S. Wagner • A. Bode • H. Brüchle • M. Brehm *EDITORS*

High Performance Computing

in Science and Engineering Garching/Munich 2016

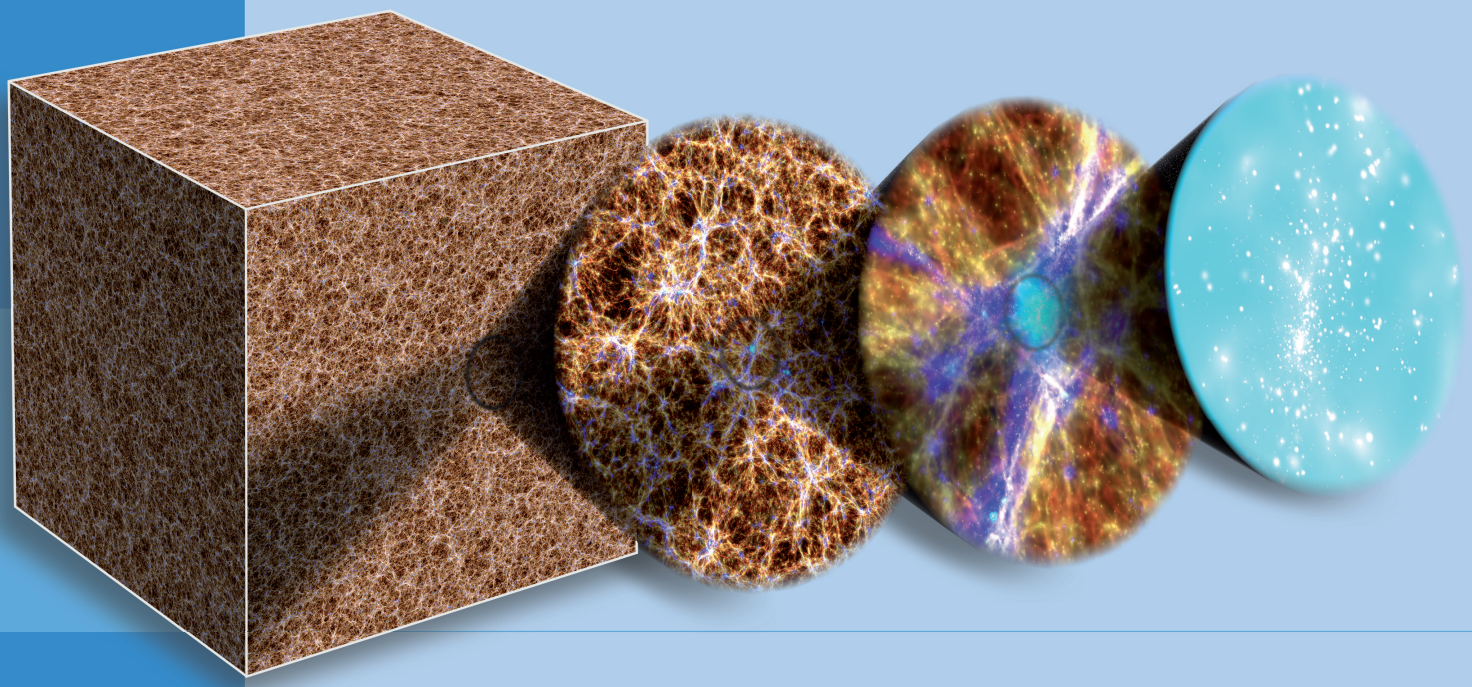

KONWIHR III

GCS
Gauss Centre for Supercomputing

Bayerische
Akademie der Wissenschaften

lrz

Impressum:

Bayerische Akademie der Wissenschaften
Alfons-Goppel-Str. 11, D-80539 München
info@badw.de, www.badw.de

Leibniz-Rechenzentrum (LRZ)
Boltzmannstraße 1, D-85748 Garching bei München
lrzpost@lrz.de, www.lrz.de

Herausgeber: Siegfried Wagner, Arndt Bode, Helmut Brüchle, Matthias Brehm
Redaktion: Helmut Satzger
Gestaltung: Tausendblauwerk, Schleißheimer Straße 21, 85221 Dachau, www.tausendblauwerk.de
Druck und Bindung: bonitasprint gmbh, Max-von-Laue-Straße 31, 97080 Würzburg

Das Titelbild zeigt die Visualisierung einer kosmologischen Simulation aus dem Magneticum Projekt.
Siehe Seite 54ff für weitere Informationen.
Bild Vorwort: Andreas Heddergott; Bild Umschlag-Rückseite: Torsten Bloth.

Das Werk einschließlich aller Abbildungen ist urheberrechtlich geschützt.
Alle Rechte liegen bei der Bayerischen Akademie der Wissenschaften.

Bezugsadresse:

Leibniz-Rechenzentrum (LRZ)
Boltzmannstraße 1, D-85748 Garching bei München

ISBN 978-3-9816675-1-6

S. Wagner • A. Bode • H. Brüchle • M. Brehm *EDITORS*

High Performance Computing

in Science and Engineering Garching/Munich 2016

KONWIHR III

GCS Gauss Centre for Supercomputing

Bayerische Akademie der Wissenschaften

lrz

Table of contents

Preface

10 *SuperMUC Phase 1 and Phase 2: The Multi-Petaflop Systems at LRZ*
Siegfried Wagner, Arndt Bode, Helmut Brüchle and Matthias Brehm

Chapter 01 – Astrophysics

14 *Near Field Cosmology*
Stefan Gottlöber
16 *Clumps in high-redshift galaxies and the Galactic Centre*
Marc Schartmann
18 *Binary neutron star merger simulations*
Bernd Brügmann
20 *Ionization Feedback in Massive Star Formation*
Thomas Peters
22 *Numerical Simulation of Binary Black Hole and Neutron Star Mergers*
W. Kastaun, L. Rezzolla
24 *Supersonic turbulence in Giant Molecular Clouds*
Lukas Konstandin
26 *The small and the beautiful: how the star formation law affects galactic disc structure*
Wolfram Schmidt
28 *Simulating and observing the formation of protostellar disks*
Daniel Seifried
30 *The world's largest turbulence simulation*
Christoph Federrath & Ralf S. Klessen
32 *Parasitic instabilities of the magneto-rotational instability*
Ewald Müller
34 *Unveiling the equation of state of nuclear matter with binary neutron stars*
F. Galeazzi, L. Rezzolla
36 *Modeling dispersive plasma waves with particle in cell simulations*
Cedric Schreiner
38 *Forming disk galaxies in magneto-hydrodynamical simulations of the Universe*
Volker Springel
40 *Mocking the Universe: Large Volume Simulations for galaxy surveys*
Gustavo Yepes
42 *Unravelling the Mechanism of Supernova Explosions*
Hans-Thomas Janka
44 *Dissipating the Energy of Space Weather Events*
Giovanni Lapenta
46 *Evolution of the corona above an active region after sunspots formed*
Hardi Peter
48 *Aurora – Simulating Cosmic Reionization*
Andreas H. Pawlik
50 *General Relativistic Simulations of Binary Neutron Star Mergers*
Bruno Giacomazzo
52 *Simulating the formation, evolution, and merging of molecular clouds*
Daniel Seifried
54 *Magneticum Pathfinder: The simulation of the evolution of the universe in an unmatched precision*
Klaus Dolag
57 *Physics of the Solar Chromosphere*
Mats Carlsson
60 *SuperCAST: Simulating the Universe*
Klaus Dolag

Chapter 02 – Plasma Physics

66 *Aacsfi-PSC: Advanced accelerator concepts for strong field interaction simulated with the Plasma-Simulation-Code*
Hartmut Ruhl
68 *Pair-dominated plasmas and radiation in ultra intense fields*
Thomas Grismayer
70 *Full-f gyrokinetic simulation of edge pedestal in Textor*
Timo Kiviniemi
72 *Riding plasma wakes: from astrophysics in the lab to twisted light*
Jorge Vieira

Chapter 03 – High Energy Physics

78 *Simulation of Interactions at the LHC*
Günter Duckeck
80 *The Chiral Condensate from Lattice QCD with Wilson Twisted Mass Quarks*
Carsten Urbach
82 *The intrinsic scale of Quantum Chromo Dynamics*
Rainer Sommer
84 *Using the complex Langevin equation to solve the sign problem of QCD*
Dénes Sexty
86 *Form factors of semileptonic B-meson decays from Lattice QCD*
Jochen Heitger
88 *The strong interactions beyond the standard model of particle physics*
Georg Bergner
90 *Lattice QCD on fine lattices*
Stefan Schaefer
92 *Hadron Physics from Lattice QCD*
Andreas Schäfer
96 *Kaon semi-leptonic form factor*
Enno E. Scholz

Chapter 04 – Engineering and Computational Fluid Dynamics

100 *Buoyant-convectively driven heat and gas exchanges*
H. Herlina
102 *Investigation of Vehicle Wheels Ventilation Moment using DoE-based Computations and Experiments*
Lu Miao
104 *Large-eddy simulation of a pseudo-shock system in a Laval nozzle*
J. F. Quaatz
106 *Shock Mach number influence in reacting shock-bubble interaction*
Felix Diegelmann
108 *Large-eddy simulation of the coaxial injection of liquid nitrogen and gaseous hydrogen under supercritical pressures*
Stefan Hickel
110 *Numerical Investigation of the Flow Field about the VFE-2 Delta Wing*
Stefan Hickel
112 *Optimal wall spacing for heat transport in thermal convection*
Olga Shishkina
114 *Initiation and evolution of subaqueous patterns in channel flow*
Markus Uhlmann
116 *Large-eddy simulation of LOx/GCH4 injection at supercritical pressures*
Michael Pfitzner
118 *Landing Aircraft and Enhanced Wake Vortex Decay*
Frank Holzäpfel
120 *Fully-resolved, finite-size particles in statistically stationary, homogeneous turbulence*
Markus Uhlmann

122 *High accuracy molecular dynamics simulation of fluids at interfaces*
Martin Horsch
124 *Investigation of three-dimensional Dynamic Stall*
Anthony D. Gardner
126 *In-situ Analysis of Flame Structures of Turbulent Non-Premixed Flames*
Christian Hasse
128 *Large Eddy Simulation of turbulent flow interacting with complex structures*
Michael Manhart
130 *Direct Numerical Simulation of an Adverse Pressure Gradient Turbulent Boundary Layer*
Vassili Kitsios
132 *Large eddy simulation of pulverized coal and biomass combustion*
Andreas Kempf
134 *Natural thermal convection at high Rayleigh numbers*
Olga Shishkina
136 *Modulation of Turbulent Properties in a Spray Flame Burning n-Heptane: Direct Numerical Simulation*
Dominique Thévenin
138 *Coupled multi-physics simulations on the massively parallel system SuperMUC*
Sabine Roller
140 *Dynamics of Shedding Partial Cavities and Erosion Aggressiveness of Cavitating Ship Propeller Flow*
Bernd Budich
142 *Specific turbulent kinetic energy transfer in very anisothermal flow*
Benoit Mathieu
144 *CFD based low-order modeling of the nonlinear flame dynamics and of swirl fluctuations*
Wolfgang Polifke
146 *Aerodynamic Investigations of Vortex-Dominated Wing Configurations with Active and Passive Flow Control*
Christian Breitsamter
148 *waLBerla – A massively parallel framework for fluid simulations*
Harald Köstler
150 *Computational acoustics of supersonic turbulent free round jets*
Jörn Sesterhenn
152 *Cavitation Erosion in Injection Systems*
Christian Egerer
154 *Direct Numerical and Large-Eddy Simulation of Wall-Bounded Flows for Aeroacoustic Source Term Identification*
Stefan Becker
156 *Coupling 4 Molecular Dynamics Codes in a Massively Parallel Molecular-Continuum Fluid Dynamics Framework*
Hans-Joachim Bungartz, Philipp Neumann
158 *Direct Numerical Simulation of the interaction of a wall-mounted cube with a turbulent boundary layer*
Christos Vassilicos
160 *Fractal flame structure due to the Darrieus-Landau instability*
Rixin Yu
162 *Relativistic effects on the Richtmyer-Meshkov instability*
Michael Dumbser
164 *Direct Numerical Simulation of Open-Channel Flow at Fully-Rough Regime*
Markus Uhlmann
166 *Towards Large-Eddy/Filtered-Density Function Simulations of Turbulent Sooting Flames*
P. Gerlinger
168 *Shock-wave/turbulent boundary-layer interaction over a flexible panel*
Vito Pasquariello

Chapter 05 – Chemistry and Material Sciences

172 *Quantum Monte Carlo and Exact Diagonalization Studies of Correlated Electron Systems*
Werner Hanke, Fakher F. Assaad
174 *Hybrid Density Functional Modeling of Transition Metal Clusters*
N. Rösch
177 *Uranyl Adsorption at Clay Mineral Surfaces*
N. Rösch
180 *Simulation of Electron Transfer and Electron Transport Processes in Molecular Systems at Surfaces*
Michael Thoss

182 *Coupling kMC and CFD in Heterogeneous Catalysis*
Karsten Reuter
184 *Theoretical investigation of photo-catalytic water splitting*
Karsten Reuter
186 *Heterogeneous catalysis from advanced molecular dynamics*
Dominik Marx
188 *Sub-Kelvin cooling with Magnetic Molecules*
Jürgen Schnack
190 *Numerical Renormalization Group studies of deposited Magnetic Molecules*
Jürgen Schnack
192 *Advanced Finite-Temperature Lanczos Method for Magnetic Molecules*
Jürgen Schnack
194 *Study of complex microstructure evolution in ternary eutectic alloys with massive parallel large-scale phase-field simulations*
Britta Nestler
196 *The inverse problem of electronic structure*
Miguel Marques
198 *The chemistry of porphyrins adsorbed on metallic surfaces*
Wilhelm Auwärter
200 *Adsorption of molecules in* giant *Metal-Organic Frameworks: An* ab initio *study*
Bartolomeo Civalleri
202 *"Hot" adatoms hopping: Equilibration dynamics from first-principles*
Karsten Reuter
204 *Shape and Catalytic Mechanism of RuO_2 Particles at CO Oxidation Reaction Conditions: First-Principles Based Multi-Scale Modeling*
Karsten Reuter
206 *New environmentally friendly thermoelectric materials*
Davide Donadio
208 *Optical Simulation of Innovative Thin Film Solar Cells*
C. Pflaum
210 *Modeling, simulation and optimization of thin film silicon solar cells on flexible Aluminum substrate*
Christoph Pflaum
212 *Tuning the electronic reconstruction at oxide surfaces and interfaces*
Rossitza Pentcheva
214 *Hydrogen-bond symmetrization and the spin transition in ε-FeOOH*
Rossitza Pencheva
216 *Phase transition based control of friction at the nanoscale*
Carlo Pignedoli
218 *Stabilization of ferroelectric properties in Hafnia and Zirconia*
Alfred Kersch
220 *Atomistic Insights into Novel Materials and Old Catalytic Problems*
Karsten Reuter
222 *Numerical simulations of topological and correlated quantum matter*
Fakher F. Assaad

Chapter 06 – Earth and Environmental Sciences

226 *Synthetic Earth Models – Taming Chaotic Mantle Flow and Exploiting Seismic Wavefield Effects to Constrain Buoyancy*
Bernhard Schuberth
228 *High resolution gravity field modeling*
Thomas Gruber
230 *EXtreme PREcipitation and Hydrological climate Scenario Simulations (EXPRESS-Hydro)*
Dieter Kranzlmüller
232 *Dynamic Rupture Simulations at Petascale*
Michael Bader, Alice Gabriel
234 *4D City – Space-time Urban Infrastructure Mapping by Multi-sensor Fusion and Visualization*
Xiaoxiang Zhu
236 *Computational Wave Propagation*
Heiner Igel

Chapter 07 – Life Sciences and Biology

248 *Dynamics of Transmembrane Domains: Impact on Complex Membrane Processes*
Dieter Langosch, Christina Scharnagl
250 *Gene mapping methods for complex diseases and integrative analysis of high dimensional 'omics data*
Konstantin Strauch
252 *Binding Specificity and Allostery of Biomolecular Interactions*
Volkhard Helms
254 *Disentangling Evolution on the SuperMUC*
Alexandros Stamatakis
256 *High Performance Methods for Computational Fluid Dynamics including Multiphysics Scenarios*
Wolfang A. Wall, Martin Kronbichler
258 *Kinetics and thermodynamics of conformational changes upon protein association studied by molecular dynamics simulations*
Martin Zacharias
260 *Simulating transition to turbulence in hemodynamics of intracranial aneurysms at extreme scale*
Sabine Roller
262 *Modulation of voltage-gated potassium channel Kv1.2 by PIP2 lipids*
Mounir Tarek
264 *The key is in the movements: allosteric activation of an oncologically relevant kinase*
Francesco L. Gervasio
266 *From Biomolecular Structures to Thermodynamic Ensembles: Cellular Logistics Controlled by Disordered FG-Nucleoporins*
Helmut Grubmüller
270 *Rapid and accurate calculation of ligand-protein binding free energies*
Dieter Kranzlmüller
272 *CAMEL Project Report*
Edward Vigmond / Gernot Plank
274 *Iphigenie/CPMD: Accurate and Efficient QM/MM Molecular Dynamics*
Gerald Mathias
276 *The Molecular Mechanism of Cooperative Activation and Control of G-Protein Coupled Receptors*
Timothy Clark
278 *Targeting FtsZ assembly for the development of new antibiotics*
Pablo Chacon

Chapter 08 – Extreme Scaling

282 *Extreme Scale-out on SuperMUC Phase 2*

Appendix

290 *The SuperMUC Multi-Petascale System*

SuperMUC Phase 1 and Phase 2: The Multi-Petaflop Systems at LRZ

Computer simulations are the well-established third pillar of natural sciences along with theory and experimentation. Particularly high performance computing is growing fast and constantly demands more and more powerful machines. To keep pace with this development, in spring 2015, the Leibniz Supercomputing Centre installed the high performance computing system SuperMUC Phase 2, only three years after the inauguration of its sibling SuperMUC Phase 1. Thereby, the compute capabilities were more than doubled. For a detailed system description of SuperMUC Phase 1 and Phase 2, please see the Appendix.

The new system was thoroughly tested by the most demanding user projects in a 30-day extreme scaling period. During that time, 14 projects had the opportunity to run their applications on the full Phase 2 system and performed stress tests on the cooling system, interconnect, processors, and storage system, as well as testing new compilers and a new parallel environment (see chapter 8 for more details).

The Leibniz Supercomputing Centre supplies its high performance computing resources to both national and international research teams. It is a member of the Gauss Centre for Supercomputing (GCS), which combines the three national centres High Performance Computing Center Stuttgart (HLRS), Jülich Supercomputing Centre (JSC), and Leibniz Supercomputing Centre (LRZ) into Germany's foremost supercomputing institution. GCS is jointly funded by the German Federal Ministry of Education and Research and the corresponding ministries of the states of Bavaria, Baden-Wuerttemberg and North Rhine-Westphalia. GCS massively contributes to European large-scale scientific and engineering research by its involvement in the Partnership for Advanced Computing in Europe (PRACE).

Proposals for computing time on SuperMUC can be submitted to the LRZ throughout the year and the projects can start immediately after they have been reviewed positively. Projects with several European partners can submit proposals via PRACE. Twice a year, GCS supports the most demanding projects through its Call for Large Scale Projects. Several PRACE and GCS Large Scale projects report about their work in this book.

Especially through the activities in the area of extreme scaling of applications, it became more and more evident that highly scalable software poses many challenges, for example message passing with several hundred thousand tasks, and massively parallel I/O. Another challenge consists of using all the different layers of parallelism, from vector instructions to multi-threading, shared memory, multi- and many core CPUs, and message passing in complex network topologies. The Bavarian Competence Network for Technical and Scientific High Performance Computing (KONWIHR) twice a year supports short to medium term projects from regional researchers that optimize existing applications to achieve better scalability. Within the scope of KONWIHR supported projects, software developers come to LRZ to directly work with application experts to profile the application, identify bottlenecks, and to develop and implement strategies for better scalability. The results of several KONWIHR projects are included in this book.

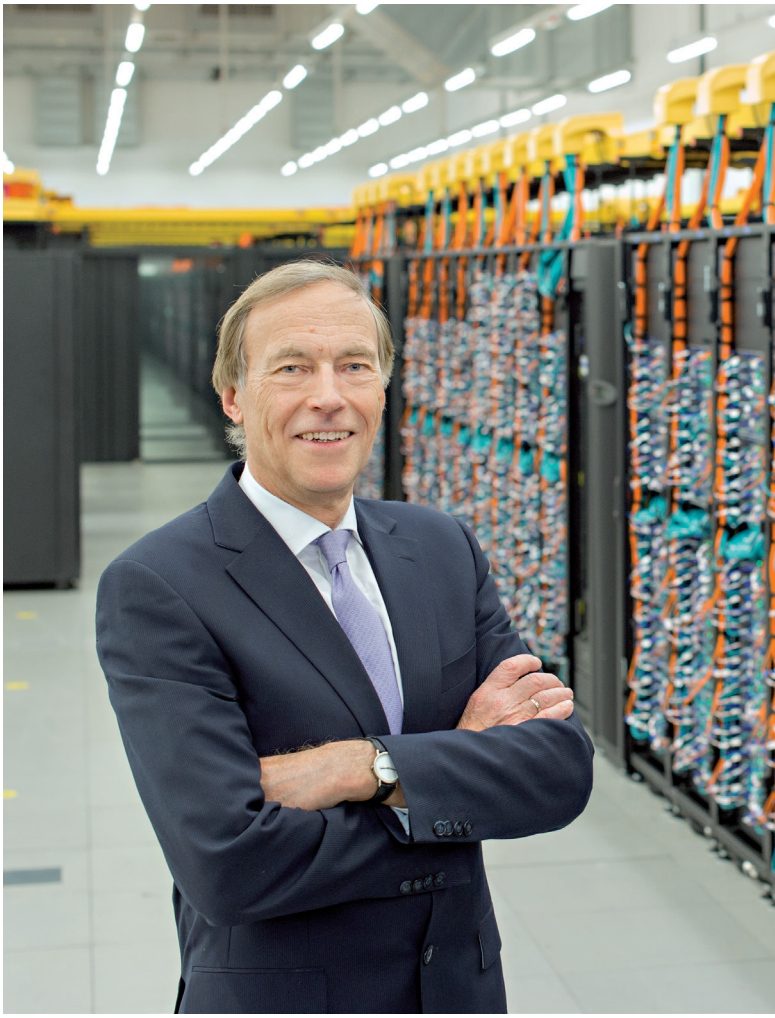

In addition, LRZ created dedicated application labs for astrophysics, big data, computational fluid dynamics, earth- and environmental sciences, digital humanities,

Figure 1: Prof. Arndt Bode, Chairman of the Board of Directors at LRZ in front of the new Phase2 nodes of SuperMUC.

and life sciences, where application experts work closely with scientists on optimization and scalability of the top applications. LRZ is also committed to the Partnership Initiative for Computational Sciences π^{CS} to address their special requirements.

This book covers the time-frame June 2014 until June 2016. Readers will find many examples of outstanding research in the more than 130 projects that are covered in this book, with each one of these projects using at least 4 million core-hours on SuperMUC.

In Spring 2016, 20% of the computing time was consumed by jobs that ran on more than 4 Islands (32,768 cores). Overall, more than 50% of jobs used at least 2,048 cores. With the start of operation of Phase 2, the batch queueing system on Phase 1 had been re-configured to allow job sizes up to 8 Islands (65,536 cores) and we see an ever increasing number of 8-Island jobs.

The largest scientific communities using SuperMUC in the last two years were computational fluid dynamics simulations, chemistry and material sciences, astrophysics, and life sciences. Reports from the largest projects in these areas are included in the book:

- Janka et al. performed three-dimensional simulations of core-collapse supernova explosions of massive stars applying neutrino hydrodynamics (page 42).
- Dolag et al. simulated the formation of the Universe in the Magneticum project, which marks the largest cosmological simulation to date (page 54).
- Adams et al. performed numerical investigations of the vortical flow field about the VFE-2 delta-wing (page 110).
- Stamatakis et al. reconstructed the phylogenetic trees for two of the largest datasets analyzed to date (page 254)
- Nestler et al. studied the complex microstructure evolution in ternary eutectic alloys (page 194)

Acknowledgments

We gratefully acknowledge the continued support of the State of Bavaria, the Bavarian Competence Network for Technical and Scientific High Performance Computing (KONWIHR), the Gauss Centre for Supercomputing, the German Research Foundation (DFG), the German Federal Ministry of Education and Research (BMBF), the Partnership for Advanced Computing in Europe (PRACE), and many other institutions promoting high performance computing. We thank the reviewers and the Steering Committees of GCS and SuperMUC for the reviews of the projects, their insights and helpful remarks. Without their efforts it would not have been possible and will not be possible in the future to sustain the high scientific quality we can see in the projects.

Garching near Munich, June 2016
Siegfried Wagner
Arndt Bode
Helmut Brüchle
Matthias Brehm

Astrophysics

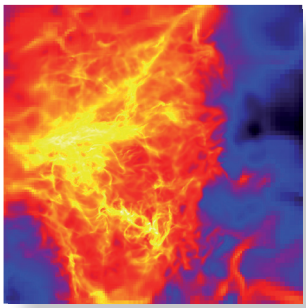

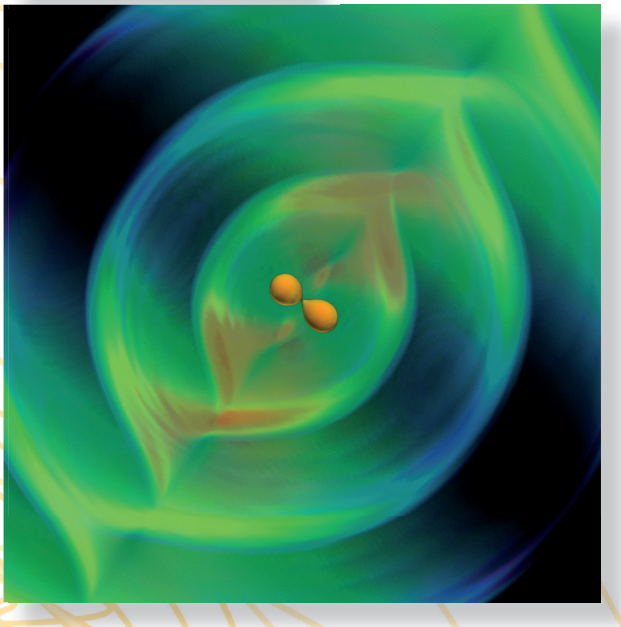

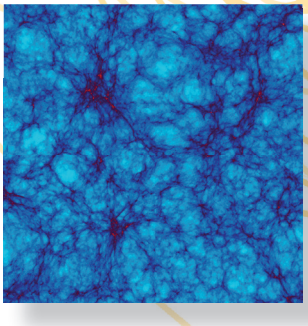

Near Field Cosmology

1

Research Institution
Leibniz-Institut für Astrophysik Potsdam

Principal Investigator
Stefan Gottlöber

Researchers
Gustavo Yepes, Alejandro Benitez-Llambay, Noam Libeskind, Jenny Sorce

Project Partners
UAM Madrid, IATE Cordoba, AIP Potsdam

SuperMUC Project ID: h009z

Introduction

Cosmology is science at the largest scale, examining the origin and fate of the Universe as a whole. By studying its evolution over the past 13.8 billion years one sees the growth of cosmic structures from an almost homogeneous medium - when the universe was in its infancy – into a vast network of filaments and voids known as the cosmic web when viewed today. Within this framework, matter collapses into halos that consist of dark and luminous components. The existence of Dark Matter is inferred by its gravitational effect on the luminous (also termed baryonic) matter, which is observed as the diverse galaxy population we see today. Cosmological studies are based on numerical simulations and observations of the evolution of cosmic structures from the first galaxies that form and then reionise the universe, to the present galaxy population and the intergalactic medium. Such simulations follow the gravitational clustering of dark and baryonic matter, hydrodynamic processes associated with the gas, the formation of stars and black holes and their back reaction known as "feedback". These computationally very expensive experiments can be performed only on the largest supercomputers.

The Local Group

One important limitation of many galaxy formation simulations is that the simulated galaxies are usually selected to be relatively isolated systems, to avoid extra computational cost. However, the environment of galaxies plays an important role in their evolution, and simulations of isolated galaxies can only offer limited insights into the more complex process of galaxy evolution, at least when it comes to the formation of the Local Group of galaxies. This system consists of two large spiral galaxies – the Milky Way and Andromeda – and several dozens smaller dwarf galaxies. In order to investigate the formation of the Local Group one either needs to simulate structure formation in a very large volume and to identify objects which look similar to the Local Group or to constrain the initial conditions of the simulations to produce the nearby galaxy distribution observed today. This is the scope of the CLUES (Constrained Local UniversE Simulations) project (http://www.clues-project.org), which provides initial conditions that aim to reproduce the environment of the Local Group. Within our project at LRZ Munich we have performed a series of simulations at different resolution that allows us to test different models of structure formation in particular on small scales which can be observed in the local environment. Figure 1 shows the simulated evolution of three different types of dwarfs in the neighborhood of the simulated Local Group. The main difference between the dwarfs is their typical time of star formation. While the upper two rows show a dwarf galaxy with early star formation (10 billion years ago) which now contains mainly old stars most of the stars in the second dwarf have been formed only 4 billion years ago. The simulation allows us to study and to understand

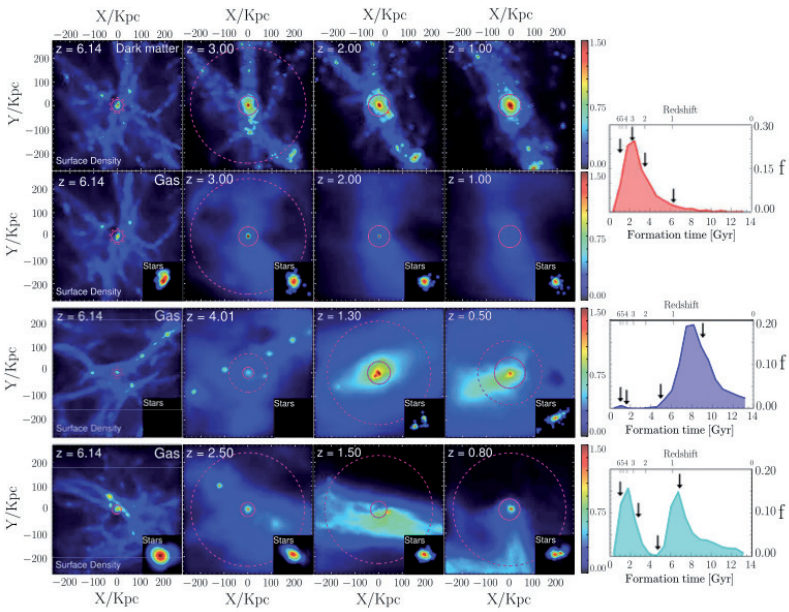

Fig. 1: Left: Dark matter (top row) and gas (second to third row) distribution of dwarf galaxies at different times. Right: Star formation rate in these dwarf galaxies as a function of time.

the processes which led to this different behavior and to detect also a third type of galaxies with a mixed population of old and young stars (fourth row).

Our neighborhood is the best observed part of the universe where the smallest satellite galaxies can be detected. The dwarf galaxies around the Milky Way and other nearby galaxies show an unexpected alignment in vast planes. In order to study the infall pattern of these satellite galaxies we analyzed a constrained simulation with 1024^3 particles within a box of roughly 300 million light years side length [3]. In this simulation we can resolve dwarf satellites down to 500 million solar masses (approximately 0.05% of the mass of our Milky Way). In the simulation one can clearly see the cosmic web in which the galaxies move. This web-like cosmic structure can be also described by the direction of the velocity flow lines along which the galaxies move. In Fig. 2 we show that the infalling satellites preferentially move towards the massive host galaxies along the axis of weakest collapse defined by the large scale cosmic web. This direction is marked in red. Orthogonal to this direction (green, blue) far fewer satellites can be detected. The four quarters of the Aitoff plot correspond to different mass ranges of the satellites with the most massive on top left. Counter clock-wise the mass decreases. With decreasing mass the signal becomes smaller. The close relation of the local infall direction to the global cosmic web can explain the unexpected alignments of dwarf galaxies seen in the Local Group.

Nearby clusters of galaxies and voids

The simulations described above were performed on initial conditions built using a so-called Wiener filter reconstruction, based on a relatively small set of observational data. We now use the much deeper CosmicFlows-2 data set [4] and an improved reconstruction algorithm [5]. Based on the new observational data and the new algorithm we are able to simulate structure formation in a much larger volume of the local universe than before. For the first time the evolution of nearby clusters of galaxies can be studied in detail. In Figure 3 we show the simulated cosmic structures around the Milky Way. The Milky Way is situated in the center of the plot and the structures are shown in a slice of a few million light years thick up to a distance of about 300 million light years. A few well-known clusters of galaxies are visible: the Coma, Shapley and Perseus-Pisces clusters. Each of

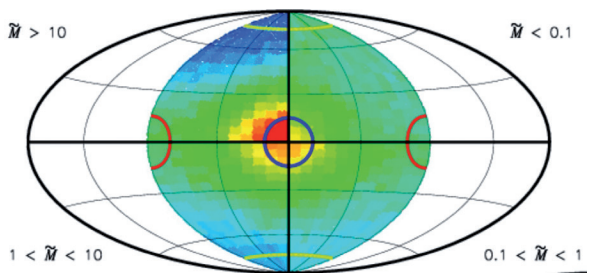

Fig. 2: Infall direction of satellites in four different mass ranges shown in an Aitoff projection. The yellow, red and blue circles define areas within 15 degrees of the cosmic web axes, e1, e2, and e3, respectively.

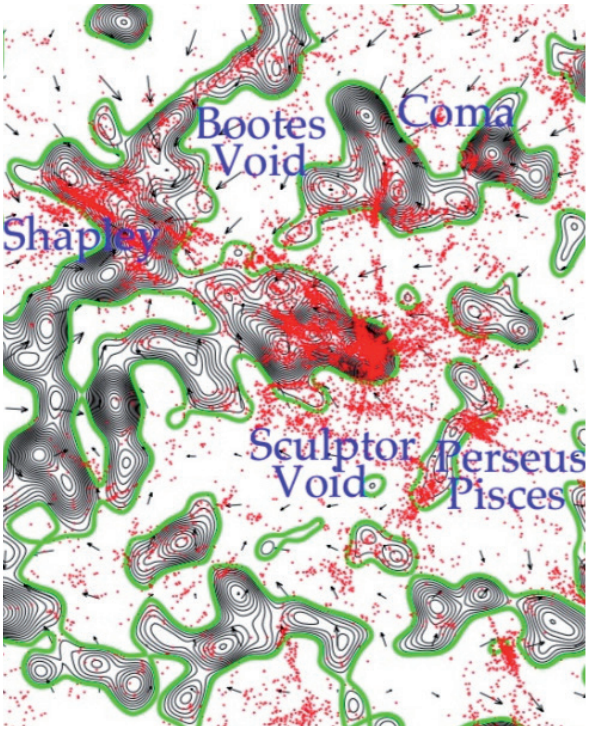

Figure 3: Simulated objects in a large volume around the Milky Way.

these clusters contains thousands of galaxies. The simulated structures can be directly compared with the observed galaxy distribution shown as red dots in this figure. Between the clusters and filaments on can see large almost empty regions in which very luminous galaxies never have been formed: the cosmic voids.

Outlook

We will use the constrained initial conditions based on the CosmicFlows-2 dataset [4] to run a high resolution simulation of the Local Volume up to a distance of a few hundred million light years as well as a set of high resolution simulations of the formation of nearby clusters like Virgo, Coma and Centaurus. All these simulations will be performed including not only Dark Matter but also baryons.

References and Links

[1] G. Yepes, S. Gottlöber, Y. Hoffman, New Astronomy Reviews 58 (2014),1
[2] A. Benitez-Llambay, J.F. Navarro, M.G. Abadi et al., MNRAS 456 (2016), 1158
[3] N. I. Libeskind, A. Knebe, Y. Hoffman, S. Gottlöber, MNRAS 443 (2014), 1274
[4] R. B. Tully et al., Astron. Journal 146 (2013), 86
[5] J. Sorce, S. Gottlöber, G. Yepes et al., MNRAS 455 (2016), 2078

The home page of the CLUES project is https://www.clues-project.org
More plots and movies can be found at
https://www.clues-project.org/cms/images-and-movies/

Clumps in high-redshift galaxies and the Galactic Centre

1

Research Institution
Max-Planck-Institute for extraterrestrial Physics / University Observatory Munich

Principal Investigator
Marc Schartmann

Researchers
Alessandro Ballone, Manuel Behrendt, Andreas Burkert, Marc Schartmann

Project Partners
Swinburne University of Technology, Melbourne (Australia)

SuperMUC Project ID: h0075

We first present a detailed simulation of the formation and evolution of clumps in a gas-rich high-redshift galaxy and then discuss our recent progress in understanding the tiny G2 cloud, which has recently passed the massive black hole (BH) in our own Galactic Centre.

Gravitational instabilities in high-redshift gas-rich discs

Introduction

Gravitational instabilities are thought to be the main driver of structure formation in high-redshift gas-rich galactic discs. In this sub-project we are interested in a deep understanding of the formation of rings and clumps by directly comparing a high-resolution hydrodynamical simulation with a newly established analytical perturbation analysis for geometrically thick discs. The self-gravitating initial disc is isothermal and in vertical hydrostatic equilibrium leading to a vertical sech^2 density distribution and is meant to resemble an observed high-redshift disc galaxy.

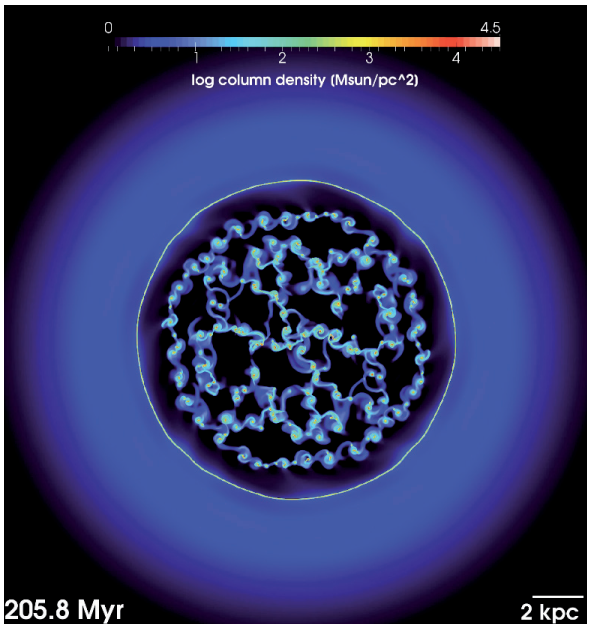

Figure 1: Column density of an evolved high-redshift, gas-rich disc. The formation of rings following gravitational instability and their subsequent break-up into clumps is shown. Taken from [2].

Results and Methods

The hydrodynamical equations together with the Poisson equation are solved with the help of the RAMSES code [1]. For this simulation, the Euler equations are evolved with the local Lax-Friedrichs scheme. The dark matter halo in which the isolated disc is embedded in is treated as an external density field added to the Poisson solver. Adaptive Mesh Refinement (AMR) allows to resolve the 48kpc size computational box with grid cell sizes between 187.5pc and 5.86pc, spanning 5 levels of refinement and ensuring that the Jeans length is resolved with at least 18 grid cells. As expected from analytical, linear stability analysis, the disc forms rings, which later break up into a large number of clouds (Fig. 1). The formation of axisymmetric overdensities with time shows the expected evolution in the linear regime: The number of simulated discrete, concentric rings and their relative position is given by the local fastest growing perturbation wavelength and the relative formation times follow the expected growth rates. In order to show this behaviour, the vertical scale heights of the discs in the simulations have to be resolved by at least 5 grid cells. These ring-like perturbations later on fragment into a large number of clumps. Being the result of ring fragmentation, their sizes are not directly related to the local fastest growing wavelength of the initial disc, as is often assumed in literature. We also find that they organise themselves into clusters of clumps, which appear as single entities at the resolution of 10m class telescopes and the substructure is able to explain many properties of observed giant clumps.

This RAMSES simulation was run on 2048 CPU-cores, consuming a total of roughly 1 Million CPU-h and producing roughly 2.4 Million files with a total of 6 TB of data.

The results of this sub-project have been published in [2].

On-going Research / Outlook

In future projects, we plan to include more physics in the simulations like cooling, the formation of stars in the dense clumps formed and the modelling of subsequent energy release by supernova explosions. We also plan to apply the expertise gained from this simulation to investigate the physics of circum-nuclear disc evolution, nuclear star cluster formation and their relation to active galactic nuclei. Due to the cooling, feedback and strong outflows, we expect higher velocities, smaller time steps and a larger volume filling factor of structure in the AMR domain. This leads to even more computationally expensive calculations, making SuperMUC or its successor system perfectly suited for this task.

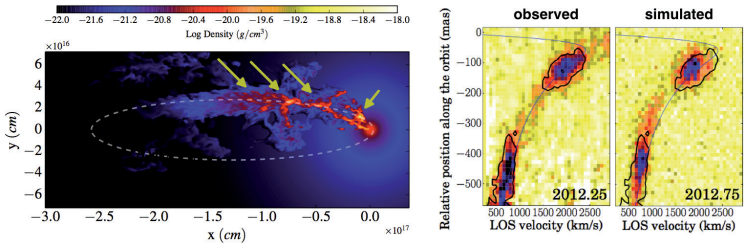

Figure 3: A mass-losing source following the observed orbit of G2 leads to the formation of a two-component system: the G2 cloud and a trailing tail component. This is visible in the density distribution in the equatorial plane (left panel) and the observable position-velocity diagrams (right panels). Taken from [5].

The Galactic Centre cloud G2: mass-losing source or condensation in a stream of gas?

Introduction

The so-called "G2" cloud has just completed its peri-centre passage around the massive BH in the Galactic Centre (GC). In this sub-project we are investigating possible scenarios for its origin and evolution in the coming years. Our strategy is to follow the two most popular scenarios for the origin of the roughly 3 Earth masses of gas, namely that G2 is part of a stream of gas pointing towards the GC or that it is made up of gas lost from a central source. By directly comparing to available high spatial resolution observations, we show that – up to now – both scenarios can be brought in good agreement with the observed appearance and kinematics of G2.

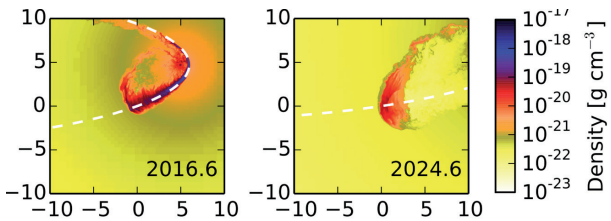

Figure 2: Evolution of the gas cloud density in the orbital plane of the observed G2 orbit. Images are centered on the nominal position of the cloud following the observed orbit. Taken from [3].

Results and Methods

For the realistic modelling of both scenarios, we employ the high-resolution shock capturing scheme PLUTO [4] and its patch-based AMR capabilities relying on the CHOMBO library within a Cartesian domain. The two-shock Riemann solver with a parabolic interpolation and the second-order Runge-Kutta time integration scheme is used. In both approaches, the domain is filled with a hot atmosphere resembling an analytical solution for the accretion flow around the central massive BH, modeled with a Newtonian point-like potential. For the case of the pure cloud simulations, we study possible origin scenarios along the observed orbit. Starting spherically symmetric in pressure equilibrium, the cloud first compresses due to the radially increasing pressure of the atmosphere, suffers from ram pressure interaction and finally stretches significantly due to the tidal shearing by the central massive BH (Fig. 2). By directly comparing to observed position-velocity diagrams, we find that a starting position close to peri-centre (before roughly the year 1900) is favoured by this setup, which places it within the discs of young stars in the GC and enables the formation of a density concentration by colliding winds in this stellar system. Similarly good agreement with PV diagrams can be reached with a cloud built-up by a central commoving source of gas. In this case we can even find a combination of velocity and outflow rate of the gas ejected by the central source, which is able to account for both, G2 itself and an observed putative trailing component. The Rayleigh-Taylor fingers, which form due to the expansion of the wind into the ambient atmosphere get partly stripped and converge behind the source to form a second distinct component in the ionised (Brackett gamma) line emission on the sky as well as in the PV diagrams (Fig. 3). Remarkably, the source properties are consistent with a T-Tauri star.

The simulations for the G2 project were run on up to 1600 CPU-cores, consuming a total of 3 Million CPU-h and producing a total of roughly 4-5 TB of data.
The work presented here is published in [3] and [5].

On-going Research / Outlook

Making the simulations more realistic includes the introduction of magnetic fields, which increases the simulation time by a factor of a few. Another step is to self-consistently build up the atmosphere from the injections of stellar winds from the young stars surrounding the observed orbit of the G2 cloud, which necessitates the simulation of the order of 100 mass-losing sources, well suitable for the next generation SuperMUC system.

References and Links

[1] Teyssier et al. 2002, A&A, 385, 337
[2] Behrendt et al. 2015, MNRAS, 448, 1007
[3] Schartmann et al. 2015, ApJ, 811, 155
[4] Mignone et al. 2007, ApJS, 170, 228
[5] Ballone et al. 2016, ApJL, 819, 28

http://www.usm.lmu.de/CAST

Binary neutron star merger simulations

Research Institution
Friedrich Schiller University Jena

Principal Investigator
Bernd Brügmann

Researchers
S. Bernuzzi, M. Bugner, T. Dietrich, T. Dörffel, E.Harms, D. Hilditch, N. K. Johnson-McDaniel, N. Moldenhauer, H. Rüter, W. Tichy, M. Ujevic Tonino, A. Weyhausen

Project Partners
Max Planck Institute for Gravitational Physics Potsdam, University of Parma, Federal University of ABC Sao Paulo

SuperMUC Project ID: h1021, pr87nu, pr48pu

Introduction

Our research focuses on the numerical tools necessary to solve Einstein's equations. In recent years we have been particularly interested in spacetimes consisting of two neutron stars in the final stages of their evolution. Because of the emission of gravitational radiation, the objects are driven together to merge; Fig. 1 visualizes the emitted gravitational wave signal. This emitted gravitational radiation carries energy and momentum away from the system and contains information about the system. Late last year the Laser Interferometer Gravitational-wave Observatory (LIGO) began searches for these gravitational wave signals at a sensitivity at which detections are expected. Although such systems can radiate a significant amount of their total mass-energy in gravitational waves, the gravitational wave signals one expects to receive on Earth are not strong, since sources of gravitational waves are often many millions of light years away. Therefore one needs accurate templates for the radiation one expects from such systems in order to be able to extract them out of the detector's noise.

Although analytical models exist for compact binary systems when the constituents are well separated, we need numerical simulation to investigate the last orbits before merger to obtain accurate templates and validate analytical approximations. Due to the strong nonlinearity of the equations and the large separation of length scales, these simulations are computationally demanding and need to be run on large supercomputers. When matter is present the computational cost as compared to pure black hole (vacuum) simulations increases even more due to the additional matter fields. But also more interesting astrophysical phenomena can happen. In fact, there is the possibility for a strong electromagnetic signal from the merger (e.g., a short gamma-ray burst or lower-energy electromagnetic signatures from the ejecta) and significant neutrino emission. Additionally, we can expect that gravitational wave observations of binaries involving neutron stars could place constraints upon the equation of state (EOS) at supranuclear densities. Neutron stars are the only places in the universe where one can probe theories of nuclear physics at greater than nuclear density in cold matter.

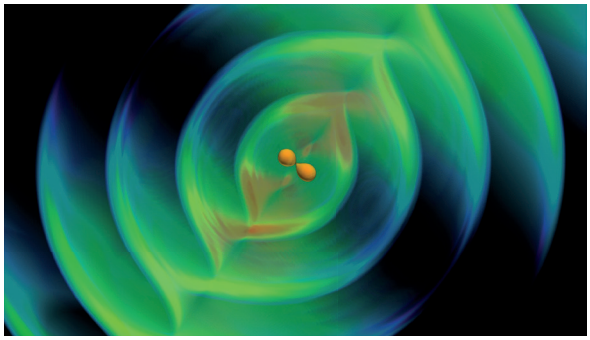

Figure 1: 3D rendering of the gravitational waves emitted from a binary neutron star system at merger. The central region (density) is stretched by a factor of ~5 for better visibility of the tidal deformation.

Results

Computational Setup

We perform our simulations with the BAM code, which combines state-of-art methods to deal with black hole spacetimes and general relativistic hydrodynamics. The code is written in C and based on the method of lines. It uses high-order finite difference stencils for the spatial discretization of the geometric variables, while high resolution shock capturing methods are used for the hydrodynamic variables. The time integration is done with an explicit Runge-Kutta method. The BAM infrastructure also supplies adaptive mesh refinement via a combination of fixed and moving boxes, as well as cubed spheres. The code is hybrid OpenMP/MPI parallelized.

For our simulations, we have to span a reasonable range in the parameter space to see the influence of individual quantities (e.g., spin, equation of state, mass-ratio) on the inspiral and post-merger dynamics. Additionally, we are forced to simulate physical setups with different resolutions to (i) show consistency, (ii) compute the convergence order, and (iii) give proper error bars for the physical quantities. Thus, the individual simulation of one physical setup using one resolution is meaningless and we can only make a meaningful scientific statement using a bundle of jobs. Depending on the resolution those jobs run on different number between ~100 and ~1000. In total, we have used ~30 million CPUh on SuperMUC over the last 2 years

within different projects. We produced ~100 million files and used a maximum of ~60TB of storage.

Scientific results
We focused on gravitational waves emitted by binary neutron star systems around merger and investigated those systems in different portions of the parameter space. Using our allocations we have published 3 Physical Review Letters, 9 Physical Review D articles, and 2 proceedings over the last 2 years.

Spinning binary neutron star simulations: One milestone was the successful computation of spinning binary neutron star mergers [2]. Before our work, almost all simulations of binary neutron stars assumed that the stars are non-rotating (irrotational). Although it is thought that many neutron stars in binaries with other neutron stars can be described as irrotational, some pulsar observations indicate that neutron stars can be spun up to periods as short as a millisecond. Consistent initial data for spinning neutron star binaries in general relativity have been only very recently computed by Prof. W. Tichy (FAU), one of our collaborators in the United States. These data allow arbitrary spins and satisfy the constraint equations of general relativity. Our simulations show that up to contact the spin-orbit interactions are possibly stronger than the tidal interactions even for astrophysically reasonable spins. This shows that a proper description of the spin is crucial to model binary pulsar systems. Depending on the initial spins, the hypermassive neutron star formed after the merger emits gravitational waves at different frequencies depending on the total angular momentum. The different centrifugal support leads also to a different collapse time leading to configurations of the final black hole and disk system with different angular momentum distribution compared to the irrotational case. These results are potentially important for astrophysical accretion disks and gamma-ray bursts as well as a possible gravitational wave detection.

High mass ratio simulations: In addition to spinning binary neutron star simulations, we also considered high mass ratio setups. We performed the first general relativistic simulation of neutron stars with a mass ratio above two. Although such high mass ratio systems are unlikely to be observed, they allow us to study several interesting physical scenarios, e.g. we find that even several revolutions before merger mass can be transferred from one constituent to the other. Furthermore, such setups eject a large amount of material in the interstellar medium and allow a careful analysis of the decomposition of the ejecta.

Tidal-effective one body formalism and quasi-universal relations: In addition to pure numerical relativity waveform modeling, we improved the effective-one-body formalism which now describes the general relativistic dynamics of neutron star binaries from the early inspiral up to the merger [3]. The effective-one-body model is a formalism to interpret the compact binary system as the motion of a particle in an effective metric. With the help of new high-resolution multiorbit numerical relativity simulations with different EOS, we found agreement

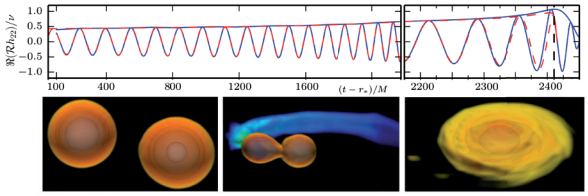

Figure 2: Gravitational wave signal (upper panel) and snapshots of the density profile (lower panels). Left: early inspiral, middle: merger, where a large amount of material is ejected (blue), right: merger remnant, surrounded by an accretion disk.

between the analytical model and the uncertainty of the numerical data; see top panel of Fig 2. Our proposed model provides the most accurate analytical representation of binary neutron star dynamics and waveforms currently available. This result also received significant press coverage. Motivated by the effective-one-body model, we found quasi-universal relations of important binary properties at merger (e.g. the mass-rescaled gravitational wave frequency and the specific binding energy). Quasi-universal means that both quantities are almost uniquely determined by some tidal coupling constants, but independent of the particular employed EOS. In the effective-one-body model, the quasiuniversality is a direct consequence of the conservative dynamics of tidally interacting bodies, but we could verify this prediction in our numerical simulations [4].

Based on this analysis, which was only valid up to the moment of merger, we performed a detailed analysis of the merger remnant. We performed a careful analysis of the main emission mode and showed that this frequency again depends on the tidal coupling constant which characterizes the binary tidal interactions during the late-inspiral-merger. The binary's total mass, mass-ratio, equation-of-state, and thermal effects do not seem to play an important role. This description opens up the possibility of developing a model of the complete gravitational spectrum of neutron star binary mergers.

Outlook

In the future we plan to extend our work on binary neutron star systems and have already made major improvements, as the implementation of higher-order flux schemes and also simulated new systems, e.g., the first precessing binary neutron star merger. Even more important is the development of a new pseudospectral code, BAMPS. This code will be the next-generation successor to BAM. Recently we implemented the first routines for general relativistic hydrodynamics in the BAMPS code within the framework of discontinous Galerkin methods.

References and Links

[1] https://www.tpi.uni-jena.de
[2] S. Bernuzzi, T. Dietrich, W. Tichy, B. Brügmann. Mergers of binary neutron stars with realistic spin. Phys. Rev. D.89, 104021.
[3] Bernuzzi, S., Nagar, A., Dietrich, T., and Damour, T. 2014. Modeling the Dynamics of Tidally Interacting Binary Neutron Stars up to the Merger. Phys.Rev.Lett. 114, (2015) 16, 161103.
[4] Bernuzzi, S., Nagar, A., Balmelli, S., Dietrich, T., and Ujevic, M. 2014. Quasiuniversal properties of neutron star mergers. Phys.Rev.Lett. 112, (2014), 201101.

Ionization Feedback in Massive Star Formation

Research Institution
Max-Planck-Institut für Astrophysik

Principal Investigator
Thomas Peters

Researchers
Ralf Klessen, Robi Banerjee, Mordecai-Mark Mac Low

Project Partners
Uni Heidelberg, Uni Zürich, Uni Hamburg, American Museum of Natural History

SuperMUC Project ID: h1343

Introduction

Understanding massive star formation is a key problem in modern astrophysics. High-mass stars, which are a factor of 10 to 100 more massive than our Sun, dominate the matter cycle and the energy budget in galaxies by their strong radiative and mechanical feedback. They also produce heavy elements, which are then released into the interstellar medium via strong stellar winds and during supernova explosions. Progress in modeling galaxy formation and evolution therefore requires knowledge of how such objects form and how their feedback impacts their surroundings. With this SuperMUC project, we have extended our previous work on ionization feedback during massive star formation in several directions. In this final report, we summarize our main research results.

Results and Methods

We use the adaptive mesh refinement code FLASH [1] for our numerical simulations. Our previous collapse simulations of high-mass star formation already included feedback by ionizing and non-ionizing radiation [2]. When the mass of the growing protostar exceeds approximately 10 solar masses, a region of ionized hydrogen forms around it. Because the ionized gas has a two to three orders of magnitude higher temperature than the surrounding molecular gas, these H II regions tend to explosively expand once they have formed. Since the density in the accretion disk around the massive protostar is particularly high, the H II region preferentially grows perpendicular to the disk plane. However, the accretion flow from which the massive protostar accretes material can become gravitationally unstable and form dense filaments that subsequently shield the ionizing radiation efficiently. When this happens, the expanding ionized gas recombines and cools down again. An example of such a molecular fountain flow is shown in Figure 1.

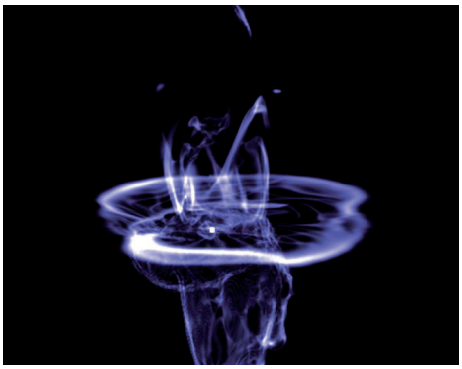

Figure 1: Volume rendering of an ionization-driven molecular outflow with the powering star shown in the center [2,6].

These ionization-driven molecular outflows have energies at the lower limit of observed high-mass outflows [3]. Furthermore, our simulations with magnetic fields suggest that magnetic launching of massive outflows via magneto-centrifugal acceleration and magnetic tower flows is difficult because fragmentation destroys the coherent azimuthal velocity structure in the disk and the ionization feedback further disrupts the magnetic field structure [4] (see Figure 2). However, the fragmentation also results in the formation of lower-mass companions around the central high-mass star [5]. We have hypothesized that magnetically-driven outflows around these companion stars could be the origin of high-mass outflows [3].

To test this proposal, we have repeated our simulation with a subgrid-scale model for the injection of protostellar outflows [7]. This subgrid model is necessary since we cannot spatially resolve the accretion disks around the low-mass stars in our simulation and self-consistently form an outflow. At each timestep, we inject a fraction of the accreted material into a cone centered on the protostar and aligned with its rotation axis. The launching speed is set to depend on the Keplerian velocity at the outflow footpoint and thus increases with protostellar mass. Carrying out this simulation with the radiative and the mechanical outflow feedback, and post-processing the resulting data, was very demanding since we have resolved the outflow over an extent of 2 pc with a grid resolution of just 100 AU. With more than 487 million grid cells, this is one of the largest star formation simulations to date. Without the generous allocation of computing time and disk space at LRZ, and in particular the very efficient tape archiving system, conducting this simulation up to this stage would have been impossible. The availability of compute nodes with a large amount of shared memory on SuperMUC was also crucial.

During this simulation, a small cluster of four stars forms within the same rotationally-flattened accretion flow. Their spin axes are all aligned. Thus, their individual outflows overlap and form a single collective outflow. The ener-

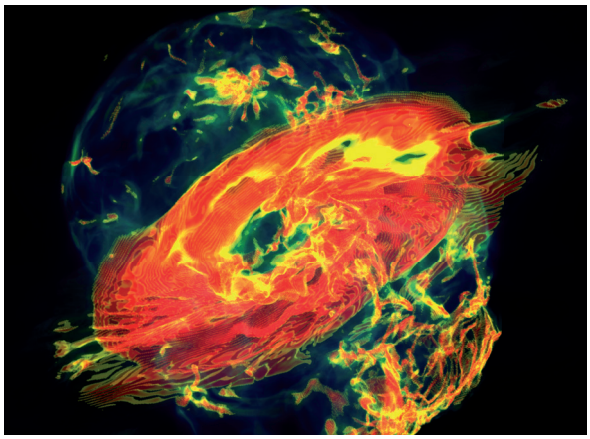

Figure 2: Volume rendering of a magnetically-driven outflow that interacts with an H II region [3,4].

getics of this collective outflow are in good agreement with observed values for high-mass outflows. Since the four stars form very close to each other, it is possible that they and their individual outflows could not be distinguished observationally. Thus, such a collective outflow could be mistakenly interpreted as driven from a single central, massive protostar. Besides the energetics, the morphology of the outflow also agrees well with observed high-mass outflows. As the protostars in the simulation grow in mass, their angular momentum vectors start to diverge. As a result, the collective outflow widens and looks less collimated. This agrees with the observation that high-mass outflows appear to be less collimated than outflows from low-mass stars. Synthetic observations of the collective outflow reproduce many features seen in Cepheus A and DR 21, two of the closest sources in which high-mass outflows can be observed. Thus, our initial results support the idea that massive outflows may be driven by the lower-mass companions of the central massive star [7].

The ionizing radiation from the forming massive star is never strong enough to stop accretion if the available gas reservoir is large enough [2]. We have demonstrated this by letting isolated stars accrete more than 100 solar masses of gas. This changes drastically when fragmentation is allowed. The massive accretion flow around the growing star is gravitationally unstable, and if it is allowed to fragment, lower-mass companions form in the vicinity of the massive star. These companion stars starve the central high-mass star of material, intercepting matter that would otherwise be accreted by it, and thereby limit its mass growth. The most massive star in this simulation reached less than 30 solar masses. Thus, it is not the radiation feedback that stops accretion onto the high-mass stars, but the fragmentation of their massive accretion flow. We have called this new process "fragmentation-induced starvation" [2,5].

Our simulations also reproduce the morphological appearances of H II regions around young stars [2,8]. These morphologies had so far been considered a characteristic of a given H II region, but their physical origin was to a large degree not understood. The simulations show that these morphologies not only change continuously during the evolution of the H II region, but that they also dramatically depend on the viewing angle onto that region. The same H II region would be classified differently when looked at from different directions. Thus, H II region morphologies appear to be nothing fundamental, but instead are simply the result of the complex interplay between the infalling gas and the ionizing radiation. Furthermore, shielding effects by the dense accretion flow lead to rapidly fluctuating H II regions. The time interval between extended and trapped phases can be as short as 10 years, and indeed changes in H II region appearance similar to the ones we find in our simulations have already been observed, but had remained unexplained [2,8].

On-going Research / Outlook

With all these breakthroughs in the theory of high-mass star formation, our SuperMUC project was a major success. Our results have spawned numerous follow-up questions. If fragmentation of the accretion flow limits the mass growth of massive stars, how do the observed 100 solar mass stars form? Does it suffice to start from an even larger initial gas reservoir? Other questions revolve around the impact of turbulence on massive star formation. How does it affect the fragmentation-induced starvation scenario? Does turbulence change the H II region morphologies or their dynamical evolution?

The next generation supercomputer will allow us to further improve our simulations. Now that we have understood our idealized simulation setups, we can start to make them more realistic. One important step will be to add turbulent velocity fluctuations to our setup. Such simulations will likely require higher numerical resolution to resolve these smaller-scale disks. Moving towards smaller spatial scales, but also towards larger clouds, is not only a computational challenge, but it also requires substantial improvements on our simulation methodology. In particular, the gas will become optically thick. Our previous approach to radiation hydrodynamics is not applicable to such situations. We will therefore use a newly developed radiative transfer scheme to propagate the stellar radiation in these optically thick media. The new scheme will further allow us to add another important feedback process to our simulations, the radiation pressure by non-ionizing radiation on dust grains. With radiation pressure included, our simulations will be among the most complete and comprehensive studies of massive star formation feedback on molecular cloud scales. They will certainly significantly improve our understanding of this important process.

References and Links

[1] Fryxell et al. 2000., Astrophys J Suppl S, 131, 273
[2] Peters et al. 2010, Astrophys J, 711, 1017
[3] Peters et al. 2012, Astrophys J, 760, 91
[4] Peters et al. 2011, Astrophys J, 729, 72
[5] Peters et al. 2010, Astrophys J, 725, 134
[6] Peters 2014, Eur J Phys, 35, 065028
[7] Peters et al. 2014, Astrophys J, 788, 14
[8] Peters et al. 2010, Astrophys J, 719, 843

Numerical Simulation of Binary Black Hole and Neutron Star Mergers

1

Research Institution
Albert Einstein Institut

Principal Investigator
W. Kastaun, L. Rezzolla

Researchers
R. Ciolfi, K. Dionysopoulou, I. Hinder, S. Hopper, D. Siegel, K. Takami

Project Partners
–

SuperMUC Project ID: pr32pi

Introduction

One of the last predictions of general relativity that still awaits direct observational confirmation is the existence of gravitational waves. Those fluctuations of the geometry of space and time are expected to travel with the speed of light and are emitted by any accelerating mass. Only the most violent events in the universe, such as mergers of two black holes or neutron stars, produce gravitational waves strong enough to be measured. Even those waves are extremely weak when arriving at Earth, and their detection is a formidable technological challenge. In recent years sufficiently sensitive detectors became operational, such as GEO600, Virgo, and LIGO. They are expected to observe around 40 events per year.

To interpret the observational data, theoretical modeling of the sources is a necessity, and requires numerical simulations of the equations of general relativity and relativistic hydrodynamics. Such computations can only be carried out on large scale supercomputers, given that many scenarios need to be simulated, each of which typically occupies hundreds of CPU cores for a week.

Our main goal is to predict the gravitational wave signal from the merger of two compact objects. Comparison with future observations will provide important insights into the fundamental forces of nature in regimes that are impossible to recreate in laboratory experiments. The waveforms from binary black hole mergers would allow one to test the correctness of general relativity in previously inaccessible regimes. The signal from binary neutron star mergers will provide input for nuclear physics, because the signal depends strongly on the unknown properties of matter at the ultra high densities inside neutron stars, which cannot be observed in any other astrophysical scenario. Besides mergers, we also want to improve the theoretical models of close encounters between black holes.

A gravitational wave detector with even higher sensitivity, the Einstein Telescope, is already in the planning stage. It is natural to ask what could be learned from the expected observations. Particularly intriguing is the possible detection of sources at cosmological distances. The fact that observed frequencies will be reduced due to the expansion of the universe could be used to determine the distance, but only if the original frequency is known. As it turns out, numerical modeling of the source might provide this information. Independent distance measures are very valuable for cosmology, improving estimates of the past expansion rate of the universe and predictions of its future fate.

Another goal of our project is to shed some light on the mystery of so called short gamma-ray bursts, intense and sudden bursts of gamma radiation that puzzled astronomers since decades. Previous simulations indicated that they are caused by neutron-star mergers. The exact emission mechanism is however unknown, and many features were completely unexplained, for example the X-ray afterglows which often accompany the main burst.

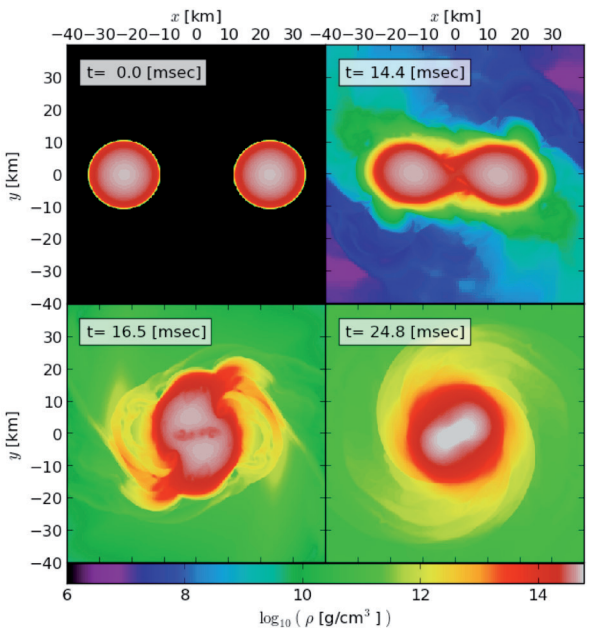

Figure 1: Snapshots of a binary neutron star merger, from late inspiral (top left), touching (top right), merging (bottom left) to the formation of a metastable hyper-massive neutron star (bottom right). The colors represent the mass density in the orbital plane.

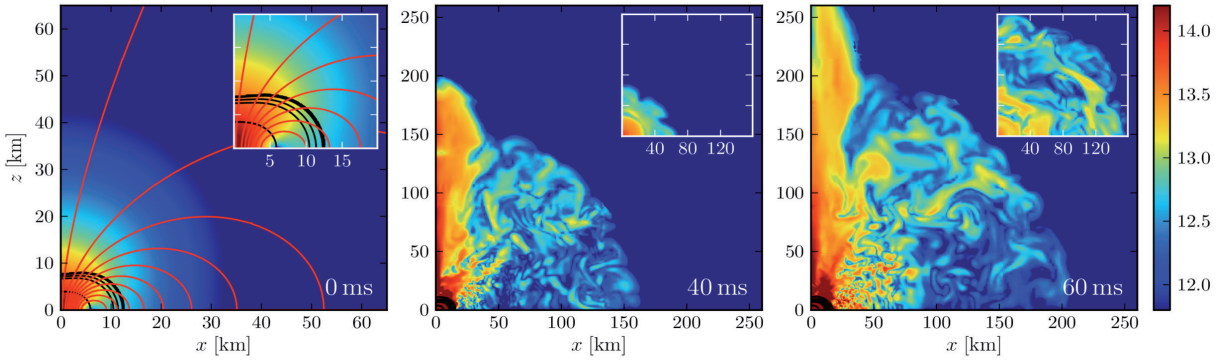

Figure 2: Evolution of a differentially rotating and strongly magnetized neutron star as produced in a neutron star merger. A dense and magnetized outflow is powered at the expense of rotational energy and generates electromagnetic emission that is compatible with the observed X-ray afterglows of short gamma-ray bursts. The panels show the color-coded magnetic field strength in Gauss (log scale) and indicate the initial magnetic field (red lines) as well as the neutron star (black lines).

Results

To achieve the project goals, we made use of a large collection of software modules based on the CACTUS computational toolkit, which provides infrastructure for large scale simulations, including MPI parallelization, mesh refinement (provided by the CARPET driver), memory management, and data storage. The equations of general relativity are integrated using the McLachlan module from the Einstein Toolkit, which is based on high order finite differencing methods. The (magneto-)hydrodynamic equations are evolved using the Whisky code. It employs modern finite volume methods for conservation laws in conjunction with various approximate Riemann problem solvers and a piecewise parabolic reconstruction algorithm. Initial data for binary neutron stars was generated using the LORENE code, a solver for elliptic PDEs based on spectral methods. In total, we used 14 million core hours for simulations of binary neutron star mergers, magnetized hyper-massive neutron stars, binary black hole mergers and close encounters. A typical run required around 200-500 cores over several days up to several weeks. The resulting data occupies more than 19 TB permanent disk storage, contained in around 1 million files.

Binary neutron star mergers
Using the above codes, we computed the inspiral and merger of binary neutron stars with 5 different masses, assuming a simple analytic model (gamma-law equation of state) for the properties of matter at high densities. An example is shown in Figure 1. For each model, we extracted the gravitational wave signal. The corresponding power spectrum contains contributions from the inspiral phase and the hyper-massive neutron star stage. It is also shifted to lower frequencies with increasing distance to the source. We demonstrated that these features could be combined to infer the absolute distance and total mass, assuming a particular equation of state of the neutron star matter [1]. We have shown that such observations of sources at cosmological distances with sufficient signal to noise ratio should be possible with the proposed Einstein Telescope.

Hypermassive neutron stars
A large part of our simulations was devoted to hyper-massive neutron stars formed during binary neutron star mergers. Besides the gravitational wave signal, the evolution of the magnetic field is of astrophysical interest. We therefore performed general relativistic magneto-hydrodynamics simulations of such models, in the limit of infinite conductivity (ideal MHD). Here, we could show that differential rotation causes magnetically driven outflows, carrying substantial amounts of energy [1]. Those winds might well be the source of the X-ray afterglows which often accompany short gamma-ray bursts. Figure 2 depicts the development of the wind.

Further, we studied the effect of large magnetic fields on the so called bar-mode instability. This instability develops only for rapidly rotating neutron stars, such as those formed in mergers, and converts rotational energy into gravitational waves, also accelerating the collapse of the star to a black hole. We found that only very high magnetic field strengths exceeding 1016 Gauss can delay or even suppress the instability [3].

Binary black hole systems
In contrast to binary black hole mergers, which are well investigated by now, little was known on the dynamics of close encounters. We performed numerical simulations of such encounters with different impact parameters and compared the results with existing analytic estimates, thus determining the accuracy of such predictions in different regimes [4]. Besides physics runs, a fraction of our resources was used to further test and develop the codes used for simulating binary black hole mergers, also with regard to numerical optimization.

On-going Research / Outlook

Currently, we are expanding our results mainly in three directions. One is the influence of the matter properties at high densities. Another is the influence of resistive effects on the evolution of magnetic fields. Finally, we are including the effects of neutrino cooling into merger simulations.

References and Links

[1] Phys. Rev. X 4, 041004, "Host redshifts from gravitational-wave observations of binary neutron star mergers"
[2] ApJL, 785, L6, "Magnetically driven winds from differentially rotating neutron stars and X-ray afterglows of short gamma-ray bursts"
[3] Phys. Rev. D 88, 104028, "Dynamical bar-mode instability in rotating and magnetized relativistic stars"
[4] Phys. Rev. D 89, 081503, "Strong-Field Scattering of Two Black Holes: Numerics Versus Analytics"

1

Supersonic turbulence in Giant Molecular Clouds

Research Institution
Institute for Theoretical Astrophysics, University of Heidelberg

Principal Investigator
Lukas Konstandin

Researchers
Philipp Girichidis, Rahul Shetty, Wolfram Schmidt, Thomas Peters, Ralf S. Klessen

Project Partners
School of Physics & Astronomy-University of Exeter, Max-Planck-Institut für Astrophysik-Garchingen, Hamburger Sternwarte-Universität Hamburg

SuperMUC Project ID: pr45si (Gauss Large Scale project), h1343

Introduction

The space in between stars is filled with a dilute mixture of charged particles, atoms, molecules and dust, called the interstellar medium (ISM). Observations reveal that the ISM is highly turbulent and characterized by complex structures on all resolvable spatial scales. The ISM is the primary galactic repository out of which stars are born and into which they deposit energy, momentum and enriched material as they die. It is the central building block of the galactic matter cycle, and as such a comprehensive understanding of it is necessary for developing an overall picture of galaxy formation and evolution. When the ISM forms a dense region that can self-shield its material against the surrounding violent radiation field, molecules form that enable the material to cool down to few Kelvins increasing the density further. These molecular clouds are the densest part of the ISM and are highly turbulent. The supersonic turbulent velocity field is likely responsible for the complex and filamentary density structures observed in molecular clouds. Turbulent motions create dense regions inside the molecular cloud that can become gravitationally unstable and collapse into dense cores, and eventually turn into new stars.

Common terrestrial flows are incompressible, in contrast to astrophysical flows, which are highly supersonic and compressible. While large improvements have been made in the understanding of incompressible turbulence, there are still open questions in our understanding of compressible turbulence. We therefore analyze super-sonic turbulent flows using three-dimensional, high performance, numerical simulations with resolutions up to 1024^3 grid cells. Our research focuses on the relation between the statistical properties of the velocity field and the density field. A detailed description of our research can be found online [1].

Results and Methods

We use FLASH4 to numerically integrate the continuity equation and the Euler equation on a uniform three-dimensional grid with periodic boundary conditions. We compute a random forcing field with a Wiener process in Fourier space such that it is smooth in space and time. With the Helmholtz decomposition we can adjust the composition of the forcing field between purely solenoidal (divergence-free) and purely compressive (curl-free) continuously. The stochastic forcing stirs the medium all the time such that the simulations reach a statistically steady state of fully developed turbulence. We performed simulations with 256^3, 512^3, 1024^3 gird cells, a Mach number of M=5, and both extreme cases - purely solenoidal and purely compressive - for the forcing field to study the influence of the different drivers on the statistical properties of the flow. These simulations used in total ~2Mio CPU/h and 10TB disk space. Since the theoretical predictions of supersonic turbulence, for e.g. the power spectrum, are reached in the time average, we run each simulation for 15 turnover times resulting in 150 files for each simulation. Figure 1 shows a cut through the xy-plane of the solenoidal (left) and compressive (right) driven simulations. The density fluctuations are more space filling with solenoidal forcing, and have smaller amplitude, while compressive forcing yields larger voids and denser regions [2, 3]. The energy flux through the scales is an important quantity, as it is assumed to be constant in incompressible turbulence theory. In our simulations the forcing routine injects the kinetic energy on large scales by creating big curls and motions on scales of half of the box size. These motions decay to smaller scales via interactions, until scales of few grid cells are reached on which dissipation effects take the kinetic energy out of the system. We measure the power spectra of the velocity field to analyze the scaling properties of the flow. We found that the compressive part of the velocity field has different scaling properties than the solenoidal one, which has in comparison more kinetic energy on small/intermediate scales [4].

Additionally, we studied the influence of the Mach number on the statistical properties of the flow. We performed therefore 15 simulations with 256^3 resolution

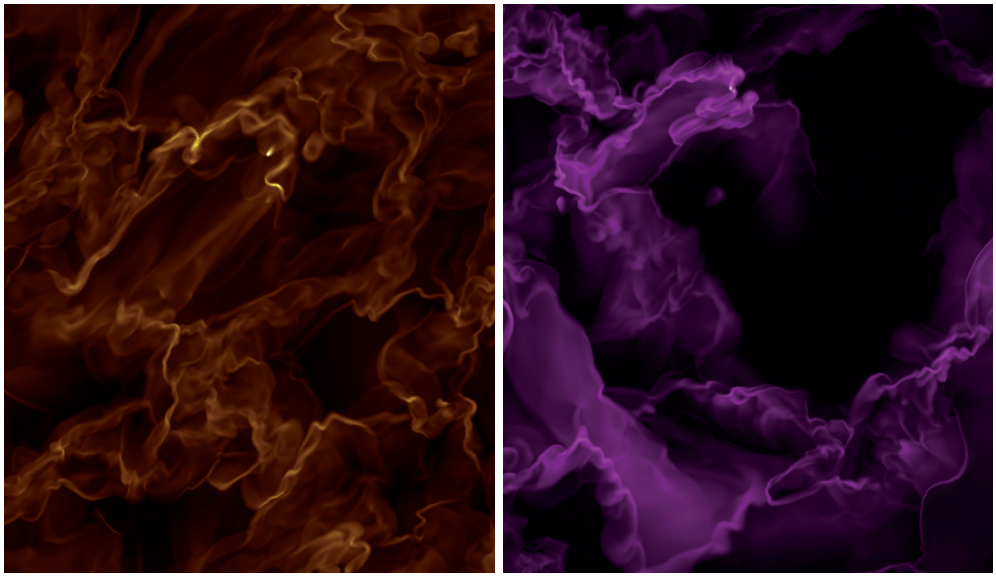

Figure 1: "The influence of the forcing on the density field": The density in the statistically steady state of fully developed turbulence, shown in a cut through the xy-plane of the simulations with Mach number M=16 driven by a purely solenoidal (left) and purely compressive (right) forcing field.

with root mean square Mach numbers between M=0.5 (slightly subsonic) and M=15 (highly supersonic), and with a mixed forcing field. We ran simulations with Mach numbers of 0.5, 5, 10, 16 and resolution of 512^3 and 1024^3 to analyze the influence of the resolution on our results. These simulations used in total ~6Mio CPU/h and 50TB disk space. The 1024^3 simulations were calculated on 512 nodes with 16 cores per node. To illustrate the influence of different Mach numbers on the flow pattern in the statistically steady state of fully developed turbulence we show in Figure 2 a cut through the xy-plane of the simulations with 1024^3 grid cells and Mach numbers of M=5 and M=16. We measured the parameters describing the scaling relation of the density power spectrum. The density power spectra follow power laws with a scaling exponent depending on the Mach number of the flow. We showed the influence of the Mach number on the space-filling factor of the flow and its fractal dimension [5]. It is likely that these parameters influence the fragmentation degree of giant molecular clouds and are therefore important for the theoretical description of the star formation process.

On-going Research / Outlook

The milestone work by Kolmogorov (1941) describing an incompressible turbulent flow is one of the few analytic results in turbulence theory. One of the most important results is the universal 5/3-scaling law for the energy power spectrum. There are many phenomenological models extending the Kolmogorov theory for the compressible case proposing a power-law behavior for the density, velocity, and combinations of these fields. We plan a systematic study of the simulations described above, measuring the influence of the Mach number and the composition of the forcing field on the scaling behavior of different mass-weighted power spectra to seek for an universal scaling law of compressible turbulence.

References and Links

[1] http://LukasKonstandin.de
[2] Konstandin L., et al., Feb. 2012, JFM, 692, 183-206
[3] Konstandin L., et al., Dec 2012, ApJ, 761-2, 149-156
[4] Konstandin L., et al., Jan. 2015, MNRAS, 446-2, 1775-1783
[5] Konstandin L., et al., submitted MNRAS, (http://arxiv.org/abs/1506.03834)

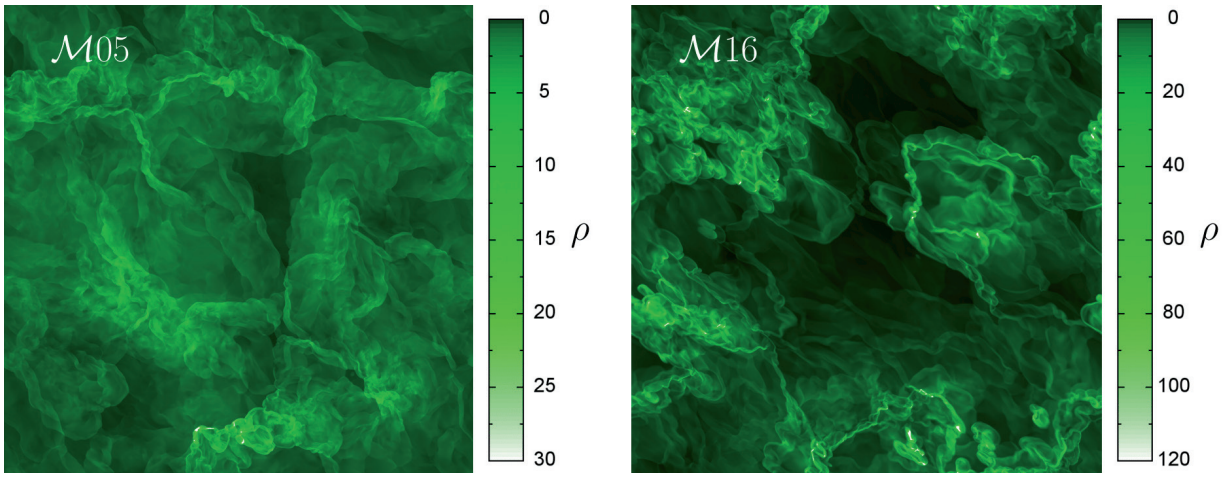

Figure 2: "The influence of the Mach number on the density field": The density in a cut through the xy-plane for Mach numbers M=5 (left) and M16 (right). Both simulations are driven by a mixed forcing field. Pictrure taken from [2].

The small and the beautiful: how the star formation law affects galactic disc structure

1

Research Institution
Institut für Astrophysik, Georg-August-Universität Göttingen, and Hamburger Sternwarte

Principal Investigator
Wolfram Schmidt

Researchers
Harald Braun

Project Partners
Computational Cosmology Center at Lawrence Berkeley National Laboratory

SuperMUC Project ID: pr47bi

Introduction

Disc galaxies are composed of gas and stars, embedded in a dark matter halo. The gas component is a highly non-uniform and dynamical medium, which is called the interstellar medium (ISM). There are several gas phases with different temperatures and densities. For star formation, the cold phase is particularly important because only in this phase the densities can reach the critical threshold for gravitational collapse, which is given by the Jeans instability criterion. Basically, this criterion states that a self-gravitating gas concentration collapses if its gravitational potential energy is sufficiently large compared to its thermal energy. A major problem is that the ISM undergoes dynamical changes on length and time scales that are extremely difficult, if not impossible to access in numerical simulations of whole disc galaxies. Although adaptive mesh refinement (AMR) enables us to increase the numerical resolution in some regions in the course of a simulation, the smallest visible details are still many lightyears across and thus larger than the size of a typical star-forming cloud. Consequently, the highest achievable resolution is too coarse to directly compute processes such as star formation and the impact of supernova explosions. This is why simulators incorporate these effects in an approximate fashion via subgrid-scale (SGS) models. For example, gas mass is converted into particles representing stars in simulations. The SGS model developed in this project controls the formation of stars by relating the star formation rate (the gas mass converted into stars per unit time) to local properties of the gas, such as the gas density and the thermal and turbulent energies, which are in turn affected by feedback from supernovae. The first part of our project was concerned with the implementation and validation of the SGS model. In the final stage of the project, we investigated how the galactic disc structure is influenced by different model assumptions for calculating the star formation rate.

Results and Methods

Our reference simulation was initialized with a rotating gaseous disc of uniform temperature in hydrostatic equilibrium. Since disc galaxies are embedded in dark matter halos, we applied a spherically symmetric gravitational potential in addition to the self-gravity of gas and stars. We used the Nyx code [2] to compute the disc dynamics resulting from the combined effects of gravity, gas cooling and heating, turbulence, star formation, and feedback due to stellar radiation and supernovae. The application of AMR allowed us to resolve the disc structure down to 30 parsecs (about 100 light years, corresponding to roughly 1/1000 of the diameter of the galactic disc). To run the simulation, we used 1024 cores (128 nodes) of the thin node segment of SuperMUC, using a hybrid parallelization scheme with 8 OMP threads per node and a single MPI process running on each node.

In the first project stage, we evolved the reference simulation over 1 Gyr (billion years) of physical time, consum-

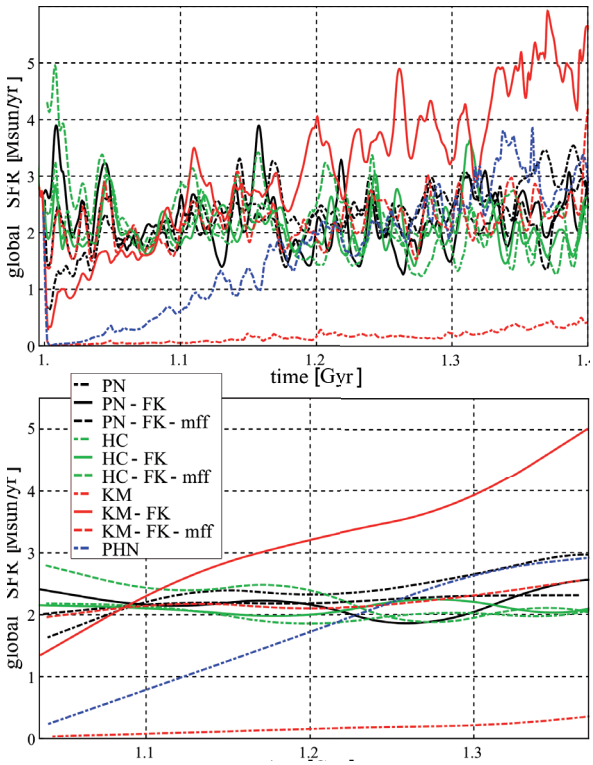

Figure 1: Time evolution of the star formation rate averaged over the whole galactic disc in units of solar masses per year. Each line corresponds to a different simulation with the star formation model indicated in the legend. For better visibility of systematic differences, the star formation rate was smoothed over a period of 0.08 Gyr in the bottom plot. Taken from [4].

ing 2 Mio. core hours. Each data dump produced about 200 GB output for postprocessing. Among the main results were predictions of star formation efficiencies (i.e. the fraction of gas turned into stars over the collapse time scale of a star forming cloud) around 1 % and a nearly linear relation between the local star formation rate and the density of molecular hydrogen [3]. In contrast to galaxy simulations performed by other groups, in which these relations are put in as empirical relations, they follow from the theoretical model of turbulence, star formation, and feedback implemented in our simulation.

However, there are various models for the calculation of the star formation rate. To understand which model-dependent changes might come about, we produced a suite of simulations by starting with the final state of the reference simulation, switching to a different star formation model and then evolving the disc for additional 0.4 Gyr. Running these simulations required roughly 1 Mio. core hours for each run.

Star Formation Models

In brief, we can distinguish three major families of models for calculating the star formation rate from local gas properties and the intensity of turbulence:

- The model of Padoan and Nordlund (PN) assumes that supersonic turbulence produces layers of shocked cold gas in the ISM, which must reach a critical density to collapse into stars. They also proposed a simplified model, which neglects the thermal energy of the gas (PHN).
- Hennebelle and Chabrier (HC) consider a hierarchy of dense structures in the turbulent ISM, whose collapse is partially inhibited by turbulence.
- Krumholz and McKee (KM) proposed a model similar to PN, however, with a different stability criterion.

Each model comes in several variants, which differ for example in the choice of numerical coefficients. For the reference simulation, we used the model labeled PN-FK in Figs 1 and 2.

Star Formation Rate and Disc Structure

An important indicator is the star formation rate averaged over the whole galactic disc (Fig. 1). For the majority of models, a nearly constant star formation rate of about 2 solar masses per year is obtained, which is typical for a Milky Way-like galaxy. In some cases, however, a steep initial drop or a drifting star formation rate can be seen, which rules out the corresponding models.

Moreover, it turns out that the gaseous disc structure is sensitive to the star formation model. Figure 2 shows contour plots of the gas density averaged over the direction perpendicular to the disc (the so-called column density) for the different simulations. Although the majority of models result in a similar, spiral-like disc structure, markedly clumpy discs are found for some models, which can also be identified as outliers in the plot of the star formation rate in Fig. 1. Although subtle differences are discernible, the PN and HC models produce sensible results, suggesting that these models capture the essential physics of star formation.

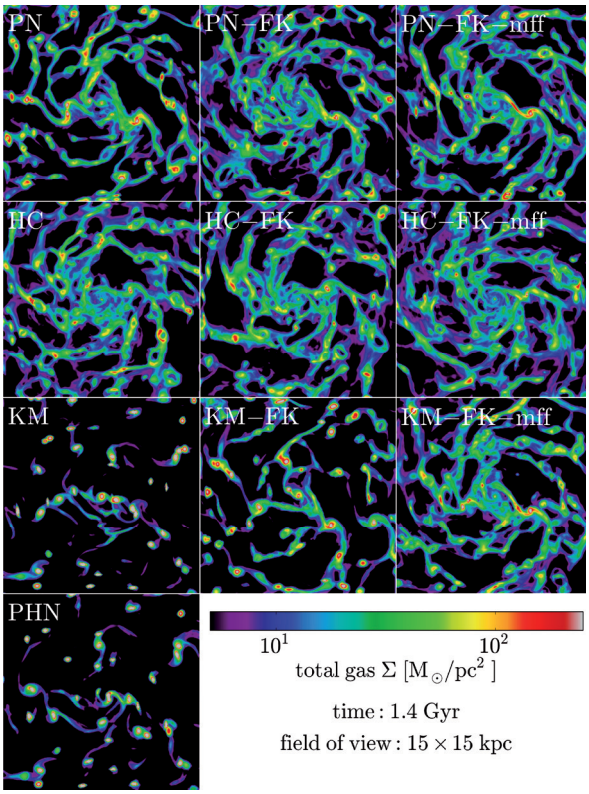

Figure 2: Gas column density for different simulations, with the star formation models listed in Fig. 1, taken from [4]. Depending on the chosen model, the disk has a more spiral-like (e.g. HC) or a very clumpy (e.g. KM) structure.

On-going Research / Outlook

The application of our SGS model allowed for the first time a comparative study of different star formation models in global simulations of galaxies.

SuperMuc is one of the few computational facilities in Germany that allows us to carry out such simulations. Apart from the speed and the available number of cores, the relatively large memory per core (about 2 GB) and the very fast interconnections are critical for communication-intensive (AMR) and memory-intensive (AMR and SGS model) applications. Furthermore, it is important that the architecture of SuperMuc is very well suited for the MPI/OMP hybrid parallelization. This allowed us to make optimal use of the load balancing implemented in the Nyx code.

We successfully finished our project in 2015. Future projects will focus on the dynamics and evolution of magnetic fields in disc galaxies, which is substantially more challenging in terms of modeling and computational requirements.

References and Links

[1] www.hs.uni-hamburg.de/index.php?lang=en
[2] ccse.lbl.gov/Research/NYX/index.html
[3] Braun, H., Schmidt, W., Niemeyer, J.C., and Almgren, A.S. 2014. Large-eddy simulations of isolated disc galaxies with thermal and turbulent feedback, MNRAS, 442, 3407
[4] Braun, H. and Schmidt, W. 2015. The small and the beautiful: how the star formation law affects galactic disc structure, MNRAS, 454, 1545

Simulating and observing the formation of protostellar disks

1

Research Institution
1. Physikalisches Institut, Universität zu Köln

Principal Investigator
Daniel Seifried

Researchers
Daniel Seifried, Robi Banerjee, Ralf Klessen, Alvaro Sanchez-Monge, Stefanie Walch

Project Partners
Hamburger Sternwarte, Universität Hamburg; Institut für theoretische Astrophysik, Universität Heidelberg; McMaster University, Hamilton, Canada

SuperMUC Project ID: pr47pi

Introduction

In our research we investigate the formation of protostars and their associated protostellar disks, the early precursors of planetary systems as well as their apperance in observations with modern telescopes like ALMA.

During the last decade simulations of collapsing molecular cloud cores have revealed the so-called catastrophic magnetic braking problem: Magnetic fields are able to transport angular momentum by means of toroidal Alfven waves. Modeling the collapse of rotating molecular cloud cores, simulations have shown that in the presence of magnetic fields with strengths comparable to observational results, the formation of rotationally supported (Keplerian) protostellar disks is largely suppressed. This is due to the fact that angular momentum is removed very efficiently from the interior of the core by the magnetic field. In previous works we could confirm this effect for the collapse of massive (100 solar masses), molecular cloud cores. This key result of the suppression of Keplerian disk formation during the earliest stages of star formation is in contrast to recent observational results which state that protostellar disks should be present already in the Class 0 stage.

Results

The results described before show up in case that highly idealized initial conditions are used for the simulations. In particular the lack of turbulent motions - frequently observed in molecular cloud cores - could have a significant effect on the formation of protostellar disks and outflows. For this reason, in our research we recently focused on the influence of turbulence on the formation of protostellar disks and outflows.

In our work we have performed a number of simulations on SUPERMUC. Each of the simulations required a computational time of a few 100 000 CPU-hours with a simultaneous use of up to 1000 CPUs per simulation. A few hundreds of files were produce for each simulation requiring a disk space of a few TB in total. The simulations are performed with the hydrodynamics code FLASH[1] written in Fortran 90. The code solves the 3-dimensional, discretized magnetohydrodynamical equations on a Cartesian grid. Making use of the adaptive-mesh-refinement (AMR) technique, only those regions which are of particular interest for us are resolved with the highest possible spatial resolution whereas other regions of minor interest are resolved more coarsely. This significantly reduces the number of calculations to be performed and hence the computational time required, thus allowing us to perform the simulations over long physical timescales.

Initial conditions
Observations of the birth places of stars show a wide range of physical quantities, in particular in their initial mass. As we do not simulate a particular region observed by astronomers but rather aim to understand the systematic influence of the initial conditions, we have to perform a number of simulations in our work covering a wide range of masses and turbulence strengths. This allows us to draw conclusions about the effect of the

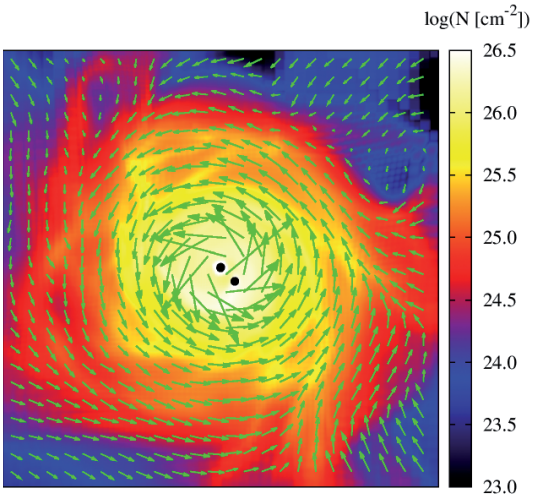

Figure 1: Protostellar disk seen from top-on. Black dots represent protostars, green arrows the velocity field.

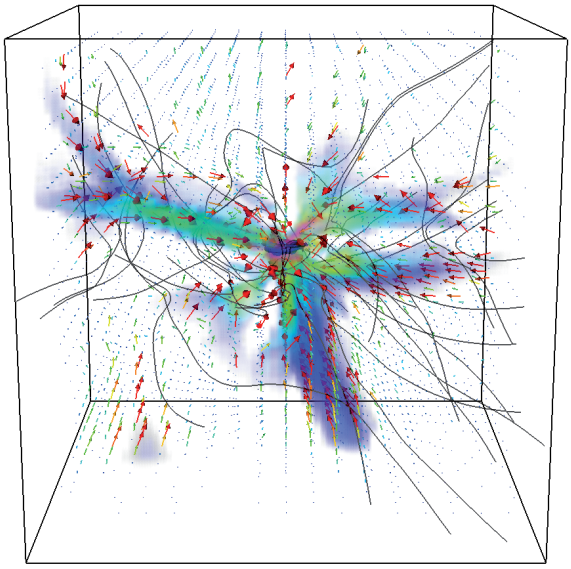

Figure 2: 3-dimensional structure of the magnetic field and gas motions around a Keplerian disk in one of our simulations.

initial conditions on the formation mechanism of stars. We modeled the collapse of molecular cloud cores with masses ranging from about 2 solar masses up to 1000 solar masses. The cores are threaded by a strong magnetic field along the z-axis and have an additional supersonic, turbulent velocity field as indicated by observations.

Turbulence-induced disk formation

We examine our simulations focusing on the question of how turbulence affects the formation of Keplerian disks[2,3]. An example result is shown in Fig. 1 showing the protostellar disk in a representative run with a core of 100 solar masses. As can be seen, in the presence of turbulence rotationally supported disks are re-obtained again, which is in contrast to the previous simulations with comparable magnetic field strengths but no turbulence. This suggests that the efficiency of the magnetic braking, which is responsible for removing angular momentum from the midplane, is reduced significantly in the turbulent case. Analyzing the surroundings of the disks we can show that this indeed the case. The turbulent surroundings of the disk exhibit no coherent rotation structure (compare Fig. 1). Therefore, an efficient build-up of a strong toroidal magnetic field responsible for angular momentum extraction is hampered. Moreover, the turbulent motions lead to a strongly disordered magnetic field which further reduces the magnetic braking efficiency. Since simultaneously the angular momentum inwards transport remains high due to the presence of local shear flows in the vicinity of the disks, there is a net inwards angular momentum transport towards the center of the disk. The combination of these effects results in the observed build-up of Keplerian disks as expected from observations. Varying the core masses (2.6 - 1000 solar masses) and the turbulence strengths does not change our findings. This clearly demonstrate that the turbulence-induced disk formation mechanism works for a wide range of initial conditions. In particular, we could show that the formation of Keplerian disks does not require an uniform rotation of the core – turbulent motions alone lead to the build-up of Keplerian disks. Moreover, we showed that even for subsonic turbulence, which is usually present in low-mass protostellar cores, the turbulence-induced formation mechanism still holds. In Fig. 2 we show the 3-dimensional structure of the magnetic field and the gas motions around a protostellar disk in one of our simulations. The magnetic field lines clearly reveal a highly complex structure being far off from well-ordered. The anisotropy of the accretion towards the disk is represented by the colored regions in the figure.

Synthetic observations

In a last step we produced so-called synthetic observations. These synthetic observations allow us to directly compare our simulation results with real observations e.g. made with ALMA. Such comparisons allow us to assess what can be inferred from observations - in particular how reliable parameters obtained from real observations are[4].

In particular we focus on the question whether discs in such an early evolutionary stage are actually detectable since they are still deeply embedded in surrounding material. We find that with ALMA it is in general possible to detect these protostellar discs and to determine their internal dynamics – a key to understand the formation of later planets. In Fig. 3 we show a synthetic observation made with ALMA for the disc shown in Fig. 2 for two different viewing angles. In both cases the disc is clearly recognisable and in the left panel even fragments can be identified.

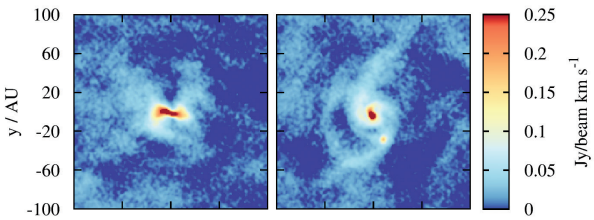

Figure 3: Synthetic observation with ALMA of one of our simulated disc seen from edge-on (left) and face-on (right).

On-going Research / Outlook

For our future research we intend to study the self-consistent ejection of protostellar outflows from the Keplerian disks formed in our turbulence simulations. In order to reach this goal we will have to redo some of our simulations with increased spatial resolution, which will require further large amounts of computing power in the future.

References and Links

[1] Fryxell, B., Olson, K., Ricker, P., et al. 2000, ASTROPHYS. J. SUPPL. S., 131, 273
[2] Seifried, D., Banerjee, R., Pudritz, R. E., & Klessen, R. S. 2012a, MNRAS, 423, L40
[3] Seifried, D., Banerjee, R., Pudritz, R. E., & Klessen, R. S., 2013, MNRAS, 432, 3320
[4] Seifried, D., A. Sanchez-Monge, S. Walch, R. Banerjee, 2016, arXiv:1601.02384

The world's largest turbulence simulation

1

Research Institution
Research School of Astronomy and Astrophysics, Australian National University;
Zentrum für Astronomie der Universität Heidelberg, Institut für Theoretische Astrophysik

Principal Investigator
Christoph Federrath & Ralf S. Klessen

Researchers
Luigi Iapichino & Nicolay Hammer

Project Partners
Leibniz-Rechenzentrum

SuperMUC Project ID: pr48pi (Gauss Large Scale project), pr32lo, pr89mu (PRACE project)

Introduction

Understanding turbulence is critical for a wide range of terrestrial and astrophysical applications. For example, turbulence on earth is responsible for the transport of pollutants in the atmosphere and determines the movement of weather patterns. But turbulence plays a central role in astrophysics as well. For instance, the turbulent motions of gas and dust particles in protostellar disks enables the formation of planets. Moreover, virtually all modern theories of star formation rest on the statistics of turbulence [1]. Especially the theoretical assumptions about turbulence behind star formation theories allow the prediction of star formation rates in the Milky Way and in distant galaxies [2]. Interstellar turbulence shapes the structure of molecular clouds and is a key process in the formation of filaments which are the building blocks of star-forming clouds.

The key ingredient for all these models is the so-called sonic scale. The sonic scale marks the transition from supersonic to subsonic turbulence and produces a break in the turbulence power spectrum from $E \propto k^{-2}$ to $E \propto k^{-5/3}$. While the power-law slopes of -2 and -5/3 for the supersonic and subsonic parts of the spectrum have been measured independently, there is no simulation currently capable of bridging the gap between both regimes. This is because previous simulations did not have enough resolution to separate the injection scale, the sonic scale and the dissipation scale.

The aim of this project is to run the first simulation that is sufficiently resolved to measure the exact position of the sonic scale and the transition region from supersonic to subsonic turbulence. A simulation with the unprecedented resolution of 10000^3 grid cells will be needed for resolving the transition scale.

Results

In the framework of a GAUSS Large Scale Project, an allocation exceeding 40 million core-h has been granted to this project on SuperMUC. The application used for this project is FLASH, a public, modular grid-based hydrodynamical code for the simulation of astrophysical flows [3]. The parallelisation is based entirely on MPI. In the framework of the SuperMUC Phase 2 scale-out, the current code version (Flash4) has been optimised to reduce the memory and MPI communication requirements. In particular, non-critical operations are now performed in single precision, without causing any significant impact on the accuracy of the results. In this way, the code runs with a factor of 4.1 less memory and 3.6 times faster than the version used for the previous large-scale project at LRZ [4], and scales remarkably well up to the full machine on SuperMUC Phase 2 (Figure 1).

The 10000^3 simulation has been nearly completed at the time of writing, and data processing is in progress. Some early impression of the forthcoming results can be seen from the highlights of the work of Federrath [4], based on the previous large-scale project on turbulence simulations (up to 4096^3 grid cells), selected as the SAO/NASA ADS paper of the year 2013.
Highly-compressible supersonic turbulence is complex, if compared to the subsonic, incompressible regime, because the gas density can vary by several orders of magnitude. Using three-dimensional simulations, we have determined the power spectrum in this regime (Figure 2), and found $E \propto k^{-2}$, confirming earlier indications obtained with much lower resolution [5]. The resolution

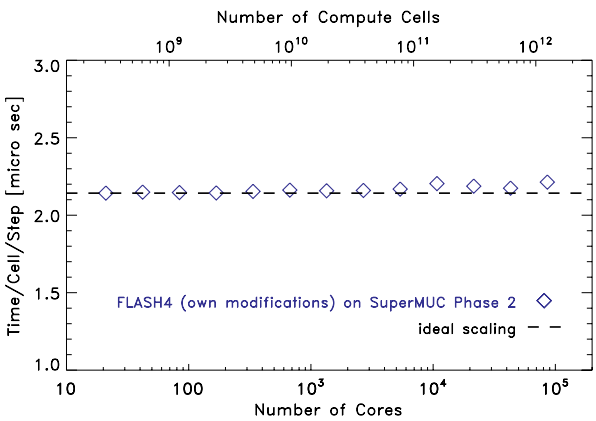

Figure 1: weak scaling of the customized version of the FLASH code, used during the SuperMUC scale-out workshop on Phase 2 in 2015. The blue diamonds indicate the scaling tests of the FLASH code, while ideal scaling is represented by the dashed line.

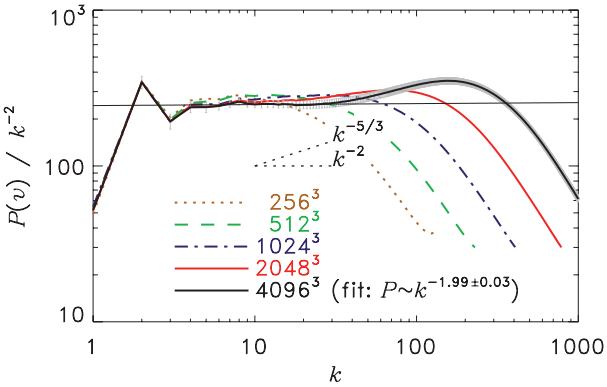

Figure 2: Power spectrum from highly-compressible, supersonic turbulence simulations (compressive driving), demonstrating a k^{-2} scaling [4].

study in Figure 2 shows that we would not have been able to identify this scaling at any lower resolution than 4096^3 cells. Extremely high resolution and compute power are absolutely necessary for the science done here.

Figure 3 displays the unprecedented level of detail in density and velocity structure achieved with our 4096^3 simulations. This visualization demonstrates the significant differences between two extreme driving modes of the turbulence, the solenoidal driving (with the stirring mechanism exciting vortices in the gas) and the compressive driving (when the gas is directly compressed by the stirring mechanisms). Simulation movies are available online (see links below).

Future Work

Turbulence has a wide range of applications in science and engineering, including the amplification of magnetic fields, star and planet formation, mixing of pollutants in the atmosphere, fuel ignition in engines, and many more. Generating the huge dataset of turbulence presented here, we have begun to reach the technical limits of what is feasible on any supercomputer in the world to date. We are currently pushing the boundaries even further by running the world's first turbulence simulation with more than 10000^3 grid cells on SuperMUC. We hope to unravel the statistics of supersonic and subsonic, magnetized turbulence in the near future, with cutting-edge supercomputing systems provided by the LRZ.

References and Links

[1] Padoan et al. 2014, Protostars and Planets VI, 77
[2] Federrath & Klessen 2012, ApJ 761, 156
[3] Fryxell et al. 2010. ApJS 131, 273
[4] Federrath 2013, MNRAS 436, 1245
[5] Kritsuk et al. 2007, ApJ 665, 416

http://www.gauss-centre.eu/gauss-centre/EN/Projects/Astrophysics/federrath_astrophysics_weltrekord.html?nn=1345700
http://www.ita.uni-heidelberg.de/~chfeder/pubs/supersonic/supersonic.shtml

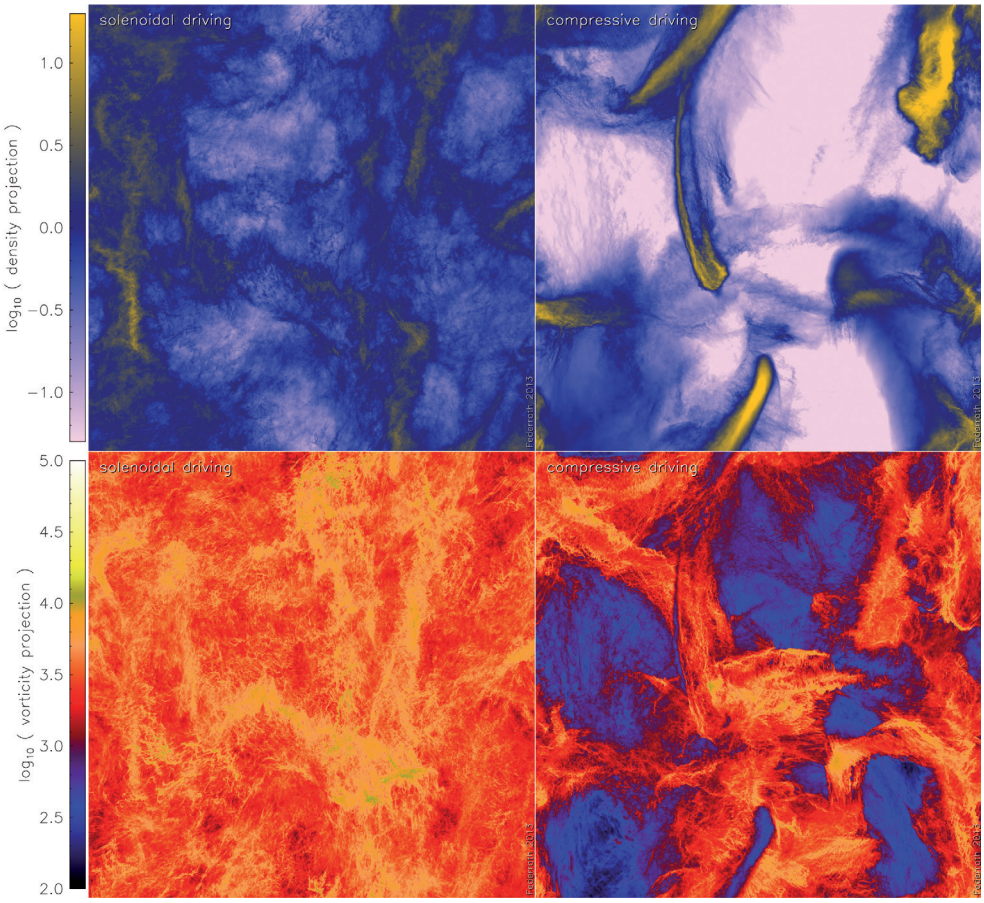

Figure 3: Projections through the three-dimensional gas density (top) and vorticity (bottom) in the simulations of supersonic turbulence with a grid resolution of 4096^3 cells [4]. Left shows solenoidal driving and right shows compressive driving of the turbulence, exhibiting extremely different statistics and star-formation activity [2].

Parasitic instabilities of the magneto-rotational instability

1

Research Institution
Max-Planck-Institut für Astrophysik

Principal Investigator
Ewald Müller

Researchers
Tomasz Rembiasz, Pablo Cerdá-Durán, Martin Obergaulinger

Project Partners
Department of Astronomy and Astrophysics, University of Valencia

SuperMUC Project ID: pr58xi

Introduction

The possible influence of magnetic fields on the dynamics of core-collapse supernovae and the evolution of nascent proto-neutron stars is largely unknown. As pointed out by [1], one of the most promising agents amplifying the magnetic fields in a core-collapse supernova explosion is the magnetorotational instability (MRI) [2]. The MRI has been studied intensively in accretion disks where it excites turbulence and provides the effective viscosity required for the accretion of the gas. Nevertheless, the saturation level of the instability is still not fully understood. It was suggested that the MRI growth is terminated by secondary parasitic instabilities (Kelvin-Helmholtz and resistive tearing-mode instabilities) growing on top of MRI modes [3,4]. Hence, the saturation level of the MRI is determined by the interplay between the MRI and the parasites. Extending our previous set of local simulations, we study this interplay in three-dimensional, non-ideal magnetohydrodynamics (MHD) simulations. Our final goal is the formulation of a simple model connecting the initial magnetic field strength, the resistivity and the viscosity of matter with the termination and saturation of the MRI.

Results and Methods

Under conditions typical for supernova cores, the MRI grows on length scales that are to small to resolve in global models of the entire core, which has a diameter of a few thousand kilometres. Hence, we performed a large set of *local* simulations that cover a small region of the core of about one kilometre extent with a very fine grid, within which we solved the equations of non-ideal magnetoydrodynamics. Within this region, we set up stellar matter in equilibrium between gravity, gas pressure, and the centrifugal forces given by a profile of differential rotation. If a weak magnetic field is included, the system is unstable, and perturbations grow in the form of *channel modes*, pairs of radial up- and down-flows threaded by layers of magnetic field of alternating polarity. They are susceptible to secondary, *parasitic* instabilities, in our case of Kelvin-Helmholtz type, whose growth rate increases along with the amplitude of the channel modes until their amplitudes become roughly equal, at which point the channels are disrupted and the MRI ceases to grow. This sequence of processes was identified in simulations and recently studied by detailed analytic theory. Our goal was to test predictions of these studies and explore the amount of field amplification can be expected before the parasites quench the MRI.

Our 3D simulations are performed with the Newtonian grid based flux-conservative finite-volume MHD code *Aenus* written in Fortran 95. The code employs a domain decomposition method and is hybrid parallelised (MPI + OpenMP). It has a good weak MPI scaling and a very good OpenMP speed up on both the FAT and THIN nodes of SuperMUC. The computational costs and the number of cores used depend on resolution. The smallest 3D simulations performed with 320 cores required ≈ 3,000 CPUhs. The largest simulations on the THIN nodes were run with 4,096 (512 MPI x 8 OpenMP) cores. The computational cost of the latter simulation is ≈ 1,000,000 CPUhs. In total, we have performed over fifty 3D simulations (not in-

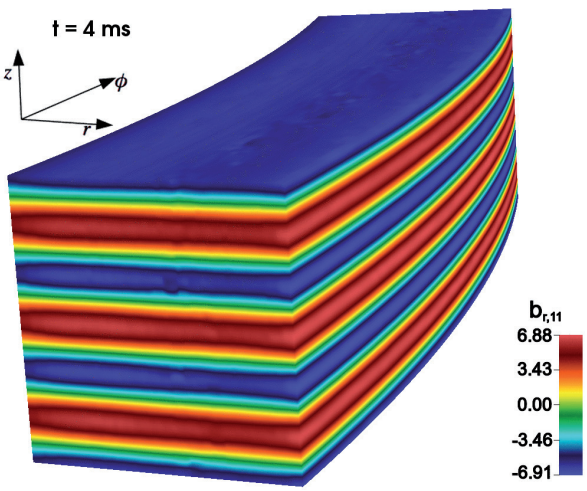

Figure 1: Radial component of the magnetic field in a 3D simulation performed in a box of a size 1 km x4 km x1 km in (r, φ, z) direction, respectively, and a resolution of 100 zones/km. 4 ms after introducing perturbations, three MRI channels are visible.

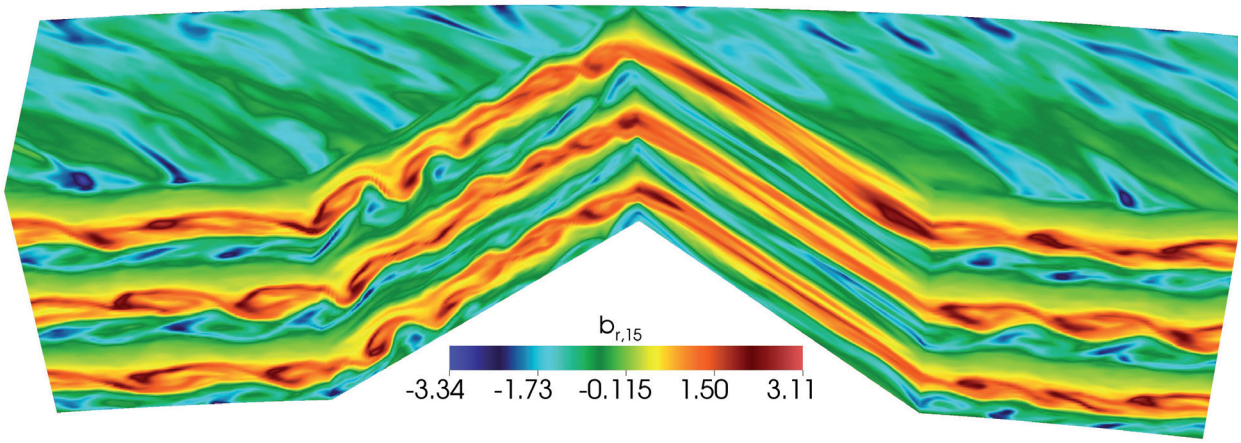

Figure 2: Radial magnetic field in the same model as shown in Fig. 1 at a time when the disruption of the MRI channels by Kelvin-Helmholtz modes (showing up as vortex rolls) is setting in.

cluding additional auxiliary tests) for which ≈ 7,000,000 CPUhs were used. More than a million files and directories were generated intermediately (post-processing reduced this number by two orders of magnitude), which occupied about 5 TB in our PROJECT directory.

Our 3D simulations, published in [5], confirm that, after channel modes such as the ones shown in Fig. 1 develop, the MRI in core-collapse supernovae is terminated by the Kelvin-Helmholtz instability (see Fig. 2). In 2D simulations, the assumed axial symmetry suppresses this secondary instability and more slowly developing tearing modes terminate the MRI growth. As a result, MRI-driven magnetic field amplification is overestimated. We conclude that the MRI termination process can be properly studied only in 3D simulations.

In a detailed analysis of our 3D models, we found a good agreement between our results and the predictions of [4] for the geometry of the parasites such as the orientation of the magnetic field and velocity vectors. In a next step, we investigated the termination level and compared our models to estimates based on the theory of parasitic modes as well as to simulations performed using a different physical framework (incompressible MHD) and numerical method (a pseudo-spectral code). When accounting for the methodological differences between the three approaches, we consistently obtain factors of amplification of the seed field of the order of a few 10. For supernova cores, the fact that we do not obtain much larger numbers seems to indicate that the MRI is not able to amplify any arbitrarily weak initial field to dynamically relevant strength during its exponential growth phase, which might limit the fraction of progenitors for which MHD effects can significantly contribute to the explosion dynamics.

On-going Research / Outlook

Supercomuting was crucial for this project mostly because of the high numerical resolution required to follow the evolution of small-scale structures developing as the MRI channels grow and, even more so, as they are disrupted by parasites. Thus, SuperMUC provided us with a unique opportunity to investigate processes that would otherwise be inaccessible to our understanding. We would in particular like to point to the last step of our project performed using SuperMUC phase 2. The two main avenues our research can now take would be, on the on hand, a closer look at the post-termination, turbulent phase of the MRI and, on the other hand, the integration of our results into global models. Both projects demand even larger computational resources because of the long time scales the simulations would have to run and because of the large computational domains to cover and the necessary inclusion of other effects such as neutrinos transport.

References and Links

[1] Akiyama, S., Wheeler, J. C., Meier, D. L., & Lichtenstadt, I., Astrophys. J., 584 (2003), 954
[2] Balbus, S. A. & Hawley, J. F., 1991, Astrophys. J., 376 (1998), 214
[3] Goodman, J. & Xu, G., Astrophys. J., 432 (1994), 213
[4] Pessah, M. E. 2010, Astrophys. J., 716 (2010), 1012
[5] Rembiasz, T., Obergaulinger, M., Cerdá-Durán, P., Mûller, E., & Aloy, M.Á., Monthly Notices of the Royal Astronomical Society, 456 (2016), 3782

Unveiling the equation of state of nuclear matter with binary neutron stars

1

Research Institution
Institut for Theoretical Physics, Goethe University (Frankfurt am Main)

Principal Investigator
F. Galeazzi, L. Rezzolla

Researchers
L. Bovard, K. Dionysopoulou, B. C. Mundim, L. J. Papenfort, K. Takami, A. Tsokaros

Project Partners
–

SuperMUC Project ID: pr84fa

Introduction

2015 marked the hundred anniversary of Albert Einstein's lecture at the Prussian Academy of Science in which he introduced, for the first time, the famous field equations which became the core of his *theory of general relativity*.

This masterpiece of 20th century science has proven extremely solid in all its predictions from the precession of the perihelion of Mercury to the observation of gravitational lensing in distant galaxies, to the more mundane time-delay corrections required by the global positioning system. One last piece of the puzzle is although still missing and comprise the direct measurement of the gravitational wave (GW) radiation emitted by any accelerating mass. These ripples in the spacetime fabric are extremely weak even when produced in the most extreme of the conditions as the ones present during the mergers of two black holes or neutron stars. For this reason they have eluded experimental scientists for almost four decades. But things are about to change, last year a new array of advanced gravitational wave detectors, namely advanced LIGO and Virgo came online in late September and they are expected to observe up to 40 events per year involving the mergers of two compact objects. Despite the high sensitivity of this generation of ground base interferometers, it is still necessary to use accurate gravitational waveforms models to extract all the information from the signal produced by the detector.

In this project we focus on the merger of two neutron stars which orbit together in a binary system. The nonlinear nature of the Einstein equations coupled with the complex microphysics behind neutron star matter requires the use of sophisticated codes which uses advanced numerical techniques to produce accurate results. By using the GW signals calculated in our numerical simulations we will be able to strongly link the properties of neutron star matter to a precise set of observable frequencies from the detector [1]. This information, together with the electromagnetic counterparts of these events, will shed some light on the engine that powers short gamma ray bursts. The properties of matter at the ultra high densities and low temperatures reached inside neutron stars cannot be observed in a conventional laboratory on Earth and for this reason accurate GW astronomy is a unique opportunity to constraint the current knowledge of the equation of state that describes these regimes. But GWs are not the only observable that can be linked to the equation of state of neutron star matter, during the violent merger of two neutron stars large amount of neutron rich material is ejected leading to the creation of heavy elements. While undergoing radioactive decay, these elements emit in near-infrared and optical bands of the electromagnetic spectrum. The characteristics of these emissions are strongly affected by the composition, temperature and total mass of the dynamically ejected material and for this reason we have developed a series of cutting-edge methods to simulate in full general relativity the inspiral, merger and collapse including relativistic hydrodynamics, the use of nuclear finite-temperature equations of state and an approximate treatment of neutrino emission and absorption. Such simulations require the use of computational facilities such as the one at LRZ where we make use of thousands of CPUs every week for each of our simulations and producing several terabytes of data. This data are processed in situ at the LRZ facility and, for a more detailed analysis, transferred to our local cluster in Frankfurt am Main (LOEWE).

Projects and results

The code that we use for our simulations is based on the CACTUS computational toolkit, which provides an efficient infrastructure for large scale simulations, including MPI and OpenMP parallelization, adaptive mesh refinement, memory and output management. The general relativistic field equations are integrated using the McLachlan module from the Einstein Toolkit, which is based on high order finite differencing methods. The general relativistic hydrodynamic equations are solved by WhiskyTHC which employs a high-order finite-difference highresolution shock-capturing method which allowed us to achieve third order convergence during the inspiral of a binary neutron star system. WhiskyTHC implements

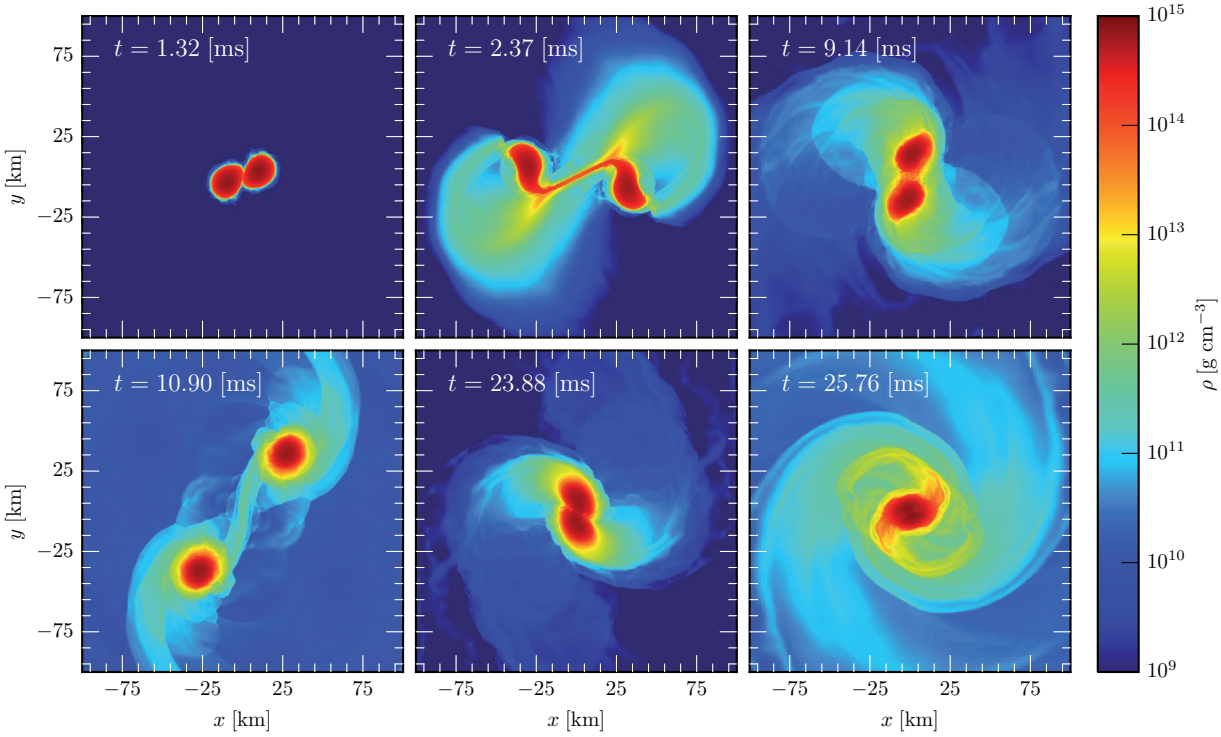

Figure 1: Snapshots of a binary neutron star collisions with an initial periastron of 14 km, from late inspiral (top left), touching (top right), merging (bottom left) to the formation of a metastable hyper-massive neutron star (bottom right). The colours represent the mass density in the orbital plane.

also a sophisticated treatment of neutrino interactions that takes into account the cooling as well as the absorption on the thick and hot neutron star matter. In total, we used 10 million core hours for simulations of binary neutron star mergers in both in quasi-circular as well as eccentric orbit using with several equations of state. A typical run required around 1000-1200 cores over several days up to few weeks of computation.

Binary neutron star collisions

We simulate the inspiral and merger of binary neutron stars configurations with 5 different values of the initial periastron radius, using a nuclear equation of state to describe the properties of matter at high densities. An example of one of these encounters is shown in Fig. 1. The dynamic of such systems is significantly more complicated than the one in a quasi-circular orbit, multiple encounters are in fact observed before the actual merger and the formation of a single fast rotating hypermassive neutron star (Fig.1 bottom right panel). Given the nonlinear nature of the system it is not possible to predict in advance the number of encounters that the binary will undergo making these simulations extremely challenging from the numerical standpoint.

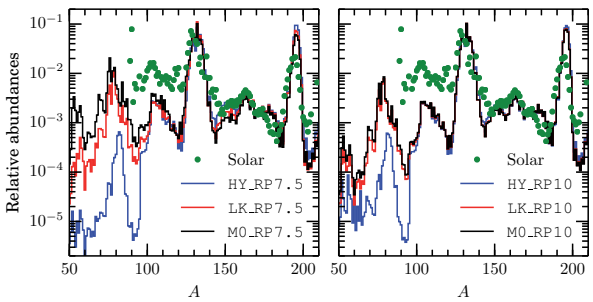

Figure 2: R-process nucleosynthesis produced by the ejected material of binary neutron star collisions, with periastron radius of 11 km (on the left) and 14 km (on the right). Different neutrino treatments are presented.

R-process nucleosynthesis from the ejecta

The most striking feature of binary neutron star collisions is the amount of material that is ejected reaching a value of about one-tenth of a solar mass. As a comparison, a similar binary in quasi-circular orbit can produce 20 times less ejected material. The ejected material is extremely neutron rich and it is thought to be responsible for the creation of a large part the heavy elements heavier than iron in the Universe. This process, know as r-process nucleosynthesis (rapid), requires a high neutron fraction environment which allow for an efficient neutron capture on the heavy seed nuclei. Our work analyses in great detail not only the effects of different orbital parameters in the production of r-process nuclei from the ejecta but also of different neutrino treatments. We have included the effects of neutrino cooling into our merger simulations and also of neutrino absorption onto the hot ejected material. Fig. 2 shows the results of the r-process nucleosynthesis calculations for different models and it will be part in the coming month of a series of publications on the topic [2].

References and Links

[1] "Spectral properties of the post-merger gravitational-wave signal from binary neutron stars" K. Takami, L. Rezzolla, L. Baiotti, 2015, PRD, 91, 064001

[2] "Dynamical Mass Ejection from Binary Neutron Star Mergers" D. Radice, F. Galeazzi, J. Lippuner and L. F. Roberts, C. D. Ott and L. Rezzolla 2016, arXiv:1601.02426 (submitted to MNRAS)

Modeling dispersive plasma waves with particle in cell simulations

Research Institution
Julius-Maximilians-Universität Würzburg

Principal Investigator
Cedric Schreiner

Researchers
Patrick Kilian, Urs Ganse, Andreas Kempf, Felix Spanier, Karl Mannheim

Project Partners
Max Planck Institute for Solar System Research, University of Helsinki, Ruhr-Universität Bochum, North-West University Potchefstroom

SuperMUC Project ID: pr84ti

Introduction

Our current SuperMUC project [1] concerns itself with an astrophysical problem which has been discussed for several decades: The Sun not only produces light, but also a constant stream of charged particles, the solar wind, which travel through the solar system. Contrary to naïve expectations, the energy spectrum of the solar wind particles does not follow a simple thermal spectrum, but an excess of highly energetic particles is measured. Therefore, an acceleration mechanism must be present which delivers energy to the particles on their way from the Sun to the Earth.

Since the possibilities for direct measurements are severely limited and mathematical approaches often need to employ drastic simplifications in order to yield analytic predictions at all, numerical simulations modeling the complex dynamics of the solar wind plasma have been on the rise over recent years. We use our Particle-in-Cell (PiC) code *ACRONYM* [2] to simulate the interaction of charged particles and electromagnetic fields self-consistently on a micro-physical scale. The focus of our research lies on the scattering of particles and plasma waves, which is believed to play an important role in the transport and acceleration of energetic particles.

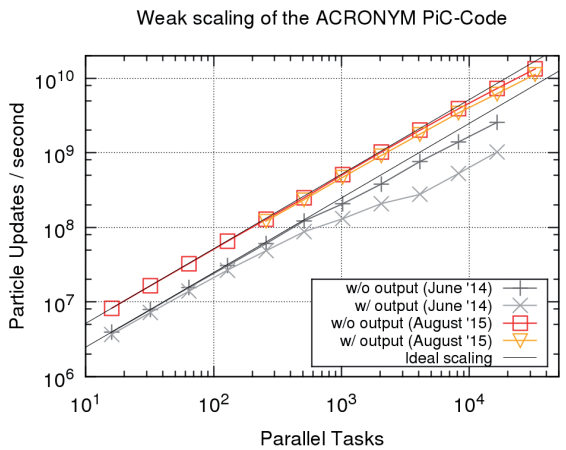

Figure 1: Scaling behavior of the *ACRONYM* PiC code on the SuperMUC system before (June '14) and after (August '15) major new output routines have been implemented. The number of particle updates per second – a measure for the performance of PiC codes – is plotted over the number of parallel tasks – or CPUs – used in the simulation.

Based on a previous project in which we investigated the scattering of energetic test particles off of a single plasma wave [3], we now want to model the interaction of fast particles and a turbulent spectrum of several waves, which presents a much more accurate picture of the processes in the solar wind. Contrary to many existing models which assume so-called Alfvénic turbulence, we focus on kinetic turbulence. While the former only considers waves which propagate with identical speeds, the latter includes dispersive waves which possess individual phase speeds for each wave, making kinetic turbulence a much more challenging problem which can hardly be tackled analytically.

Results and Methods

During the first months of the project we successfully applied for the AstroLab Support Call offered by the LRZ. AstroLab aimed at improving the performance of existing codes on the SuperMUC system. In the course of this project we were able to improve our output routines significantly, thus clearing one major bottleneck detected in earlier use of our *ACRONYM* code. Together with some other modifications of our code, the results from the AstroLab project helped to increase the scalability and overall performance of our simulations. Compared to test runs performed during the Extreme Scaling Workshop 2014, our updated code now runs twice as fast on the SuperMUC system and suffers far less from writing data to disk (see Fig. 1).

Besides the work on the code, a few test simulations have been carried out to investigate the parameter ranges and computational resources required for production runs concerning kinetic turbulence and particle

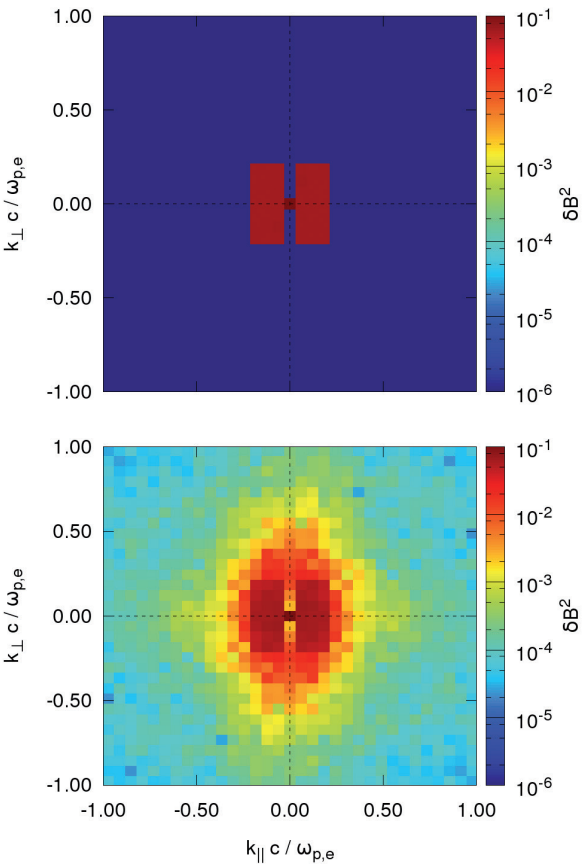

Figure 2: Two-dimensional power spectrum showing the magnetic field strength of waves with different wave numbers and traveling in different directions. The initial distribution of energy is plotted in the top image. During the simulation, a cascading spectrum of waves is established, as can be seen in the bottom image. The energy decreases with increasing wave number.

transport. In order to do so, the theoretical foundation underlying the initialization of our simulations has been expanded.

To create turbulence in a PiC simulation, two different mechanisms can be used: In both cases, a thermal, magnetized plasma is used as a starting point. Then, one can either initialize streaming particles which trigger plasma instabilities, thus creating plasma waves and eventually establishing a turbulent cascade. Or one can already initialize a few plasma waves, which then develop a cascade by wave-wave interaction. We chose the latter approach, since it is simpler – no additional population of streaming particles is needed – and faster – one does not need to wait for the instability to develop.

For first tests, a setup described by [4] is reproduced. In this setup, a two-dimensional simulation box containing thermal protons and electrons is initialized. Then, several plasma waves are excited by deploying the electromagnetic fields accordingly throughout the simulation box. To do this correctly, the dispersion relations, i.e. the relations between the wave's frequency and its wave length, as well as the polarizations of the waves, i.e. the directions of the electric and magnetic field vectors with respect to each other and the direction of the wave's propagation, have to be computed analytically during the initialization of the simulation.

After the initialization the simulation runs self-consistently, with waves propagating and interacting with each other. Theory predicts that waves with long wave lengths produce new waves with shorter wave length by their interaction. This produces a cascade of waves, starting with waves at large spatial scales and developing structures at smaller and smaller scales.

It is convenient to study this process in Fourier space, meaning that waves are not characterized by their wave length, but by its inverse – the so-called wave number k. Figure 2 shows the two-dimensional space of wave numbers, where the axes of the plot are given by the directions parallel and perpendicular to a static background magnetic field in the simulation. The two plots in Fig. 2 show data from the initialization (top) and a later point in time in the simulation (bottom), where the turbulent cascade has already been developed. It can be seen that the energy density decreases from small to large wave numbers, as expected from theory. The spectrum ends at some point due to damping mechanisms which dissipate the electromagnetic fields of the waves and transfer it to the particles, eventually leading to plasma heating.

The setup includes so-called whistler waves, which are right-handed dispersive waves. The mechanisms responsible for the dissipation of those waves are electron cyclotron and Landau damping at sufficiently large wave numbers. Damping affects mostly the thermal electrons in the background plasma, but an additional population of energetic electrons will be able to interact with waves at smaller wave numbers as well.

On-going Research / Outlook

Still in the first phase of the project, we have not investigated the scattering of energetic particles and turbulent waves, but we expect to find resonant interactions of fast electrons and turbulent whistler waves. Our next steps will be to further study the evolution of a turbulent cascade from a few initial waves in the simulation. It is especially interesting to test whether the dimensionality of the simulation box has an influence on the turbulent cascade and the evolution time scales.

After having obtained a better understanding of the influences of physical and numerical parameters on the spectrum of waves, we plan to include energetic test particles in the simulations and to study their transport in the turbulent plasma.

References and Links

Example:
[1] http://www.lrz.de/projekte/hlrb-projects/0000000000F43805.html
[2] P. Kilian, T. Burkart, F. Spanier, in: W. E. Nagel, D. B. Krner, M. M. Resch (Eds.), High Performance Computing in Science and Engineering '11, Springer, 2012, p. 5
[3] C. Schreiner, F. Spanier, Computer Physics Communications 185 (2014), p. 1981
[4] S. P. Gary, C. Smith, Journal of Geophysical Research 114 (2009), A12

1

Forming disk galaxies in magneto-hydrodynamical simulations of the Universe

Research Institution
Heidelberg Institute for Theoretical Studies (HITS)

Principal Investigator
Volker Springel

Researchers
Robert Grand, Federico Marinacci, Rüdiger Pakmor, and further collaborators

Project Partners
Massachusetts Institute of Technology, Kavli Institute for Astrophysics and Space Research

SuperMUC Project ID: pr85je (Gauss Large Scale project)

Introduction

One of the most important and interesting class of objects populating the Universe are so-called disk galaxies. They are comprised of stars, dust and interstellar gas, and this material is organized in an approximately disk-like shape that spins around its symmetry axis and is embedded in a dark matter halo. The most prominent example of this type of objects is our own galaxy, the Milky Way. For decades, the formation of disk galaxies has been a puzzle for cosmologists – numerical simulations invariably yielded far too massive galaxies with an overly large central bulge and only a very small disk component. In recent years, however, substantial progress has been made that much improved the physical realism of the simulation predictions.

The difficulties in modeling galaxy formation ab initio in numerical simulations stem from the fact that the problem is intrinsically complex, due to its multi-scale and multi-physics nature. To obtain realistic results, simulations ought to be able to model physical processes acting on very small scales – much smaller than the typical size of a galaxy – that nevertheless have a large impact on the global galaxy properties. Moreover, there is a large variety of physical processes that have just started to be investigated in theoretical studies but can still play a key role in galaxy evolution. This includes the dynamics of magnetic fields within the full cosmological context. The latter poses additional mathematical and numerical challenges compared to plain gas dynamics, which has typically been employed for computing galaxies so far. Our project pr85je on SuperMUC represents the first successful attempt to include a realistic modeling of magnetic fields in cosmological simulations of disk galaxy formation.

Results and methods

Within our project pr85je on SuperMUC, we have used the massively parallel simulation code AREPO, developed at HITS, to study galaxy formation within a fully cosmological setting. AREPO uses a comprehensive treatment of galaxy formation physics designed for large-scale cosmological simulations, and now also includes an accurate MHD scheme capable of following the intricate dynamics of magnetic fields as the Universe evolves. What distinguishes AREPO from other astrophysical codes is its moving and fully dynamics mesh. The code does not partition the simulated universe with a fixed grid but rather uses a movable and deformable mesh, which allows a very accurate processing of the vastly different size and mass scales occurring in cosmological simulations. The moving mesh approach combines the accuracy of traditional grid codes in modeling gas dynamics with the high adaptivity of particle-based codes, allowing it to follow the clustering of matter particularly well.

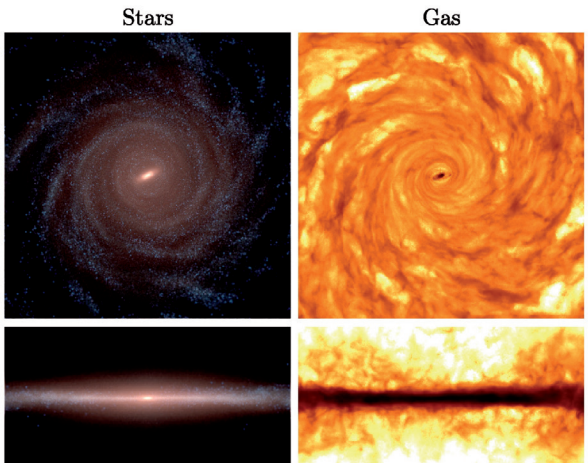

Figure 1: Example of one of the simulated Auriga galaxies (Au16) at one of the fine resolution levels in our set of magneto-hydrodynamical cosmological simulations. The figure shows stellar (left column) and gas (right column) projections at the end of the simulation, after ~14 billion years of evolution, in face-on and edge-on projections. Note how the gas, although spatially more extended, tracks the underlying stellar distribution well.

In the context of project pr85je, we succeeded in computing a set of fully cosmological simulations of 30 halos hosting realistic disk galaxies. In these calculations, collectively named the Auriga Project, all the simulated objects include magnetic fields, evolved self-consist-

ently from the initial conditions up to the final simulation time. In addition to magnetic fields, another distinctive trait of the project is the large number of simulated objects, allowing us to probe how the formation of disk galaxies is linked to the assembly history and the properties of the hosting halos. Also, a subsample of the full simulation set has been simulated at different resolutions without changing the galaxy formation model parameters, thus allowing us to investigate whether the galaxy properties are robust with respect to large (a factor of 64 in mass) variations of the numerical resolution. It is worth mentioning that at the finest resolution level, our runs improve by a factor of 8 in mass over our previous achievements and are among the highest resolved cosmological simulations of galaxy formation to date.

High resolution is a key aspect of our calculations because it allows us to resolve galactic sub-structure and the associated dynamical influence on the evolution of the stellar disks (which has previously been limited to isolated simulations with simplified initial conditions) in systems representative of a globally successful galaxy formation model in a fully cosmological setting. In particular, the galaxies develop well-formed bars and spiral arms, whose associated dynamical influence are thought to play an important role in shaping the structural properties of the disk.

In some of our recent work, we have analyzed the vertical structure of the stellar disks with the aim of finding the dynamical mechanisms responsible for scattering stars into vertically thickened distributions, which is currently a topic of major debate in the galactic astronomy community. From the Auriga simulations, we were able to establish that the stellar bar is the most ubiquitous vertical star scattering source, increasing the vertical velocity dispersion of star particles gradually over long timescales. Significant satellite interactions increase the dispersion dramatically on short timescales, but are less prevalent. Interestingly, we found that thick disks in the outer galaxy seem unlikely to originate in high velocity dispersion stars from the inner galaxy moved there through a process termed 'radial migration', which has been previously believed to be a likely contributor to thick disk formation. In addition, we found that in nearly all cases the vertical structure of the disk is dominated by the formation of new stars, which are born progressively closer to the mid plane of the disk over time, thereby creating a vertically thin disk. This process is in many cases dominant over the heating of preexisting stars.

On-going research

SuperMUC has been instrumental to carry out the Auriga simulations, which have extended what it is currently possible in terms of resolution and faithful modeling of the physical processes (for instance by the inclusion of magnetic fields) in cosmological simulations of the formation and evolution of galaxies like our own Milky Way. Still higher resolution is required to further improve our

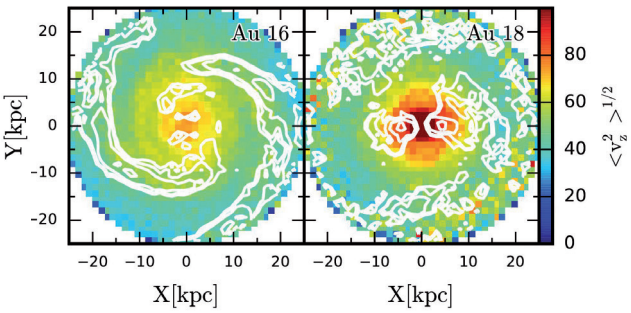

Figure 2: Face-on maps of the vertical velocity dispersion of stars for a galaxy with no bar (left) and a strong bar (right). Over-density contours are overlaid in white. Note that the velocity dispersion is much larger in the strong bar case than that of the galaxy with no bar (adapted from Grand et al. 2016).

theoretical understanding of galaxy formation and evolution, because only this will finally enable us to study many aspects of this complex problem with the required accuracy.

One aspect, the formation of spiral arms, is directly related to stellar dynamics within the disk, which can now be sampled with an appropriate number of star particles. Simulations of disk galaxies have long been known to produce transient, winding spiral arms that rotate at the same speed as the stars and have a lifetime roughly comparable to the dynamical time of the galaxy. Such properties, though in agreement with some observational studies of external galaxies, are at odds with the most widely accepted theory of spiral arms: spiral density wave theory, which predicts that spiral arms should be long-lived features that rotate rigidly around the disk and therefore never wind up. With their unprecedented resolution, our simulations are an ideal testbed for studying the formation of spiral arms in a variety of cosmological environments, ranging from quiescent to more active evolutionary histories.

Other areas in which we expect substantial progress from the analysis of the Auriga simulations are the study of the gas circulation between the disk and the halo, and the analysis of the properties of satellite galaxies. For the former, which is fundamental for understanding how disk galaxies can power star formation throughout a Hubble time, high resolution is essential to capture the intrinsic multiphase nature of these gas flows. For the latter, high resolution not only probes the properties of satellite galaxies to lower masses, but it will also allow us to investigate how the evolution of such objects is shaped by their interactions with the central galaxy and the surrounding environment.

References and Links

[1] V. Springel, MNRAS 411 (2010) 1525
[2] R. Pakmor, F. Marinacci, V. Springel, ApJL 783 (2014) L20
[3] F. Marinacci, R. Pakmor, V. Springel, MNRAS 437 (2014) 1750
[4] R. J. J. Grand, V. Springel, F. A. Gómez, F. Marinacci, R. Pakmor, D. R. Campbell, A. Jenkins, MNRAS in press (2016), http://arxiv.org/abs/1512.02219
[5] A. Monachesi, F. A. Gómez, R. J. J. Grand, G. Kauffmann, F. Marinacci, R. Pakmor, V. Springel, C. S. Frenk, MNRAS in press (2016), http://arxiv.org/abs/1512.03064

1

Mocking the Universe: Large Volume Simulations for galaxy surveys

Research Institution
Universidad Autónoma de Madrid

Principal Investigator
Gustavo Yepes

Researchers
S. Gottlöber[2], K. Riebe[2], A.Klypin[3], P. Behroozi[4], E. Jullo[5]

Project Partners
[2]Leibniz Astrophysical Institute Potsdam, [3]New Mexico State University,
[4]UC Berkeley, [5]LAM Marseille

SuperMUC Project ID: pr86bu (PRACE project), pr87yi

Introduction

More than 96% of the total energy content of the universe consists of an still unknown form of dark energy and dark matter. Understanding the physical nature of these two components is one of the great challenges for the (astro)physics of the 21st century. Both dark energy and dark matter are only detectable through their gravitational interaction. Together they give rise to the observed large scale structure in the universe. Thus understanding large scale structure formation is a prerequisite to understand dark energy and dark matter. Due to the non-linear nature of gravitational clustering computer simulations play a very important role in cosmology. Research in computational cosmology started four decades ago with the first N-body simulations having a few hundred or thousand particles. Nowadays, large simulations with billions of particles are rather easy to do. They are used to address numerous aspects of the evolution of fluctuations and formation of structures on the vast range of relevant scales. Meanwhile numerical simulations provide extraordinary accuracy for important statistics such as the mass function and the correlation function of objects, the biases between dark matter and luminous matter distribution. Thus large dark matter simulations are an essential tool to interpret the observational data. In the next years there will be a substantial number of observational projects that will be gathering photometric and redshift data from galaxies at different epochs in the universe and covering a large area of the sky (e.g BOSS, DES, KIDS, DESI, LSST, Euclid, eROSITA, etc). Galaxies and the matter distribution on larger scales will be mapped more accurately and, consequently, this will put strong limits on the cosmological model of the Universe. But, in order to derive cosmological constraints from these data, they must be compared with the theoretical predictions of the different cosmological scenarios. Due to the intrinsic non-linear nature of the physics behind the formation and evolution of galaxies and larger structures, the only reliable tool to perform this task is the numerical simulation of the gravitational evolution of fluctuations in computational boxes that are large enough to accommodate the sampled volumes of the different observational surveys. State of the Art N-body simulations are now capable of dealing trillions (10^{12}) of particles. Moreover, these large volume simulations would need to be created for different models of dark energy (i.e. Cosmological constant, dynamical dark energy with an evolving equation of state, dark energy coupled with dark matter, etc) and also in other alternative cosmological scenarios such as those with modified gravity, non-homogeneous or isotropic models, non-gaussian initial conditions or non standard particles such as sterile neutrinos. All these alternatives to the standard cosmological model would need to be properly modelled numerically in order to produce mock galaxy catalogues to the different observational surveys, so their predictions on clustering, abundance of galaxies, weak lensing maps, etc can be reliably compared with their observational counterparts. Therefore, this would require that the simulations will have enough mass resolution to properly resolve the dark matter objects which host the galaxies detected in observations.

Simulations

The PRACE pr86bu project was run in Supermuc during 2013 and 2014. It produced a set of N-body simulations in sufficiently large volumes to accommodate the volume sampled by the recently finished BOSS [5] redshift survey with enough resolution to be able to resolve all the dark matter halos which host the same galaxies as those detected in the survey. This imposed a severe constraint in the minimum number of particles that had to be used in the simulations. A total of more than 56 billion (3840^3) were used in the various runs performed at LRZ [2]. Since 2015 we have continued running simulations within the LRZ pr87yi project. We have completed our simulation dataset with two extra boxes, one larger and another one smaller than those made in the previous project. In Figure 1 we show the different box sizes that we have produced in these two project, all of them run with the same number of particles and using the cosmo-

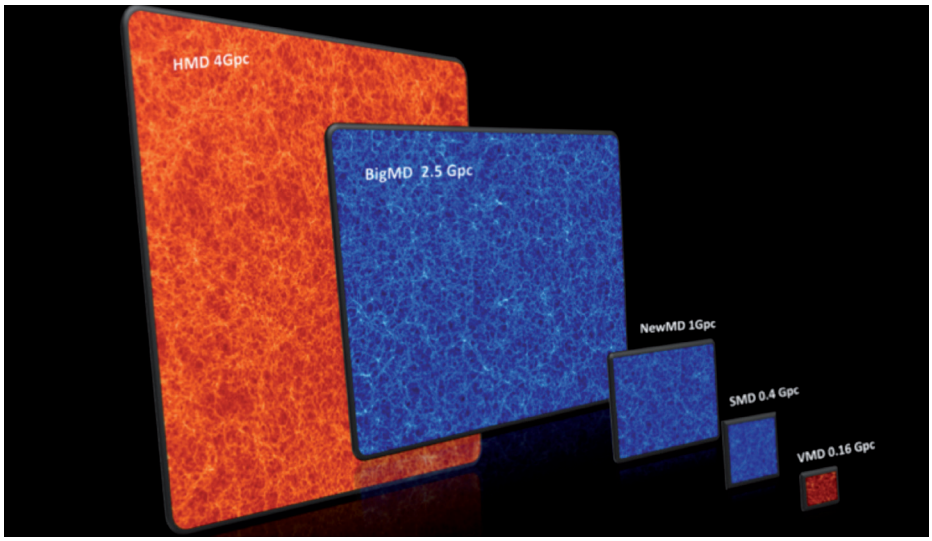

Figure 1: The density of a dark matter slice from the 5 different simulated boxes. The blue ones correspond to the PRACE pr86bu project. The two red boxes were computed within the LRZ pr87yi Project. Each box contains the same number of particles, 56.6 billion (3840^3).

logical parameters derived from ESA's Planck satellite. In total, we produced more than 3 PBYTES of raw data that have to be analysed in order to derive galaxy catalogues which mimic those from the BOSS survey. This is a very time consuming process due to the large amount of data that have to be processed. In total, we have generated more than 60 billion dark matter objects (halos), which are suppose to host galaxies, from all these simulations. In the end, our goal is to generate a distribution of object in redshift which closely resembles those from the observed galaxies. From these 3D datasets, we can then compare their clustering properties and derive conclusions about the nature and content of the dark matter and dark energy components of the universe.

Our simulations are also being used to produce weak gravitational lensing maps. That is, to simulate the bending of the light emitted from early galaxies due to the intervening dark matter distribution [3]. Observational projects like CHFTLens [4] are measuring the galaxy distortions due to this effect . A detailed comparison with the predictions of cosmological models, it is only possible though these large volume simulations (see Figure 2).

All the mock catalogues and other deliverables from our simulation dataset are publicly available through the COSMOSIM database [1]

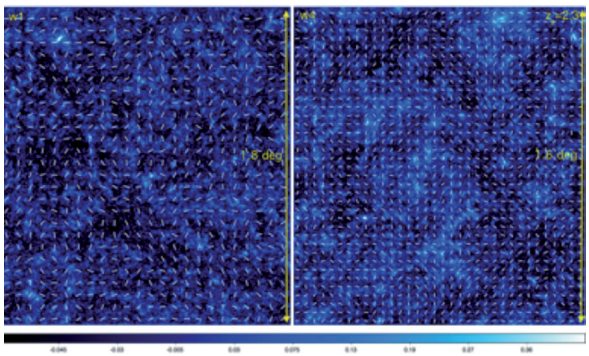

Figure 2: Maps of weak lensing shear fields for two equivalent areas of the VIPERS survey (see [3] for further details).

Future work

We plan to continue the analysis of the dataset of the simulations performed in the course of the PRACE pr-86bu and LRZ pr87yi projects. There is still a fair amount of work to do until we finish the production of halo catalogues and the derivation of light-cone samples. We are currently working on the production of halo merger trees to be used in semi-analytical models of galaxy formation in order to be able to assign galaxies to our dark matter objects. In this regard we will be able to make predictions for the clustering of galaxies of the upcoming galaxy redshift surveys (DESI [6], Euclid [7]) which are dealing with galaxies with special observational properties whose dark matter halo hosts are not yet well known. Our simulations will shed light on those dark objects.

The large dynamical range, both in mass and spatial resolution (4 orders of magnitude) covered by our simulations suite will also allow us to study other interesting problems. For instance, using the smallest computational volume, which has the highest mass resolution, we are going to study the epoch of formation of the first stars and galaxies which were responsible of reionizing the primordial gas of the Universe. Experiments like Square

Kilometer Array (SKA) are going to probe this early epoch and will give us observational information on the distribution of neutral hydrogen that can be directly compared with the simulation predictions.

References and links

[1] http://www.cosmosim.org
[2] Kypin et al 2016, MNRAS, 457, 4340
[3] Giocoli, Jullo, et al, 2015, arxiv:1511.08211
[4] http://www.cfhtlens.org/
[5] http://cosmology.lbl.gov/BOSS/
[6] http://desi.lbl.gov/
[7] http://sci.esa.int/euclid/

1

Unravelling the Mechanism of Supernova Explosions

Research Institution
Max Planck Institute for Astrophysics

Principal Investigator
Hans-Thomas Janka

Researchers
Tobias Melson, Alexander Summa, Bernhard Müller, Andreas Marek

Project Partners
Max Planck Computing and Data Facility (MPCDF)

SuperMUC Project ID: pr86la (PRACE project), pr85ja (Gauss Large Scale project), pr84ne (PRACE project), pr48ra (Gauss Large Scale project)

Introduction

Latest three-dimensional (3D) computer simulations are closing in on the solution of a decades-old problem: how do massive stars die in gigantic supernova explosions? Since the mid-1960s, astronomers thought that neutrinos, elementary particles that are radiated in huge numbers by the newly formed neutron star, could be the ones to energize the blast wave that disrupts the star. However, only now the power of modern supercomputers has made it possible to actually demonstrate the viability of this neutrino-driven mechanism.

Supernovae are among the brightest and most violent explosive events in the Universe. They are not only the birth sites of neutron stars and black holes; they also produce and disseminate heavy chemical elements up to iron and possibly even nuclear species heavier than iron, which could be forged during the explosion. Understanding the explosion mechanism of massive stars is therefore of fundamental importance to better define the role of supernovae in the cosmic cycle of matter.

Already in the 1960's it was speculated that neutrinos might be involved. Myriads of these high-energy elementary particles are radiated by the extremely hot, newly formed neutron star. If less than one percent of them gets absorbed in the matter behind the stalled shock, a healthy supernova explosion will be the consequence.

However, many aspects of the relevant physics in previous computer models were still too crude and too approximate to be realistic. In particular, nature has three spatial dimensions, but only very recently the increasing power of modern supercomputers has made it possible to perform supernova simulations without any artificial symmetry constraints. A new level of realism in such simulations is thus reached and brings us closer to the solution of a 50 year old problem.

The stellar collapse group at the Max Planck Institute for Astrophysics (MPA; [1]) plays a leading role in the worldwide race for such models. With all relevant physics included, in particular using a highly complex treatment of neutrino transport and interactions, such computations are at the very limit of what is currently feasible on the biggest available computers.

Methods

In this project the stellar core collapse, neutron star formation, and the onset of the supernova explosion are studied with the PROMETHEUS-VERTEX code for multi-dimensional hydrodynamical simulations including a highly sophisticated description of three-flavor neutrino transport and neutrino-matter interactions with full energy dependence. While the former is treated with an explicit, higher-order Godunov-type scheme, the latter is

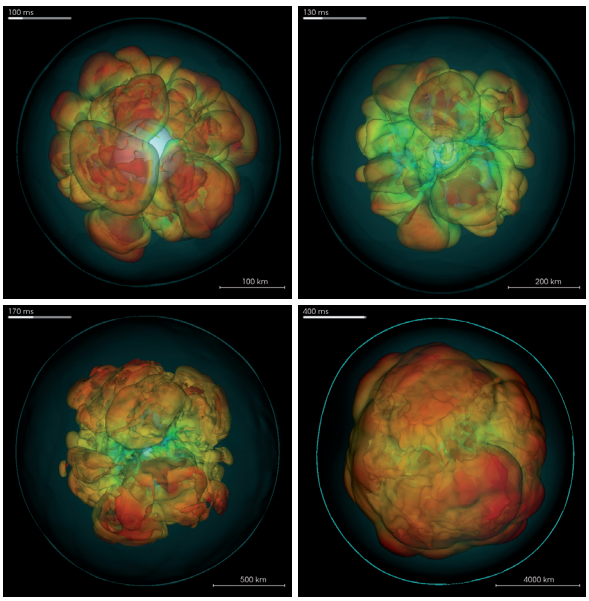

Figure 1: Development of the explosion in the interior of a 9.6 solar-mass star: within fractions of a second, the supernova shock inflates to many times its initial volume. The snapshots show that the explosion is far from symmetric and that convective buoyancy and turbulence play an important role. The colour coding indicates the speed of the ejected material. The thin bluish line shows the position of the shock front [2,3]. (Visualization: Elena Erastova and Markus Rampp, Max Planck Computing and Data Facility (MPCDF); copyright (2015) by American Astronomical Society)

solved by an implicit integrator of the neutrino energy and momentum equations, supplemented by a state-of-the-art set of neutrino-reaction kernels and a closure relation computed from a simplified model-Boltzmann equation. Consistent with the basically spherical geometry, the equations are discretized on a polar coordinate grid, and the computational efficiency is enhanced by the use of an axis-free Yin-Yang implementation and a time- and space-variable radial grid. The thermodynamics and changing chemical composition of the stellar medium are determined by high-dimensional equation-of-state tables and nuclear burning at non-equilibrium conditions.

Employing mixed MPI-OpenMP parallelization, our ray-by-ray-plus approximation of multidimensional transport allows for essentially linear scaling up to tested processor-core numbers of more than 130,000. The model simulations reported here were performed on about 16,000 cores in parallel (achieving 10–15% of the peak performance), which is the largest share of SuperMUC that the MPA team was granted access to, but constrains the models to currently two degrees angular resolution. Nevertheless, one full supernova run, conducted over an evolution time of typically half a second, consumes up to 50 million core hours and takes more than 1/2 year of project time to be completed.

Results

The enormous effort has payed off! The MPA team has recently been able to report a first successful 3D explosion for a 9.6 solar-mass star (Figure 1; [2,3]) and has now also obtained a 3D explosion of a 20 solar-mass progenitor (Figure 2; [4,5]). Based on the presently most advanced description of the neutrino physics in collapsing stellar cores worldwide, these results are a true milestone in supernova modeling. They confirm the viability of the neutrino-heating mechanism in principle, applying our currently best knowledge of all processes that play a role in the center of dying stars, whose extreme conditions in temperature and density are hardly accessible by laboratory experiments on Earth. Since not all aspects of the complex neutrino reactions in the newly formed neutron star are finally understood, the 3D models demonstrate that within existing uncertainties neutrinos can indeed transfer enough energy to revive the stalled shock. As known from previous models in two dimensions, violent non-radial fluid flows must provide crucial support to relaunch the blast wave and will function as seeds of the later, large-scale asymmetries that are observed in supernova explosions (Figures 1, 2).

On-going Research and Outlook

Further work on the theoretical models is necessary. So far the successful 3D simulations could only be done with rather coarse resolution, because bigger numbers of computing cores would be needed to perform more refined supernova calculations. Moreover, a wider range of stellar masses must be investigated, varying the initial conditions in the pre-collapse cores and testing still

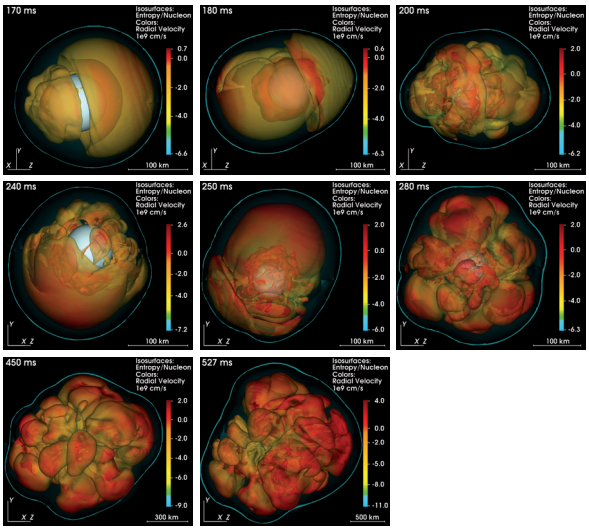

Figure 2: Sequence of volume-rendered images showing the violent non-spherical mass motions that drive the evolution of the collapsing 20 solar-mass star towards the onset of a neutrino-powered explosion. The whitish central sphere indicates the newly formed neutron star, the enveloping bluish surface marks the supernova shock [4,5]. (Visualization: Elena Erastova and Markus Rampp, Max Planck Computing and Data Facility (MPCDF); copyright (2015) by American Astronomical Society)

uncertain aspects of physics like the amount of anglar momentum in the collapsing star and different models of the incompletely understood equation of state of hot neutron-star matter. A final confirmation of our theoretical picture of the explosion mechanism and the role of neutrinos, however, can only come from observations. On the one hand this demands a closer link of the explosion models to observable supernova properties, requesting simulations over longer evolution times, on the other hand much hope rests on a next supernova that will occur in our Milky Way galaxy. Such a nearby event will flood the Earth with 10^{30} neutrinos, of which several thousand to tens of thousands will be captured in huge underground experiments like Super-Kamiokande in Japan and IceCube at the South Pole. Neutrinos (besides gravitational waves) will thus serve as unique messengers: since they escape from the center of the supernova they will bring us information directly from the very heart of the explosion.

Acknowledgments
This project was partly funded by the European Research Council through grant ERC-AdG No. 341157-COCO2CASA. Computing time was kindly provided by PRACE on SuperMUC (GCS@LRZ, Germany), Curie TN (GENCI@CEA, France), and MareNostrum (BSC, Spain), and by the Gauss Centre for Supercomputing on SuperMUC (GCS@LRZ, Germany).

References and Links

[1] http://mpa.iwww.mpg.de/220337/Modeling-Stellar-Collapse-and-Explosion
[2] http://www.mpa-garching.mpg.de/112431/hl201504
[3] T. Melson et al., ApJL 801 (2015) L24
[4] http://www.mpa-garching.mpg.de/208528/hl201508
[5] T. Melson et al., ApJL 808 (2015) L42

Dissipating the Energy of Space Weather Events

Research Institution
KU Leuven

Principal Investigator
Giovanni Lapenta

Researchers
Vyacheslav Olshevsky, Jan Deca, Jorge Amaya, Maria Elena Innocenti, Emanuele Cazzola, Lorenzo Siddi, and Emmanuel Chané

Project Partners
–

SuperMUC Project ID: pr87di (PRACE project)

Introduction

Space is permeated by magnetic fields. They serve different purposes: solar magnetic field provides storage and transport for its enormous internal energy. Magnetic fields in the solar wind guide the propagation of charged particles and solar ejecta. Earth's magnetic field shields us from the deadly impact of high-energy particles from space. Often, during interaction of different magnetic fields, a tremendous amount of energy is released. The most spectacular manifestations of these events are polar lights. However, practical value of such space weather [1] events are even broader, as energetic particles and magnetic field perturbations endanger space vehicles (especially those flying outside the magnetosphere, e.g., to Moon or Mars), GPS, aircrafts and even ground-based facilities such as power grids.

Plasma-astrophysics group at KU Leuven [2] studies the fundamental properties of solar and interplanetary magnetic fields and their interaction. At our disposal are present-day massively parallel codes which we use to model magnetized plasma behavior. Although space plasma is very rarified (only a few particles per cubic cm), the regions of interest are so big that individual description of each plasma particle is far beyond the scope of any computer in the coming decades. Therefore different simplifying approximations such as single-fluid and multi-fluid magnetohydrodynamics (MHD) are used in the numerical models. Availability of the computing power in the last decade allowed practical applications of the more advanced modeling techniques, namely the kinetic Particle-in-Cell (PIC) [3]. These models combine the advantages of the kinetic (individual) description of plasma particles with the 'traditional' evolution of electromagnetic field on a regular (or irregular) computational grid. Of course, PIC is much more demanding than any conventional fluid model (it is impossible to run a reliable simulation on, say, less than 1000 cores), but it rises our physical insight to the qualitatively new level. Physical effects we found using the kinetic PIC code iPic3D [4] would have never been revealed by any fluid model of space plasma.

Results and Methods

Our project aims at breaking the long-standing paradigm in the physics of space plasma, the so-called magnetic reconnection. This concept predicts that primary interaction of magnetic fields happens where the opposite-polarity fields approach each other. These locations often referred to as X-points and are present, for instance

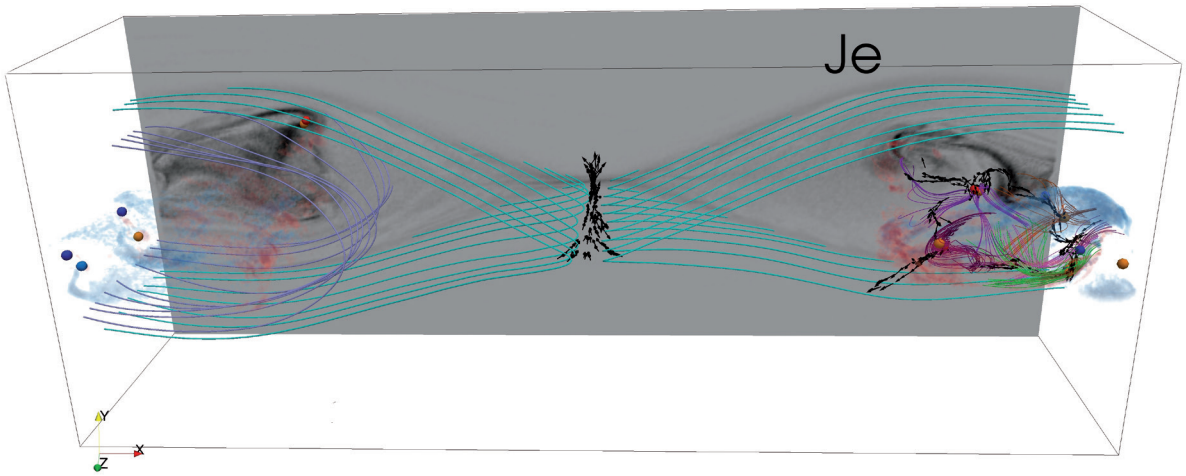

Figure 1: Breakage the classical concept: magnetic reconnection is believed to happen at the central X-points where black arrows depict electron currents. However massive energy dissipation (blue and red fog) happens far from it, in the flux ropes of reconnection exhaust.

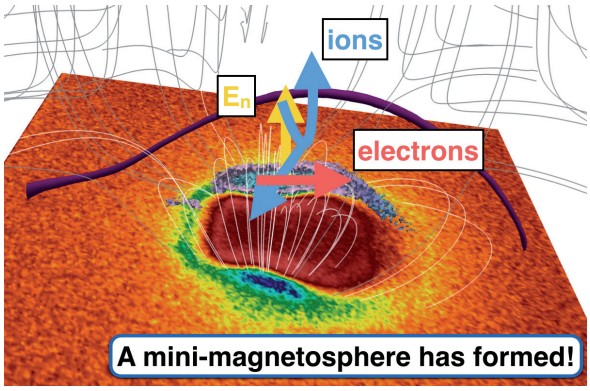

Figure 2: A mini-magnetosphere of a lunar magnetic anomaly protects its surface from energetic particles of the solar wind: lunar base should be underneath its shield!

in the Earth's magnetotail, or in magnetopause. Simulations in large three-dimensional boxes representing the magnetotail topology pointed us to the other regions where intensive energy dissipation is happening, the so-called secondary reconnection sites [5]. Further simulations have confirmed our initial guess: no magnetic reconnection happens near Lunar Magnetic Anomalies (Figure 2) despite the presence of X-points. The dedicated simulations of turbulent plasma have revealed the importance of magnetic flux ropes and secondary instabilities (Figure 3). **Magnetic energy dissipation in space plasma is not associated with specific topological features such as X-points, it takes place everywhere, and is closely tied with interacting magnetic flux ropes.**

As noted above, PIC simulations are demanding to computational power because they are using large number of particles. For instance, a simulation of magnetic flux ropes (Figure 3) used 65536 million particles! As an upside, PIC codes scale very well, and we were able to use up to 64000 cores on SuperMUC. Regular runs employed 8-16 thousand cores, and ran for a few days (with restarts). Typical simulation needed five restarts (5 days) and 2 million CPU hours. In total all work packages of our SuperMUC allocations have consumed 25 million core hours.

Great challenge related to PIC simulations is data handling. iPic3D uses parallel hdf5 output to dump the data of each process. Particle and grid quantities are saved separately. In this project we mainly addressed field quantities (which are much less data). Still, we are talking about terabytes per simulation. Parallel post-processing routines allow conversion to the more convenient formats such as VTK, directly on the cluster. Only the necessary data, already converted and compressed, is transferred from the HPC facility. We plan to use the improved output capabilities in our new, more ambitious project devoted to acceleration of plasma particles.

On-going Research / Outlook

Modeling plasmas with PIC methods is at the frontier of modern physics. Correspondingly, it demands the most advanced supercomputing facilities. SuperMUC's thin nodes are ideal for PIC simulations because of large amount of RAM per node. Our study was just the first insight that brought out a new concept of energy transformation in the turbulent space plasmas. It opens the way for larger, realistic simulations and, finally the comparison with real observations of the new and existing space missions such as Cluster, MMS and THOR.

Memory-demanding, PIC pays back with almost ideal scalability, and is a perfect application for the new generation of shared-memory hybrid architectures. Our code is being actively developed and adjusted for use on such clusters. We are looking forward for real exascale simulations on the "SuperMUC Next Generation".

References and Links

[1] http://www.spaceweather.com
[2] http://wis.kuleuven.be/CmPA
[3] https://www.particleincell.com
[4] Markidis, Lapenta & Rizwan-udding, MATH COMPUT SIMULAT 80 (2010)
[5] Lapenta et al., NAT PHYS 11 (2015)

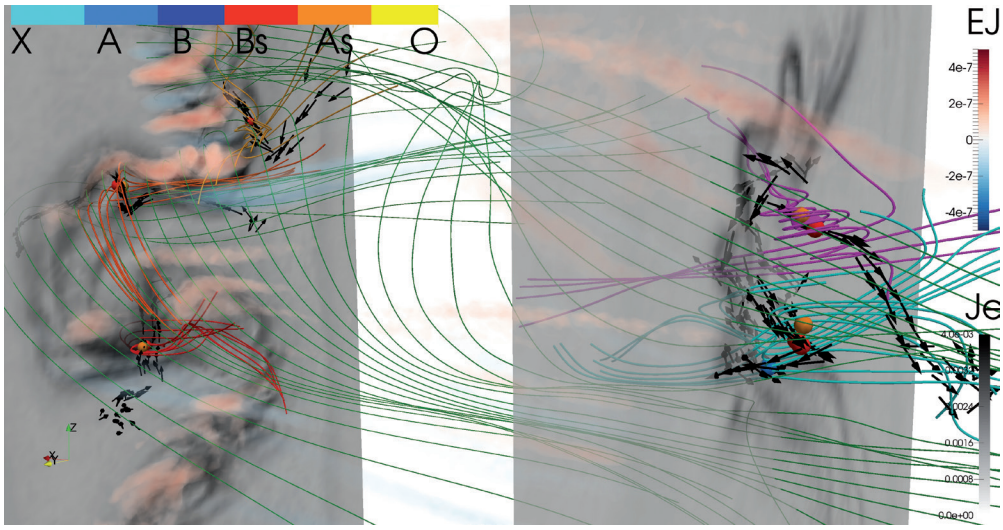

Figure 3: Proposed mechanism of magnetic energy dissipation is linked to plasma instabilities produced in the kinked magnetic flux ropes embedding current channels (black arrows) and magnetic nulls (color spheres).

Evolution of the corona above an active region after sunspots formed

1

Research Institution
Max Planck Institute for Solar system research
Principal Investigator
Hardi Peter
Researchers
Sven Bingert, Feng Chen
Project Partners
–

SuperMUC Project ID: pr87la (PRACE project)

Introduction

The corona is the outer atmosphere of the Sun, which can be brilliantly seen during total solar eclipses. Since 1940s, it has been clear that the corona is composed of tenuous and million K hot plasma, which is more than 100 times hotter than the solar surface. What heats the corona is a great mastery in astrophysics and plasma physics. Modern observations of the corona in X-rays and extreme ultraviolet (EUV) show astonishing loop-like structures (coronal loops) that arch over the solar surface with their apex reaching several tens Mm. Coronal loops outline the magnetic field lines, because the magnetic energy is so overwhelming in the corona, that the plasma is confined by the magnetic field lines.

Theoretical models suggest that the magnetic field plays a key role in heating the corona. In general, the granular motions at the surface of the Sun can tweak and tangle the magnetic field lines, which induces currents in the coronal magnetic field. The dissipation of the currents (i.e. magnetic energy) is supposed to be sufficient to heat the coronal plasma to over 1 MK. To test these theoretical ideas in a "realistic" setup that can properly account for the complexity of the real corona, a numerical model must solve a three dimensional magnetohydrodynamical (MHD) problem with the processes (e.g. energy loss through optically thin radiation and highly anisotropic heat conduction) that are substantially important for the thermal dynamics of the coronal plasma. Thus, realistic models for the corona above an active region are very computationally demanding.

With the help of the computation resource of SuperMUC, the simulation in our previous project (pr86bi) became the first realistic coronal model that allows us to study how the hot corona is formed during the emergence of magnetic flux from the solar interior to the surface[1]. The simulation in the present project is a continuation of the successful model in the previous project. We evolved a high resolution run in the previous project for more than half an hour solar time, which gives us a rich mine of data to study the evolution of the corona after the magnetic flux emerged.

Results and Methods

Model setup
To simulate the evolution of the corona in response to the energy input from the solar surface, our coronal simulation is coupled to a numerical experiment that simulated the formation of a pair of sunspots though the emergence of a magnetic flux tube from the solar interior to the surface. We use several consecutive horizontal layers at the surface of the flux emergence simulation as a time-dependent bottom boundary of the coronal simulation. Thus the coronal simulation is driven by the flux emergence at the solar surface.

The coronal model uses a 1024x512x256 grid to resolve a 147x74x50 Mm^3 domain. The horizontal grid spacing of 144 km is sufficient to resolve the granular motions and small-scale magnetic features at the solar surface, which contribute a substantial part in the energy input to the corona. The vertical grid is stretched so that the grid spacing is 32 km near the bottom to match that of

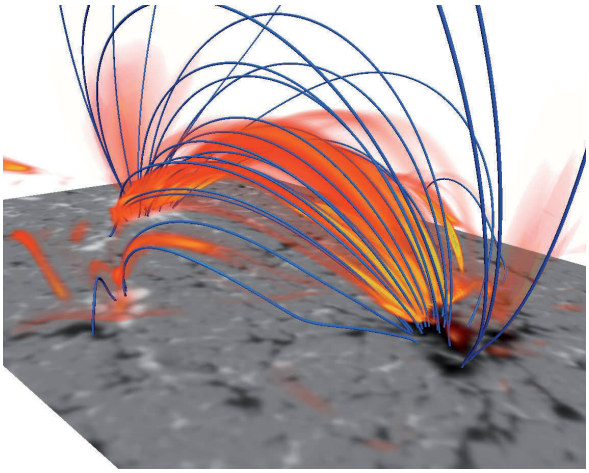

Figure 1: An overview of the model corona. The map in gray scale at the bottom is the vertical magnetic field at the solar surface, where the large patches of white and black are the two sunspots. Blue lines show the magnetic field lines anchored in the sunspots. Reddish structures represent the synthesized EUV emission from million K hot plasma. This figure is adapted from Chen (2015) [3].

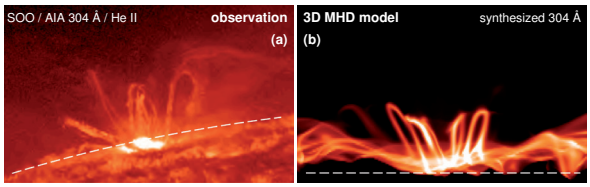

Figure 2: Loop structures arching over the solar surface. Left: An active region at the solar limb observed by SDO / AIA in the 304 channel, which is dominated by emission from the He II spectral line. It mainly samples the plasma of about 5x105 K. The dashed line indicates the limb of the solar disk. Right: The emission synthesized from the model according to the response function of the AIA 304 channel. The point of view is equivalent to the observation in the left panel. This figure is taken from Chen et al. (2015) [4].

the flux emergence simulation, 192 km in the majority of the coronal part, and gradually increases near the top boundary, which is not the region of interest.

Numerical methods

The numerical experiment is done using the Pencil-code[3]. The spatial derivatives are evaluated by a 6th order central difference scheme, and the time integration is done by a 3rd order Runge-Kutta scheme. In particular, we apply a super-time-stepping scheme for the heat conduction terms in an operator-splitting manner to improve the speed of the calculation, since the heat conduction is the major term limiting the time step. We run the simulation with 4096 cores. In combination with the previous project, we consumed about 20 million CPU-hours to evolve the model for about 2 hours solar time. This generates more than 3 TB of data.

Reproducing a real corona

The Atmospheric Imaging Assembly (AIA) on the Solar Dynamics Observatory (SDO) observes the solar coronal in several EUV channels. Each channel samples the EUV emission from the plasma within a particular temperature range. We use the density and temperature in our simulation to synthesize EUV emission according to the temperature response functions of the AIA channels, and compare the synthesized emission with real AIA observations.

Coronal loops become visible in the synthesized EUV images after the sunspots start to form through the coalescence of the small scale magnetic flux at the solar surface. The coronal loops are inline with the magnetic field lines as in the real corona, and their foot points are co-spatial with the enhanced upward magnetic energy flux, which is mostly found at the edge of the forming sunspots (Fig. 1). This is consistent with the scenario we found in our previous simulations. However, much more fine coronal structures are formed in the new simulation, because with a high resolution we are able to account for the small-scale magnetic features that are important for coronal structuring.

The synthesized AIA observations give a count rate of more than 1000 DN/pixel/s in a coronal loop [1], which is consistent with real AIA observations of active regions in the similar size. Given that the EUV emission of the corona is a convolution of its density and temperature, the consistency in the count rate implies that the model is able to reproduce a density and temperature similar to the real corona. This is also confirmed by comparing the density (10^8 to 10^9 cm^{-3}) and temperature (a few MK) in the model corona with those deduced from observations.

The thermal properties of the corona are closely related to the heat input. Therefore, the result above clearly suggests that the model produces a correct amount of heating in the corona. Furthermore, the synthesized observation also shows the coronal structures that are astonishingly similar to the real Sun, which are the numerous thin and elongated loops arching over the solar surface, as shown in Fig 2. The structuring of the corona essentially reflects the inhomogeneity of the heating in the corona, because only the magnetic field lines that have sufficient heat input can become visible coronal loops.

A successful model for the coronal heating must provide not only the correct amount of heating to maintain the temperature of the corona, but also the spatial distribution of the heating that gives rise to the observed coronal structures. Previous models have mainly focused on the first point, while our model sheds new light on coronal heating problem. To the best of our knowledge, it is the first model that self-consistently produces both the right amount of heat input to heat the coronal plasma to over 1 MK, and the coronal structuring that are highly consistent with the real corona. All together, this strongly implies the heating in the model is a good representative of the heating in the real corona.

On-going Research / Outlook

The analysis of the data has focused on the basic thermal dynamics of the coronal loops [1], the relation of EUV loops to magnetic field lines [4], and the oscillation in coronal loops [5]. These lead to several publications in leading journals in astrophysics, and the dataset still has a great potential for further studies.

In the following, we will try to reveal how the heating results in some fundamental characteristics of the observed EUV structures, such as the width and cross-section profile of a coronal loop. We will also put efforts on comprehensive comparisons between observations and the model to find out where the model may fail to reproduce the observational properties. The comparison is very helpful to explore the missing factors in present models, and leads us to more realistic models and better understanding of the solar corona.

References and Links

[1] F. Chen, H. Peter, S. Bingert, and M. C. M. Cheung, 2014 A&A, 564, A12
[2] https://github.com/pencil-code/pencil-code
[3] F. Chen, PhD thesis, 2015, Coronal dynamics driven by magnetic flux emergence, University of Goettingen
[4] F. Chen, H. Peter, S. Bingert, and M. C. M. Cheung, 2015 Nature Physics, 11, 492
[5] F. Chen and H. Peter, 2015 A&A, 581, A137

Aurora – Simulating Cosmic Reionization

Research Institution
Max Planck Institute for Astrophysics, Garching, Germany

Principal Investigator
Andreas H. Pawlik

Researchers
C. Dalla Vecchia, M. Jeon, B. Oppenheimer, A. Rahmati, A. Richings, J. Rosdahl, J. Schaye

Project Partners
Leiden University, University of Texas at Austin, IAC Tenerife

SuperMUC Project ID: pr87ne (PRACE project), pr83le (Gauss Large Scale project)

Introduction

The ionizing radiation emitted by the first stars and galaxies photoionized and photoheated the cold and neutral cosmic gas that filled space shortly after the birth of the universe in the Big Bang. This started the epoch of reionization, which took place in the first billion years, and which left galaxies embedded in a warm and ionized intergalactic medium (Figure 1). A new generation of telescopes such as the Low Frequency Array and the space-based James Webb Space Telescope is currently underway to help scientists unravel the astrophysics of reionization, creating exciting opportunities to test and improve the understanding of the universe. In this report we summarize how we used two multi-million compute hours allocations by the Gauss Centre for Supercomputing and by PRACE, the Partnership for Supercomputing in Europe, both on SuperMUC at the Leibniz Supercomputing Centre (LRZ), to carry out Aurora, a new set of radiation-hydrodynamical simulations of galaxy formation during reionization.

Understanding reionization is important because of its strong coupling with structure formation caused by ionizing radiative feedback, which strongly impacts the assembly of galaxies, including our own galaxy, the Milky Way. The absorption of ionizing photons by neutral gas photoionizes the gas and raises its temperature to a few ten thousand degrees Kelvin. The associated increase in gas pressure increases the Jeans scale, i.e., the scale below which pressure prevents the collapse of gas into gravitationally bound objects, thus smoothing out density fluctuations in the cosmic gas. Because the rate at which the ionized gas recombines to return to its neutral state is proportional to the clumpiness of the gas between galaxies, photoionization heating makes it easier to keep the gas ionized and therefore facilitates reionization (positive feedback on reionization). However, the increase in the gas pressure also boils out gas from small galaxies and impedes the accretion of fresh gas from the intergalactic medium onto them. This reduces the fuel for star formation, which reduces the rate at which galaxies emit ionizing photons and therefore impedes reionization (negative feedback on reionization).

The net effect of photoheating on reionization is difficult to predict short of treating it with full cosmological simulations that explicitly follow the relevant physical processes, primarily the transport of ionizing photons and their interaction with the cosmic gas. While radiative transfer simulations of reionization have become feasible in the last decade thanks to advances in numerical techniques and the advent of high-performance supercomputing facilities, such as SuperMUC at LRZ, simulating reionization remains a computationally highly demanding task.

The prime challenge is the need to transport ionizing photons in large cosmological volumes with structure on a large range of scales and with many stars and galaxies emitting radiation. Aurora addresses this challenge by using a novel spatially adaptive radiation transport technique that solves multi-scale problems efficiently and has a computational cost independent of the number of radiation sources. A major improvement on many previous reionization simulations is that the transport of radiation is carried out fully coupled to the gas dynamics. In combination with the very high resolution exclusively enabled by our spatially adaptive radiation transport technique and the generous allocations of compute time awarded to us on SuperMUC, this lets us investigate the impact of photoionization heating on the formation of structure in exquisite detail. In addition to ionizing radiation, Aurora tracks the explosion of stars as supernovae and the enrichment of the universe with metals synthesized in the stars. The Aurora simulations thus make detailed predictions of reionization by galaxies that will be tested in comparisons against observations by the upcoming telescopes.

Methods and Results

The Aurora simulations make use of the novel, accurate and well-tested radiation transport technique TRAPHIC [1] developed by us. TRAPHIC is coupled to the galaxy formation code GADGET [2]. Our simulations thus capture the physics both of reionization and galaxy formation, as well as their tight interplay. In the following, we provide a brief description of the main simulation techniques used.

GADGET is a massively MPI parallel ANSI-C code which simulates the dynamics of the two matter components of the universe, the dark matter and the baryons (gas and stars). The well-tested scalable standard version of GADGET computes N-body gravitational forces using a fast Particle-Mesh/Tree solver and tracks the dynamics of the gas us-

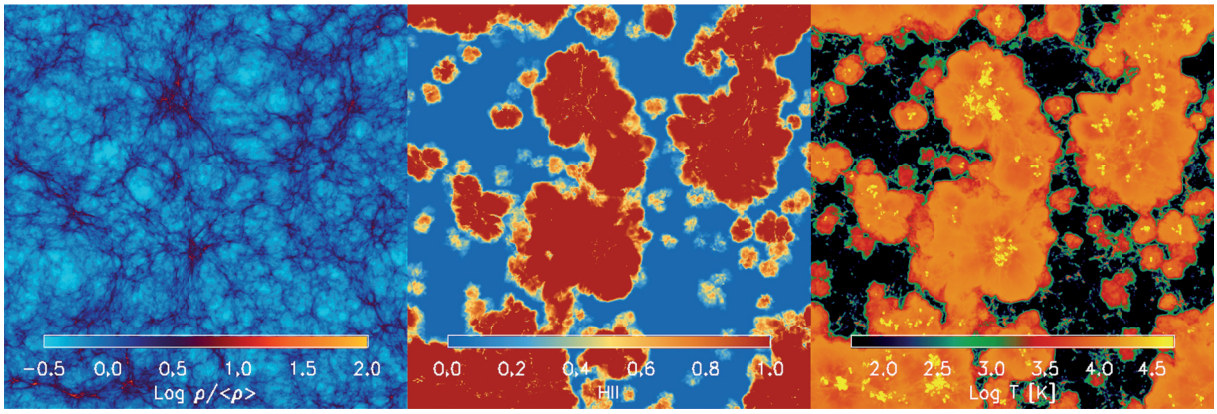

Figure 1: Snapshot of the Aurora reionization reference simulation carried out on SuperMUC at LRZ. From left to right, panels show the gas density (blue is underdense, red/yellow is overdense), ionized fraction (blue is neutral, red is ionized), and temperature (black is cold, orange/yellow is warm/hot) in a thin central slice through the simulated cubical volume of linear size 35 Mpc, just a few hundred million years after the Big Bang. Galaxies form in the densest regions and emit radiation that photoionizes and photoheats the gas surrounding them. Reionization is complete once the individual ionized regions overlap, which happens about one billion years after the Big Bang. The localized bright yellow regions of very high temperatures are caused by the explosions of stars as supernovae.

ing the Smoothed Particle Hydrodynamics (SPH) technique. SPH is a spatially adaptive, Lagrangian simulation technique that discretizes the mass in the simulated volume using particles and follows the evolution of these particles subject to the Euler equations of fluid flow. Our customized version of GADGET enables us to follow the key astrophysical processes relevant for the formation of galaxies in the early universe: the collapse of dark matter in halos hosting the first galaxies, the accretion of gas by these halos, leading to the formation of stars, as well as the explosion of stars in supernovae and their death as black holes, and the chemical enrichment of the intergalactic medium. Crucially, it also contains an MPI ANSI-C implementation of our multiscale and scalable radiation transport method TRAPHIC.

TRAPHIC solves the time-dependent radiative transfer equation by tracing photon packets emitted by the ionizing sources at the speed of light and in a photon-conserving manner through the simulation box. The photon packets are transported directly on the spatially adaptive, unstructured grid defined by the SPH particles, which lets us use the full dynamic range of the hydrodynamical simulation. For comparison, in most previous reionization simulations the radiation transport was carried out on an external and often uniform grid superimposed on the cosmological simulation. The use of an external (uniform) grid typically implies a substantial reduction in dynamic range in comparison with the underlying spatially adaptive gas simulation and impedes the coupling of the radiation and the dynamics of the gas. In fact, because of the large computational expense, the gas dynamics has often been ignored altogether even in some of the most sophisticated works, assuming it traces that of the dark matter, thus preventing an accurate treatment of the feedback from photoionization heating.

The Aurora simulations are designed to capture small-scale structure in the ionized gas as well as the local radiative impact of stars in the first galaxies in cosmological, representative volumes. Our reference simulation resolves all galaxies above the dwarf scale, the main drivers of reionization, in a box of size 35 comoving Mpc. This is accompanied by simulations in larger and smaller boxes and at higher and lower resolution to investigate numerical convergence. In other simulations we varied parameters to gain further insight in our results by comparison with the reference simulation. E.g., in some simulations we turned off photoheating to isolate the impact of this physical process. All simulations were run on the Thin Nodes partition of the SuperMUC supercomputer at LRZ. In total, about 2x30 million compute hours, awarded to us by the Gauss Centre for Supercomputing and by PRACE, were used. A typical simulation was run on 4096 cores over the course of several months and generated more than 10 Tb of data.

The analysis of the simulations is still ongoing. Exciting initial results from Aurora include that the coupling of galaxy formation and reionization by photoheating is very strong, and that an accurate treatment of feedback from photoheating and supernovae is key to reproducing observable properties of the early universe [3]. The simulations provide a rich source of information both for scientists studying the first galaxies and their evolution and for scientists studying the reionization of the universe.

On-going Research / Outlook

Efficient access to SuperMUC, one of the largest existing supercomputers, has been critical to the success of the Aurora project. Future upgrades of SuperMUC will enable still larger and more comprehensive simulations than presented here. Executing these simulations will require substantial algorithmical changes to our simulation codes, both to incorporate additional physical processes and to take advantage of the new parallel computing architectures, such as MIC processors and GPUs. We look forward to actively engage in these new directions of computational research.

References and Links

[1] Andreas H. Pawlik, Joop Schaye. 2008. TRAPHIC – Radiative Transfer for Smoothed Particle Hydrodynamics. MON.NOT.R.ASTRON.SOC. 381 (2008), 651.

[2] Volker Springel. 2005. The cosmological simulation code GADGET-2. MON.NOT.R.ASTRON.SOC. 364 (2005), 1105.

[3] Andreas H. Pawlik, Joop Schaye, Claudio Dalla Vecchia. 2015. Spatially adaptive, radiation-hydro simulations of galaxy formation during reionization. MON.NOT.R.ASTRON.SOC. 451 (2015), 1586.

General Relativistic Simulations of Binary Neutron Star Mergers

1

Research Institution
University of Trento

Principal Investigator
Bruno Giacomazzo

Researchers
L. Baiotti, R. Ciolfi, W. Kastaun, T. Kawamura, R. Perna, D. Siegel

Project Partners
Osaka University, Columbia University, Stony Brook University

SuperMUC Project ID: pr94bo (PRACE project)

Introduction

Neutron stars are the remnants of supernova explosions (the spectacular deaths of massive stars) and the most dense objects in the universe besides black holes. A typical neutron star concentrates more than the mass of the Sun within a radius of only around 10 km. Because of their extreme gravity a proper description of neutron stars requires Einstein's theory of General Relativity. Investigating neutron star properties can shed light on the behavior of matter at very high densities, which is not yet understood well by nuclear physics.

Two neutron stars can also bind together in a binary system, and orbit around each other for millions of years with a smaller and smaller separation. Eventually, the two merge together in an instant (just a few milliseconds) resulting either in a black hole or in a rapidly rotating super-heavy neutron star, which can still collapse to a black hole later on. Binary neutron star mergers are among the most violent astrophysical events and can be observed in several ways. First, they emit gravitational waves, perturbations in the geometry of spacetime that propagate at the speed of light and manifest themselves, e.g., as tiny periodic changes in the distance between free-floating objects. Gravitational wave emission is responsible for the decrease in orbital separation and becomes particularly strong towards the end of the long inspiral phase and the time of merger. Several modern laser interferometers such as LIGO and Virgo are searching for those signals, and have reached a sensitivity that might allow to detect them from binary neutron star mergers occurring more than 100 million light-years away.

Moreover, mergers are the most likely scenario to explain the so-called short gamma ray bursts, which are powerful outburst of gamma rays originating in distant galaxies. According to the leading model, when the merger results in a black hole surrounded by a hot and massive accretion disk endowed with a strong magnetic field, a powerful energy outflow can be generated, resulting in a burst of gamma-rays. Finally, the matter ejected during such mergers offers the extreme conditions necessary to synthesize elements heavier than iron and, together with supernova explosions, seems to be the only explanation for the abundance of such elements in the universe.

In this project, we perform computer simulations of binary neutron star mergers, taking into account general relativity, magnetic fields, and different models for the description of very dense matter. We consider different masses and mass ratios of the neutron stars, as well as different initial configurations of their magnetic fields, and we predict gravitational wave emission and mass ejection. In addition, we investigate the conditions for the emission of short gamma-ray bursts.

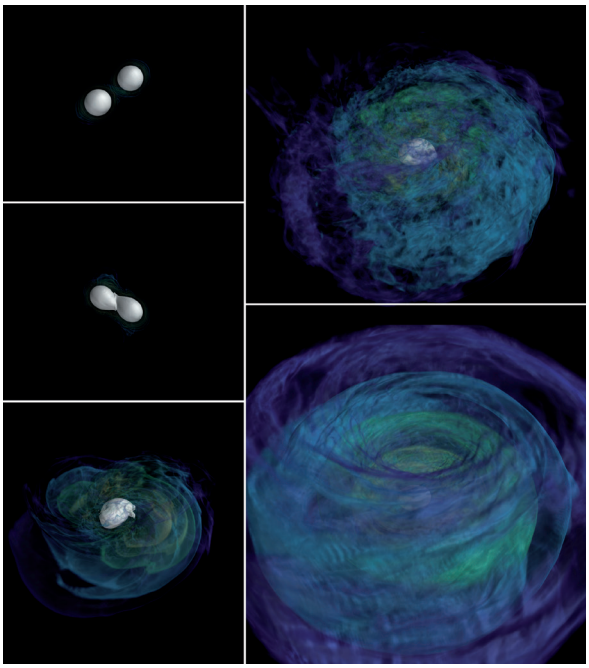

Figure 1: Typical merger event. The two neutron stars (drawn in white) spiral around each other (top left panel), touch (middle left), and merge into one heavy neutron star while ejecting debris matter (bottom left). Most of the ejected matter forms a torus orbiting around the remnant (right panels), while the rest escapes.

Results and Methods

Our simulations were carried out using the Whisky code for general relativistic magneto-hydrodynamics [2] in conjunction with the publicly available McLachlan code for the evolution of the spacetime geometry (responsible for gravitational forces and gravitational waves). The code makes use of the Einstein Computational Toolkit [3], which provides MPI based parallelisation and mesh refinement. On top of that, our code uses OpenMP. A typical simulation runs on 1024-2048 cores for around 20 days. So far, we evolved 15 different setups, resulting in around 100 TB of scientific data.

Figure 1 shows snapshots for a merger that resulted in a long-lived heavy neutron star surrounded by a dense and thick accretion disk. This an interesting case for gravitational wave astronomy, since the signal carries also an imprint of the strongly oscillating remnant. Gravitational wave signals for all our simulations will be made available on our webpage [1]. We also investigate in detail the structure of long-lived remnants, their mass distribution and their rotation profile, which represents key information for the interpretation of future observations. For larger masses, the merger remnant quickly collapses into a black hole. In this case, we investigate in detail the evolution of the magnetic field, in order to understand if the conditions support the formation of a relativistic jet, which can lead to a short gamma-ray burst. The magnetic field evolution is shown in Figure 2 for one of our models. Main features are the development of a strong toroidal field in the orbital plane (see also Figure 3), followed by the emergence of a strong conical field at the edges of the accretion disk, with a structure that is reminiscent of a tornado.

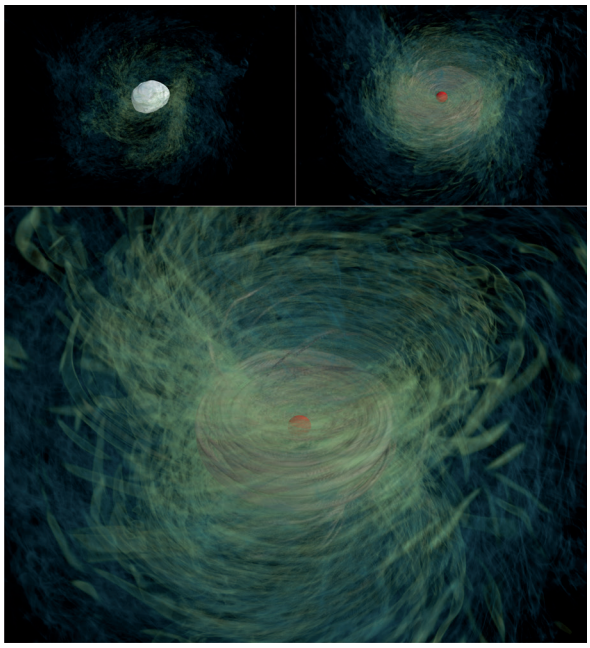

Figure 2: Development of the magnetic field. Top left panel: shortly after the merger, a short-lived neutron star (white) is surrounded by a relatively weak and unordered magnetic field. Top right: the star has collapsed to a black hole (red) and the field in the equatorial plane has become stronger. Bottom: a conical, tornado-shaped field is forming at a later stage.

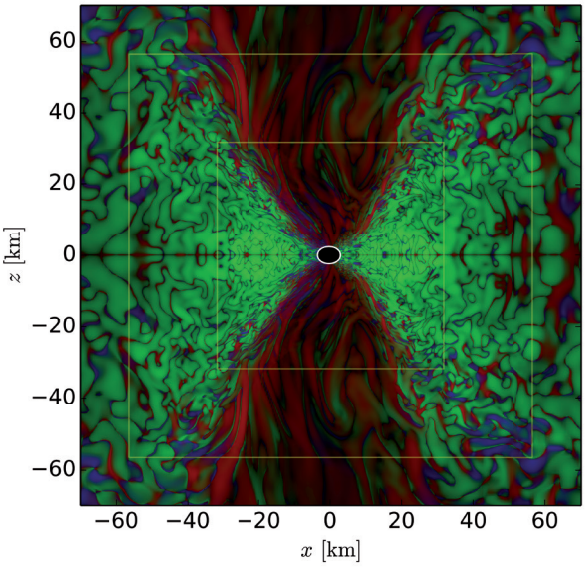

Figure 3: Distribution of magnetic field in a plane aligned with the rotation axis. The color intensity is proportional to the logarithmic field strength, while the color indicates the field orientation: red means up or down, green pointing in or out of the picture plane, blue means left or right. The white line marks the black hole horizon.

More importantly, we find that the magnetic field emerging inside the open conical funnel is mainly radially oriented, as observed in earlier simulations of the same model with less sophisticated methods and lower resolutions, but it remains too weak to be compatible with the launch of a relativistic jet. In addition, we determined for the first time the impact of using different orientations for the initial magnetic fields in the two neutron stars.

On-going Research / Outlook

Currently, we are running additional simulations to investigate additional models for the properties of dense matter. Furthermore, we are using remote visualization resources provided by LRZ to produce movies showing 3D visualizations of our simulations, which will be available soon on the webpage of our group [1]. One of the main challenges for our simulations is the fact that some important effects leading to magnetic field amplification happen on small length scales (compare Figure 3). This makes it very difficult to resolve them numerically. In order to further improve the accuracy, we proposed a follow-up study in which we will evolve one or more models with very high resolution and then use the results to calibrate a so-called sub-grid model [4], which is designed to capture the field amplification on scales not resolved with the lower, more affordable resolutions. Once calibrated, the sub-grid approach will allow to investigate a large number of models without the need for very high resolutions.

References and Links

[1] http://www.brunogiacomazzo.org
[2] B. Giacomazzo et al, Phys. Rev. D 83, 044014 (2011)
[3] http://einsteintoolkit.org/
[4] B. Giacomazzo et al, Astrophys. J. 809, 39 (2015).

Simulating the formation, evolution, and merging of molecular clouds

1

Research Institution
1. Physikalisches Institut, Universität zu Köln

Principal Investigator
Daniel Seifried

Researchers
Daniel Seifried, Stefanie Walch, Philipp Girichidis

Project Partners
MPA Garching

SuperMUC Project ID: pr94du

Introduction

In our research we investigate the formation and evolution of molecular clouds by means of high-resolution zoom-in simulations of stratified galactic disks. The simulations are a follow-up of the work performed in the Large-Scale Gauss Project pr45si carried out on the general purpose supercomputer SuperMuc, where the long-term evolution of different galactic disks was modeled with significantly lower spatial resolution[1,2]. By zooming in with a smart adaptive mesh refinement technique, we center on individual molecular clouds while these are forming and evolving within a realistic environment. We thus can explore the impact of e.g. supernova explosions on the clouds. In particular we are interested in the chemical evolution of the clouds as well as the internal dynamics and structure.

Results

In our work we have so far performed a number of simulations on SUPERMUC. Each of the simulations required a computational time of a about 1 - 2 Mio. CPU-hours with a simultaneous use of up to 1000 CPUs per simulation. A few hundreds of files were produce for each simulation requiring a disk space of about 20 TB in total. The simulations are performed with the hydrodynamics code FLASH[3] written in Fortran 90. The code solves the 3-dimensional, discretized magnetohydrodynamical equations on a Cartesian grid.

Making use of the adaptive-mesh-refinement (AMR) technique, only those regions which are of particular interest for us are resolved with the highest possible spatial resolution whereas other regions of minor interest are resolved more coarsely. This significantly reduces the number of calculations to be performed and hence the computational time required, thus allowing us to perform the simulations over long physical timescales. Furthermore, we use a chemical network designed for astrophysical problems which allows us to model the formation of molecular hydrogen and CO, and non-equilibrium cooling and heating effects.

Initial conditions
The simulations we carry out follow the evolution of the multi-phase ISM in a $(500\ \mathrm{pc})^2 \times \pm 5$ kpc region of a galactic disk, with a gas surface density of 10 M_{sun}/pc^2. We include an external potential, self-gravity, magnetic fields, heating and radiative cooling, as well as time-dependent chemistry. We explore SN explosions at different rates in located either in high-density regions, in random locations, in a combination of both, or clustered in space and time. We select zoom-in regions from these runs where we know that molecular clouds are forming and follow their evolution with a significantly higher spatial resolution of 0.06 pc.

Structure of molecular clouds
We first would like to point out that we are still at the start of analyzing the simulation data, which is why the results presented here are only preliminary.

We investigate the formation of a couple of molecular clouds in our simulations. In Fig. 1 we show the projection of two clouds formed. Even from these simple projection a couple of interesting results emerge. First, the clouds investigated are far off from a spherically symmetric structure but rather show a highly fragmented and filamentary shape. It appears also that the individual clouds show large morphological differences between each other.

The fragmentary and filamentary structure is also reflected by the so-called fractal dimension of these clouds. The fractal dimension describes how volume-filling a structure is: a homogeneous medium would have a fractal dimension of 3, a sheet-like structure 2, and a thin string-like filament a fractal dimension of 1. By analyzing the individual molecular clouds we find that in general the clouds tend to have a fractal dimension around 2.5, which is in excellent agreement with actually observed molecular clouds.

Time evolution
We also analyze the time evolution of global properties of the clouds. In general we find that the cloud masses increase over time reaching masses up to about 100 000

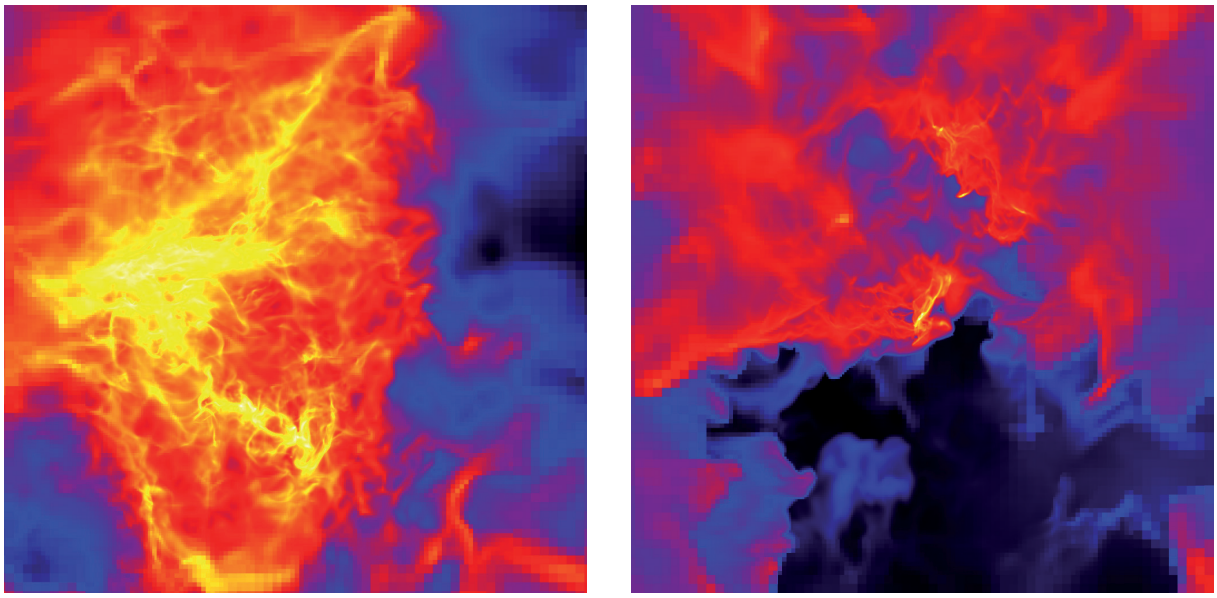

Figure 1: Projection of two molecular, star forming clouds in our simulations. The color reflects the column density of the gas.

Msun. Moreover, we find the cloud radii – which we compute from the occupied volume under the assumption of spherical symmetry – to be in the range of about 10 parsec. These values are in accordance with typical observed clouds in the Milky Way.

An interesting results from the analysis done so far, arrives from the comparison of global cloud properties derived from projected data (see Fig. 1) or when using the full 3D information present in the simulation data. In general we find that depending on the method (2D vs 3D) as well as on the angle under which the cloud is observed, even basic quantities like the mass and radius can differ quite significantly.

Dynamics of molecular clouds

It is a long-standing debate in which state observed molecular cloud are. Are they gravitationally bound or unbound, are the collapsing and how strong are the internal turbulent motions. With the present simulations we can investigate these question in detail.

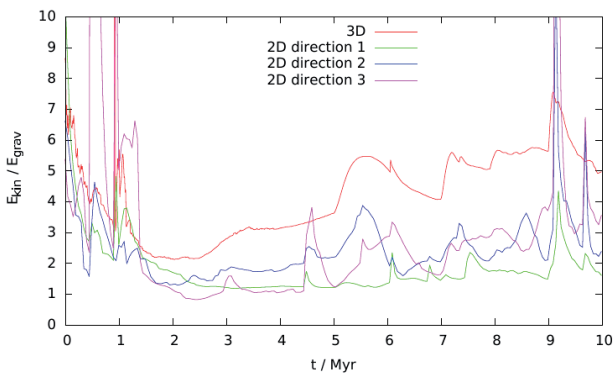

Figure 2: Time evolution of the ratio of kinetic to gravitational energy in one of our simulated molecular clouds (see right panel of Fig. 1 for a graphical representation of the cloud at t = 10 Myr)

In Fig. 2 we investigate the evolution of the ratio of turbulent to gravitational energy in the molecular cloud shown in the right panel of Fig. 1. It can be seen that the ratio is in general somewhat larger than 1, thus the cloud appears to be only marginally gravitationally bound which is in agreement with recent observations.

On-going Research / Outlook

For our future research we intend to study the evolution of further molecular clouds. In particular we will focus on the impact of magnetic fields not discussed here so far. This will require further large amounts of computing power in the future.

References and Links

[1] Walch, S., Girichidis, P., Naab, T., et al. 2015, MNRAS, 454, 238
[2] Girichidis, P., Walch, S., Naab, T., et al. 2015, MNRAS, 456, 3432
[3] Fryxell, B., Olson, K., Ricker, P., et al. 2000, ASTROPHYS. J. SUPPL. S., 131, 273

1

Magneticum Pathfinder: The simulation of the evolution of the universe in an unmatched precision

Research Institution
Universitäts-Sternwarte München, Fakultät für Physik der Ludwig-Maximilians-Universität

Principal Investigator
Klaus Dolag

Researchers
Veronica Biffi, Sebastina Bocquet, Stefano Borgani, Andreas Burkert, Nicolay J. Hammer, Michaela Hirschmann, Joseph Mohr, Margarita Petkova, Antonio Ragagnin, Rhea-Silvia Remus, Alexandro Saro, Felix Schulze, Lisa Steinborn, Adelheid Teklu

Project Partners
C2PAP Universe Cluster, LRZ

SuperMUC Project ID: pr83li (Gauss Large Scale project), pr86re

Introduction

Within modern cosmology, the Big Bang marks the beginning of the universe and the creation of matter, space and time about 13.8 billion years ago. Since then, the visible structures of the cosmos have developed: billions of galaxies which bind gas, dust, stars and planets with gravity and host supermassive black holes in their centres. But how could these visible structures have formed from the universe's initial conditions?

To answer this question, theoretical astrophysicists carry out large cosmological simulations. They transform our knowledge about the physical processes which drive the formation of our universe into models and simulate the resulting evolution of our universe across a large range of spacial scales and over billions of years. To be comparable to ongoing and future cosmological surveys, such theoretical models have to cover very large volumes, especially to host the rarest, most massive galaxy clusters expected to be the lighthouses of structure formation detectable already at early times (e.g. at high redshifts). While the Universe makes its transition from dark matter dominated to dark energy dominated (i.e. accelerated expansion), the objects which form within it make their transition from young, dynamically active and star formation driven systems to more relaxed and equilibrated systems observed at late time (e.g. low redshifts). Especially here theoretical models in form of complex, hydrodynamical cosmological simulations are needed to disentangle the internal evolution of clusters of galaxies with respect to the evolution of the cosmological background. Such simulations will be essential to interpret the outstanding discoveries expected from cosmological surveys.

In cooperation with experts of the Excellence Cluster Universe's data centre C2PAP and of LRZ, the world's most elaborated cosmological simulation of the evolution of our universe was accomplished. The most comprehensive simulation within the Magneticum Pathfinder project pursues the development of a record number of 180 billion tiny spatial elements – each representing the detailed properties of the universe and containing about 500 bytes of information – in a previously unreached spatial length scale of 12.5 billion light years (see table 1).

Results and Challenges

To perform such simulations, we incorporated a variety of physical processes in the calculations, among them three are considered particularly important for the development of the visible universe: first, the condensation of matter into stars, second, their further evolution when

Resolution	[NParticles]	Medium	High	Ultra high	Extremely high
Box 0	[2688 Mpc/h] ^{3}	***2 x 4536^3***			
Box 1	[896 Mpc/h] ^{3}	2 x 1536^3			
Box 2b	[640 Mpc/h] ^{3}		***2 x 2880^3***		
Box 2	[352 Mpc/h] ^{3}	2 x 594^3	2 x 1584^3		
Box 3	[128 Mpc/h] ^{3}	2 x 576^3	2 x 576^3	2 x 1536^3	
Box 4	[48 Mpc/h] ^{3}	2 x 81^3	2 x 576^3	2 x 576^3	
Box 5	[18 Mpc/h] ^{3}		2 x 81^3	2 x 216^3	2 x 576^3

Table 1: Overall Magneticum Pathfinder simulation set listed by box size (volume) and resolution (number of particles). The gray shaded simulations where performed within previous simulation campaigns, whereas the two largest simulations emphasized by the orange shaded entries, are added by this project.

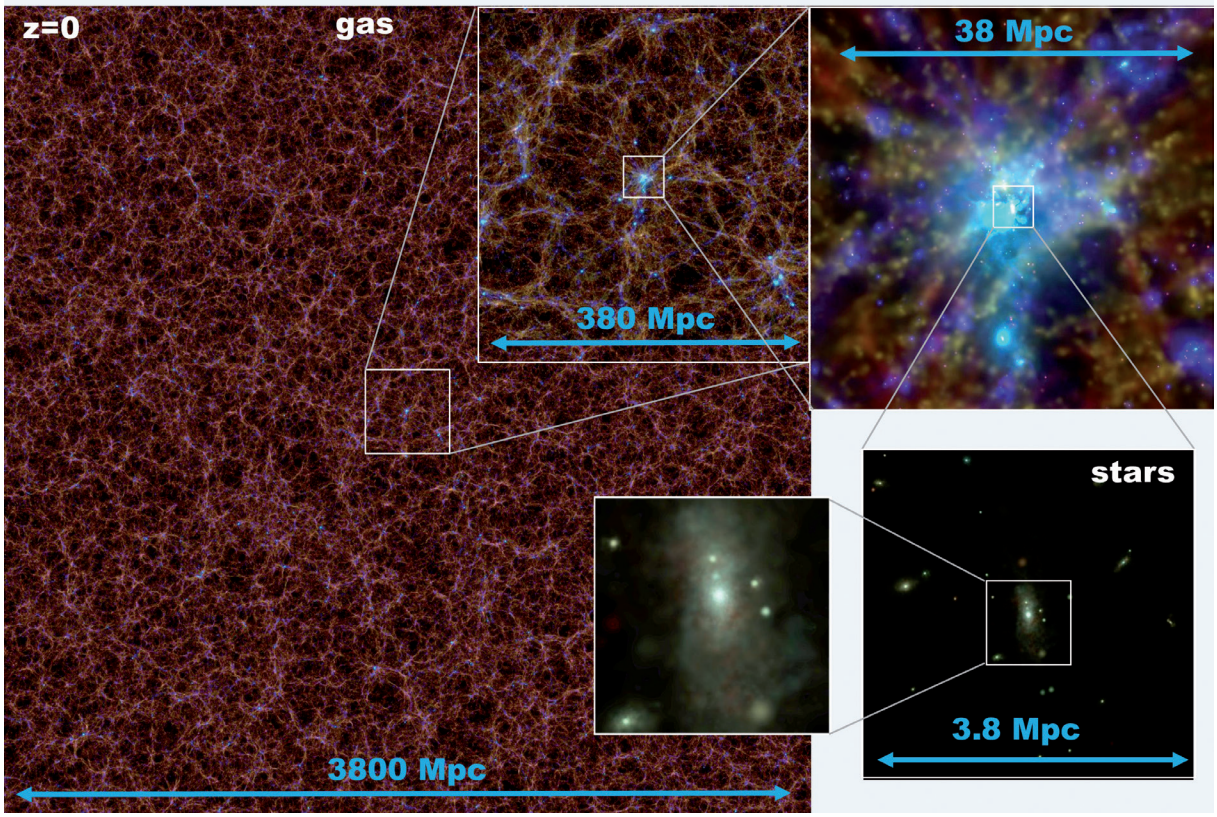

Figure 1: Visualization of the large scale distribution of the gas and stellar component within Box0 at the end of the simulation, zooming onto the most massive galaxy cluster, where individual galaxies gets visible, see Magneticum web page [1].

the surrounding matter is heated by stellar winds and supernova explosions and enriched with chemical elements, and third, the feedback of super-massive black holes that eject enormous amounts of energy into the universe.

For the first time, these numerous characteristics of the simulations performed (see figure 1 and 2) make it possible to compare cosmological simulations in detail with large-scale astronomical surveys. Astronomical surveys from space telescopes like Planck or Hubble observe a large segment of the visible universe while sophisticated simulations so far could only model very small parts of the universe, making a direct comparison virtually impossible. Thus, Magneticum Pathfinder marks the beginning of a new era in computer-based cosmology.

This achievement is preceded by more than ten years of research and development accompanied through support by HPC centers, especially experts from the Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences. One of the biggest challenges for such a complex problem is to find the right balance between optimizing the simulation code and the development of the astrophysical modeling. While the code permanently needs to be adjusted to changing technologies and new hardware, the underlying models need to be improved by including better or additional descriptions of the physical processes that form our visible universe.

To perform these largest simulations of the Magneticum Pathfinder project took about two years, including initial preparation and testing works. The research group was supported by the physicists of the data centre C2PAP which is operated by the Excellence Cluster Universe and located at the LRZ. Within the framework of several one-week workshops, the Magneticum Pathfinder team got the opportunity to use the LRZ's entire highest-performance supercomputer SuperMUC for its simulation.

Overall, the Magneticum Pathfinder simulation of Box0 utilized all 86,016 computing cores and the complete usable main memory – about 155 out of a total of 194 terabytes – of the expansion stage "Phase 2" of the SuperMUC which was put into operation last year. The entire simulation required 25 million CPU hours and generated 320 terabytes of scientific data.

On-going Research / Outlook

The Magneticum research collaboration will continue to analyze the large amount of data produced within this project, see for example [2], [3], [4], [5] and [6]. Furthermore, these data will be made available for interested researchers worldwide via a public web service, which is currently already in the testing phase. The Munich-based astrophysicists are already engaged in further projects: Among others, Klaus Dolag is currently collaborating with scientists from the Planck collaboration to compare observations of the Planck satellite with the results of the Magneticum simulations.

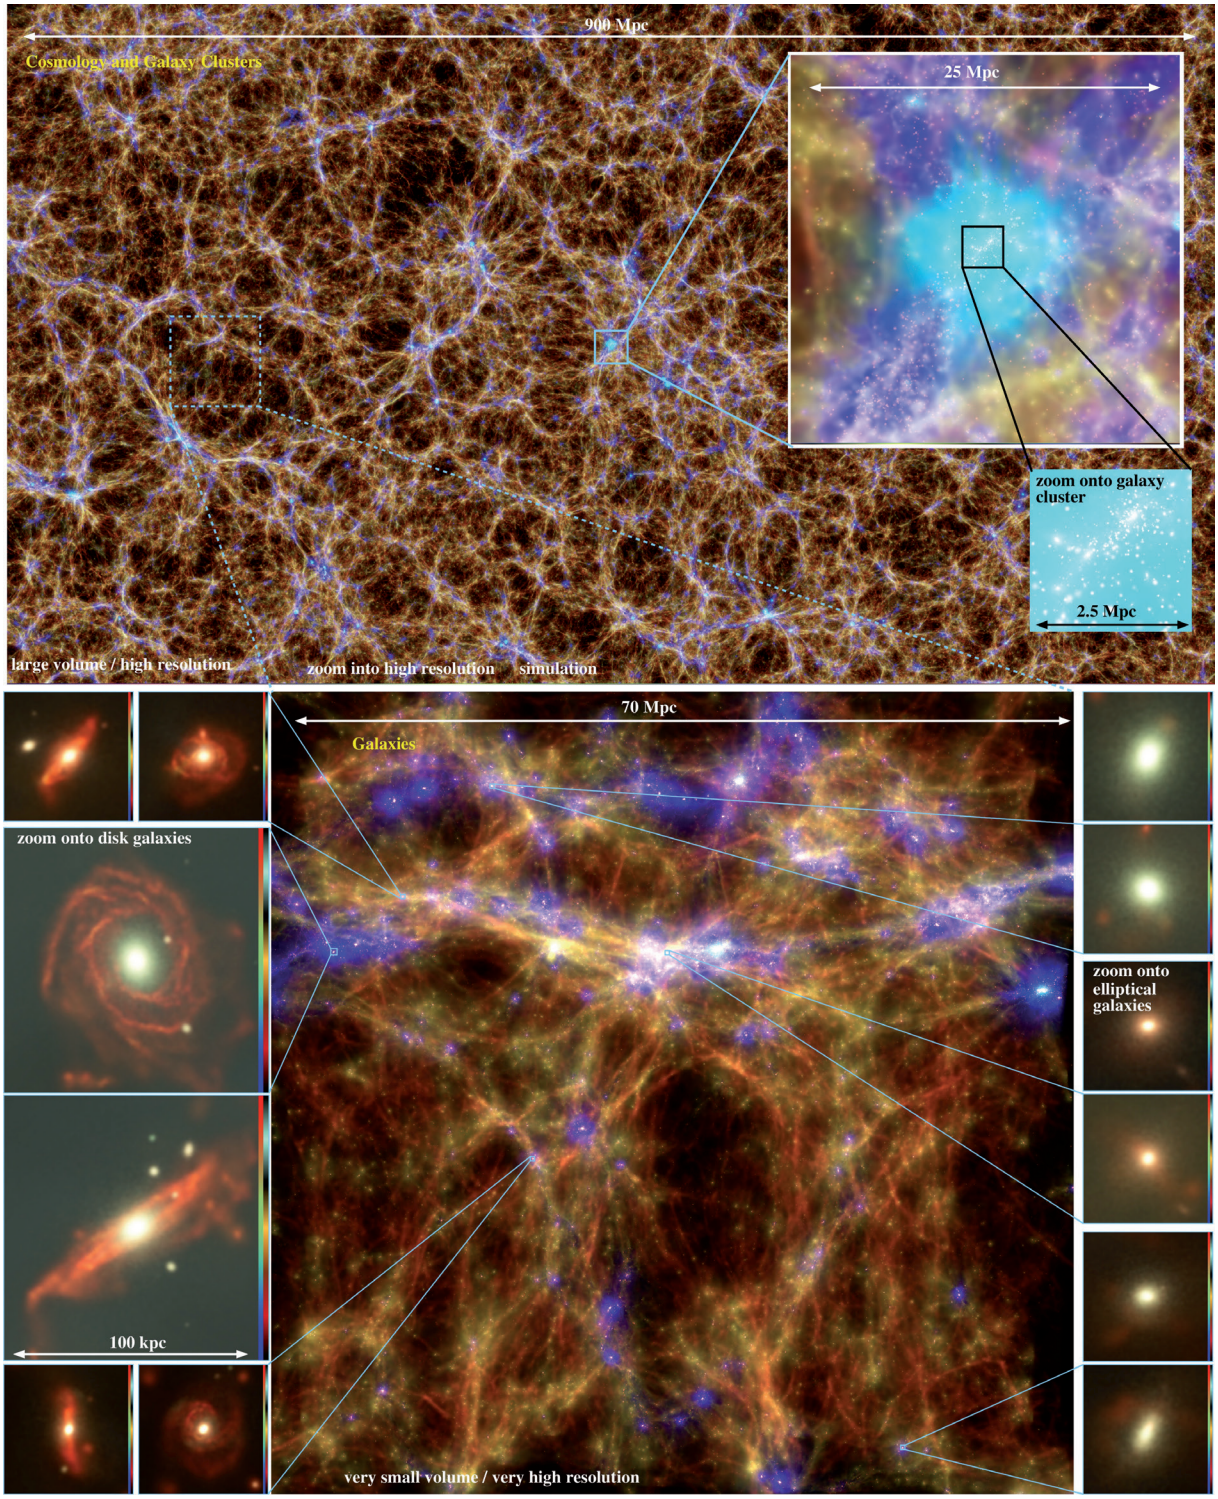

Figure 2: The shown region of Box2b/hr spans a total size of ~900 Mpc. Shown is the gas which fills the space between the galaxies (color coded according to its temperature from cold/brown to hot/light blue) together with the galaxies and stars forming in the simulation (colored in white). The zoom onto the galaxy clusters revile the ability of this simulation to resolve up to thousands of individual member galaxies within massive clusters, each resolved with hundreds up to even ten thousands of stellar particles. In the smaller simulation volumes, as the example of Box4/uhr shows, the resolution is finally large enough that the morphology of galaxies can be resolved within the simulations, reproducing the two observed, fundamental galaxy types in our universe, elliptical and spiral galaxies.

References and Links

[1] www.magneticum.org

[2] Dolag, Gaensler, Beck & Beck: Constraints on the distribution and energetics of fast radio bursts using cosmological hydrodynamic simulations, 2015, MNRAS 451, 4277

[3] Teklu, Remus & Dolag et al.: Connecting Angular Momentum and Galactic Dynamics: The complex Interplay between Spin, Mass, and Morphology, The Astrophysical Journal 2015, 812, 29

[4] Remus, Dolag & Bachmann et al.: Disk Galaxies in the Magneticum Pathfinder Simulations, 2015, International Astronomical Union Symposium, Volume 309,145-148

[5] Dolag, Komatsu & Sunyaev: SZ effects in the Magneticum Pathfinder Simulation: Comparison with the Planck, SPT, and ACT results, 2015, arXiv:1509.05134

[6] Bocquet, Saro, Dolag & Mohr: Halo mass function: baryon impact, fitting formulae, and implications for cluster cosmology, 2016, MNRAS, 456, 2631

Physics of the Solar Chromosphere

Research Institution
Institute of Theoretical Astrophysics, University of Oslo, Norway

Principal Investigator
Mats Carlsson

Researchers
Boris Gudiksen, Viggo Hansteen

Project Partners
–

SuperMUC Project ID: pr85wo (PRACE project)

Introduction

This project aims at a breakthrough in our understanding of the solar chromosphere by developing sophisticated radiation-magnetohydrodynamic simulations in order to interpret observations from the NASA SMEX mission Interface Region Imaging Spectrograph (IRIS) that was launched in June 2013.

The enigmatic chromosphere is the transition between the solar surface and the eruptive outer solar atmosphere. The chromosphere harbours and constrains the mass and energy loading processes that define the heating of the corona, the acceleration and the composition of the solar wind, and the energetics and triggering of solar outbursts (filament eruptions, flares, coronal mass ejections) that govern near-Earth space weather and affect mankind's technological environment.

Small-scale MHD processes play a pivotal role in defining the intricate fine structure and enormous dynamics of the chromosphere, controlling a reservoir of mass and energy much in excess of what is sent up into the corona. This project targets the intrinsic physics of the chromosphere in order to understand its mass and energy budgets and transfer mechanisms. Elucidating these

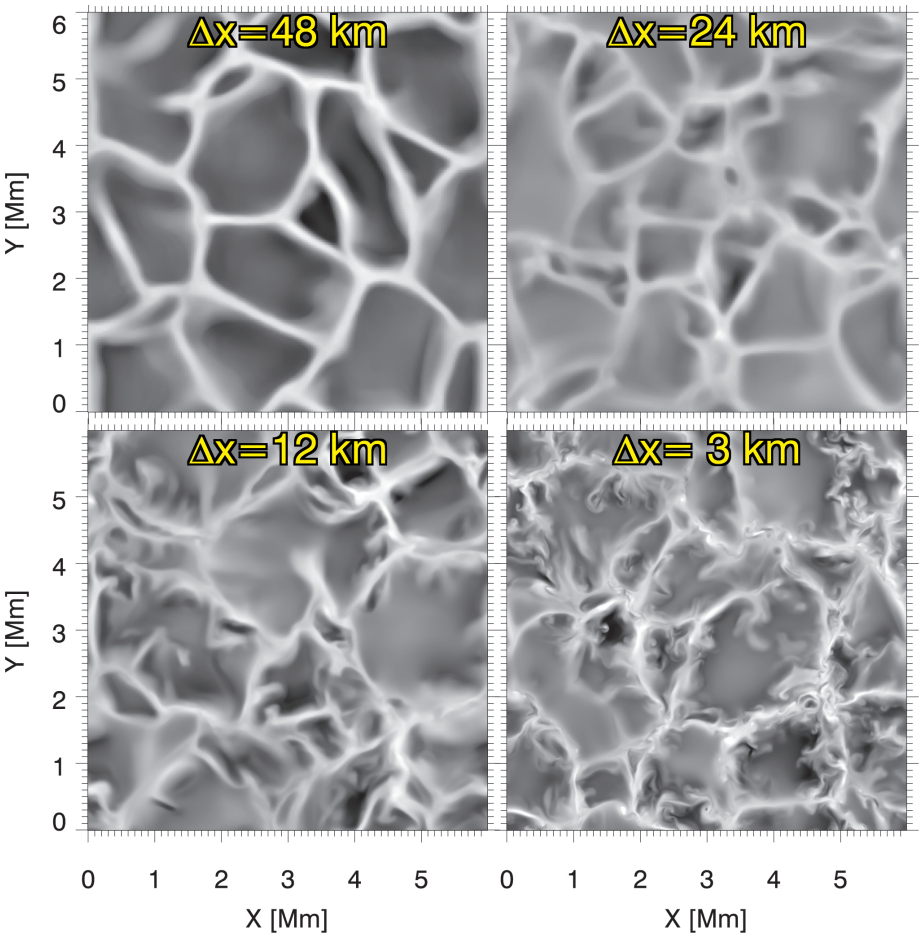

Figure 1. Solar convection shown as vertical velocity (bright is downflow) at the visible surface showing that increasing the resolution from a grid size of 48 km to 3 km dramatically increases the small scale structure. This has a corresponding effect on how the convective energy can be transported to the outer solar atmosphere. The region shown is about half the diameter of the Earth on the side.

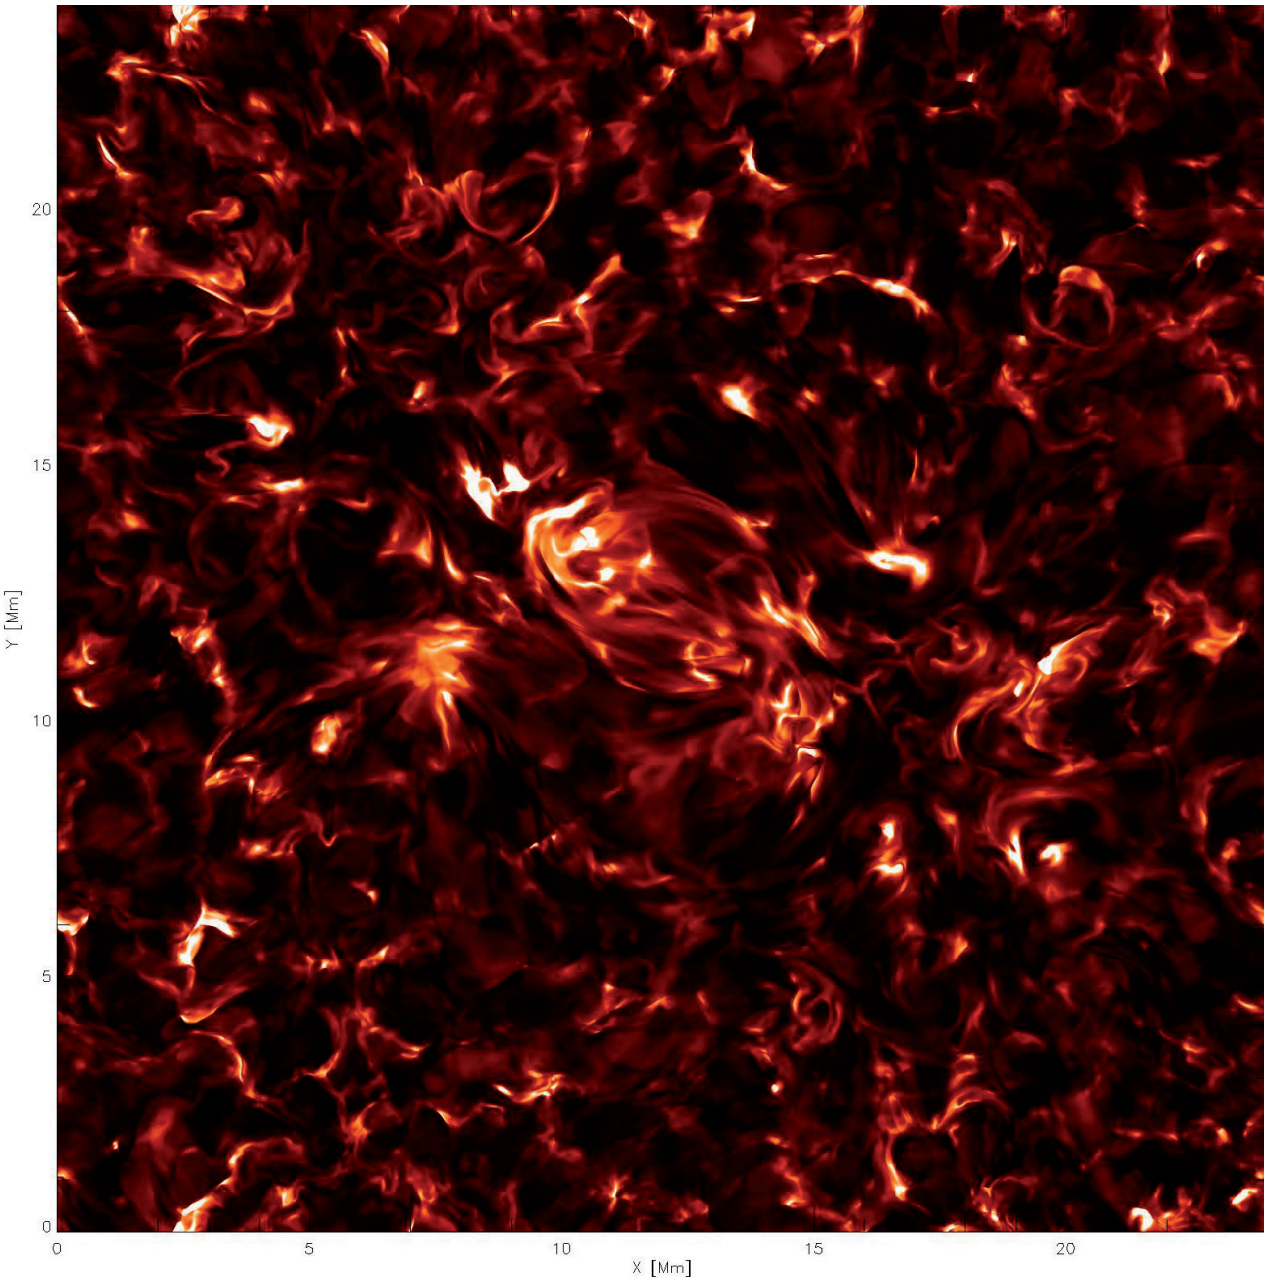

Figure 2: the Sun as viewed in a narrow band filter showing light emitted in the solar chromosphere at a temperature of about 10,000 Kelvin. The synthetic image is derived from numerical simulations that reveal how the Sun's magnetic field structures its atmosphere on fine scales. The region shown is twice as large as the Earth on the side.

is a principal quest of solar physics, a necessary step towards better space-weather prediction, and of interest to general astrophysics using the Sun as a close-up Rosetta-Stone star and to plasma physics using the Sun and heliosphere as a nearby laboratory.

Our group is world-leading in modelling the solar atmosphere as one system; from the convection zone where the motions feed energy into the magnetic field and all the way to the corona where the release of magnetic energy is more or less violent. The computational challenge is both in simplifying the complex physics without loosing the main properties and in treating a large enough volume to encompass the large chromospheric structures with enough resolution to capture the dynamics of the system. We have developed a massively parallel code, called Bifrost, to tackle this challenge. The resulting simulations are very time-consuming but crucial for the understanding of the magnetic outer atmosphere of the Sun.

Results and Methods

The models are constructed by using the Bifrost code developed by us [2]. The core of this code employs a high-order (6th order derivatives, 5th order interpolations) finite difference scheme on a staggered grid that includes high-order artificial viscosity and resistivity

("hyper-diffusion") in order to maintain numerical stability on a grid of finite resolution. These terms are also the source of magnetic and viscous heating. The equations are stepped forward in time using either the third order predictor-corrector procedure of Hyman, modified for variable time steps, or a third order Runga-Kutta scheme. When modeling the solar atmosphere the code is usually set up with periodic boundary conditions in the horizontal directions and characteristic boundaries in the vertical direction, such that the flux of energy and mass into and out of the computational domain can be controlled. The code is built up in a modular fashion, a number of extra modules are included in order to model the solar atmosphere as realistically as possible: Non-gray, optically thick radiative losses including the effects of scattering in the chromosphere are included, as are non-LTE radiative losses in the chromosphere, optically thin radiative losses in the upper chromosphere and corona and thermal conduction along the magnetic field lines in the corona. We include a realistic equation of state including partial ionization

Earlier chromospheric simulations performed by our group indicate that the numerical resolution is a critical factor for the quantitative behavior of the simulated chromospheric plasma.

In this project we have concentrated on one simulation encompassing a computational volume on the Sun of 6x6x3 Mm^3 at much higher resolution (3 km grid-size in a 2048x2048x512 box compared with the 16 km resolution in earlier simulations). This simulation shows that higher numerical resolution allows for more dynamical events in the simulated volume. These more dynamic events are closer to the observed behavior than was the case in previous simulations. The increased resolution allows for more vorticity in the narrow downdrafts in the solar convection leading to both torsional waves entering the upper atmosphere and increased winding of flux-tubes. The simulation was run on 8192 cores and used 30 M core hours. Towards the end of the allocation, the upper atmosphere was added to the box (extending it to 1628 points). We have started on the analysis of the energy balance in the simulation with special focus of the heating of the magnetic chromosphere.

On-going Research / Outlook

Thanks to the SuperMUC project we have been able to perform simulations of the solar chromosphere with twice the spatial resolution of our previous highest resolution models. To finish the simulations in a reasonable time (6-8 months) we need to run on more cores (8192) than is available to us through our national resources (now maximum 4096 cores). After the analysis of the energy balance in the magnetic chromosphere in the current simulation we will study the most promising magnetic field configurations from lower resolution simulations with similar high resolution as the SuperMUC simulation. We will also make simulations for larger computational domains including a full solar active region.

References and Links

[1] http://www.mn.uio.no/astro/english/research/projects/solar-atmospheric-modelling/
[2] Gudiksen B.V, Carlsson M., Hansteen V.H., Hayek W., Leenarts J., Martinez-Sykora J., 2011, Astronomy & Astrophysics, 531, 154

SuperCAST: Simulating the Universe

Research Institution
Universitäts-Sternwarte, Fakultät für Physik, Ludwig-Maximilians Universität München

Principal Investigator
Klaus Dolag

Researchers
Christian Alig, Alexander Arth, Alexander Beck, Andreas Burkert, Rhea-Silvia Remus, Lisa Steinborn, Jens Stücker(*), Adelheid Teklu

Project Partners
Computational Center for Particle and Astrophysics, Excellence Cluster Universe

SuperMUC Project ID: pr86re

Introduction

It is now well accepted that the observed structure of our universe is best reproduced in the presence of cold dark matter and dark energy, within the framework of LCDM cosmology, in which structures form in a hierarchical bottom up fashion. In such hierarchical picture of structure formation, small objects collapse first and then merge in a complex manner to form larger and larger structures. Astronomical instruments nowadays have opened the so-called era of precision cosmology, in which we need to understand the formation of structures in the universe with high precision, e.g. it requires that we understand the complex, non-gravitational, physical processes which determine the evolution of the cosmic baryons. The evolution of each of the underlying building blocks – where the baryons fall into the potential well of the underlying dark matter distribution, cool, and finally condense to form stars – within the hierarchical formation scenario will contribute to the state and composition of the inter-galactic and intra-cluster media (IGM and ICM, respectively), and is responsible for energy and metal feedback, magnetic fields, and high-energy particles. Depending on their origin, these components will be blown out by jets, winds or ram pressure effects and finally mix with the surrounding IGM/ICM. Some of these effects will be naturally followed within hydrodynamic simulations (like ram pressure effects), others have to be included in simulations via effective models (like star formation and related feedback and chemical enrichment by supernovae). Further components like black holes and their related AGN feedback need additional modeling of their formation and evolution processes, and must also be self consistently coupled with the hydrodynamics.

SuperCAST

Research in the Computational Astrophysics Group (CAST) at the University Observatory Munich ranges from the theoretical investigation of star formation to studies of processes on cosmological scales. A variety of different, well known numerical codes (like RAMSES, GADGET, PLUTO, SEREN) are used. Primary investigations regard the relation between turbulence and phase transitions in the multiphase interstellar medium (ISM), energetic feedback processes, molecular cloud and star formation in galaxies as well as cosmological structure and galaxy formation and the interplay between feedback processes, AGN and galaxy evolution and their imprint on the intergalactic medium (IGM) or intercluster medium (ICM). It is now clear that small-scale processes like the condensation of molecular clouds into stars and large-scale processes like gas infall from the cosmic web into galaxies are intimately coupled and have to be investigated in a concerted effort. The various projects cover a link between the various scales and contribute to our understanding of crucial aspects of the formation and evolution of central black holes and AGNs [1,2], star-forming regions and the ISM, galaxies and their IGM [3,4], galaxy clusters and the ICM [5,6,7,8] as well the large scale structures in the universe [9,10,11]. It also drives the continuous effort to develop and to apply new numerical methods and the next generation of multi-scale codes within the framework of numerical astrophysics within our group of students and young Post-Docs.

Simulating the Universe

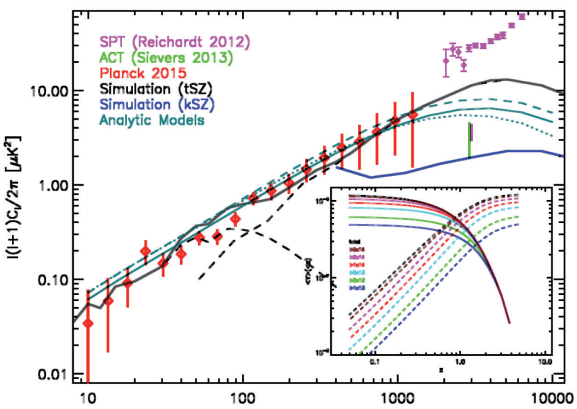

Figure 1: The observed angular power spectrum at 150 GHz of the CMB from PLANCK (red data points) and SPT (pink data points, with only the primary signal subtracted). The black solid line shows the thermal SZ signal obtained from our simulation assuming CMB cosmology. The light blue solid line shows the analytic prediction assuming the mass function from Bocquet et al. (2015), while the dashed line uses the mass function from Tinker et al. (2008). The inlay shows the build up (solid lines) of the mean Compton Y parameter with declining redshift and the contribution of halos with different mass to it (color coded).

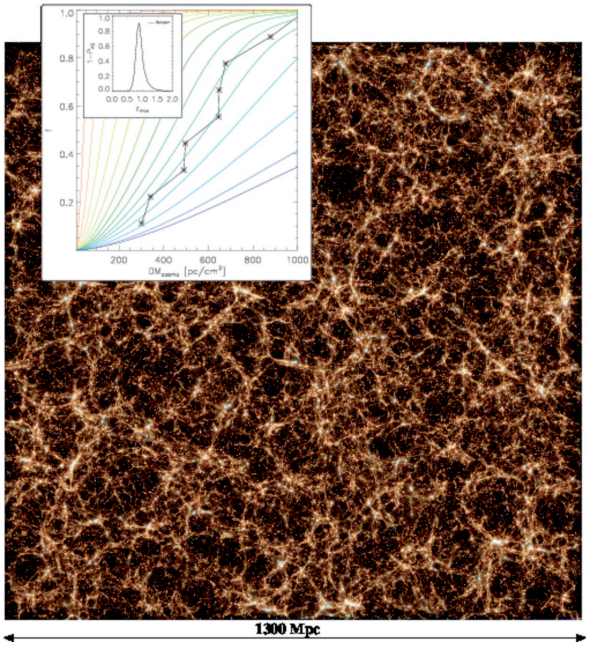

Figure 2: Shown is a visualization of the baryonic matter distribution of a large, cosmological box. The inlay shows the inferred DM distribution of sources originating at different redshifts with observations. The simulations match the observations for a redshift around 1.

Recent observations by PLANCK revealed a tension between the cosmological parameters inferred from the primary cosmic microwave background (CMB) power spectrum and the cosmological parameters inferred from the secondary signal of galaxy clusters, which leave their imprint in form of the so called Sunyaev-Zel'dovich (SZ) effect on the CMB. We demonstrate that assuming cosmological parameters inferred from the CMB, the thermal SZ power spectrum as observed by PLANCK is well matched by the deep light-cones constructed from our set of cosmological hydrodynamic simulations (*Magneticum*). The thermal SZ prediction from the full SZ maps are significantly exceeding previous templates at large multipoles l (e.g., l > 1000) and therefore predict a significantly larger contribution to the signal at l = 3000 compared to previous findings. The excess of positive values within the probability distribution of the thermal SZ signal within the simulated light-cone agrees with the one seen by PLANCK. This excess signal follows a power law shape with an index of roughly -3.2. The bulk of the thermal SZ signal originates from clusters and groups which form between z = 0 and z ≈ 2 where at high redshift (z > 1) significant part of the signal originates from proto-cluster regions which are not yet virialized. The simulation predicts a mean fluctuating Compton Y value of 1 .18 x 10^{-6}, with a remaining contribution of almost 5 x 10^{-7} when removing contribution from halos above a virial mass of 10^{13}M⊙/h. [11].

The Quest for FRBs

Fast radio bursts (FRBs) are a newly identified and as-yet-unexplained class of transient objects. The ten currently known FRBs in the literature are characterized by short (~ ms), bright (> 1 Jansky) bursts of radio emission; none have been seen to repeat, and all but two occurred at high Galactic latitude, |b| > 20°. The implied all-sky event rate is enormous, around 10000 per day. The radio signals from FRBs experience a frequency-dependent dispersion delay as they propagate through ionized gas, just as is routinely seen for radio pulsars. However, for most observed FRBs, the very high dispersion measures (DMs), in the range 400-1100 pc cm^{-3}, are more than an order of magnitude larger than the DM contribution expected from the interstellar medium (ISM) of the Milky Way in these directions. We have used the *Magneticum* set of cosmological hydrodynamic simulations to investigate the contribution from cosmological large-scale structure and found that the foreground-subtracted DMs are consistent with a cosmological origin, corresponding to a source population observable to a maximum redshift z ≈ 0.6-0.9. We considered models for the spatial distribution of FRBs in which they are randomly distributed in the Universe, track the star-formation rate of their host galaxies, track total stellar mass, or require a central supermassive black hole. Current data do not discriminate between these possibilities, but the predicted DM distributions for different models will differ considerably once we begin detecting FRBs at higher DMs and higher redshifts. We additionally consider the distribution of FRB fluences, and showed that the observations are consistent with FRBs being standard candles, each burst producing the same radiated isotropic energy. The data imply a constant isotropic burst energy of ≈ 7 x 10^{40} erg if FRBs are embedded in host galaxies, or ≈ 9 x 10^{40} erg if FRBs are randomly distributed. These energies are one to two orders of magnitudes larger than had previously been inferred. Within the constraints of the available small sample of data, our analysis favors FRB mechanisms for which the isotropic radiated energy has a narrow distribution in excess of 10^{40} erg [9].

Accretion and star-formation in a Galaxy

We conduct high resolution simulations of self-regulated isolated Galactic Discs. Galaxies form stars from a reser-

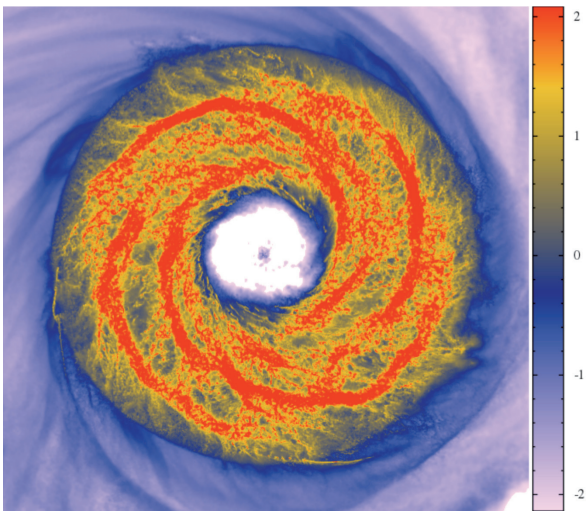

Figure 3: Shown is the logarithmic column density of a galactic disk with a high infall rate of 250 solar masses per year at 100 Myrs into the simulation.

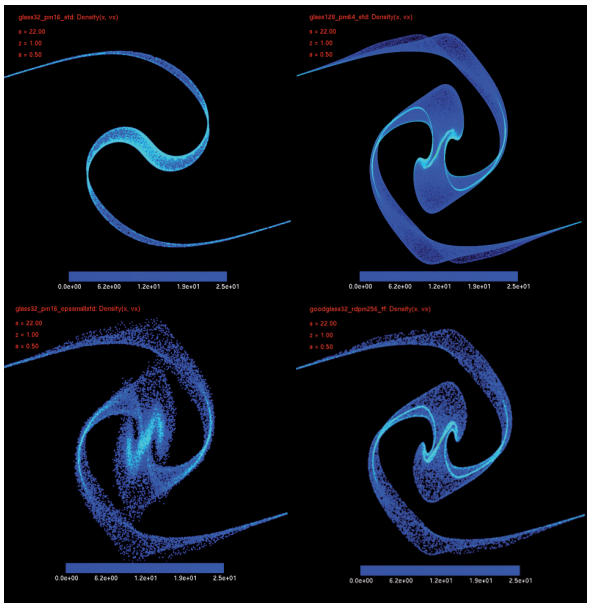

Figure 4: Phase-space diagrams of the antisymmetrically perturbed plane wave (Valinia 1997) at a=0.5 for different numerical setups. By ε we denote the softening and by d the mean particle separation. Top Left: Standard N-Body with 32^3 particles and ε ≈ 2d: The force resolution is insufficient to follow the distortions of the dark matter sheet. Bottom Left: Standard N-Body with 32^3 particles and ε ≈ 0.2d: The force resolution is higher than the mass resolution. The sheet gets destroyed by two-body collisions which leads to the formation of artificial haloes along the z-axis. Top right: high resolution 128^3 N-Body with ε ≈ 2d. Bottom right: Anisotropic softening for 32^3 particles and ε_o ≈ 2d. By adapting the softening to the distortions of the sheet, anisotropic softening resolves all forces correctly without any two-body scattering. It is converged already at 64 times less particles than the case on the upper right [12](*).

voir of cold molecular gas. For star formation to continue, this reservoir has to be fueled with extragalactic material or else the gas depletion stalls star formation. In this project we use the Smoothed Particle Hydrodynamics (SPH) code GADGET3 to simulate a spiral galaxy that is constantly fed with new gas at a given rate. This feeding naturally occurs in cosmological simulations. However, due to the large scales considered in cosmological simulations individual galaxies are rather poorly resolved.

Our simulations combine the effect of the large scale gas infall together with the ability to resolve the galaxy in much greater detail. This allows us to study the impact of external feeding on the evolution of the galaxy. It is possible to trace the position of the accreted gas and follow the spread throughout the disc. Simulating different types of infall will help us to better understand the impact of feeding and distinguish between the effects caused by feeding. Also, the impact on star-formation can be studied, which according to theory should adjust to the accretion rate in order to achieve equilibrium. A simple star-formation feedback model allows us to study the formation and destruction of structure inside the disc. The large amount of produced data can also be used for future student projects.

Warm Dark Matter Simulations

For simulations that exhibit a thermal cutoff in the power spectrum, the standard N-Body technique suffers from the fragmentation of filaments into artificial haloes. This problem does not appear if the softening ε is chosen similar to the mean particle separation d. However, such simulations suffer from far to low force-resolution. We develop the new simulation technique *anisotropic softening*. Anisotropic softening is based on the potential of an ellipsoidal density distribution with Epanechnikov-kernel which is evaluated for every softened pair interaction. The ellipsoids are reshaped according to the distortion tensor which follows the deformations of an infinitesimal volume element around each simulation particle (Vogelsberger et al., 2008). Using anisotropic softening we can match mass- and force-resolution precisely also in situations of highly anisotropic collapse, and thereby avoid artificial fragmentation while keeping the force resolution high [12](*).

Project related Publications

[1] Hirschmann, Dolag, Saro, Bachmann, Borgani & Burkert, 2014, MNRAS, 442, 2304
[2] Steinborn, Dolag, Hirschmann, Prieto & Remus 2015, MNRAS, 448, 1504
[3] Remus et al. 2013, ApJ 766, 71
[4] Teklu, Remus, Dolag, Beck, Burkert, Schmidt, Schulze, Steinborn 2015, ApJ, 812, 29
[5] Planck Collaboration, 2013, A&A, 558, 2
[6] McDonald et al 2014, ApJ 794, 67
[7] Bocquet, Saro, Dolag & Mohr 2015, 2016, MNRAS 456, 2361
[8] Arth, Dolag, Beck, Petkova & Lesch 2014, MNRAS submitted, arxiv14126533
[9] Dolag, Gaensler, Beck & Beck 2015, MNRAS, 451, 4277
[10] Marulli, Veropalumbo, Moscardini, Cimatti & Dolag 2015, arXiv150501170
[11] Dolag, Komatsu & Sunyaev 2015, MNRAS submitted, arxiv150905134
[12](*) Jens Stücker, Masterarbeit
[13] Beck, Murante, Arth, Remus, Teklu, Donnert, Planelles, Beck, Foerster, Imgrund, Dolag & Borgani 2015, MNRAS, 455, 2110
(*) Master/bachelor thesis can be downloaded from http://www.usm.uni-muenchen.de/CAST/student_projects.html

Project related Proceedings

- Dolag, Remus & Teklu 2015, IAU, Symposium 317
- Dolag, Komatsu & Sunyaev 2015, IAU, FM 5
- Remus, Dolag, Bachmann, Beck, Burkert, Hirschmann & Teklu 2014, IAU Symposium 309
- Teklu, Remus, Dolag & Burkert 2014, IAU Symposium 309
- Remus, Dolag & Burkert 2014, IAU Symposium 311

Plasma Physics

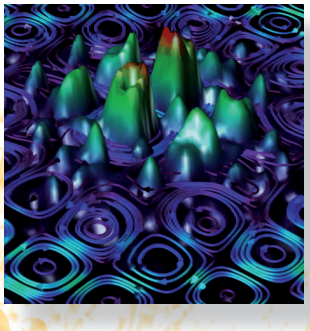

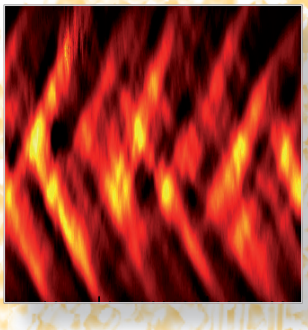

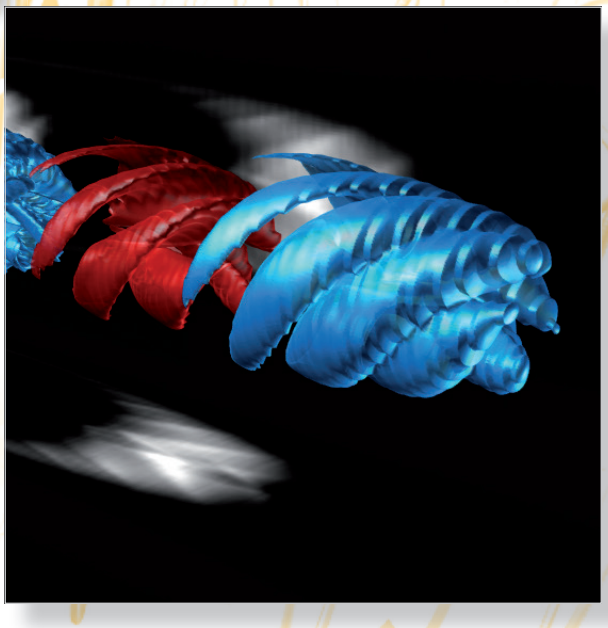

Aacsfi-PSC: Advanced accelerator concepts for strong field interaction simulated with the Plasma-Simulation-Code

Research Institution
University of Munich, faculty of physics, chair for computational and plasma physics

Principal Investigator
Hartmut Ruhl

Researchers
Karl-Ulrich Bamberg, Nils Moschüring, Viktoria Pauw

Project Partners
AWAKE Collaboration (CERN), Max Planck Institute of Quantum Optics (MPQ), ELI-NP

SuperMUC Project ID: pr84me (Gauss Large Scale project)

Introduction

Due to the AWAKE project at CERN there is a great interest in proton driven wakefield generation, as they are considered to be a promising way of accelerating electrons to the TeV scale. This is investigated using carefully planned Particle-In-Cell (PIC) simulations at the cutting edge of high performance computing technology. At the same time the ELI-NP project is gaining momentum. ELI-NP promises to become an ultra-high field laser facility where novel high-field experiments can be carried out. Also in this project the interaction of short and ultra-intense laser pulses with ultra-thin foils and nano targets in cooperation with J. Schreiber and L. Veisz at the Max Planck Institute of Quantum Optics is studied. The resulting generation of fast particles in the MeV up to GeV range as well as strong electro-magnetic radiation can be employed in fundamental science and medical applications. The facility ELI and the AWAKE project at CERN have budgets of about €1B and €15M, respectively. In both cases a detailed understanding of the expected physics, backed up by in-silico experiments, is needed and can strongly increase the return on the investment. Further results in laser driven proton acceleration can potentially help to cut costs in the order of €100M for cancer therapy centers by a factor of ten.

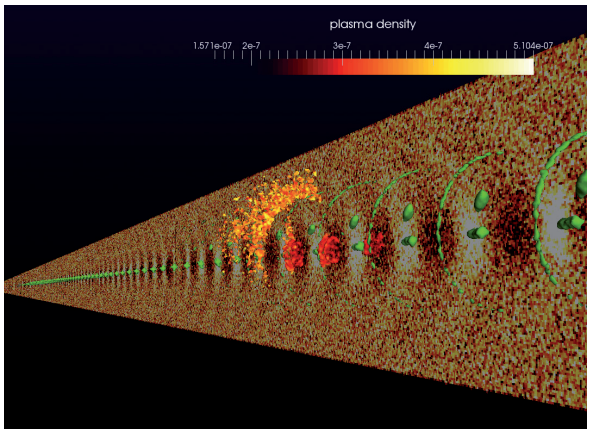

Figure 1: Visualization of the 450 GeV ion beam (green) inside the plasma (2D slice) and the accelerated electron witness beam (red/yellow). After 5m micro-bunching and a strong wakefield (black/white pits) is clearly visible.

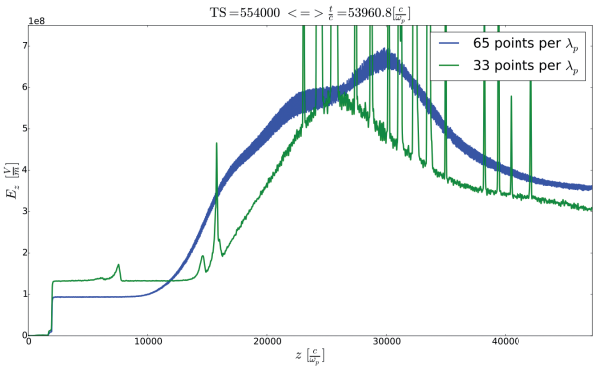

Figure 2: The maximum acceleration field of two runs, resolved differently, clearly demonstrating the need for high resolutions.

Results and Methods

The Plasma-Simulation-Code (PSC) is a general purpose framework to solve the extended Maxwell-Vlasov-Boltzmann system of equations via the PIC approach [1]. Recent extensions comprise the self-field effects of radiation and electron-positron pair production in strong fields. The original FORTRAN version evolved to a modern modularized C simulation framework supporting bindings to FORTRAN as well as C/CUDA and features selectable field and particle pushers. The PIC approach is well-known for its good scaling capability via configuration space parallelization. A Hilbert-Peano space filling curve is used for efficient, dynamic and adaptive load and memory balancing allowing for complex and dynamic geometries.

AWAKE

In the AWAKE project the interaction of a 450 GeV proton beam of the SPS pre-accelerator at CERN with a 10 m long plasma is studied (Fig. 1). The required resolution is on the μm scale and therefore 735 billion grid cells are necessary (Fig. 3). Moving window technology allows for reducing the active memory footprint and the costs to

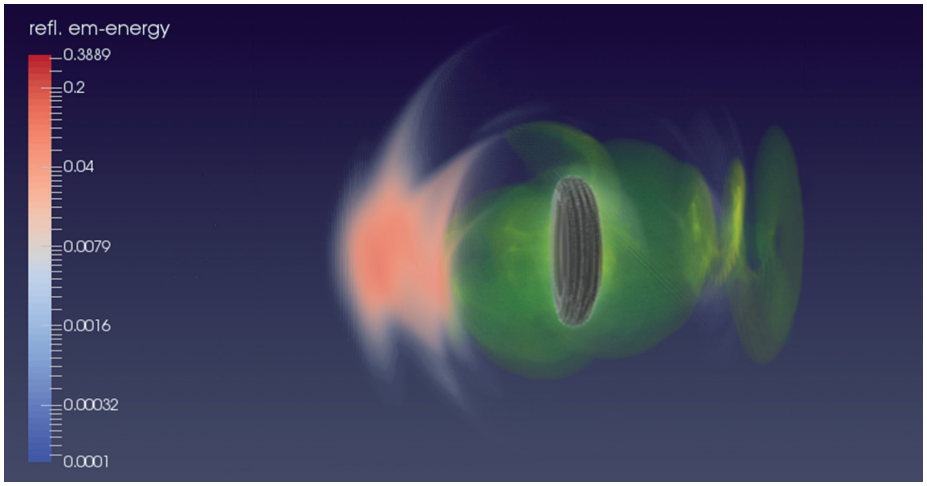

Figure 3: A 3D visualization of a preliminary simulation for an ultra thin foil interacting with a short pulse laser. The circularly polarized plane wave laser (10^{25} W/m², 2 cycles) passed the foil. In the reflected light (blue-red), the longitudinal oscillation of electron sheets (green-red) created a multi-layer-structured modulation giving rise to very short pulses of high frequency circularly polarized laser light (AXP). Ion background colored in black and white.

about 3% of the total simulation, but it still takes several weeks on a large fraction of SuperMUC. A minimum of 10 full output steps produce around 300 TB. We were able to demonstrate via a full 3D kinetic simulation that a maximum acceleration field of 700 MV is generated (Fig. 2), in accordance with 2D cylindrically symmetric codes [3], and therefore showing that PIC represents the physics correctly while opening research possibilities to important effects like beam filamentation.

Ultra-thin foils and nano targets

This project part aims for new particle as well as light sources. To produce atto-second X-ray pulses (AXP) via high harmonic generation (Fig. 3) of few cycle laser pulses enormous resolutions are necessary. E.g. the main simulation requires 2 *nm* resolution for a (16 μ*m*)³ box. Therefore 512 billion simultaneous grid cells (half a trillion) are necessary, that need up to 120 TB of main memory. The simulation took 30 hours on 16 islands (131.072 cores) of SuperMUC phase 1, i.e. four million core hours. Despite heavy inline data reduction still 100 TB of output data were generated and the frequent checkpoints of 30 TB were written with 105*GB/s* average I/O-throughput.

For Ion acceleration low emittance and high conversion efficiency from laser energy to fast ions are desirable. As the focus is not to resolve the electro-magnetic high harmonics, lower resolutions and therefore 10^5 core hours per shot are sufficient to cover the relevant physics.

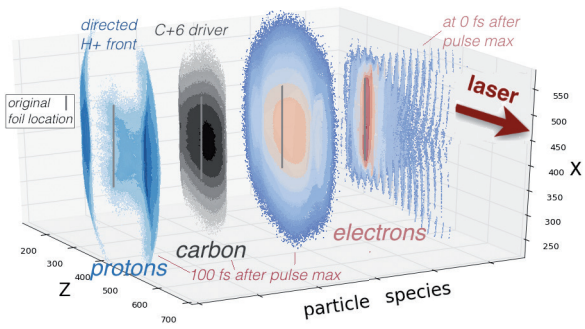

Figure 4: Polystyrene platelet irradiated by a relativistic laser pulse (20 cycles). Shown are the displacement of the electrons by the laser and the highly directed subsequent coulomb explosion of the ions.

The removal of electrons in the form of fs-bunches was studied as well as the subsequent Coulomb explosion, producing fast ions (up to 90 MeV for protons, 120 MeV for carbon, Fig. 4) with a very directed angular distribution. However, to develop better experimental designs a larger series of runs was carried out. Each run used 7168 cores for 12 hours and 5-10 TB disk space.

On-going Research / Outlook

Since the installation of SuperMUC phase 2 the 9216 nodes of phase 1 are more easily available for large scale runs allowing for the thin foil and AWAKE simulations. Besides phase 2 could be used in parallel for high throughput of the ion acceleration simulations.

Challenging to our project were the full-volume checkpoints required by PIC that strained the I/O-subsystem of SuperMUC to its limits. New approaches considered for the next generation system, like burst buffers could overcome this bottleneck. Additionally, as the FDTD solver in PIC is strongly bandwidth bound, PSC will benefit profoundly from high-bandwidth memory (HBM) that most likely will be available in future HPC machines. This will be of great advantage as in 2018 phase II of AWAKE should begin, with a longer plasma channel further increasing the need for additional computing resources.

Last but not least, it is expected that our methods used in plasma physics (many body interaction with radiation) will be more and more adapted for medical diagnostics and treatments. For this research field we expect centimeter sized volumes with necessary resolutions of tens of micro meters resulting in boxes of >10^{12} voxels (100-200 TB) on a regular basis. In consequence the demand for computing time and especially for data storage and data handling capacities will also increase significantly.

References and Links

[1] www.plasma-simulation-code.net
[2] Pauw, Ostermayr, Bamberg et al 2016, NIMA Section-A (10.1016/j.nima.2016.02.012)
[3] Caldwell & Lotov 2011, Phys. Plasmas 18, 103101

Pair-dominated plasmas and radiation in ultra intense fields

Research Institution
GoLP/IPFN, Instituto Superior Técnico

Principal Investigator
Thomas Grismayer

Researchers
Marija Vranic, Joana Martins, Jorge Vieira, Ricardo Fonseca, Luís Silva

Project Partners
–

SuperMUC Project ID: pr84yi (PRACE project)

Introduction

Relativistic electron-positron pair plasmas are tightly related to extreme astrophysical objects such as pulsar magnetospheres or gamma-ray bursts. Due to the inherent difficulties of studying these remote objects it is extremely desirable to study dense pair plasmas in the laboratory, both for fundamental purposes and astrophysical applications. The recent spectacular rise in laser intensities accompanied by the ongoing construction of new laser facilities such as ELI [1] or the Vulcan 20 PW will place intensities above 10^{23}W/cm^2 within reach. The magnitude of these lasers electromagnetic fields overlaps with the estimated fields of milliseconds pulsars. Producing pair plasmas in ultra strong fields may demonstrate that we can mimic the conditions appropriate to these astrophysical environments in terrestrial laboratories. The pair creation in such energy density environments is caused by the decay of gamma rays in intense fields. This process usually leads to quantum electrodynamics (QED) cascades, as the pairs created re-emit hard photons that decay anew in pairs, eventually resulting in electron-positron-photon plasmas.

It is known that the radiation emitted by relativistic particles can change in the presence of sufficiently strong fields. Quantum effects then start to play a role and emission of hard photons becomes significant at fields on the order of the Schwinger field in the particle rest frame. With the aim of studying the change in the radiation emission properties in the transition from the classical to the QED regime, we have incorporated quantum corrections into the radiation spectrum calculation, which are due to the electron recoil. On the other hand, most advances regarding coherent light amplification are reached by exploiting a narrow set of fundamental laser properties. The orbital angular momentum (OAM) is a new fundamental degree of freedom that can be used to reach new laser-plasma interactions regimes.

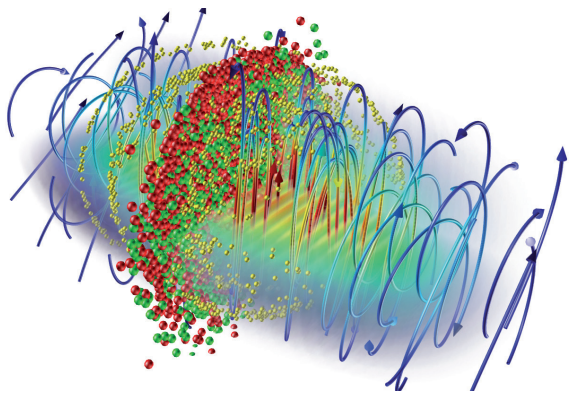

Figure 1: 3D simulation of a 2-lasers QED cascade showing the electric field lines and the pairs and photons

Results

QED cascades with counter propagating laser pulses
Our work relies on a QED module, part of our PIC code OSIRIS 3.0, which includes real photon emission from an electron or a positron, and decay of photons into pairs. The exponential growth of the number of PIC particles in cascades is sorted out with the use of a novel particle-merging algorithm [2] that resamples the 6D phase space with different weighted macro-particles. A configuration was recently proposed, comprising two counter propagating laser pulses with some seed electrons in the interaction region to initiate the cascade. In fact, analytical predictions are difficult to achieve for such complicated scenarios and the extreme multiplicity of pairs and photons in the cascade development requires tremendous computational power to allow multi-dimensional simulations.

We studied the development of electron-positron cascade shown in Fig.1 for 2 colliding laser as well as the laser absorption and the output radiation in the gamma-ray range. The conditions to achieve substantial laser absorption were investigated and the main consequence is that the ratio between the absorption time and the pulse duration is the relevant parameter for the laser depletion. We also confirmed that the significant laser energy conversion to hard photons could be achieved at lower intensities than when using four colliding lasers instead of two. The different laser polarization directions

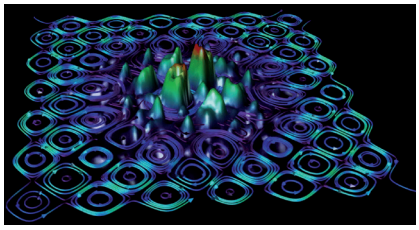

Figure 2: 2D simulations of a 4-lasers QED cascade showing electric field vortices and pair density.

result in the variation on the field microstructure, that in turn changes the way the field interacts with the plasma. Consequently, the laser absorption levels vary according to the polarization direction, as well as the directions of dominant gamma-ray emission. Fig. 2 shows the electric field structure of a standing wave formed by 4 lasers as well as the shape of the density distribution of the created plasma. A typical simulation box was 300.0 c/w_p wide and 300.0 c/w_p long, had 30 000 x 30 000 cells, and was followed for 30 000 iterations. This makes for approximately 0.5 million cpu hours for each 2D simulation of 4-laser cascades performed with OSIRIS-QED.

Classical and quantum Radiation

This quantum correction was introduced in our JRad post-processing tool and used to explore the effects of recoil in the radiation emission for very intense laser fields. Several scenarios close to the classical regime were studied in a setup where electrons were counter-propagating with a circularly polarized laser pulse and with a linearly polarized laser pulse. Using increasing intensities it was possible to observe the changes in the spectrum due to recoil and associated quantum corrections. An example of spectra obtained for such simulations is shown in Figure 3. It depicts the spectrum over a line for the collision between an ultra-relativistic electron and a circularly polarized laser. The trajectory was calculated including the radiation damping. Quantum corrections associated with the change in emission process lead to a reduction of energy radiated at higher frequencies. Furthermore, the total radiated energy from the detectors in the circularly polarized case showed that only the inclusion of the quantum corrections would allow the total energy captured in the detector be consistent with the energy lost by the particle as measured from its trajectory. This indicates that in addition to the inclusion of radiation damping in the particle trajectory, quantum corrections need to be taken into account in the spectrum calculation to obtain consistent

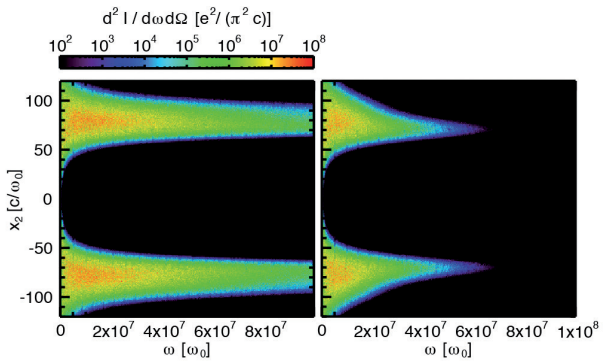

Figure 3: Left / right scattering spectra represent the results without / with quantum corrections.

results. Many simulations were performed for several different scenarios, with and without the corrections, with typical grid resolutions of 2000 x 98304 cells, few thousands of cores and tens of hours. OSIRIS QED simulations were also performed for the same setups for comparison. The results were reported in [3].

We have performed 3D simulations that show that stimulated Raman backscattering in plasmas can be used to generate and amplify orbital angular momentum lasers to petawatt intensities. Fig. 4 illustrates a laser pulse with orbital angular momentum after being Raman amplified in the plasma. These intense laser pulses could be used as a driver for compact plasma based accelerators, exciting nonlinear plasma waves suitable for positron acceleration. Typical simulations have box sizes with 650x2400x2400 cells with up to 4 particles per cell. Each simulation requires more than 50000 time steps, requiring 50,000-200,000 CPUhours to run in SuperMUC.

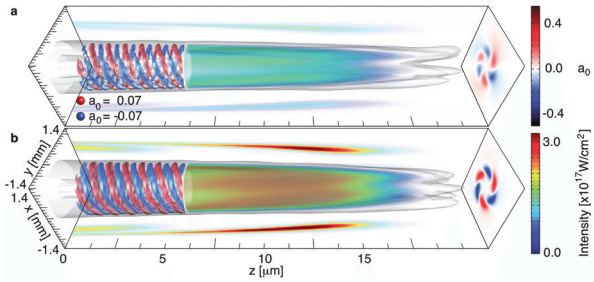

Figure 4: Illustration of a 3D Osiris simulation result showing the fields of a twisted laser pulse during (a) and after (b) the amplification. Then braided red-blue field structures are an indication of the orbital angular momentum of the laser.

SuperMIC Intel Xeon Phi Cluster

We used this system to further our efforts on deploying the OSIRIS code on this architecture, focusing not only on algorithm changes and single card performance but specifically on scalability using multiple cards in an MPI environment. Our benchmarks show code performance of ~ 600(2D) / 300(3D) million particle pushes per second on a single board, and above 74% (strong)/ 94% (weak) scaling efficiency up to 32 boards. Our tests also highlighted a bottleneck when writing simulation output, mainly due to the (well known) limitations of NFS access through the coprocessor OS, that we are currently working on overcoming to allow production work to be deployed on these systems.

On-going Research / Outlook

SuperMUC enabled to make significant advances in pair plasma creation and radiation in the laboratory in conditions for which analytical models are currently unavailable. We managed to perform very large simulations, which would not have been possible in smaller supercomputers. Besides publications in top journals (e.g. Nature Communications) we also expect several additional papers to appear in Physical Review Letters (submitted) as a scientific result from this grant.

References and Links

[1] http://www.eli-laser.eu
[2] M. Vranic et al, Comp Phys Comm 191, 65-73 (2015)
[3] J. Martins et al, PPCF 58, 014035 (2016)
[4] J. Vieira et al, Nature Communications 7, 10371(2016)

Full-f gyrokinetic simulation of edge pedestal in Textor

Research Institution
Aalto University

Principal Investigator
Timo Kiviniemi

Researchers
Paavo Niskala, Susan Leerink, Salomon Janhunen, Tuomas Korpilo, Jukka Heikkinen

Project Partners
VTT, Finland

SuperMUC Project ID: pr86go (PRACE project)

Introduction

Negligible greenhouse gas emissions and high power density make thermonuclear fusion an attractive long-term energy solution. The progress towards the significant benefits has been hampered by challenging practical execution. Magnetically confined tokamaks represent the most mature fusion energy technology at the moment. The cost and viability of a tokamak fusion reactor depend on the efficiency of the magnetic confinement, which is limited by energy losses. Applying sufficient external heating on the plasma can spontaneously increase the confinement performance of a tokamak by a factor of two. The sudden improvement in confinement is understood to be a product of the interplay between two fundamental phenomena: plasma turbulence and flows.

The research activities carried out on GCS supercomputer SuperMUC of LRZ Garching shed some light on the possible role of the radial derivative of a time-varying electric field in triggering this transition using the ELMFIRE turbulence simulation code [1,2] which investigates these phenomena with a so-called first principal computer model. This model tracks individual particles providing information on the complex interplay between the magnetic field, the electric field and the particle trajectories. ELMFIRE is even more computationally demanding than other turbulence codes as it does not assume a Maxwellian distribution of particles but instead simulates the full distribution of electrons and ions. While this approach enables more realistic simulations especially near the plasma edge, a very large number of simulation particles is needed for good statistics. So far excellent results for the small FT-2 tokamak have been obtained and the GCS supercomputing resources, made available through the Partnership for Advanced Computing in Europe (PRACE), were used to investigate phenomenology of confinement improvement in the mid-sized tokamak TEXTOR.

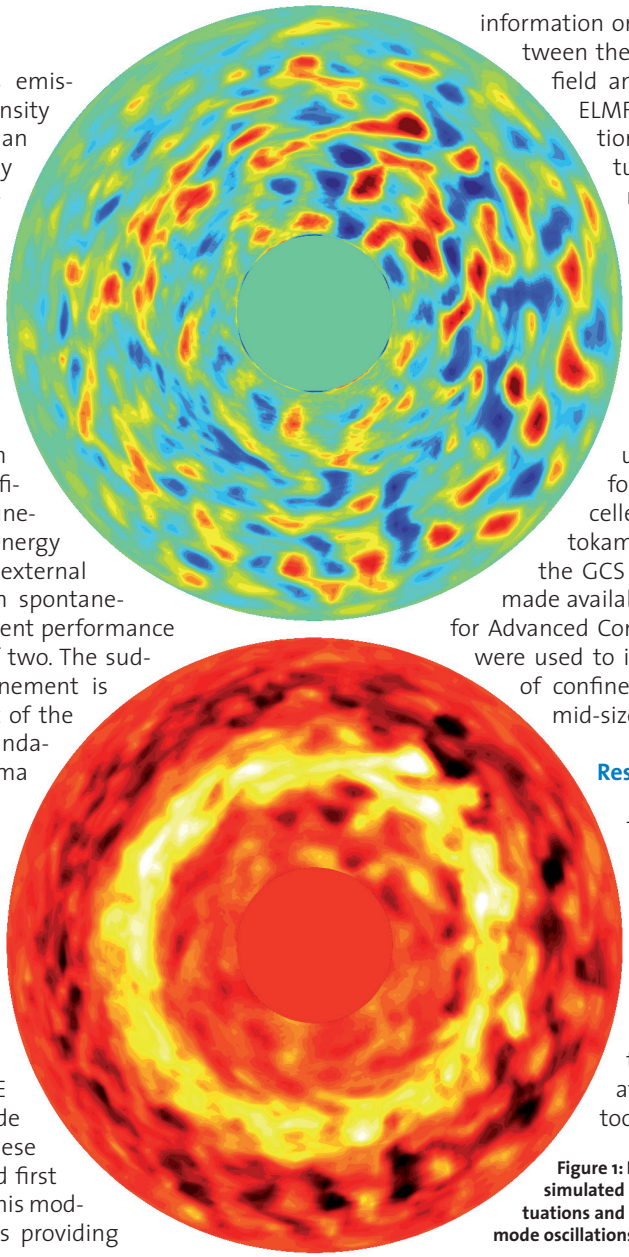

Figure 1: Poloidal cross section from a simulated plasma shows a) density fluctuations and b) typical geodesic acoustic mode oscillations in potential.

Results

The main effort in the project was a scan over local parameters like temperature and density starting from experimental Textor parameters [3]. Most of the runs were run with 4096 cores as it was noticed that scaling got worse after that. Each of the cases took about 400 000 CPUh (to-

tal of 30 MCPUh was allocated for the project) and produced about 200 GB data of results and 120 temporary snapshot files, which were used to continue the runs. At the end of the project overall storage used was totalling up to 10TB of which 4TB was transferred to our local computer for future analysis.

Onset of turbulent transport and radial electric field, E_r, oscillations with geodesic acoustic mode (GAM) characteristics was observed when the radial density profile was steepened. Typical turbulent structures of density can be seen in Fig. 1a and structure of potential oscillation during the GAM can be seen in Fig. 1b. Analysis of the radial electric field dynamics was conducted for the steep density profile with dominating trapped electron mode turbulence. Scaling of radial electric field oscillations was studied while temperature, density, magnetic field, and ion species were varied. All of the simulations showed radial electric field oscillations with GAM characteristics. Frequency and wavelength of the oscillations scaled similarly to theoretical estimates. Frequency values were generally higher than those given by the analytical estimate and no clear radial trends were obtained, possibly due to limited resolution. Wavelength increased roughly as a function of ion Larmor radius and temperature gradient scale length, although the results were again constrained by grid size.

Different scalings were observed for GAM amplitude and radial propagation speed. In the simulations, both of these increased linearly as a function of temperature and temperature gradient. Moving from hydrogen to deuterium and tritium decreased radial propagation speed but produced an increase in the amplitude. Overall, the isotope effect of GAMs seemed to be weak compared to the temperature dependency. Results indicate that collisionality could affect especially the amplitude of the oscillations.

In Fig. 2a, transport quantities can be seen to oscillate at GAM frequency in ELMFIRE simulations with evident geodesic acoustic mode activity. These modulations are visible in both particle and energy or heat fluxes as ballistically propagating avalanche-like structures and by extension similar oscillations naturally exist for density. A strong gradient in the middle of the simulation region drives the mean E_r, dividing it to positive and negative E_r shear regions which may explain the propagation of avalanches even though these structures are clearly associated with regular GAM oscillations shown in Fig. 2b and not intermittent avalanches. Part of the resources was also used to finish similar work for FT-2 [4].

From allocations we have received via others we have got excellent results for small (diameter d=1.1 m) FT-2 tokamak including first ever multi-scale benchmark of gyrokinetic simulations to experimental measurements. With the present PRACE resources we are able to investigate phenomenology of confinement improvement in TEXTOR (d=3.5 m), which is one of the largest tokamaks in Europe. With PRACE resources our group is the only one in Europe which is able to do direct comparison of

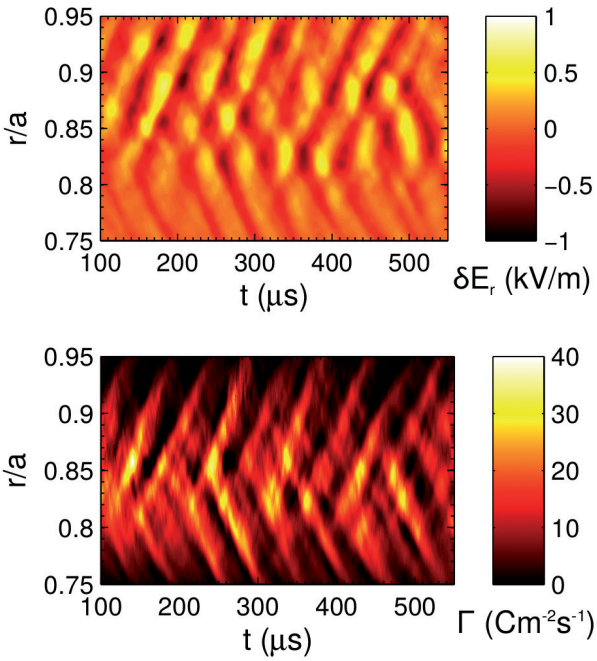

Figure 2: a) GAM activity in E_r is typically associated with corresponding b) fluctuations of transport.

first principles simulations to multi-scale experimental data for a tokamak close to reactor relevant size. Present PRACE project is a highly important and necessary step on a way towards simulating ITER (d = 12 m).

On-going Research / Outlook

In ongoing simulations we have noticed that change phase angle between electric field and density oscillation may be important for changes in particle transport for different isotopes which could explain part of the so-called isotope-effect. Even the present database from the PRACE simulation (about 20 cases and some 4TB of data) can still be further explored for this as the 3D data for both electric field and density exists. After finishing the PRACE project the code has been updated to include scrape-off-layer (SOL) which has opened several possibilities for future research.

References and Links

[1] http://www.elmfire.eu
[2] J. A. Heikkinen, S. J. Janhunen, T. P. Kiviniemi and F. Ogando, Journal of Computational Physics 227 (2008) 5582.
[3] P. Niskala, T. P. Kiviniemi, S. Leerink, T. Korpilo. Nucl. Fusion 55 073012 (2015).
[4] A. D. Gurchenko, E. Z. Gusakov, P. Niskala, A. B. Altukhov, L. A. Esipov, T. P. Kiviniemi, D. V. Kouprienko, M. Yu. Kantor, S. I. Lashkul, S. Leerink, A. A. Perevalov and T. Korpilo, European Physical Letters 110 (2015) 55001

Riding plasma wakes: from astrophysics in the lab to twisted light

Research Institution
GoLP/IPFN, Instituto Superior Técnico

Principal Investigator
Jorge Vieira

Researchers
P. Alves, U. Sinha, R.A. Fonseca, L. O. Silva, P. Muggli, A. Caldwell

Project Partners
–

SuperMUC Project ID: pr89wo

Introduction

Plasmas can sustain electric fields that can be orders of magnitude higher than material breakdown thresholds. Thus, even in the presence of extremely high electric fields, a plasma in an equilibrium will simply readjust in order to reach a new equilibrium state, without loosing any of its properties. This unique feature has been exploited towards the development of novel, more compact, plasma-based accelerators and light sources.

Plasma accelerators have the potential to either enhance or to assist conventional particle acceleration techniques towards the construction of increasingly powerful accelerators. There are many university scale laboratories where plasma based acceleration experiments are routinely performed, and several experiments are being planned in national and international facilities around the world. One of the most exciting of such projects is the proton driven plasma wakefield accelerator, currently under preparation at CERN. In this project, a long proton beam from the Super-Proton Synchrotron will excite plasma waves that will be used to accelerate electrons to high energies. In order to realize this concept, a large collaboration has been set up, called the AWAKE collaboration [1]. AWAKE could be a first step towards a single stage plasma accelerator towards the energy frontier. A plasma based linear collider based on this scheme would also require the acceleration of positrons to ultra-relativistic speeds. Positron acceleration in plasmas, however, is challenging, and the ideal technique is still under discussion. A plasma based electron-positron collider could be used to study the most fundamental constituents of matter. Electrons and positrons co-propagating in the plasma, on the other hand, could be used to recreate astrophysical scenarios in controlled laboratory settings. This can thus an exciting test bed to tackle some of the greatest questions in astrophysics, related to the generation and amplification of magnetic fields, and to the origin of gamma ray bursts.

In this report we will discuss our main contributions related to these topics. Our research was conducted using the massively parallel three-dimensional particle-in-cell (PIC) code Osiris [2], which fully exploited the tremendous computing capabilities of SuperMUC.

Results

Towards stable proton driven plasma wakefield acceleration

In the AWAKE experiment, the plasma will turn the long proton bunch into in a sequence of smaller proton beamlets. Ensuring the stable propagation of these beamlets for arbitrarily long propagation distances is crucial for the success of AWAKE. We performed three-dimensional (3D) simulations of AWAKE relevant scenarios, which demonstrated a set of parameters where the (long) bunch dynamics could stabilized for very long propagation distances (See Fig. 1 for an example simulation) [3]. Typical simulations ran for roughly $5x10^4$ core-hours, pushing $6.4x10^8$ particles for more than $2x10^5$ time-steps.

Astrophysics in a palm-top: large scale magnetic fields in a plasma accelerator

The proton driven plasma wakefield accelerator could also be geared towards the study of astrophysical scenarios. We have then performed 3D simulations that can be related with the origin and amplification of galactic or extra galactic magnetic fields in regimes that could also be attained in AWAKE relevant configurations.

In collaboration with an experimental team at *Laboratoire d'Optique Appliquée*, we found a new mechanism for the generation of large-scale, persistent magnetic

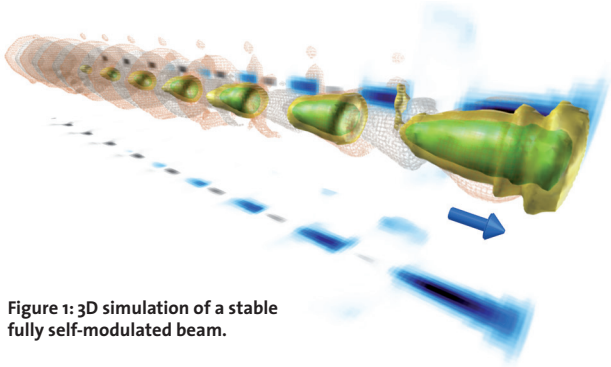

Figure 1: 3D simulation of a stable fully self-modulated beam.

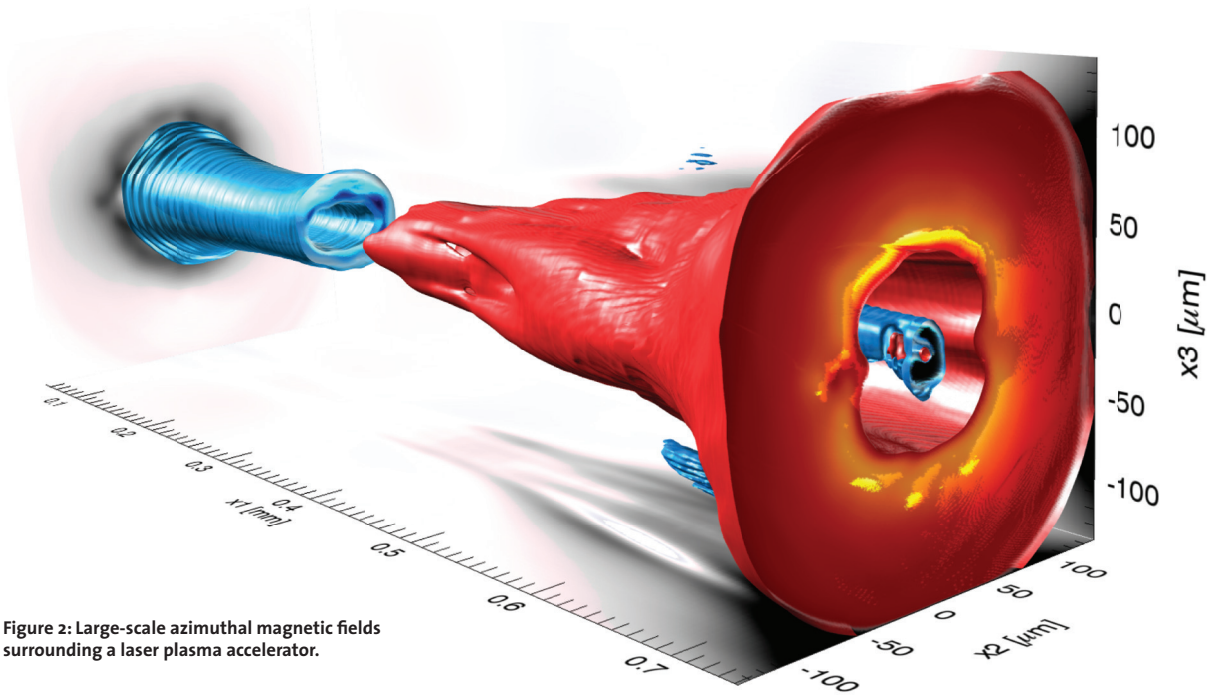

Figure 2: Large-scale azimuthal magnetic fields surrounding a laser plasma accelerator.

fields that could occur in astrophysical scenarios [4] (Fig. 2). Our three-dimensional simulations revealed that the magnetic fields form when energetic electrons cross the boundary from the plasma to the neutral region.

To observe the magnetic field generation mechanism, numerical calculations included ionization effects. The entire >1 mm long gas jet was also initialized in the simulation box. We have derived analytical estimates that indicate that this mechanism can generate the magnetic fields present in astrophysical scenarios. The simulations used a laser pulse driver, but similar findings could also be found by using proton beam drivers. Given the simulation size, access to SuperMUC was critical for the success of our endeavors. Typical 3D simulations ran for $>3\times10^5$ CPU hours, pushing $>3.3\times10^{10}$ particles during roughly 3.5×10^4 timesteps.

These large-scale magnetic field seeds could be amplified via electromagnetic plasma instabilities resulting from the propagation of electronically neutral electron-positron fireballs in plasmas. Fireball beams can be created after the collision of an ultra-relativistic electron bunch with a solid, high-Z material. The process produces Bremssrahlung radiation, and some of the high-energy Bremssrahlung photons may then decay in electron-positron pairs. Under the appropriate conditions, the fireball is electronically neutral. We have shown that these neutral fireball beams, created from laser-plasma accelerators, can greatly amplify ambient magnetic fields in plasmas [5].

We then showed that the propagation of neutral fireball beams is unstable. During the propagation, initial seed azimuthal magnetic fields start by separating the fireball electrons from the fireball positrons. The resulting electrical currents enhance the initial magnetic fields. This mechanism, illustrated in Fig. 3, is also known as the current filamentation instability, and it can amplify initial seed magnetic fields in astrophysics. A typical 3D run took 5×10^4 CPU hours, pushed 9.5×10^{10} particles during $>5\times10^4$ time-steps.

Creating optical tornadoes in plasmas

Using our SuperMUC allocation, we proposed to use intense Laguerre-Gaussian beams to produce ultra-relativistic positron bunches in plasmas [6]. These twisted laser modes are characterized by doughnut shaped intensity profiles, and orbital angular momentum. Producing twisted light at low intensities, below the damage thresholds of optical media, is straightforward. In addition, the amplification of Laguerre-Gaussian modes to very high intensities is also possible by using conventional means. However, since plasmas support very high fields, we proposed to use a stimulated Raman backscattering plasma amplifier to generate and amplify ultra-intense twisted beams or optical tornadoes.

Raman backscattering is a laser-plasma instability that occurs spontaneously when a pump laser pulse propagates in the plasma. The interaction produces a plasma wave and a scattered counter-propagating electromagnetic wave. We can stimulate this process by sending a counter propagating light wave with the same frequency of the scattered light wave. In this setup, there will be efficient energy transfer from the pump laser to the scattered wave.

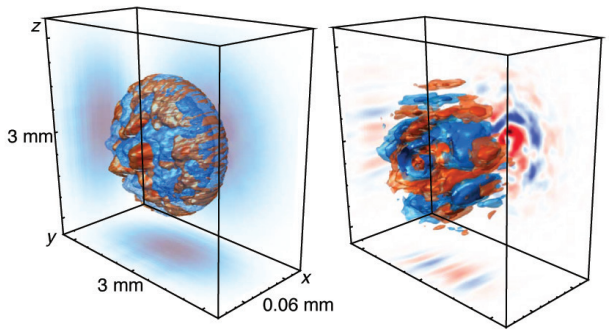

Figure 3: Current filamentation instability of an electrically neutral electron-positron fireball beam (left) and corresponding azimuthal magnetic field structure (right).

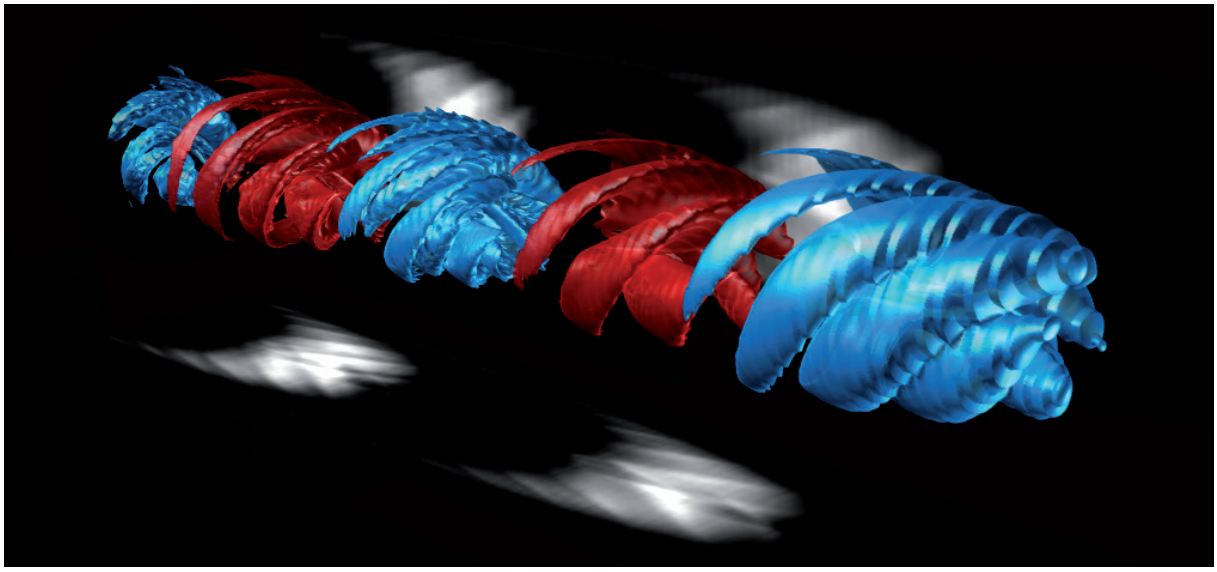

Figure 4: Plasma wave with orbital angular momentum created during stimulated Raman backscattering.

We have proposed to use this scheme to produce and to amplify twisted light in plasmas. Since Laguerre-Gaussian modes are inherently three-dimensional, we performed 3D particle-in-cell simulations to illustrate the process. Numerical calculations showed that the plasma wave excited through Raman scattering also processes orbital angular momentum (Fig. 4). In addition, we demonstrated that stimulated Raman backscattering could be used to create and amplify light beyond the Petawatt frontier. Figure 5 shows an example of an amplified twisted light beam [7]. Each run took nearly $5x10^4$ CPU hours, pushing almost 4x109 particles during 105 timesteps.

On-going Research / Outlook

Our allocation in SuperMUC opened new exciting research paths in both plasma accelerators, where we are currently investigating the properties of plasma waves driven by exotic beams, and in astrophysics, where we are studying the radiation properties associated with beam-plasma instabilities.

References and Links

[1] http://awake.web.cern.ch/awake/
[2] R.A. Fonseca, J. Vieira, et al., Plasma Phys. Control. Fusion, 55 124011 (2013).
[3] J.Vieira, et al., Phys. Rev. Lett. 112 205001 (2014)
[4] A. Flacco, J. Vieira, et al. Nat. Physics 11, 409 (2015)
[5] G. Sarri et al., Nat. Comms. 6 6747 (2015)
[6] J. Vieira, et al., Phys. Rev. Lett. 112 205001 (2014)
[7] J. Vieira et al., Nat. Comms. 7 10371 (2016).

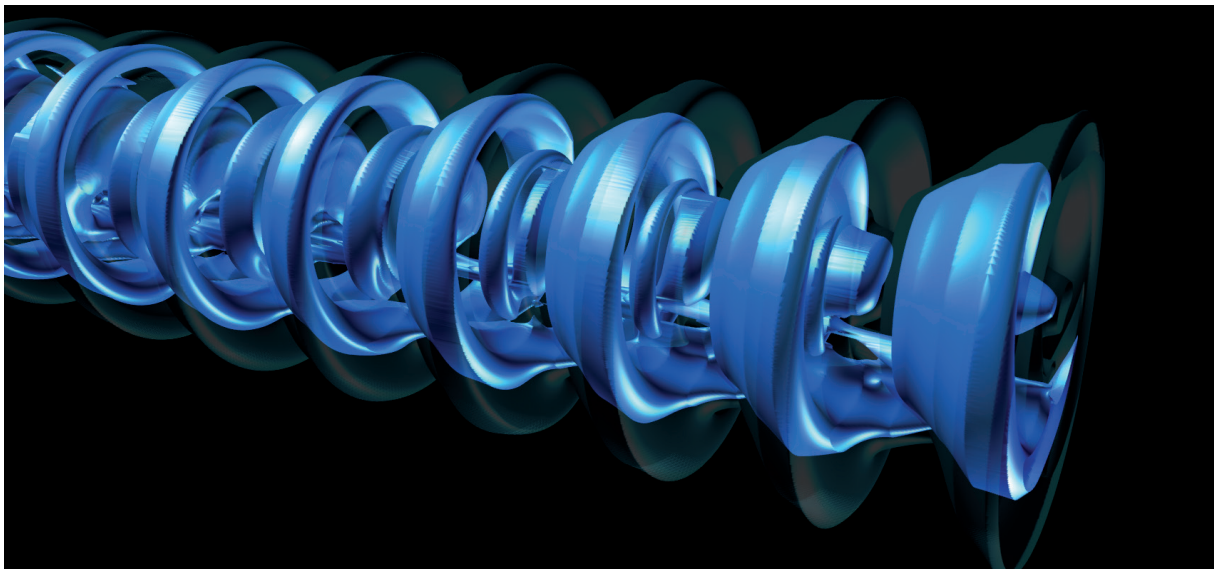

Figure 5: Amplified laser pulse with orbital angular momentum in stimulated Raman backscattering.

High Energy Physics

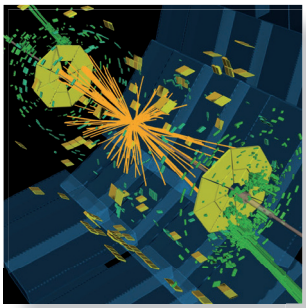

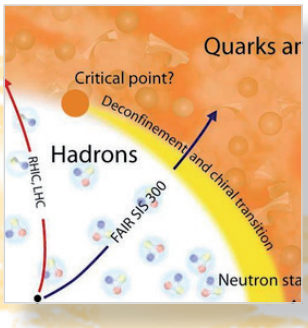

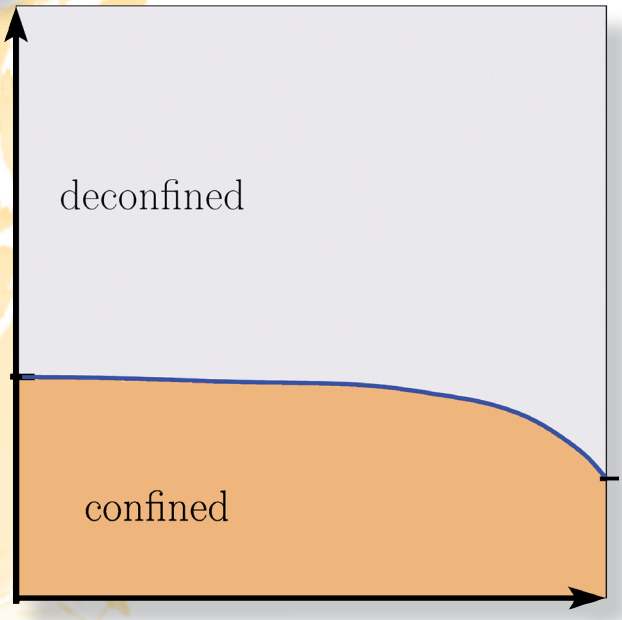

Simulation of Interactions at the LHC

Research Institution
LMU Munich, Physics Faculty

Principal Investigator
Günter Duckeck

Researchers
Rodney Walker, ATLAS Collaboration

Project Partners
MPP Munich

SuperMUC Project ID: pr58be

Introduction

Our project (ATLMUC) runs simulations of high energy proton-proton collisions in the large hadron collider (LHC) at CERN combined with a simulation of the ATLAS detector response on SuperMUC. LHC started its 2nd phase (LHC run-2) in 2015 operating at a centre-of-mass energy of 13 TeV. The ATLAS experiment[1] is one of two multi-purpose experiments at the LHC designed to record large numbers of these proton-proton collision events. Figure 1 shows an example event from the recent data taking. The ATLAS collaboration has already published more than 400 journal articles including the celebrated discovery of the Higgs boson. In searches for new phenomena, as well as for precise measurements, simulations of proton-proton collisions, based on theoretical predictions, combined with a detailed simulation of the detector response are indispensable. These simulations are computationally expensive, e.g. the complete simulation of a complex collision event takes up to 1000 seconds on a single CPU core.

The ATLAS experiment records about 10 billion collision events per year. The detailed analysis of this data requires at least the same amount of simulated events for the standard processes in order to perform the baseline optimizations and background corrections. Detailed searches for contributions from 'New Physics' processes – the main purpose of the LHC program – require additional samples of simulated events for these processes,

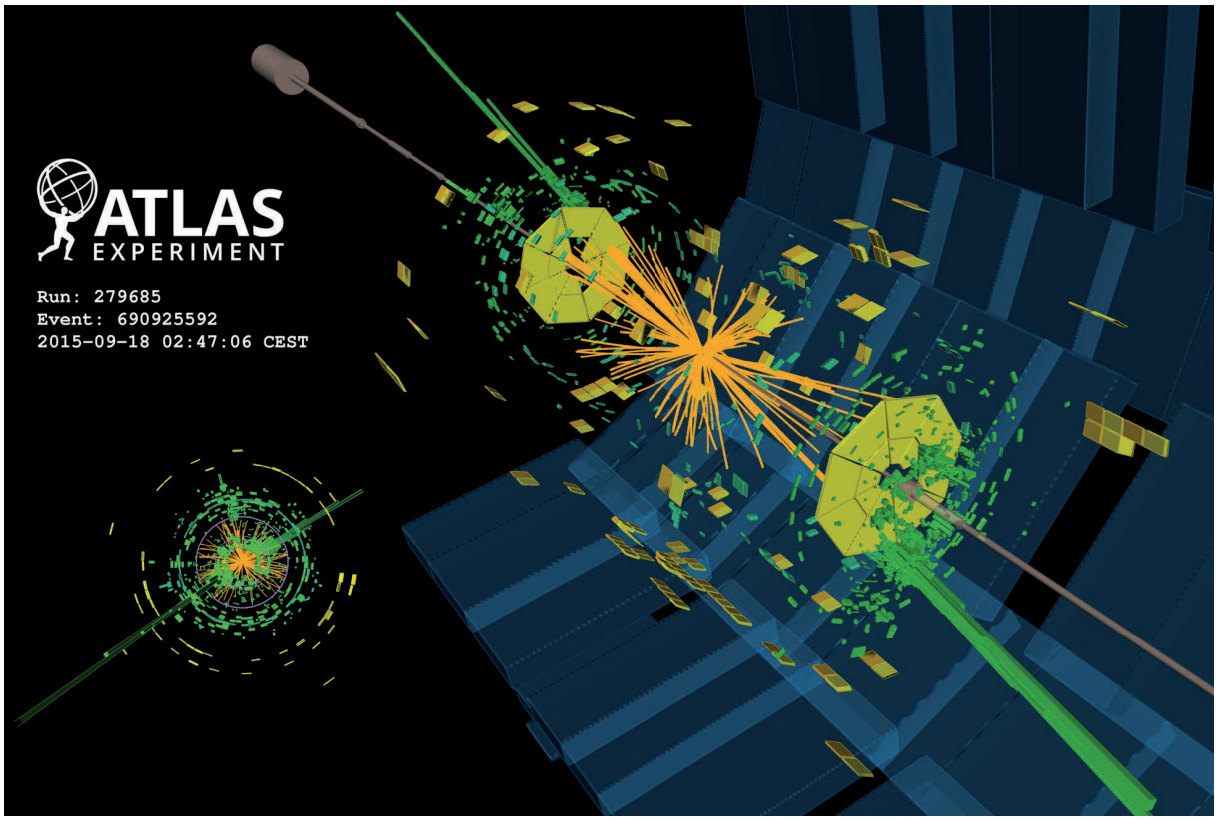

Figure 1: Example of a proton-proton collision event in ATLAS LHC Run-2. This event was collected in September 2015 and characterized by two particular high energetic jets which have an invariant mass of 8.8 TeV,

typically for multiple settings of parameters specific for these models. These measurements provide stringent constraints on theoretical models beyond the Standard Model of particle physics, as illustrated in Figure 2 for Super-Symmetry.

In many cases the scientific output of the ATLAS collaboration is not limited by the capacity to process and reduce the data but by the capacity to produce all necessary simulations. Therefore using CPU resources at HPC systems such as SuperMUC/LRZ is a crucial extension of the worldwide LHC computing grid resources which primarily focus on data storage and reconstruction of LHC events.

Results and Methods

SuperMUC was integrated into the ATLAS production system to run a cpu-intensive part of the Monte Carlo simulation of LHC events in the ATLAS detector. The integration required a gateway service to receive job requests, stage-in input data, submit into the batch system, and stage-out the output data. Due to the large number of jobs submitted automized submission procedures are required. The gateway is provided by an ARC CE[2] running on a remote node with key-based ssh access to the SuperMUC login nodes. Submission into the batch system and subsequent monitoring proceeds using commands run via ssh. The GPFS files systems are fuse-mounted (ssh-fs) and therefore available for stage-in and out of data. Several technical problems were identified and solved in doing this. The remote ARC CE via ssh is remarkably stable. Submission is strictly controlled via X509 certificate. The workloads are a well-defined subset of ATLAS central production workflows, namely detector simulation based on Geant4[3]. Geant4 is a toolkit for simulation of the passage of particles through matter. It is cpu-limited and dominated by integer arithmetic operations. Although the passage of particles through matter is serial by nature, ATLAS developed a means to usefully use multiple cpu cores. After some initialization, the process forks to N sub-processes using a copy-on-write shared memory. Each process then processes a stream of independent events, before merging at the end. This enables the efficient use of whole-nodes and thereby fulfills the basic SuperMUC requirement.

The workloads are deliberately defined to be short (<4hrs) in order to maximize backfill potential. The project was accepted on the basis of backfill with pre-emptable jobs. In lieu of check-pointing, upon which we are working, the short jobs ensure little work is lost in case of pre-emption.

The initial 10M core-hours allocation was consumed by October 2015, after which 10M core-hours were added. The main problem encountered was the poor GPFS client performance on phase-1 compute nodes. It leads to a halving of the cpu efficiency, due to delays in file access. The identical problem on the RZG Hydra machine was solved by a simple client reconfiguration. Phase 2 nodes have no such problem. A partial solution was found by using Parrot-cvmfs[4] for the software access. The cvmfs part caches file metadata and leads to reduced GPFS lookup.

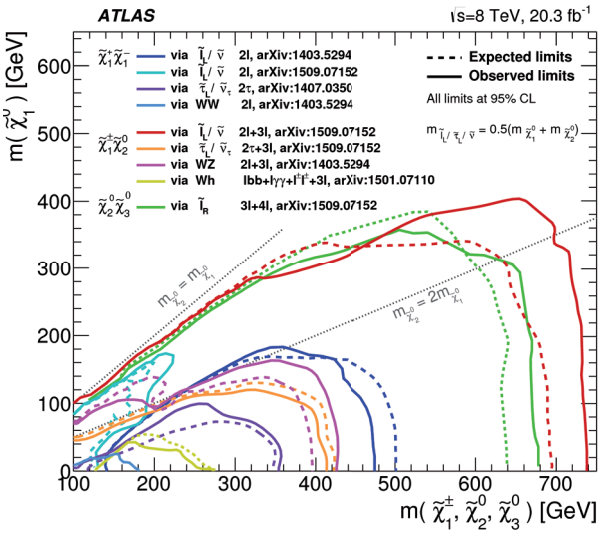

Figure 2: Example of ATLAS results on searches for super-symmetric particles predicted in extensions of the Standard Model theory. proton-proton collision event in ATLAS LHC Run-2. The plot shows 95% confidence level exclusion limits for the production of charginos and neutralinos.

GPFS scratch space is used for caching input files. We can turnover the cache to remain inside some limit, but currently rely on the system cleanup. The work GPFS is used to store the software, in a format required by the cvmfs client, and also for work directories of active jobs. The 11 TB quota is adequate for the foreseeable needs.

On-going Research / Outlook

The LHC Run-2 is planned to continue until end of 2018 and should increase the data volume by at least a factor 5 compared to Run-1. A corresponding increase of the simulated data volume is required in order to analyze and interpret the recorded data. This will allow us to determine with much better precision the properties of the Higgs Boson and either find new particles as predicted by 'New Physics' theories or further increase the constraints on these models. Using SuperMUC to simulate events will be a crucial component to reach these goals. Active development of the simulation software is ongoing in order to make the workflow more flexible and better parallelizable for smaller work-units. Adapting the software for Intel/Mic architectures is an important goal, though presumably more in the long-term after LHC Run-2 (Run-3 is planned to start in 2021). We would hope that "SuperMUC Next Generation" provides Intel/Mic architecture extensions.

References and Links

[1] The ATLAS experiment, http://atlas.ch/
[2] "Advanced Resource Connector middleware for lightweight computational Grids". M.Ellert et al., Future Generation Computer Systems 23 (2007) 219-240.
[3] Geant4, Nuclear Instruments and Methods in Physics Research A 506 (2003) 250-303,
[4] http://cernvm.cern.ch/portal/filesystem/hpc

The Chiral Condensate from Lattice QCD with Wilson Twisted Mass Quarks

Research Institution
HISKP (Theorie), Universität Bonn, Nussallee 14-16, 53115 Bonn

Principal Investigator
Carsten Urbach

Researchers
Krzysztof Cichy, Elena Garcia-Ramos, Karl Jansen, Konstantin Ottnad, Falk Zimmermann

Project Partners
NIC DESY

SuperMUC Project ID: pr63po

Introduction

Symmetries are important concepts in all fields of physics. In particular, symmetries are fundamental for the three forces building the Standard Model of Particle Physics, the electromagnetic, the weak and the strong force: all three forces are described by so-called gauge theories. The underlying gauge symmetry gives rise to the force carriers, so-called gauge bosons, with the photon being the most known example for the electromagnetic interaction.

Symmetries are not only interesting when they are exact, but also when they are broken. This breaking might be either generated by external fields, or dynamically from the theory itself. A well-known example is the magnetisation in ferromagnetic systems. Due to rotational symmetry the magnetisation is expected to be identically zero, however, by applying an external magnetic field, a finite magnetisation can be generated. Moreover, the system dynamically generates a non-zero magnetisation below a certain temperture, the Curie temperature, even without an external field.

Quantum Chromodynamics (QCD), the theory of strong interactions exhibits a similar phenomenon. QCD contains gluons as gauge bosons and quarks as fermionic matter. Chiral symmetry is formally the invariance of QCD in the limit of vanishing quark masses under the exchange of massless left- and right-handed quarks. A non-zero value of the masses of the quarks plays in rough analogy to the ferromagnet the role of an external field. But even at vanishing quark masses chiral symmetry is broken spontaneously leading to a non-zero vacuum expectation value of the chiral condensate Σ.

QCD is a strongly coupled theory at low energies and, hence, perturbation theory cannot be used to investigate low energy properties of QCD, like chiral symmetry. Therefore, a non-perturbative method needs to be applied, like lattice QCD. In lattice QCD the space-time is discretised with a small but finite lattice spacing a. For finite lattice spacing, QCD can then be solved in Euclidean space-time

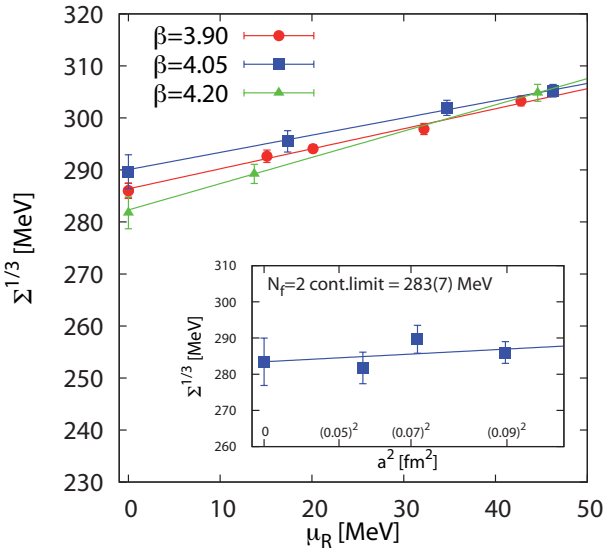

Figure 1: chiral extrapolations of the chiral condensate Σ as a function of the renormalised quark mass for the 2 flavour ensembles and three values of the lattice spacing. The lines are extrapolations to the chiral limit, linear in the quark mass. Inset: continuum extrapolation of the chirally extrapolated chiral condensate versus the lattice spacing squared.

using Markov-Chain Monte-Carlo methods. However, the discretisation we are using, the so-called Wilson twisted mass formulation of lattice QCD, breaks chiral symmetry at finite lattice spacing values, which makes it necessary to devise special methods to measure the chiral condensate. Such a method was invented by the authors of Ref. [1]. Their so-called *spectral projector* method consists in stochastically evaluating the number of eigenmodes (mode number) of the lattice Dirac operator below some threshold value M. In a certain range of M values, the dependence of this mode number on M is linear and its slope is proportional to the chiral condensate Σ. Σ needs to be renormalised and can, therefore, only be quoted in a given renormalisation scheme at a given scale. We will quote all results in the modified minimal subtraction scheme at a renormalisation scale of 2 GeV. The spectral projector method can also be used to compute the topological susceptibility χ, which is a quantity expressing the topological fluctuations of the QCD vacuum.

The simulations are performed at finite values of the

lattice spacing, hence the continuum limit has to be performed. Moreover, we simulate at finite values of the quark mass μ_R and, hence, the limit to vanishing quark mass (chiral limit) has to be performed as well.

We have investigated the chiral condensate for QCD with two light flavours (2 flavours) and for QCD with 2 light and a strange and a charm quark (2+1+1 flavours). In both cases we have studied three values of the lattice spacing and a wide range of quark mass values. The lattice spacing a is measured in units of fm.

Results

We have computed the chiral condensate Σ for QCD with 2 and 2+1+1 dynamical quark flavours, in both cases for three values of the lattice spacing [2]. The result for the 2 flavour case is shown in Figure 1 and the corresponding 2+1+1 results in Figure 2. It is visible that we control the chiral extrapolation of the renormalised chiral condensate Σ in the quark mass μ_R and the continuum extrapolation (shown in the inset) very well. The final results are $\Sigma^{1/3}$ = 283(14) MeV for the case of 2 flavours and $\Sigma^{1/3}$ = 280(12) MeV for 2+1+1 flavours (the reported errors are larger than in Figures 1 and 2, because additional sources of systematic errors were considered, in particular the choice of the range of M values in the mode number computation). This is the first evaluation of the chiral condensate using dynamical *up, down, strange* and *charm* quarks. However, within the precision reached we observed no difference in between the 2 flavour and 2+1+1 flavour results. This finding suggests that the influence of the heavier strange and charm quarks is not very large for this observable.

We remark in closing this section that we cannot report in this document here all the results obtained within this project. In particular, let us mention our computation of the topological susceptibility with the spectral projector method [3]. This was also performed for a wide range of lattice spacings and quark masses, with 2 and 2+1+1 dynamical quark flavours. The quark mass dependence of χ can be described in the framework of an effective theory, the so-called chiral perturbation theory. This theory predicts that the above mentioned dependence is governed by the chiral condensate (and other low energy constants). Thus, it is possible to extract Σ in an alternative way. We only mention this alternative extraction leads to a compatible value of Σ in comparison with the values given above, but the precision of this method is worse than of the direct spectral projector calculation.

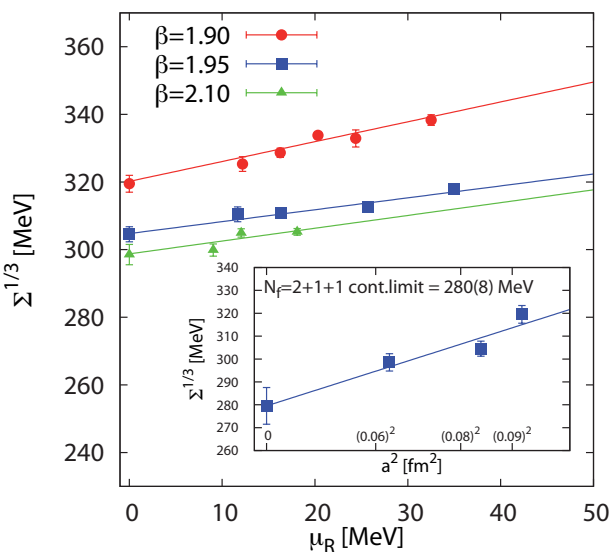

Figure 2: same as Figure 1, but for the 2+1+1 flavour ensembles.

Finally, we can confront the computation of pseudo-scalar flavour singlet meson masses [4] in 2+1+1 flavour QCD with the topological susceptibility in the so-called quenched approximation. The connection is provided by the famous Witten-Veneziano formula. The latter formula provides a link between the so-called axial anomaly in QCD and the quenched topological susceptibility. We could confirm this relation for the first time non-perturbatively [5].

On-going Research / Outlook

Lattice QCD is a very computer time demanding scientific application. Only with the computer time made available on supercomputers like SuperMUC significant progress, like the one reported here, can be reached.

Moreover, the computing resources made available by LRZ are used to reduce the systematic uncertainties in our results even further: in another project we are generating ensembles with physical values of the quark masses, such that a chiral extrapolation is not needed anymore.

References and Links

[1] L. Giusti, M. Lüscher, JHEP 0903 (2009) 013
[2] K. Cichy, E. Garcia-Ramos, K. Jansen, JHEP 1310 (2013) 175
[3] K. Cichy, E. Garcia-Ramos, K. Jansen, JHEP 1402 (2014) 119
[4] C. Michael, K. Ottnad, C. Urbach, Phys. Rev. Lett. 111 (2013) 181602
[5] K. Cichy, E. Garcia-Ramos, K. Ottnad, C. Urbach, JHEP 1509 (2015) 020

For the simulation software see
https://github.com/etmc/tmLQCD

The intrinsic scale of Quantum Chromo Dynamics

Research Institution
Neumann Institute for Computing, DESY

Principal Investigator
Rainer Sommer

Researchers
ALPHA Collaboration: Felix Bahr, Debasish Banerjee, Mattia Bruno, John Bulava, Mattia Dalla Brida, Michele Della Morte, Patrick Fritzsch, Jochen Heitger, Tomasz Korzec, Bjoern Leder, Alberto Ramos, Stefan Schaefer, Hubert Simma, Stefan Sint, Christian Wittemeier, Ulli Wolff

Project Partners
Humboldt Univeristät zu Berlin, Universidad Autonoma de Madrid, Universität Münster, Trinity College Dublin, CERN, University of Southern Denmark

SuperMUC Project ID: pr84mi (Gauss Large Scale project)

Introduction

Decades of research have lead to the Standard Model of particle physics. It is a theory which describes the structure of matter at length scales below the diameters of nuclei to an astonishing level of precision. Equivalently it successfully predicts particle decays and scattering cross sections of high-energy processes up to the energies reached at the Large Hadron Collider (LHC) at CERN in Geneva. Indeed, we also believe that physics at larger length scales, i.e. nuclear and atomic physics, would emerge from the fundamental equations of the Standard Model, if we were able to solve them directly.

A very attractive feature of the Standard Model is that it has very few free parameters. These are – as far as we presently know – fundamental parameters of Nature. Their precise determination is thus an important part of particle physics and physics in general. It is also essential in order to put the Standard Model to tests of ever increasing precision. Such tests are especially motivated by observations that go beyond the physics described by the Standard Model, such as the existence of dark matter or the degree of matter - anti-mattter asymmetry in the universe. Thus, despite its tremendous success, the Standard Model must be incomplete! In the quest for a more complete theory precision tests of the Standard Model complement direct searches for dark matter candidates and other effects of "new" physics at the LHC and other laboratories.

Both the determination of the fundamental parameters and the precision tests of the theory require precision experiments on the one hand and a precise solution of the theory (as a function of the fundamental constants) on the other hand. For many processes, the theory can be accurately solved as a series expansion in the couplings of the theory. An exception are processes affected or dominated by the "strong" interactions part of the theory. This part is called Quantum Chromo Dynamics (QCD). Its most important feature is that it describes the structure of the smallest nucleus, the proton, in terms of constituents called quarks. These are bound inside the proton by the strong force, which is mediated by the exchange of quanta called gluons. The analogous phenomenon on the atomic level is the binding of the electrons by the exchange of photons with the nucleus. However, while the Coulomb force between electron and nucleus falls off rapidly with the distance (and is relatively weak altogether, characterized by a small fine structure constant α=1/137), the force between quarks remains strong at arbitrarily large distances and leads to the phenomenon of confinement: quarks are always bound. They do not exist as true particles by themselves. How do we know that confinement is indeed a property of the theory? It is only due to its "simulations" on a *space-time grid* on super-computers such as SuperMUC. We put "simulations" in quotation marks, since these are not simulations of how particles move in space-time, but rather are stochastic solutions of *Feynman's path integral*, which provides the quantization of the fundamental fields, the quarks and the gluons. The stochastic solution of the path integral is possible independently of any series (perturbative) expansions. It thus provides non-perturbative predictions of the theory. The stochastic evaluation of the path integral on a grid is called a *lattice QCD simulation*.

Despite its strong interactions, one may define α_s, the analogue of the fine structure constant in QCD. This is the coupling of the theory and a perturbative treatment means the series expansion in this coupling. A simple, physically appealing definition of the coupling is the force between static (infinitely heavy) quarks multiplied by the square of the distance. The aforementioned property of confinement means that at distances around a proton radius this coupling is much larger than one and a series expansion makes no sense. However, at smaller distances also the QCD force becomes Coulomb-like and the distance-dependent ("running") coupling, α_s, becomes weaker and weaker. Perturbation theory then predicts that α_s vanishes at small distances, r, like $-1/\log(r\Lambda)$. The constant Λ characterizes the coupling uniquely and is the fundamental intrinsic energy scale of the theory. Once it is known, perturbative predictions, valid at short distances or, equivalently, high energies become parameter-free. This is important for tests of the theory at LHC, which provides the necessary high energies.

Results and Methods

A precision determination of the Λ-parameter is a challenge. A physical observable has to be evaluated which simultaneously has three properties: 1) it is a short distance (high energy) quantity; 2) it can be obtained with high precision and 3) its perturbative expansion is known to high order. Lattice QCD, once its free parameters are determined from low energy experimental data, can provide such observables, but in addition to the above, care has to be taken that 4) lattice spacings are small compared to the physical short distance involved.

Our collaboration has developed a systematic strategy to cope with all challenges, in particular 4). After applications to simplified theories, we are now finalising a very precise result in the three-flavor theory with u,d,s quarks. It can be connected perturbatively to the physical five-flavor number [3]. We thus determine Λ in a controlled fashion, from experimental input at the lowest energy: masses and decay constants of Pion and Kaon.

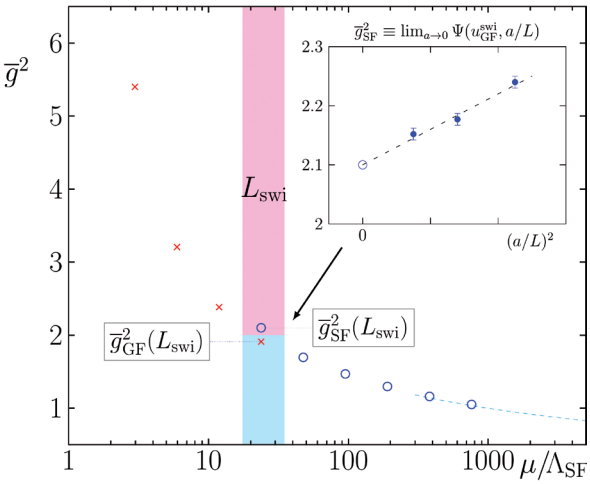

Figure 1: A sketch of the strategy for iteratively covering large scale factors and reaching the perturbative domain. See [1] for details. Here the relation $g^2=4\pi\,\alpha_s$ is used.

Our strategy (Figure 1) to connect large distances L=1/μ and small distances involves couplings defined deliberately in a finite small volume, where very small grid spacings can be simulated. Two couplings with complementary properties at larger and smaller distances are used. Their μ-dependence and their connection are both computed non-perturbatively by simulations of lattice QCD and extrapolations to vanishing grid spacing a. In the weak coupling region, the non-perturbative results are compared to the perturbative expression in terms of Λ and that fundamental parameter is determined. Here the precision of perturbation theory is at the % level because a coupling of α_s=0.1 is reached and the unknown perturbative corrections are proportional to its square.

In our *SuperMUC project*, we deal with the numerically most challenging part, the connection of the coupling to the low energy, non-perturbative scales of the theory. This is very difficult because now physically large volumes have to be simulated and different grid spacings need to be considered in order to take a continuum limit by extrapolation.

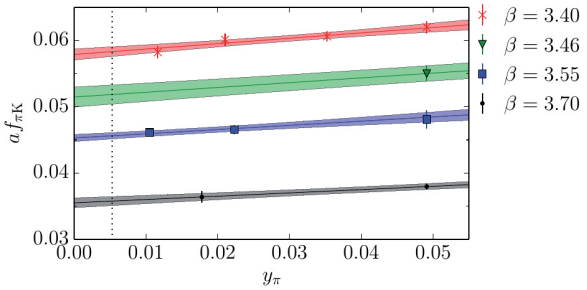

Figure 2: Linear combination of decay constants of Pion and Kaon as a function of the light quark mass for four different lattice spacings a. Graph from the Ph.D. thesis of M. Bruno, see also [2].

We use the decay constants of Pion and Kaon as low energy scales. These can be computed precisely (see Figure 2) and are also known well experimentally from the weak decays of these two mesons.

Intermediate results are shown in Figure 2 for four different grid spacings a. They have to be extrapolated to the physical quark mass indicated by the dotted line.

On-going Research / Outlook

We are presently checking that the necessarily finite size of the simulated grid does not affect the decay constants at the level of our precision and we are connecting the coupling at the smallest scale μ to the decay constants through simulations at matching grid spacings.

In the summer 2016 we will be able to put the SuperMUC results (such as Figure 2) and the analysis of the running coupling (indicated in Figure 1) together and present our high quality result at the summer conferences on particle physics. It will represent a milestone in lattice QCD: the scales μ reached are an order of magnitude higher than ever before in the three-flavour theory. Consequently the α^2 correction is truly small for the first time. In addition there is full control of the continuum limit.

The large volume simulations were carried out in a GAUSS project on both SuperMUC and Juqueen, using the most suitable architecture for each grid size. The many smaller volume simulations were done at HLRN with a much smaller number of cores per simulation. The combination of these supercomputing resources is essential for carrying out such a challenging project.
Once there is again a jump in compute resources by a factor of order 10, we would like to simulate the 4-flavour theory in a way where the decoupling of the heaviest quark from the low-energy physics is used [3].

References and Links

[1] M. Dalla Brida et al, http://archiv.org/abs/arXiv: 1511.05831
[2] M. Bruno et al, J High Energy Phys, 1502 (2015) 043
[3] M. Bruno et al, Phys. Rev.Lett. 114 (2015) 10, 102001

Using the complex Langevin equation to solve the sign problem of QCD

Research Institution
Bergische Universitaet Wuppertal

Principal Investigator
Dénes Sexty

Researchers
Szabolcs Borsányi, Zoltán Fodor, Sándor Katz, Csaba Török

Project Partners
Leibniz Supercomputing Centre, SuperMUC cluster

SuperMUC Project ID: pr84to (Gauss Large Scale project)

Introduction

The theoretical (as well as experimental) exprolation of the phase diagram of QCD matter is one of the big outstanding problems of theoretical (and experimental) physics today.

At zero fermionic density the standard non-perturbative tool of QCD, lattice discretization allows the calculation of observables of interest. However, at nonzero density such calculations are hampered by the sign problem of QCD, which is a consequence of non-zero chemical potential rendering the measure of the ensemble complex. This invalidates naive Monte Carlo calculations as the interpretation of the measure as probability density is not valid anymore. Many ideas have been invented trying to evade the sign problem such as Taylor extrapolation from the zero density ensemble, reweighting which proposes a positive measure and tries to recover the correct ensemble in the observables and extrapolation from imaginary chemical potentials. It turns out that these methods work only for small chemical potentials, that is for small fermionic densities.

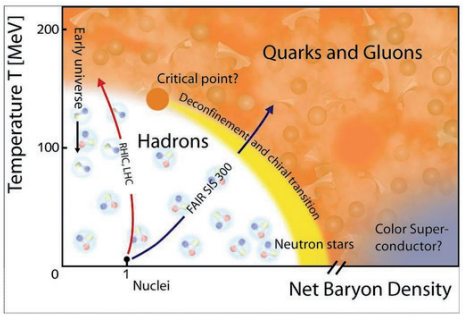

Thus the phase diagram is only easily explorable close to the temperature axis, and very little is known about it elsewhere. Many interesting phases have been proposed at high densities such as color superconductor and quarkyonic phases, and many important questions remain unanswered such as the existence and position of the critical point, the compressibility of nuclear matter etc.

In this project [1] we test an idea which tries to evade the sign problem by using analyticity: the manifold of the variables is complexified using the complex Langevin equation (CLE) [2]. The enlargement of the manifolds from SU(3) to SL(3,C) leads to instable behavior in the case of gauge theories as the gauge copies have now infinite volume. Recently it has been shown that a procedure called gauge cooling allows stabilization of the complexified process [3] allowing to perform simulations of full QCD with light quarks [4]. We compare the results of CLE to those calculated with multiparameter reweighting [5].

Results and Methods

For this study we have used two kinds of calculations: first the complex Langevin equation is solved for the link variables on a space-time lattice describing a gauge field configuration. We use adaptive step sizes to suppress runaways and thus ensure stable behavior. We also use gauge cooling to stabilize gauge degrees of freedom. The calculation of the drift term of the Langevin equation involves a contribution from the fermionic degrees of freedom for which we use the iterative Conjugate Gradient algorithm to invert the fermion matrix. We have used lattice sizes from 8^3*4 up to 16^3*8. We typically used Langevin step sizes from 10^(-5) up to 5*10^{-5} with a total run of 10-100 Langevin time. We also use the real Langevin equation to generate configurations for the starting point of the reweighting as well as the measurement of the lattice spacings and hadron masses.

The second kind of calculation involves calculating the reweighting of the ensemble to include effects of non-zero density of fermions. For this one must calculate the determinant of the fermion matrix for various chemical potentials. This is time-consuming as the size of the matrix is proportional to the space-time volume of the lattice. We have diagonalized the matrices with LAPACK routines after an initial reduction with some steps using Gauss elimination which reduces the size of the problem to be proportional to the space volume only.

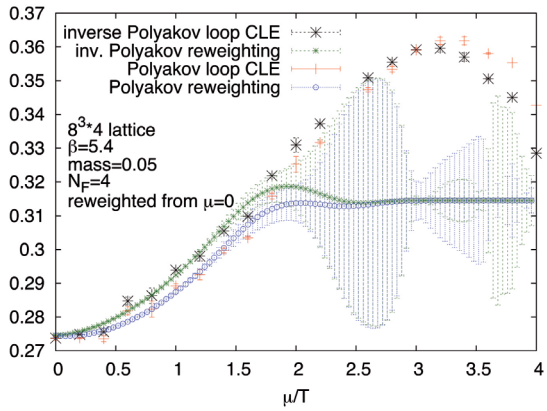

Figure 1: Comparison of reweighting and CLE results as a function of the chemical potential at a fixed temperature.

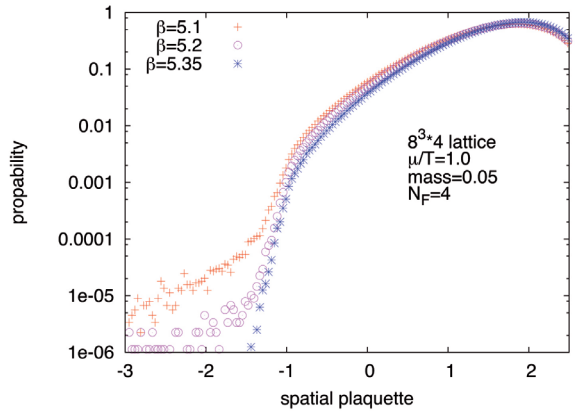

Figure 2: Histogram of the spatial plaquettes in a CLE simulation at several different temperatures.

We have used some 32 million core hours in the project. A typical job used 224 cores. The space-time lattice is split between different nodes and the boundaries are communicated using MPI. The local lattice is updated parallel using OMP parallelization. In total we generated some 50000 files totaling 1.6 TB of data. We have determined the lattice spacings and the pseudoscalar masses such that we stay on the line of constant physics as we go closer to the continuum limit. We have used N_F=4 staggered fermion flavors suche that no rooting is necessary. The bare quark masses are chosen to give a pseudoscalar mass in the vicinity of the transition temperature m_pi/T_c =2.2.

We observed good agreement between reweighting and Complex Langevin in the region where they are both trustworthy, see in Figure 1.

They have different limitations: the reweighting method fails if the chemical potential is too large, meaning that we can't get very far from the original ensemble because overlap and sign problems become strong. One alsnotices this by the jack-knife errors of the reweighting growing large. The CLE in contrast is stable at large chemical potentials as well, all the way to the saturation, where all fermionic modes on the lattice are filled. At low temperatures however, the gauge cooling becomes ineffective to counter the instability of the gauge degrees of freedom. One observes long tailed distributions of various observables, see e.g. the histogram of plaquettes in Fig.2. Similar behavior was observed before in HDQCD where the fermions are made static, and this difficulty could be overcome by getting closer to the continuum limit, i. e. increasing the lattice size. Here we observed that the lattices up to N_t=8 are not fine enough to get to the confined phase, see in Fig. 2, where we studied the system for to different pseudoscalar masses.

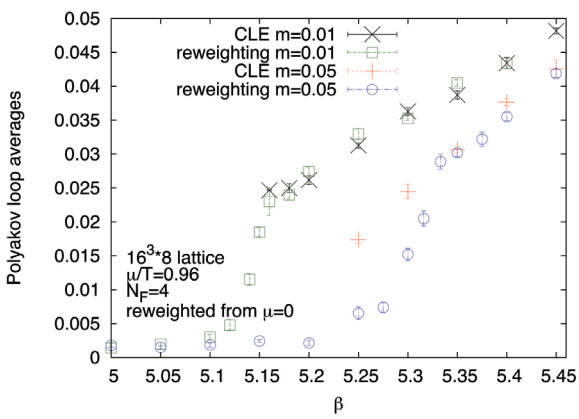

Figure 3: Comparison of reweighting and CLE results as a function of the temperature at a fixed chemical potential on N_T=8 lattices.

We have also investigated systems with N_F=2 fermion flavors to study the effects of rooting as well as reweighting from the phase-quenched ensemble. We have found good agreement between reweighting and CLE in these cases as well.

In summary we have seen that the CLE and reweighting results agree where they should and their validity can be also assessed without relying on the comparison to the other method.

On-going Research / Outlook

Using the resources of SuperMUC we have been able to calculate the reweighting results and compare them to the CLE for lattice sizes up to N_t=8. This did not allow the exploration of the phase transition line. It's an open question whether increasing the lattice size will allow us to go to smaller temperatures. The cost of larger lattices is of course increasing, especially the reweighting becomes much more expensive at larger volumes, as it's cost is proportional to the spatial volume cubed. An other important open question is the question of the poles: the fermionic drift term has singularities on the complex manifold, which in some cases can lead to the breakdown of the method, but it is unknown what its effect is on QCD, especially at low temperatures.

References and Links

[1] Z. Fodor, S.D. Katz, D. Sexty, C. Torok, Phys.Rev. D92 (2015) 9, 094516
[2] G. Parisi, Phys. Lett.131 B (1983) 393
[3] E. Seiler, D. Sexty, I.-O. Stamatescu, Phys.Lett. B723 (2013) 213-216; Aarts, Bongiovanni, Seiler, Sexty, Stamatescu Eur.Phys.J. A49 (2013) 89
[4] D.Sexty Phys.Lett. B729 (2014) 108-111
[5] Z. Fodor and S. D. Katz Phys. Lett. B534 (2002) 87

Form factors of semileptonic B-meson decays from Lattice QCD

Research Institution
Westfälische Wilhelms-Universität Münster

Principal Investigator
Jochen Heitger

Researchers
Felix Bahr, Fabio Bernardoni, Michele Della Morte, Anosh Joseph, Piotr Korcyl, Hubert Simma, Rainer Sommer, Christian Wittemeier

Project Partners
Neumann Institute for Computing, DESY, Medizinische Fakultät der TU Dresden, University of Southern Denmark, DAMTP University of Cambridge, Universität Regensburg, Humboldt Universität zu Berlin

SuperMUC Project ID: pr85ju

Introduction

The goal of this project is to contribute to the search of signals for New Physics beyond the Standard Model (SM) of elementary particle physics.

The SM is very successful in describing a wide class of phenomena that originate from the electromagnetic, the weak and the strong interactions, over a wide range of energy scales. While the SM is holding up superbly against extensive experimental scrutinies, it is well established that it can only be seen as an effective theory, valid up to some energy cut-off. Amongst others, it does not contain a candidate for dark matter, and it does not explain the observed baryon asymmetry of the universe. The direct search of new particles is the main approach to establish New Physics at the LHC, CERN, whereas it is possible that first signals will emerge from indirect searches. In this case one looks for deviations from the SM theoretical predictions in precise (low-energy) measurements. In this respect, the flavour physics sector plays a key role. Flavour-changing neutral current processes, for example, are highly suppressed in the SM and therefore very sensitive to New Physics. However, some of them are still large enough to be studied with high statistics in dedicated experiments.

Flavour transitions in the SM are parameterized by the Cabibbo-Kobayashi-Maskawa (CKM) matrix that contains information on the strengths of flavour-changing weak decays involving quarks. While these transitions have their origin in the weak interaction, the particles appearing in the physical processes, as an effect of the confinement phenomenon, are the hadrons, not the quarks. Consequently, the decay rates are written in terms of hadronic matrix elements, where the associated low-energy processes are mediated by the *theory of the strong interaction*, called *Quantum Chromodynamics (QCD)*. However, precise tests of such QCD processes are still rare, because the traditional perturbative approach is accurate in predicting short-distance effects only, but not applicable to the strongly coupled low-energy regime of the theory, where the hadrons and their decay matrix elements live. Therefore, computations in this regime have to be performed non-perturbatively.

Lattice QCD is the natural, genuinely non-perturbative ab-initio method that allows such computations, without relying on any model-dependent assumptions, and where systematic errors can be fully controlled. It starts from a discretization of space and time and puts the fundamental (quark and gluon) field variables on the sites and links of a lattice, resulting in a finite (but still very large: ~$O(10^8)$) number of degrees of freedom. The definition of physical observables in the lattice approach rests upon the Euclidean version of *Feynman's path integral representation* of the QCD partition function as well as of expectation values derived from it. In practice, these calculations are realized as *"computer experiments"*: Stochastic evaluation of the expectation values by numerical *Monte Carlo simulations of lattice QCD*, employing "importance sampling" methods to estimate the multidimensional integrals involved. In the last years, realistic lattice QCD simulations, i.e., with dynamical light sea quarks with masses close to nature and in large volumes at fine lattice spacings, have become customary, but are still very challenging. Hence, extensive large-scale numerical simulations on super-computers such as *SuperMUC* represent valuable analysis inputs to achieve a precision for many interesting physical quantities that now is comparable to the one reached in experiments.

One very interesting case, where one expects beyond SM effects to be quite large (if present), is constituted by *semileptonic weak decays of mesons containing a b-quark, which are promising processes for indirect searches of New Physics*. The SM prediction for the decay rate requires, as mentioned above, the non-perturbative computation of a hadronic matrix element and is proportional to one entry of the CKM-matrix. With this project we aim at a lattice QCD computation of the hadronic matrix element relevant for the semileptonic process $B_s \to K \ell \nu$. Combining this result with the experimental decay rate, one is able to determine the element $|V_{ub}|$ of the CKM-matrix.

Comparing this exclusive determination with others coming from $B \to \tau \nu$, $B \to \pi \ell \nu$ and an inclusive one, possible discrepancies between these different determinations would hint at the presence of New Physics. The occurrence of the b-quark, however, represents an additional difficulty, because its mass is very large. As a consequence, it cannot be simulated as a relativistic quark with present computing power. Thus the b-quark is treated using effective theories of QCD. Our collaboration employs the *Heavy Quark Effective Theory (HQET)*, where the hard degrees of freedom of the order of the mass of the b-quark are integrated out through an expansion in the inverse b-quark's mass such that large discretization effects in hadronic quantities are suppressed when the theory is regularized on the lattice. Within our determination of the $B_s \to K \ell \nu$ decay matrix element, which involves the computation of the form factor f+, we also intend to study and implement the best techniques for the treatment of excited states as well as for the chiral and continuum extrapolation of the simulation results, later to be applied to the even more demanding computation of the $B \to \pi \ell \nu$ form factor, too. Note that an experimental measurement of the $B_s \to K \ell \nu$ decay rate is within the potential experimental progamme of LHCb, Belle and planned Super-B factories. A convincing theoretical treatment is thus expected to stimulate further experimental effort to carry out the necessary analysis.

Results and Methods

In our form factor computations, we work with HQET at next-to-leading order in the inverse heavy quark mass. The underlying *strategy* is explained in [1,2,3], illustrated in Figure 1 and splits into two parts: (i) the determination of the HQET parameters appearing in the Lagrangian and in the full set of components of the heavy-light (axial and vector) quark currents via a non-perturbative matching of HQET to QCD in small volume, and (ii) the calculation of HQET energies and matrix elements in large volume, where the once determined HQET parameters enter to extract the physical hadronic matrix elements of the semileptonic decay process in question.
Within our ongoing SuperMUC project we so far have completed the finite-volume matching calculations and

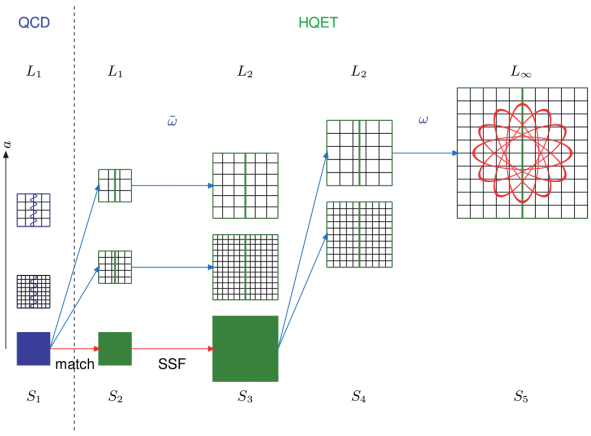

Figure 1: Strategy to determine the parameters of HQET through matching it to QCD [1]. The bottom row denotes the continuum limit obtained by extrapolating the results in each column. Graph from [2] with more details.

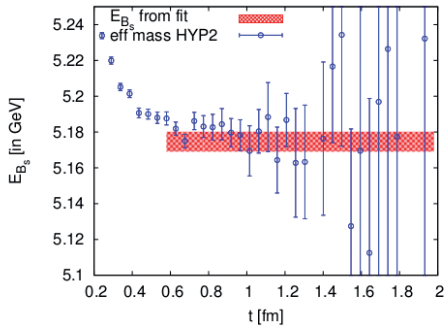

Figure 2: Example of the determination of the ground-state mass of the B_s-meson based on computations of the project. See [4] for details.

investigated the level of precision that can be obtained in large-volume calculations, taking the continuum limit at fixed pion mass and squared momentum transfer q^2 of the decay. In particular, we evaluated the relevant form factor $f_+(q^2)$ for $B_s \to K \ell \nu$ decays for two (not yet physical) values of the pion mass, approximately 270 and 330 MeV, and with a fixed q^2. To enable a continuum extrapolation, the computation has been repeated for each pion mass on three large-volume ensembles with lattice spacings of about 0.05 to 0.08 fm. Given the computational challenges from excited states contaminations at small time separations and rapidly growing statistical noise of correlators with static quarks, the high-quality data obtained in this project, together with sophisticated analysis techniques, leads to a promising status of results, see Figures 2 and 3 [4]. It is for the first time that the continuum limit of semileptonic decay matrix elements in lattice HQET with controlled errors is taken.

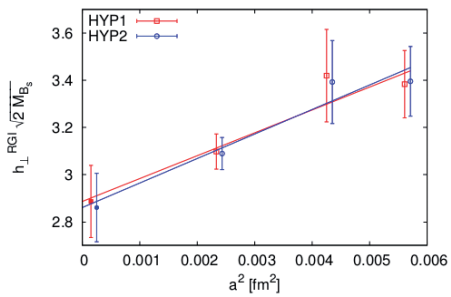

Figure 3: Continuum extrapolations linear in a^2 (from [4]) for the (properly renormalized) QCD matrix element of the spatial vector current between B_s- and K-meson states, which enters the semileptonic form factors of interest.

On-going Research / Outlook

With the remaining part of our total SuperMUC grant of 41 M core-hours we want to reduce the overall uncertainties of our results to a phenomenologically relevant precision. This amounts to add a second q^2-value, data at lighter pion mass to stabilize the chiral limit, and to include the next-to-leading order HQET corrections.

References and Links

[1] J. Heitger and R. Sommer, JHEP 0402 (2004) 022
[2] B. Blossier et al., JHEP 1209 (2012) 132
[3] M. Della Morte et al., JHEP 1405 (2014) 060
[4] F. Bahr et al., Phys. Lett. B757 (2016) 473

The strong interactions beyond the standard model of particle physics

Research Institution
Insitute for theoretical physics, University of Münster

Principal Investigator
Georg Bergner

Researchers
Pietro Giudice, Gernot Münster, Istvan Montvay, Stefano Piemonte

Project Partners
–

SuperMUC Project ID: pr85so

Introduction

The standard model of particle physics is an extremely successful theory: it contains all known fundamental particles and forces, except gravity. The discovery of the Higgs particle has completed the experimental search for the constituents of the standard model. However, the real nature of this particle is so far unknown and its mass is unnaturally light. Moreover, astronomical observations have revealed that only a small fraction of the matter in the universe consists of the particles of the standard model. It is hence essential to find consistent extensions of this theory. The most promising theoretical concepts for a solution of these open issues are based on additional symmetries, compositeness, or extra dimensions. Our project is related to all of these approaches: in supersymmetric theories an additional symmetry leads to a natural Higgs sector, in a Technicolour theory the Higgs emerges as a composite state of a new strong dynamics, and in gauge theories with extra dimensions the Higgs is protected by the gauge principle of the higher dimensional theory. The extra dimensions are not directly observable since they are compactified and therefore there is only a constraint dynamic in the corresponding direction. These approaches are based on new theories with strong interactions. Interesting phenomena, like the phase transitions and the bound state spectrum of these theories, are so far not well understood. This project is concerned with these challenging questions.

Supersymmetry is one interesting new concept for the extensions of the standard model, favored by many different theoretical considerations. This symmetry connects two completely different classes of particles, fermions and bosons. The former class consists of the matter particles, like the electrons, the latter of the mediators of the forces, the gauge bosons, and the Higgs particle. The standard model is extended by the supersymmetric partner particles as shown in Figure 1. The extended theory includes a supersymmetric version of the strong interactions of the standard model. The supersymmetric gluodynamics describes the interactions of gluons and their superpartners the gluinos.

Results and Methods

Strong interactions and the lattice
Quantum chromodynamics (QCD) describes the strong interactions responsible for the binding of quarks into proton and neutrons and for the formation of nuclear matter. The mechanisms and effects in strongly interacting theories are so far not analytically calculable. For example, the analytical understanding of the confinement mechanism is regarded as one of the most challenging problems, the "Millennium Prize Problems", by the Clay Mathematics Institute. This mechanism is the fundamental explanation of why color charges are not as directly observable as their electromagnetic counterparts: particles interacting by this force are strongly bound into neutral bound states and can not be isolated. In our project we also aim to offer a new perspective for the understanding of this mechanism based on the investigations of new strongly interacting theories.

So far the only method for an investigation of strongly interacting theories are numerical simulations on a discretised space-time, the lattice. Efficient algorithms have been recently developed for the simulations of QCD leading to a remarkable agreement between numerical data and experiment. In this project these methods are applied in the investigations of the new strongly interacting theories beyond the standard model of particle physics.

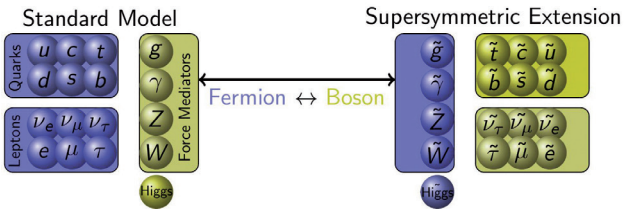

Figure 1: Simplified picture of the supersymmetric extension of the standard model: the bosonic particles get additional fermonic partners and vice versa. In the standard model the interactions between the quarks and leptons are mediated by the gauge bosons: the strong interactions by the gluons, the electromagnetism by the photons, and the weak interactions by the W and Z bosons.

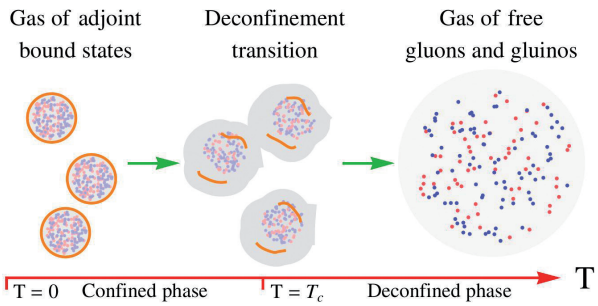

Figure 2: Illustration of the confinement transition in strongly interacting theories. The bound states at small temperatures are dissolved into a free gas of fundamental particles at high temperatures.

The spectrum of bound states
At low energies, the observable particles are bound states of the fundamental gluons and fermionic particles. In supersymmetric theories these should form multiplets of fermions and bosons with the same mass. We were able to show, for the first time, the expected degeneracy of the non perturbative particle spectrum [1].

In Technicolour theories the Higgs particle emerges as a bound state of the strong dynamics and additional bound states correspond to new resonances to be observed in experiments. We have been able to observe patterns in the particle spectrum of these theories, which are, as required for a viable Technicolour theory, considerably different from QCD.

The phase transitions
At low temperatures strongly interacting theories are confined, but at high temperatures they behave like a free gas of gluons and fermions, see Figure 2. The deconfinement transition separates these two phases. In addition, the fermions condense at low temperatures. The analysis of these transitions is important to understand the phenomenological implications of the extensions of the standard model. We have done intense numerical investigations to measure the transition temperatures.

At finite temperature the different nature of fermions and bosons becomes apparent: they obey different statistics and supersymmetry gets broken. In quantum field theory, non-zero temperature is realised by a compactified dimension with different boundary conditions for fermions and bosons. If the same boundary conditions are instead applied for fermions and bosons, their contributions cancel out and the deconfinement transition disappears, as shown in Figure 3. We were able to verify

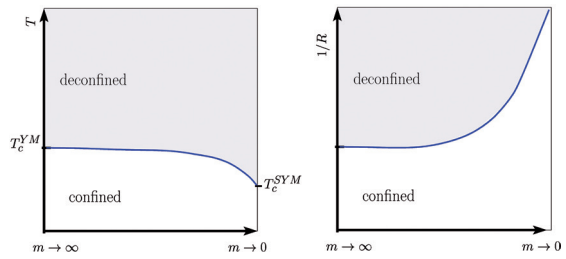

Figure 3: Theoretical predictions of the different phases in supersymmetric Yang-Mills theory at finite temperature T and at finite compactification radius R. The numerical confirmation of this effect can be found in [3,4].

for the first time this effect by numerical simulations [2-4]. It is observed, for example, in the vanishing derivative of the partition function for periodic boundary conditions, see Figure 4. The disappearance of the deconfinement phase transition allows to formulate an analytically tractable asymptotic expansion were strong interactions can be understood in their non-pertubative nature. The possibility for a tuning of the phase transition is a particular feature of the supersymmetric theory.

Typical requirements for the lattice simulations
Lattice simulations need a large amount of computing resources. In particular, the consideration of new theories requires a careful tuning of the parameters and advanced methods are needed for the measurements of the relevant observables. The typical requirements are quite different. The zero temperature investigations usually require the consideration of rather large lattices, while for the investigations of the phase transitions smaller simulations, but a scan of a large parameter range, is necessary. Further a large amount of configurations is required to sample correctly all possible phases, typically at least 5000 configurations are produced for simulations near a phase transition. Our project has around 20 Mio. core-hours. The current zero temperature runs use ~3072 cores and the runs at finite temperature ~54 cores. Our C++ code uses an efficient hybrid MPI + openMP parallelisation with vectorisation in the most time consuming parts.

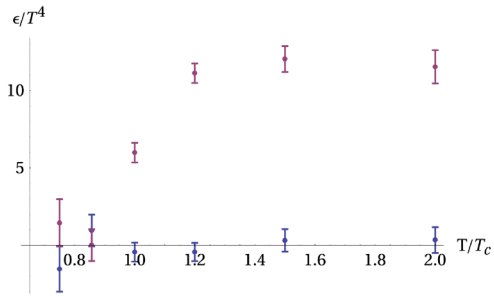

Figure 4: The quantity shown in this plot represents a derivative of the partition function for thermal and periodic boundary conditions [2]. In case of periodic boundary conditions it is the sum of the differences between bosonic and fermionic energy states. It is consistent with zero which proves the cancellation between bosonic and fermionic contributions as required by supersymmetry.

On-going Research / Outlook

SuperMUC is one of the most convenient high performance machines for our project since it offers a high performance and flexibility regarding different applications. This is of particular importance for investigations of new theories, where on the one hand the parameters and systematic uncertainties have to be estimated in smaller simulations and on the other hand a large computational performance is needed for the estimations of the scale at zero temperature. Our project is just the first investigation of the new physics beyond the standard model of particle physics and we hope to proceed with our studies towards more involved Technicolour candidates, supersymmetric QCD, and extended supersymmetry.

References and Links

[1] G. Bergner, P. Giudice, I. Montvay, G. Münster, S. Piemonte, JHEP **1603** (2016) 080.
[2] G. Bergner, P. Giudice, G. Münster, S. Piemonte, arXiv: 1510.05926 [hep-lat].
[3] G. Bergner, S. Piemonte, JHEP **1412** (2014) 133.
[4] G. Bergner, P. Giudice, G. Münster, S. Piemonte, D. Sandbrink, JHEP **1411** (2014) 049.

Lattice QCD on fine lattices

Research Institution
Neumann Institute for Computing, DESY

Principal Investigator
Stefan Schaefer

Researchers
M. Bruno, D. Djukanovic, G.P. Engel, A. Francis, G. Herdoiza, H. Horch, P. Korcyl, T. Korzec, M. Papinutto, E.E. Scholz, J. Simeth, H. Simma, W. Söldner

Project Partners
Helmholtz Institute Mainz, Università Milano-Bicocca, INFN, Universidad Autónoma de Madrid, Universität Mainz, Humboldt Universität zu Berlin, Sapienza Università di Roma, Universität Regensburg

SuperMUC Project ID: pr85we (PRACE project)

Introduction

Quantum Chromodynamics is the theory of quarks bound together by gluons to form the hadrons like the protons, neutrons and pions which we observe in Nature. It is a prime example of a strongly interacting theory such that specialized numerical calculations need to be performed to compute quantities of interest: masses of hadrons, matrix elements needed in the interpretation of high energy physics data.

These computations consist of performing a high-dimensional integral by Monte Carlo methods which require the generation of field configurations on which the observables are then evaluated. Since this is the computationally most expensive part of a lattice computation, these lattices, once generated, can be used in many physics projects.

In this project we started the large scale generation of such a set of gauge field configurations, at parameters which are chosen such that many physics questions can be answered during the years to come with high accuracy.

In particular we include the full dynamical effect of the three lightest quarks: called the up, down, and strange. In our choice of parameters we pay particular attention to the *continuum limit* of our results: in a lattice computation space-time is discretized on a lattice with lattice constant a. Each computation needs to be repeated on finer and finer lattices and the result then needs to be extrapolated to zero lattice spacing.

This process is very costly. In fixed volume the number of lattice points rises with the forth power of the inverse lattice spacing 1/a, and so does the computational effort. Furthermore, algorithms face the problem of critical slowing down, i.e. one needs to increase the number of steps of the algorithm as the continuum is approached. Here it is expected to add an additional $1/a^2$ to the scaling. In the traditional setup a significantly worse behavior has been observed, however, which could be traced back to the formation of topological sectors as the continuum is approached. In the past this problem has essentially prevented reliable simulations below lattice spacings of *a=0.05 fm*. For the first time in a QCD simulation we have used a recently proposed solution to this problem, a set-up with open boundary conditions in time such that no sector formation even in the continuum occurs [1].

Results and Methods

The algorithmic details of the publicly available **openQCD code** are described in detail in Ref.[2] and the parameters chosen in the present simulation in Ref.[3]. Most of the computational effort in these computations goes into the solution of the Dirac equation, a sparse system of linear equations. The Dirac matrix had up to $6 \cdot 10^8$ rows and collumns in our case. A key feature of the openQCD code is the locally deflated (two grid) solver for the this equation. It bascially eliminates the increase of the solver iterations as the condition number of the Dirac matrix rises.

Furthermore, we used the so-called twisted mass reweighting, which limits the condition number of the Dirac matrix and stabilizes the field generation algorithm. The combination of these improvements proposed during the last decade has been employed for the first time during these simulations. While they made the present simulations possible, an important outcome of

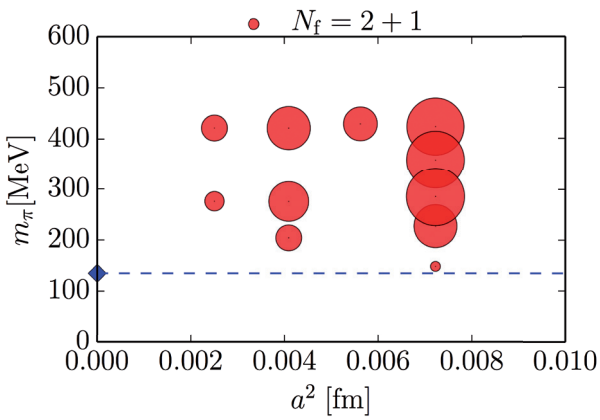

Figure 1: Parameters of the generated ensembles. The size of the circles indicates the statistics.

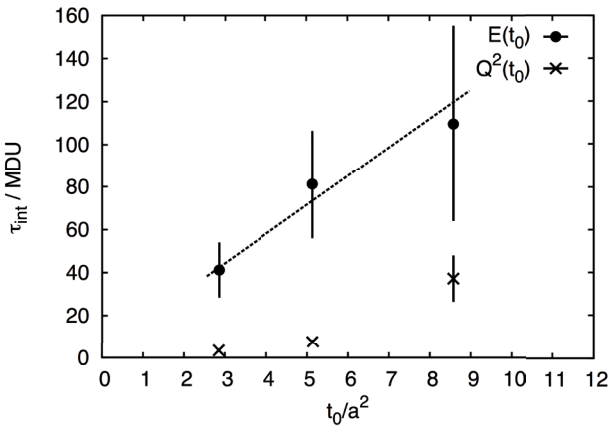

Figure 2: Simulation time needed to generate an independent field configuration as a function of the square of the lattice spacing for two particularly slow observables: the smoothed action density E and the square of the topological charge Q^2. No exceptional slowing down is observed.

this project was also the prove that these techniques work in this challenging setting.

The openQCD code is written in C and is parallelized using standard MPI. This allowed us to run on SuperMUC with jobs from 1024 to 4096 cores and generate a new field configuration every 10min to 30min.

Lattice field configurations at three different values of the lattice spacing ranging from *a=0.05 fm* to *a=0.09 fm* have been generated, with a pion mass ranging from *420 MeV* to close the physical point of *140 MeV*. They are on large volumes such that the results do not depend noticeably on their size.

In Figure 1 we give an overview of the ensembles generated. We extensively studied the autocorrelations in the simulations and scaled the statistics with the slowest relaxation times in our Markov Chain, for which we investigated certain smoothed observables known to be particularly sensitive to slow modes.

In Fig. 2 the critical slowing down observed in the simulations is shown. The integrated autocorrelation time, a measure for how many update steps it takes to generate an independent measurement, is shown as a function of the lattice spacing. The theoretical expectation of a^{-2} is confirmed even for the particularly slow quantities given here. Since during production we took into account this increase in the number of update steps, the achieved statistics is always large enough for controlled physics analysis.

First physics results also confirm that the goal of results with percent accuracy could be reached. In Figure 3 we show the approach to the continuum of the pion decay constant f_π in units of the scale t_0. ≈ 0.42 fm. Here we consider the line on which all quark masses are degenerate, with a mass of the pion at 420 MeV.

In the setup we are using, the leading deviation from the continuum result is expected to be proportional to a^{-2}, which is compatible with the observation.

On-going Research / Outlook

These configurations are currently in use in many on-going projects carried out by researchers throughout Europe. In particular this data will serve as an essential input into the computation of the coupling constant of QCD, where some of the simulations are still on-going. But also projects computing the masses of hadrons and investigating their structure are underway as well as activities in the physics of heavy quarks.

As this initial project of gauge field generation has been successful, it is worthwhile to extend the currently available ensembles with further points in parameter space.

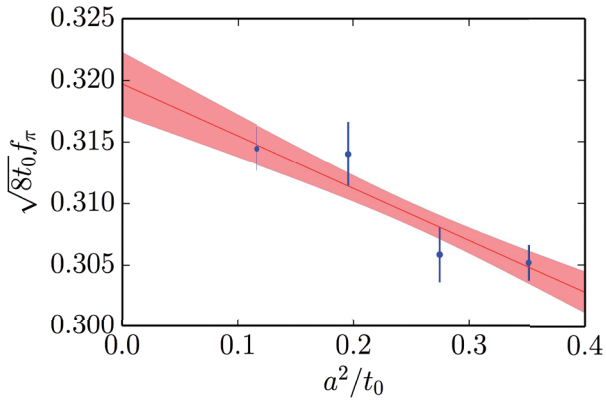

Figure 3: The decay constant of the pion in units of the scale parameter t_0 a function of the lattice spacing a.

These will allow to further study and control systematic effects like the ones introduced by the finite volume, the non-physical quark masses and the finite lattice spacing. In particular, as is obvious from Figure 3, certain compromises have still been made in the region where pion masses and lattice spacing are both small. This is because physical pion masses require larger lattices to keep the effects of the finite volume under control. At light pion masses, a precise control of the continuum extrapolation is therefore difficult, but certainly a main goal of future simulations. To reach this goal, algorithmic developments as well as faster hardware will be needed.

References and Links

[1] M. Lüscher, S. Schaefer, J High Energy Phys, 1107 (2011) 036
[2] M. Lüscher, S. Schaefer, Comput Phys Commun 184 (2013) 519-528
[3] Bruno et al, J High Energy Phys ,1502 (2015) 043

Hadron Physics from Lattice QCD

Research Institution
Institute for Theoretical Physics, Regensburg University

Principal Investigator
Andreas Schäfer

Researchers
G. Bali, V. Braun, S. Collins, B. Gläßle, M. Göckeler, B. Lang, B. Musch, J. Najjar, R. Rödl, D. Richtmann, R. Schiel, J. Simeth, W. Söldner, A. Sternbeck

Project Partners
–

SuperMUC Project ID: pr85xi, pr84qe, pr86te (all Gauss Large Scale projects)

Introduction

This is the combined report for the projects:

pr85xi "Excited state artifacts in calculations of hadron 3-point functions"
pr84qe "The strangeness content of the nucleon"
pr86te "Nucleon and meson matrix elements close to the physical point"

Particle physics experiments at modern high luminosity particle accelerators achieve orders of magnitude higher count rates than what was possible ten or twenty years ago. This extremely large statistics allows to draw far reaching conclusions even from minute signals, provided that these signals are well understood by theory. This is, however, ever more difficult to achieve. Presently, technical and scientific progress in general and experimental progress in particle physics in particular, shows typically an exponential growth rate. For example, data acquisition and analysis are, among many other factor, driven by the development of ever more efficient computers and thus by Moore's law. Theory has to keep up with this development by also achieving an exponential increase in precision, which is only possible using powerful computers. This is true for both types of calculations, analytic ones as, e.g., in quantum field perturbation theory, and purely numerical ones as in Lattice QCD. As stated above such calculations are absolutely indispensable to make best use of the extremely costly large particle physics experiments. Thus, it is economically reasonable to invest a certain percentage of the cost of accelerators and experiments in related theory efforts.

The basic ideas behind Lattice QCD simulations are the following: Because quarks and gluons can never be observed individually but are always "confined" into colorless hadrons, like the proton, all quark-gluon states can be expressed in two different systems of basis states, namely in a quark-gluon basis and the basis of hadron states. The proton, e.g., is an eigenstate of the latter, a specific quark-gluon configuration is part of the former. In the quark-gluon basis a physical hadron, like a proton, is given by an extremely complicated multi-particle wave function containing all effects of quantum fluctuations. This state is so complicated that it is basically impossible to model it with the required precision. However, quantum field theory has a very important fundamental property, which allows to make progress: When the variable "time" is analytically continued to imaginary time (in the sense of square root of minus one) it gets mapped onto thermodynamics and statistics and questions in quantum field theory are transformed into purely statistical problems, which can be solved numerically by Monte Carlo techniques. While there might be more to it, this can be seen as just a mathematical trick. This trick does not only make numerical simulations of quantum field theories possible, but it solves at the same time the problem alluded to above: Within QCD any quark-gluon model which is simple enough that one can use it for practical calculations, fails to describe a real hadron. More precisely a simple quark-gluon state, which can easily be described within QCD corresponds to an infinitely complicated superposition of hadronic states. *However*, if such a superposition is propagated in imaginary time in the right manner all components except the lowest mass physical hadron, e.g. the proton, get exponentially suppressed. Thus the exact many particle wave function of the physical proton is obtained with which one can then calculate all physical quantities one is interested in, with one constraint: Because time has lost its meaning, only time-independent quantities can be obtained. Consequently, Lattice QCD has nearly always to be combined with real time treatments, most prominently perturbative QCD, to obtain physical predictions.

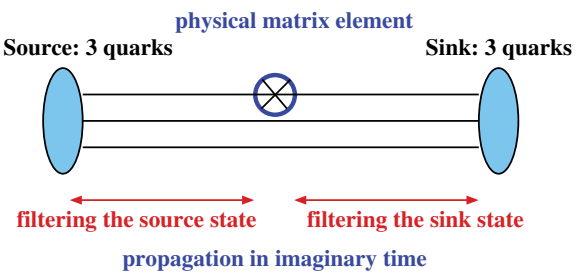

Figure 1: The calculation of, e.g. proton matrix elements on the lattice. At the source and sink a combination of three quark fields with the correct quantum numbers is inserted. The Green's function propagates these in imaginary time, filtering out all excited states. Then, all physical observables which are local in time can be directly evaluated.

The schematic structure of hadron structure lattice calculations is illustrated in Fig. 1. Because source, sink and matrix element define three points in space-time such amplitudes are called "3-point functions".The Greens function on the lattice is just the inverse of a large sparse matrix. This inversion is one of the computationally most expensive tasks in lattice QCD calculations. To propagate a three quark state simply three such Greens functions have to be convoluted.

As we discussed, the unwanted heavy hadronic states, the so-called "excited states" get exponentially suppressed by the propagation in imaginary time. However, in parallel the signal of interest is also reduced, which faces any Lattice QCD practitioner with a serious optimization problem: If he chooses a very strong exponential suppression of unwanted states also the signal is lost, if he chooses a weak suppression, his signal is distorted by excited state artifacts. Therefore, many ingenious ideas were developed to improve this filtering process. One of the most efficient ones is to increase the overlap of source and sink with the hadron of interest by a process called "smearing" .

In practice, to cope with this problem, the careful analysis of many different statistical ensembles using a multitude of analysis strategies, e.g. different smearing prescriptions, is needed to obtain a realistic systematic uncertainties for all analyzed quantities. To perform such analyses for a large selection of hadronic observables was the purpose of this project. Some of these were already discussed in last years report. Therefore, we are concentrating here on some new ones. As the theoretical challenges are very similar for all of them, many applications are equally well suited for such a demonstration.

Nucleon coupling constants

We will present a few data for nucleon coupling constants and for the so called nucleon sigma term which parameterizes the coupling to many forms of the hypothetical dark matter. Dark matter search experiments form presently one of the most active large scale research activities and these couplings are needed to obtain exclusion bounds from experimental results or, in the case of a positive signal, the properties of the dark matter particle(s).

Let us start by showing in Fig.2 some of the ensembles of lattice configurations we used. As stated above, having a large set of ensembles is crucial to achieve a trustworthy extrapolation to the physical point. Fig.3 illustrates for one ensemble and one observable the optimization process with respect to suppression of excited states and statistical accuracy.

The result of a combined analysis for quantities which are experimentally known is then either satisfactory (Fig.4) or still not optimal (Fig.5). This determines the accuracy which one can claim for quantities, for which there is no experimental knowledge like certain moments of Generalized Parton Distribution (GPDs) and the nucleon sigma term.

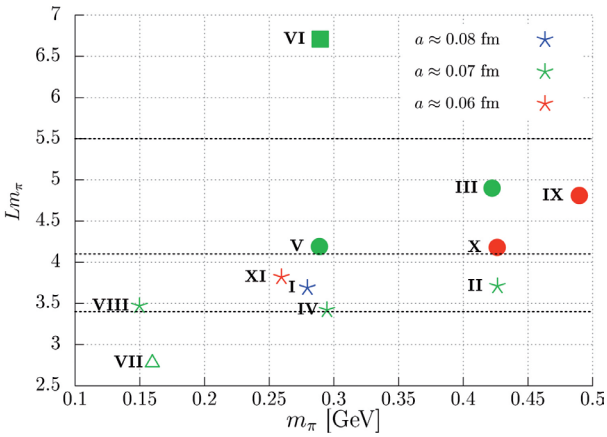

Figure 2: The different Lattice ensembles used in this analysis, plotted according to the pion mass at which the simulations are performed (the physical value is 0.14 GeV) and the ratio of lattice size and pion Compton wave length. The lattice constant of each configuration is indicated by color.

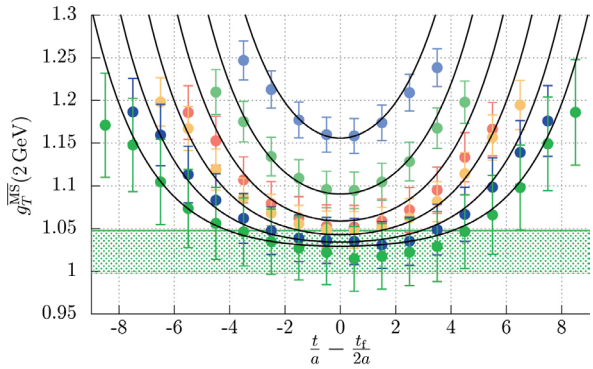

Figure 3: Results for one specific coupling constant, using dimensional renormalization at a momentum scale of 2 GeV. If the excited states are not very strongly sup-pressed (blue points) the error bars are smallest but one encounters severe lattice artifacts. If they are strongly suppressed (lowest data points) the artifacts are much reduced, but the statistical error increases. The green band shows the value and systematic uncertainty we deduced from many such analysis. These data used are from just one ensemble, see [1] for details.

Moments of GPDs

Today the understanding of hadron structure is so advanced that one is interested in dozens of functions already just for the proton. A large class of these can be grouped together in a unified formulation called Generalized Parton Distributions (GPDs). State-of-the-art calculations of these requires to perform very many different analysis and to histogram the results to obtain a reliable estimate for the systematic uncertainties, as illustrated in Fig. 6 and 7. The example presented is a particularly interesting one because the total angular momentum of a quark in a nucleon cannot be measured directly in an experiment (only the spin part can) but is needed to understand how all the different angular momenta in a proton (orbital and spin, quark and gluon) add up to one half.

The nucleon sigma term

The nucleon sigma term parameterizes how many virtual strange-antistrange quark pair quantum fluctuations

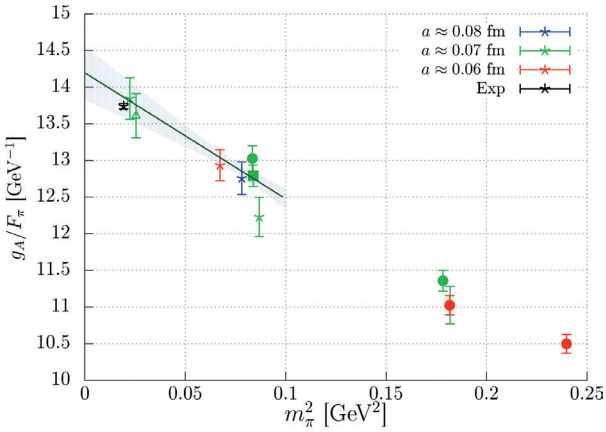

Figure 4: Our results for the ratio of the iso-vector axial coupling constant of the nucleon and the pion decay constant, a case which works out perfectly (the experimental values is shown as black star).

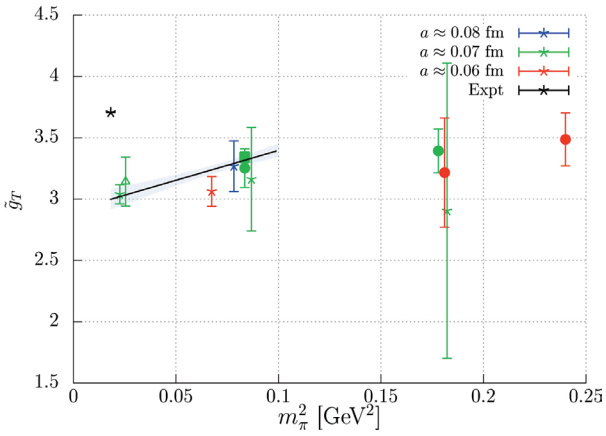

Figure 5: Our results for a specific tensor coupling. This is our worst case, in which we face a twenty percent discrepancy.

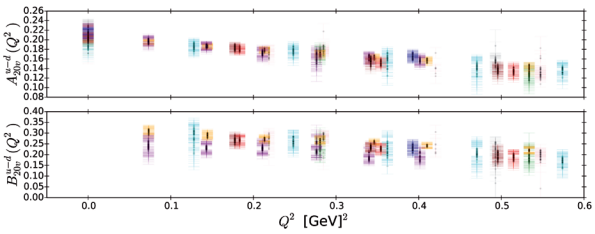

Figure 6: The two relevant generalized form factors are fit for different momenta in many different ways to the lattice data. Each point presents the result of one such a fit.

exist on average in a proton. This quantity is obtained from the evaluation of so called "disconnected" 3-point functions, as illustrated in Fig.6. This name is actually a misnomer cause by the fact that gluon lines are simply not shown. Fig. 1 and 6 represent in reality the coherent sum over infinitely many amplitudes, containing arbitrarily many additional gluon lines.

In recent years the interest in the proton sigma term has greatly increased not only because it parameterizes the coupling of detector material to dark matter but also because it became more and more clear that one encounters a severe conflict between results of effective field theory, which describes the low energy limit of QCD, and Lattice QCD. This discrepancy is a matter of grave concern, because both frameworks are accepted as stringent formulations which therefore have to agree. If they do not this could indicate in the worst case, that one of them has a fundamental, yet unknown problem and thus could invalidate many previous results. We improved on previous lattice calculations be evaluating the disconnected graphs explicitly in [2] but our result, 35.0 (6.1) MeV, agrees perfectly with other lattice determinations, and is much smaller than the value suggested by effective field theory, which in turn might be in better agreement with experimental suggestions from LHC, though this connection can only be made within rough models. Thus the problem of the proton sigma term is rather accentuated than resolved by our results.

A more efficient way of calculating hadron matrix elements ?

Lattice QCD calculations involve very many technical problems most of which we have not even hinted to so far. One of them is that because the quantity one evaluates has to be local in (imaginary) time one cannot calculate directly most quantities one is interested in, like GPDs, but only special integrals of them, called moments.

The standard procedure is to assume some parameterization for, e.g., a GPD, to then calculate the moments in terms of the parameters and to choose the latter such that the lattice results for the moments are reproduced. It would be much better to have a direct method to get the full functional form of, e.g. GPDs, directly from the lattice. A way to do so was proposed in [3]. The idea is basically to calculate on the lattice only what is easy to calculate, namely correlators local in imaginary time

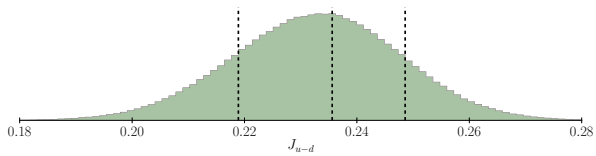

Figure 7: Resulting histogram of thousands of fits for the total angular momentum carried by up quarks minus that carried by down quarks in a nucleon.

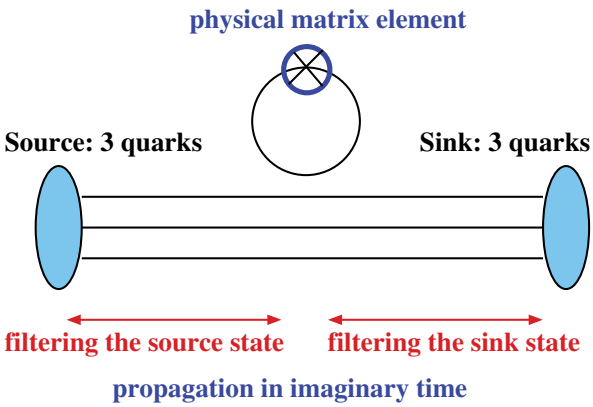

Figure 8: A disconnected contribution which describes, e.g., the strangeness content of the proton.

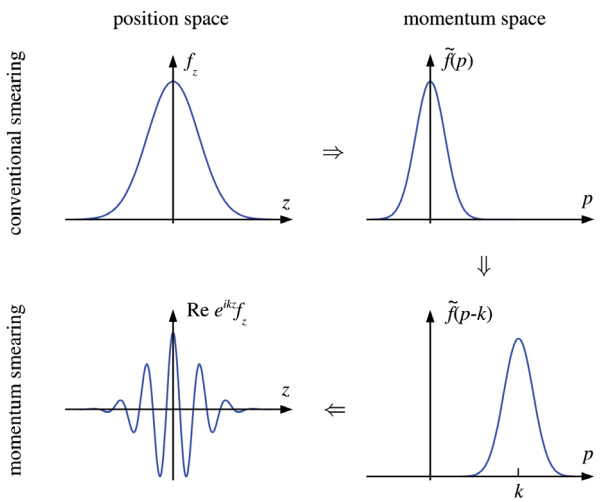

Figure 9: The basic idea to realize large spatial hadron momenta. Conventional smearing is illustrated in the first row. In the new method we use smearing with a super-imposed plane wave.

but non-local in space, and to relate it to what is needed, namely correlators non-local in time and space, by perturbative QCD in the continuum theory. The technical challenge to make this work in practice is to realize large spatial hadron momenta on the lattice. Recently we managed to do so, see Fig. 9, 10 and 11. We hope that this will greatly improve the quality with which hadron properties can be studied on the lattice, in future.

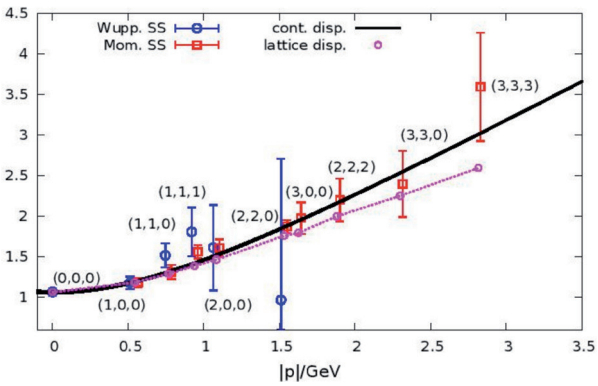

Figure 10: The resulting dispersion relation for the proton. With the conventional method (blue data points) one cannot achieve momenta larger than 1 GeV, which is insufficient . With the new method we can easily realize 3 GeV, which should be sufficient. (The momentum squared must be significantly large than 1 GeV squared.)

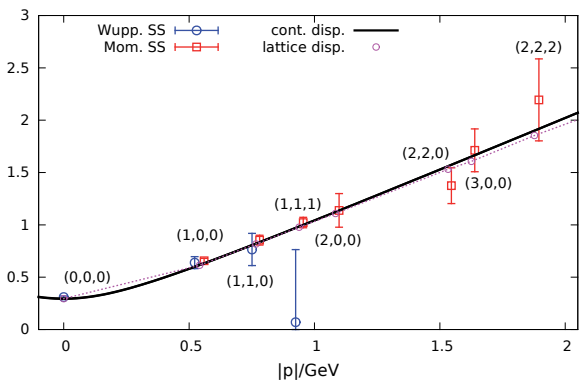

Figure 11: The same as Fig. 8 but now for the pion.

The upshot of our investigations is, that once you have really large statistics, as we have, one encounter a lot of problems which before were kind of hidden by the error bars, but that up to now all of these could be solved.

References and Links

[1] G.S. Bali et al., Phys.Rev. D91 (2015) 054501
[2] G.S. Bali et al., arXiv:1603.00827
[3] V. Braun and D. Müller, Eur. Phys. J. C55 (2008) 349
X. Ji, Phys. Rev. Lett. 110 (2013) 262002
[4] G. S. Bali et al., arXiv 1602.05525

Kaon semi-leptonic form factor

Research Institution
Institut für Theoretische Physik, Universität Regensburg

Principal Investigator
Enno E. Scholz

Researchers
Gunnar Bali, Vladimir Braun, Alessio Burello, Benjamin Gläßle, Rudolf Rödl, Wolfgang Söldner

Project Partners
–

SuperMUC Project ID: pr89ti (PRACE project)

Introduction

In the Standard Model (SM) of particle physics the interactions between the elementary particles are mediated by the strong and electro-weak forces. The strong forces are described by the theory of Quantum Chromodynamics (QCD). QCD forces are mediated by gluons between the quarks, which are the fundamental building blocks of all hadrons, especially the proton and the neutron, and therefore of most matter around us. Six different quark flavors grouped in 3 generations have been observed in Nature (up and down, charm and strange, top and bottom), which differ, e.g., in their masses but also their electric charges. The electro-weak forces are described by another gauge field theory, the electro-weak theory of which for example Quantum Electrodynamics mediating the electric forces via photons is a subgroup. Since the quarks carry electric charges, the electro-weak forces act via photons between different quarks. But there are also the W- and Z-bosons as further force carriers in the electro-weak theory being responsible for more interactions between the elementary particles, also including the leptons (electron, muon, tau and their respective neutrinos).

The eigenstates of the two theories, QCD and electro-weak theory, do not coincide. The electro-weak mixing matrix, commonly referred to as the Cabibbo-Kobayashi-Maskawa (CKM) matrix, describes the mixing between the mass and electro-weak eigenstates of the different quark flavors in the SM. Originally formulated for two generations of quark flavors in 1963 by Cabibbo, later Kobayashi and Maskawa realized that in accordance with the observed violation of charge-parity (CP) symmetry at least a third generation of quark flavors had to be added. The existence of such a third generation was shown by the discovery of the bottom-quark in 1977. Eventually, the completeness of the third generation was verified experimentally with the discovery of the top-quark at Fermilab's Tevatron accelerator in 1995. Over time, the elements of the CKM-matrix have been measured and extracted with increasing precision. They are particularly interesting, because unitarity violations of the matrix may give hints for physics beyond the SM. Particle collider experiments like the Large Hadron Collider (LHC) at CERN or the Belle2 experiments presently are or will in the near future produce improved measurements to determine these matrix elements.

In such experiments decay rates for certain processes involving the elementary particles or bound states of those, like mesons or hadrons, are measured. To extract the CKM-matrix elements from such measurements, one has to know the different contributions of the QCD and electro-weak interactions in these processes. Due to their nature, unlike the electro-weak contributions which can be (at least in principle) calculated analytically order-by-order in a perturbative expansion to the desired precision, for the strong forces such a perturbative expansion is not possible. But it is possible to simulate QCD numerically on a discretized, finite space-time lattice. Basically, one first generates a statistical ensemble of gauge configurations, which sample the correct distribution of the gluon fields according to the QCD action. Then one "measures" the interactions between quarks and the hadronic operators of interest by calculating propagator functions on the generated gauge fields and combine them with certain operators. Numerically this requires the inversion of large sparse matrices to obtain the propagators between certain quarks. These lattice-QCD simulations have matured over the last decades such that by today reliable results from simulations, e.g., for the mass spectrum of the hadrons and certain hadronic matrix elements are available.

Currently, the biggest uncertainty in the first row unitarity relation $|V_{ud}|^2+|V_{us}|^2+|V_{ub}|^2=1$ originates from the CKM-matrix element $|V_{us}|$, which is our motivation to address it in this project. Given the available experimental data, the most precise value for this matrix element can be obtained from the semi-leptonic kaon decay $K\rightarrow\pi l\nu$ (Kl3). From those measurements, one obtains the combination of the matrix element $|V_{us}|$ and the vector form factor $f_+^{K\pi}(q^2)$ at zero momentum transfer $q^2=0$. We will extract the form factor for the Kl3-decay from lattice measurements of the matrix element of the weak vector current $V_\mu=\bar{s}\gamma_\mu u$ between a kaon and pion state with momenta p_i and p_f, respectively. The vector current matrix element is related to the semi-leptonic form factors at momentum transfer $q = p_i - p_f$ via

$$\langle\pi(p_f)|V_\mu|K(p_i)\rangle = \\ = f_+^{K\pi}(q^2)\,(p_i+p_f)_\mu + f_-^{K\pi}(q^2)\,(p_i-p_f)_\mu .$$

High precision results for the vector form factor can be obtained from lattice simulations of QCD since leading lattice artifacts are proportional to the kaon-pion mass difference $m_K - m_\pi$ and apply only to the deviation of the form factor from unity in the limit $m_K = m_\pi$. As one is ultimately interested in form factors at zero momentum transfer, it is necessary to reach zero or near zero momentum transfer in the lattice simulations as well. As outlined below, in this study this is achieved by employing so called twisted boundary conditions.

Results

For the analysis Stout-link Non-Perturbative Clover (SLiNC)-fermion gauge configurations with $N_f = 2+1$ flavors are used, which were generated by the QCDSF-collaboration [1, 2] and are available at the ILDG. We used four ensembles with pion masses in the range of approximately 210 to 340 MeV, see Table 1. These ensembles start from the symmetric point at $\varkappa_l = \varkappa_s$ and in all ensembles the average quark mass $(2m_l+m_s)/3$ has been kept constant. For each of the three ensembles at lattice sizes of $32^3 \times 64$ we had around 1800 to 2000 configurations available, while for the ensemble with the lightest pion mass at lattice size of $48^3 \times 96$ only 600 configurations were available.

Table 1: SLiNC-ensembles at $\beta = 5.50$

$\varkappa_l$	$\varkappa_s$	$L^3 \times T$
0.121040	0.120620	$32^3 \times 64$
0.121095	0.120512	$32^3 \times 64$
0.121145	0.120413	$32^3 \times 64$
0.121166	0.120371	$48^3 \times 96$

To measure the form factor at zero or near zero momentum transfer, we induce additional momentum either to the pion or the kaon state by utilizing twisted boundary conditions in the valence quarks, meaning the boundary conditions are only applied during the calculation of the quark propagators (valence quarks) and were not applied during the generation of the gauge configurations. Applying twist angles θ_k to the spatial directions k to the quark fields

$$\psi(x_k+L)=e^{i\theta_k}\psi(x_k)$$

induces a momentum $|\vec{\theta}|/L$. Here we follow the strategy outlined by the RBC/UKQCD-Coll. [3, 4] for achieving zero momentum by implementing either

$$|\vec{\theta}_K|=L\sqrt{\left(\frac{m_K^2+m_\pi^2}{2m_\pi}\right)^2-m_K^2}\,, \quad \vec{\theta}_\pi=0$$

or

$$|\vec{\theta}_\pi|=L\sqrt{\left(\frac{m_K^2+m_\pi^2}{2m_K}\right)^2-m_\pi^2}\,, \quad \vec{\theta}_K=0\,.$$

The necessary two- and three-point correlation functions are then calculated with the twists applied accordingly. For the calculation of the three-point correlator the sequential source technique is used. The computations are performed using the publicly available Chroma software package for lattice QCD [5]. The inverter we used and optimized for running on the SuperMUC architecture is a multi-grid domain-decomposition solver. Typically, for the lattice sizes used in this project, we run on partition sizes of 1024 nodes, but scaling of our inverter code would have allowed the use of larger partition sizes as well.

We first started by estimating the pion- and kaon masses from two-point correlators using untwisted fermion fields for each ensemble. From these masses, we determined the twist angles for zero momentum transfer according to the above formulae. In the following step the two- and three-point applying the twist angles were generated and stored to disk. Here we also generated data for additional twist angles leading to small but non-zero momentum transfer. In that way we will be able to interpolate the form factors at several q^2 to zero momentum transfer. The form factor itself can be obtained from suitable ratios of the correlation functions, cf. [3,4].

We acknowledge the generation of the gauge configurations used in this work by the QCDSF-Collaboration and resources of the SFB-TRR 55 "Hadron Physics from Lattice QCD" [1,2]. The principal investigator is supported by the DFG SFB-TRR 55 and EU grants PITN-GA-2009-238353 (ITN STRONGnet) and PIRG07-GA-2010-268367.

References and Links

[1] W. Bietenholz et al. (QCDSF-Coll.), Phys. Lett. B690 (2010) 436, arXiv:1003.1114 [hep-lat].
[2] W. Bietenholz et al. (QCDSF-Coll.), Phys. Rev. D84 (2011) 054509, arXiv:1102.5300 [hep-lat].
[3] P.A. Boyle et al. (RBC-UKQCD-Coll.), J. High Energy Phys. 05 (2007) 016, arXiv:hep-lat/0703005.
[4] P.A. Boyle et al. (RBC-UKQCD-Coll.), Eur. Phys. J. C69 (2010) 159, arXiv:1004.0886[hep-lat].
[5] R.G. Edwards (LHPC-Coll.), B. Joo (UKQCD-Coll.), The chroma software system for lattice QCD, Nucl. Phys. Proc. Suppl. B140 (2005) 832.

Engineering and Computational Fluid Dynamics

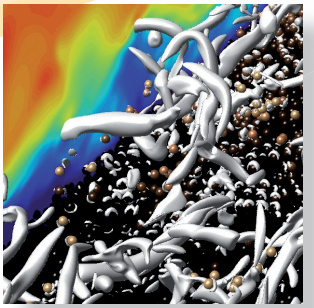

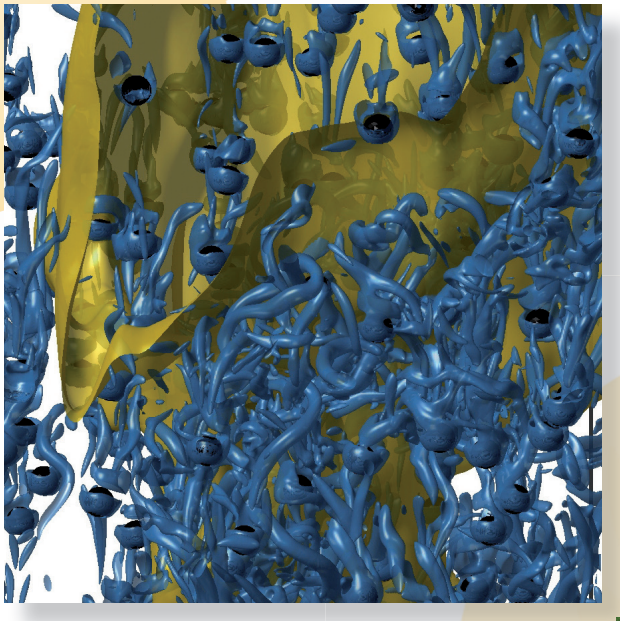

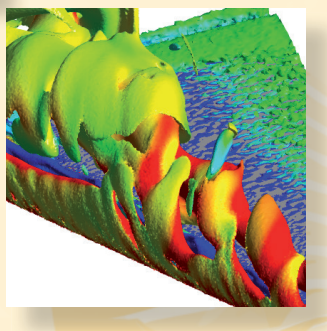

Buoyant-convectively driven heat and gas exchange

Research Institution
Institute for Hydromechanics, Karlsruhe Institute of Technology (KIT), Germany

Principal Investigator
H. Herlina

Researchers
H. Herlina, J.G. Wissink

Project Partners
Dept. Mechanical, Aerospace and Civil Engineering, Brunel University London, UK

SuperMUC Project ID: pr28ca

Introduction

The study of heat and atmospheric gas exchange across the air-water interface received increasing interest in the last decades because of its important role in the global heat and green-house gas budget. Most estimations of the heat / gas flux only take wind-shear into account, which often results in an over- or under-prediction of the actual flux, particularly during low wind conditions. Equatorial ocean regions and sheltered lakes are examples of regions with typically low wind speed. Recent studies have reported that in such conditions buoyancy-driven flux cannot be neglected and may even dominate. Hence, to improve the accuracy of the prediction of heat and gas fluxes an in-depth understanding of the buoyant-convective mixing process in deep waters is required. To address this, a series of direct numerical simulations (DNS) of interfacial gas transfer driven by a surface-cooling-induced buoyant convective instability were performed [1, 2].

Results and Methods

The DNS builds on earlier experiments performed at KIT. Compared to the size of the experimental domain (50x50x42cm³), only a small part adjacent to the water surface was modeled (Table 1). As in the experiments, the buoyant instability is caused by sudden cooling of the water surface. The initial conditions for the velocity and concentration fields were set to zero and the temperature field to T=1 (normalized so that T=0 and T=1 are the coldest temperature at the surface and the initial temperature of the warmer bulk water, respectively).
As a result, a thin thermal boundary layer of cool water formed adjacent to the air-water interface, with an even thinner layer of gas-saturated water at the top. At t=9.6s small random disturbances were added to the temperature field to trigger the instability resulting in the formation of thin falling sheets of cold water. Consequently, warmer water from the bulk starts to move upward resulting in the formation of convection cells. At intersections of three or more convection cells –where the added quantity of cold water falling down resulted in an increase in the local sink velocity – mushroom like plumes are formed (Figure 1). These cold sinking plumes transport high gas-saturated fluid deep into the bulk as shown in Figure 2.

At the grid plane immediately beneath the surface a typical net-like pattern is visible which can be seen as the footprint of the convection cells in the water.

Table 1: Overview of the simulations.

Case	Sc	Base mesh (grid points)	Domain (cm³)	Ra_L
BC3	20; 500	800x800x512	10x10x10	34000
BC4	20; 500	800x800x352	10x10x5	34000
BC5	20; 500	400x400x256	5x5x5	23000

Snapshots of the net-like pattern show that in time the convection cells tend to grow. The high resolution of the present DNS allowed a quantitative prediction of the average convection cell-size near the surface L_C. L_C, together with the velocity fluctuations at the surface u_{surf} were found to be relevant choices for the length and velocity scales, respectively, to be used in the large-eddy model for the estimation of the transfer velocity K_L.

As illustrated by the isolines for the scalar concentration with Schmidt number (Sc) = 500 (relevant for the transfer of oxygen in water) in Figure 2, the transfer of high Sc atmospheric gaseous into water is characterized by an extremely thin gas boundary layer and instantaneous occurrence of steep gradients in other regions. To accurate-

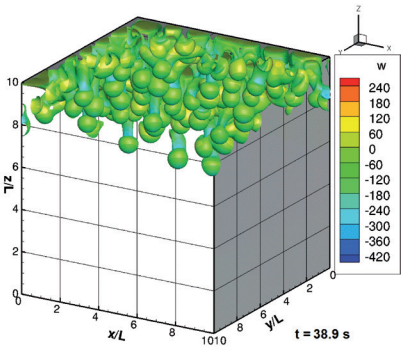

Figure 1: Isosurface of T=0.75 flooded with the vertical velocity that is normalized by U=ϰ / L=0.0158 mm/s, where ϰ is the thermal diffusivity of water at 298.15K.

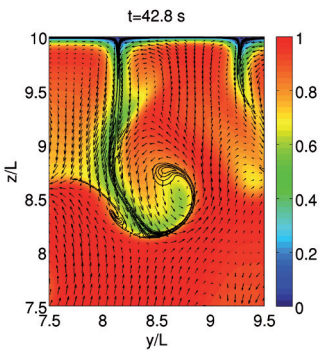

Figure 2: Correlation between temperature field (colour) and the Sc=500 scalar (isolines : C=0.1 to 0.6).

ly simulate this problem, we used a specifically designed numerical scheme capable of resolving the scalar field without having under- and/or over-shoots of the scalar quantity. In this code, the 5th-order WENO-scheme [3] for scalar convection, combined with a 4th-order central method for scalar diffusion, was implemented on a staggered and stretched mesh. In the simulations, up to five convection-diffusion equations for the scalar and temperature were solved simultaneously (further details in [2, 4]).

To optimize the load balancing, in each simulation the computational domain was divided into a number of blocks of equal size, each of which was assigned to its own processing core (Table 2). The standard Message-Passing-Interface (MPI) was used for the exchange of data between processes. MPI-I/O routines were used for the input and output of the large amount of data that need to be periodically stored to be able to continue runs. Memory usage was minimized by allowing certain scalars to be solved on the coarse mesh while others were solved on the refined mesh.

As implied above, to resolve high Sc transport processes, a very fine mesh for the scalar is needed. Because of the special numerical scheme, the usage of a refinement factor of 3 for the Sc=500 scalar field was found to be sufficient in the present runs. The required computing resources were still significant but feasible (cf. Table 2). Each simulation typically generated twelve files. The overall storage needed (including the on-going large scale simulation described below) is about 19TB.

Table 2: Refined mesh size and used resources.

Case	Refined scalar mesh	Simulated time (s)	Cores	CPUh
BC3	2400x2400x1536	130	16,384	6.9×10^6
BC4	2400x2400x1056	73	11,264	2.9×10^6
BC5	1200x1200x768	67	4,096	0.4×10^6

On-going Research / Outlook

Running on SuperMUC allowed us to perform very large scale, highly accurate numerical simulations of gas transfer across the air-water interface at high Schmidt numbers. The installation of phase 2 made it possible to further increase the resolution in our simulations so that we are currently able to run simulations at higher turbulent Reynolds numbers (R_T). These DNS calculations, to our knowledge, are the first that allow a detailed study of the gas transfer driven by high intensity turbulence diffusing from below (case ITfs1 : R_T up to 2000). Such highly turbulent simulations are crucial to study the role of small vortical structures in the gas transfer mechanism.

To be able to run our simulation on such a fine mesh we had to exchange the traditional I/O (using one result file per processing core) by MPI I/O. Also, we needed to minimize the memory requirements by using dynamic memory allocation for all arrays and (as mentioned above) allowing certain scalars to be calculated on the coarse (base) mesh while only a selected few are calcu-

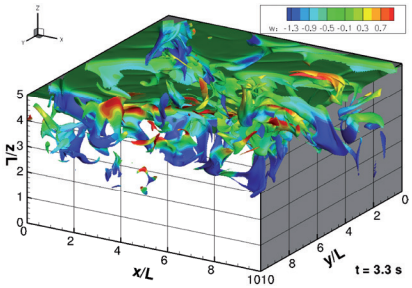

Figure 3: Simulation ITfs1 : Isosurface of 25% concentration saturation of Sc=20 and R_T=2000.

lated on the fine mesh. Restrictions that we encountered were mainly related to available memory, both internal memory and disk space. While solving the memory problem we cooperated with Mohamed Shaheen from IBM to help solve some issues with intel MPI that occurred when running large jobs.

A snapshot of the concentration iso-contours from this large ITfs1 simulation (Figure 3) clearly shows the dynamics of the eddies impinging on the surface affecting the gas transfer. For this particular case 20,992 cores were employed, a total of 5×10^6 CPUh were used to achieve a well-developed turbulent flow-field and another 11×10^6 CPUh for resolving the interfacial gas transfer (Sc=20 and 500). It should be noted that at some occasions SuperMUC had experienced stability problems (e.g. hardware failures) that cost us a non-negligible amount of computing time. At present, the results already reached a statistically quasi steady behavior. Approximately another 2.5×10^6 CPUh would be needed for recording a series of snapshots of the 3D flow and concentration fields that allow us to explore features (e.g. 3D vortical structures) important to gas transfer.

In the simulations described above, idealized conditions without surface contamination (free-slip boundary condition) were used. In reality, water surfaces are rarely clean. It is therefore our aim to perform further DNS computations of gas transfer by adding complexities as occurring in nature such as surface contaminations and also the combined effect of buoyancy-induced and isotropic-turbulence-induced flow.

References and Links

[1] http://www.ifh.kit.edu/26_1776.php
[2] Wissink, J.G. and Herlina, H. 2016. J. Fluid Mech. 787, 508-540.
[3] Liu X. D., Osher S., Chan T. 1994. J. Comput. Phys. 115, 200–212.
[4] Kubrak B., Herlina H., Greve F., Wissink J.G. 2013. J. Comput. Phys. 240, 158-173.

Investigation of Vehicle Wheels Ventilation Moment using DoE-based Computations and Experiments

Research Institution
Institute of Aerodynamics and Fluid Mechanics

Principal Investigator
Lu Miao

Researchers
Thomas Indinger

Project Partners
–

SuperMUC Project ID: pr42re

Introduction

With constantly growing fuel prices and toughening of environmental legislation, the vehicle industry is struggling to reduce fuel consumption and decrease emission levels for the new and existing vehicles. One of the ways to achieve this goal is to improve aerodynamic performance by decreasing aerodynamic resistance. Rotating wheels together with the wheelhouses can produce up to 25% of the total aerodynamic drag. Furthermore, there are power losses associated with the resistance moments acting on the wheels; originating from the relative movement of the wheels in the air. In order to better understand those losses a closer look at forces acting on the rotating wheel is required.

The improvement of vehicle aerodynamics requires the tools of aerodynamics development to perform at ever-increasing levels of accuracy. Computational Fluid Dynamic (CFD) is very important due to the complexity of problems and accuracy required. Requirements of CFD are high process integration to keep pace with vehicle development cycle and the accuracy of results must be reliable, especially where no experiments are available. The Open-Source CFD is chosen in this project, because the commercial environment for CFD codes are truly viable for productive use with limited insight or the black-box approach. The simulations of full car in a wind tunnel have been studied with rotating wheels and different parameters. And the forces and moments were calculated through an user-defined utility in OpenFOAM.

The simulations were carried on on the super computer of Munich (supermuc) in Leibniz Supercomputing Center.

Methods

The DrivAer model was selected for the project. The DrivAer model is developed at TU München in cooperation with the automotive industry companies BMW and Audi. The experiments reported in this project were executed in the Wind Tunnel A of the Institute of Aerodynamics and Fluid Mechanics at Technical University Munich, a 1:2.5 model wind tunnel with a blockage ratio 8%. The test section is 4.8m long, the cross section of nozzle exit is 4.32m2. Vortex generators are installed at the nozzle exit to reduce the pressure fluctuations induced by the developing shear layers. The maximum wind speed is 65m/s. Four different Setup experiments were studied, which would be used to validate the simulation results.

For the numerical investigation, the open source code OpenFOAM was chosen. The customizability of open-source software, along with the absence of licensing restrictions, is increasing its presence in the engineering and research environments [1]; the user has the choice of technology provider. Full transparency of technology permits complete analysis and solves problems, which is very flexible for calculating the ventilation moment of the rotating wheels. The model is 1:1, according to the experiment. To allow for a sufficient development of the turbulent characteristics of the flow, the model was positioned in the detailed wind tunnel model, with

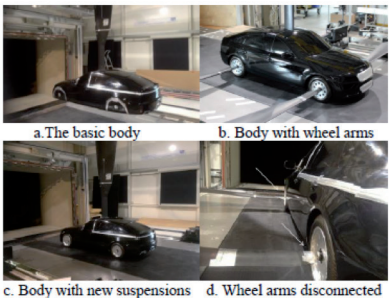

Figure 1: Four different Setups in wind tunnel

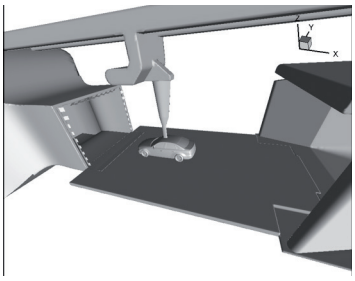

Figure 2: Simulation Model in wind tunnel

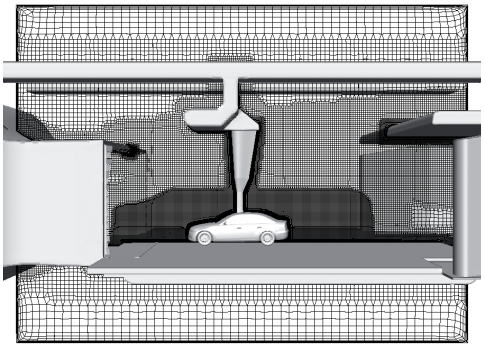

Figure 3: Volume mesh with refinement regions

moving belt, boundary layer scoop, which could make the most agreement between experiments and simulations.

The meshing was performed with the commercial software Spider. The meshes produced by Spider consist prevalently of hexaheder cells, some prisms, pyramids and tetraheders. To minimize the computational resources, the refinement boxes were adjusted to the body geometry and also the the shear layer and scoop and wheel and wheelhouse were also refined. The final mesh size depending on the case had around 95 million. The volume mesh qualities were optimized in OpenFOAM with tools to manipulate a mesh,see Fig3.

The pressure-velocity coupling in the present work is realized using the SIMPLE algorithm implemented in simpleFoam. For the Reynolds Averaged Navier Stokes Equations (RANS) simulations, the k-ω-SST (Shear-Stress-Transport) model following Menter [2] was chosen. The ω-equation has significant advantages near adverse pressure gradient flows, leading to improved wall shear stress. The SST model combines the k-ω model near the wall and the k-ε model away from the wall as a unified two-equation turbulence model. It was developed for external aerodynamic flow simulation and has shown to be superior to other two-equation models in view of separation, lift and drag prediction.

The simulations were conducted at the velocity u = 41.2 m/s with rotating wheels, the turbulent intensity is 0.4% and the turbulent length scale is 1.5mm, which are in line with the inlet conditions during the wind tunnel experiment. To prevent backflow into the domain, the velocity boundary condition at the outlet was set to an inletOutlet condition. A no-slip boundary condition was enforced at the walls and a symmetry boundary condition was chosen for the wind tunnel buffer. The rotation of the wheels is approximated by imposing a rotatingWallVelocity boundary condition. The chosen wall function for the viscosity term imposes a continuous vt profile near the wall based on the velocity, as proposed by Launder and Spalding [3]. To solve the transport equations, a basic second-order scheme linearupwind discretization was implemented for the divergence terms. To compute the velocity components from the divergence term the Gauss bounded linearUpwindV scheme has been chosen. The Gauss linear scheme, a basic second-order gradient scheme using the Gauss theorem and face interpolation, was chosen as the base gradient scheme for the simulations. Gauss limited was chosen as the snGradSchemes based on the mesh information, normally when orthogonality is bigger than 60, and cellimited Gauss linear scheme was used for gradSchemes, in order to avoid unphysical oscillations.

The results of simulation and wind tunnel experiment were compared, and then the project would focus on the forces and moments of the wheels. There are pressure and shear forces acting on any surface affected by the airflow. In the case of the rotating wheel, the surfaces affected are the tyre, the rim in wind tunnel, see Figure 4. The forces mentioned produce a rotational moment of

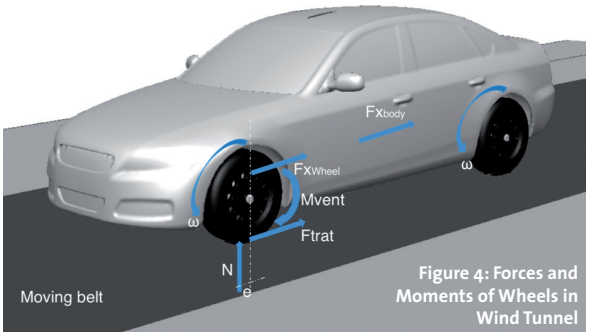

Figure 4: Forces and Moments of Wheels in Wind Tunnel

resistance around the wheel axle of rotation - the ventilation torque, Mvent. It can be calculated in OpenFOAM through an user-difined utility: Integrating normal and tangential stresses over the surface of rotating wheels, to calculate the ventilation moment around the axle of rotation [4].

Table 1: CPU hours of simulations

	No. of Cells	No. of CPU core	CPU hours
Mesh1	44Mio.	144	1664
Mesh2	96 Mio.	528	11088

The simulations were running on the superMUC. in Leibniz Supercomputing Center. Table1 shows the needed resources, for one simulation with 44Mio cells and 96 Mio cells respectively, the CPU hours of Mesh 2 are 11088 hours. For the RANS simulations of the project, 20 cases of different parameters of cars and Reynolds numbers are needed, as many as 0.2 Mio CPU hours would be needed for the RANS simulations for this single project and more CPU hours would be needed for the DES simulations in the future. One case need storage of about 31G, the overall storage needed in project should be around 350G at least.

References and Links

[1] OpenFOAM Foundation, "Features of OpenFOAM", OpenFOAM, URL: http://www.openfoam.org/features
[2] Menter, F., and Egorov, Z.,2005. "A scale adaptive simulation model using two-equation models". AIAA(2005-1095).
[3] Launder, B., and Spalding, D., 1974: "The numerical computation of turbulent flows". Computer methods in applied mechanics and engineering, 3(2), pp.269-289.
[4] Vdovin. A, Lofdahl. L, Sebben. S "Investigation of Wheel Aerodynamic Resistance of Passenger Cars" SAE Technical Paper 2014-01- 0606,doi:10.4271/2014-01-0606.

Large-eddy simulation of a pseudo-shock system in a Laval nozzle

Research Institution
Institute of Aerodynamics and Fluid Mechanics, Technical University of Munich

Principal Investigator
J. F. Quaatz

Researchers
M. Giglmaier, S. Hickel, N.A. Adams

Project Partners
–

SuperMUC Project ID: pr45tu

Introduction

We investigate a novel apparatus that employs gasdynamic processes to produce nanoparticles in industrial scale and of high quality. The device consists of two subsequent chocked Laval nozzles. Sonic speed is reached in the primary nozzle throat and the flow is then further accelerated to supersonic speed. The static temperature and the static pressure of the hot pressurized gas decrease during the acceleration. A precursor – the source material which is burnt to provide the combustion products for the formation of the nanoparticles further downstream – is injected into the flow below ignition conditions. The secondary Laval nozzle forces the flow to return to subsonic conditions ideally by forming a single normal gasdynamic shock. Passing the shock, pressure and temperature instantaneously jump above ignition conditions of the precursor and the particle growth starts at a well-defined location and under specified thermodynamic conditions.

In real operation, where the process is affected by viscous boundary layers at the channel walls, subsonic conditions are not fashioned by a single normal shock but by a complex system of oblique shocks and expansion waves, referred to as pseudo-shock system. Thus, deceleration to subsonic speed is achieved no longer instantaneously. Moreover, flow separation at the walls leads to recirculation zones and adds inhomogeneity to the flow field. The focus of the present work is the numerical investigation of the mean flow pattern and the transient behavior of this pseudo-shock system.

Methods

We perform highly resolved Large-Eddy Simulations (LES) with our in-house flow solver INCA to investigate the topology and real-time dynamics of pseudo-shock systems in the divergent part of the primary Laval nozzle based on the setup of a reference experiment [1]. These LES incorporate, for the first time, the full 3-D geometry of the rectangular Laval nozzle duct including parallel sidewalls, such that the influences of secondary flow features and flow separation at all channel walls can be analyzed. Our numerical discretization method and LES subgrid-scale model ALDM is specifically designed for compressible turbulence and shock-turbulence interaction [2]. Sophisticated numerical techniques such as locally refined grids, an immersed boundary method, and the recycling-rescaling method at the domain inlet are employed to facilitate these unprecedented, computationally challenging simulations.

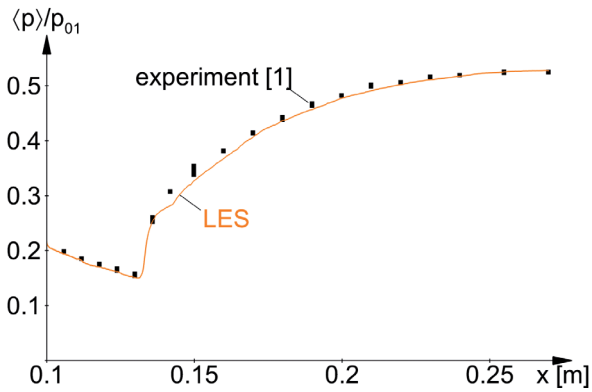

Figure 1: Validation of the wall pressure distribution predicted by the simulations (solid line) against experimental data (squares) provided by [1]

Comprehensive statistical data is generated that will also help to validate and further improve simpler approaches such as Reynolds-averaged simulations (RANS) [3]. Within this paper we give an overview of the progress during the recent project phase.

Flow field analysis

Figure 1 shows the wall pressure distribution predicted by the LES (solid line) compared to experimental data (squares). Location of the pressure jump caused by the primary shock as well as extent of the pressure rise is in excellent agreement. A detailed validation against experimental results and further information about the set-up can be found in [4].

Figure 2 shows transient snapshots of the axial density gradient predicted by the LES at different points in time.

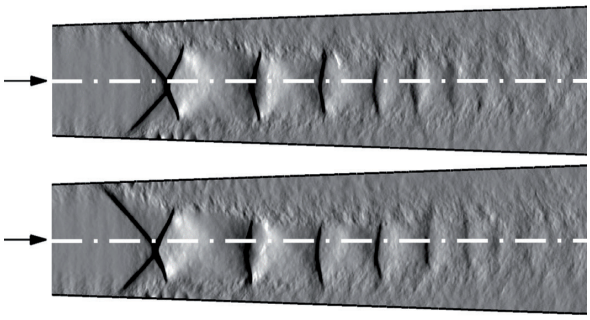

Figure 2: Numerical schlieren pictures showing the axial density gradient at two different instances. Top: Symmetric shape. Bottom: Asymmetric shape.

Shocks occur in black, expansions in withe. It depicts the complex structure of the pseudo-shock system consisting of a bifurcated primary shock followed by expansions and several curved shocks at both instances.

Although the channel geometry is symmetric, Gawehn et al. [1] showed that the shocktrain is not permanently symmetric along the centerline of the channel. It shows a stochastic asymmetry and bends to the channel walls. This behavior can be reproduced by unsteady RANS simulations, but the underlying mechanism has not been completely understood yet. During the recent project period, we observed this asymmetry also in our LES as presented in Figure 2. While located along the centerline in the upper picture, the pseudo-shock system is bent to the lower channel wall at the later instance (bottom). The detailed transient data provided by the time-resolved LES allows for better insight into this mechanism, which may lead to strategies for controlling and avoiding it.

Shock-induced precursor mixing

Another objective of this work is the investigation of the distribution of the injected precursor over the channel cross section while passing the pseudo-shock system. A homogeneous mixture is very important for the quality of the produced nanoparticles. Therefore we analyzed the turbulent transport and mixing of passive scalars injected at four representative locations within the inlet plane.

Figure 3 shows the cumulated distribution of the passive scalars based on our LES results. For visualization of the flow topology, the Mach number distribution at the mid-plane is shown. The centrally injected scalar passes mostly normal shocks and does not spread over the channel cross section, whereas the scalars injected near the nozzle walls are successively transported into the core flow. This is an important finding because sedimentations within the hot recirculation zones are reduced and particle growth within the core flow is more homogenous.

Outlook

These high-fidelity simulations clearly require a high performance computing system. The LRZ supercomputing center provided the resources for this investigation, which consumed in total around $30 \cdot 10^6$ CPUh on SuperMUC Phase 1 and 2. The crucial limitation is the available memory per compute node. This issue has been sustainably improved during the past project phase by intensive review and optimization of our flow solver. The fat node island was used for grid generation and post-processing of the huge datasets. In the next project phase, we have to extend the transient data set enabling a detailed investigation of the flipping process of the pseudo-shock system and collect more data for statistical analysis. Moreover, more information about the unsteady behavior will explain the interaction between the pseudo-shock system and the recirculation zones at the channel walls in more detail.

References

[1] Gawehn, T., Gühlhan, A., Al-Hasan, N.S., Schnerr, G.H. (2010). Experimental and numerical analysis of the structure of pseudo-shock systems in Laval nozzles with parallel side walls. Shock Waves 20(4), pp. 297-306. doi: 10.1007/s00193-010-0263-1

[2] Hickel, S., Egerer, C.P., Larsson J. (2014). Subgrid-scale modeling for implicit Large Eddy Simulation of compressible flows and shock turbulence interaction. Physics of Fluids 26, 106101 doi:10.1063/1.4898641.

[3] Giglmaier, M., Quaatz, J.F., Gawehn, T., Gühlhan, A., Adams, N.A. (2014). Numerical and experimental investigations of pseudo-shock systems in a planar nozzle: impact of bypass mass flow due to narrow gaps. Shock Waves 24(2), pp 139-156. doi: 10.1007/s00193-013-0475-2

[4] Quaatz, J.F., Giglmaier, M., Hickel, S., Adams, N.A. (2014) Large-eddy simulation of a pseudo-shock system in a Laval nozzle. International Journal of Heat and Fluid Flow 49, pp. 108-115. doi: 10.1016/j.ijheatfluidflow.2014.05.006

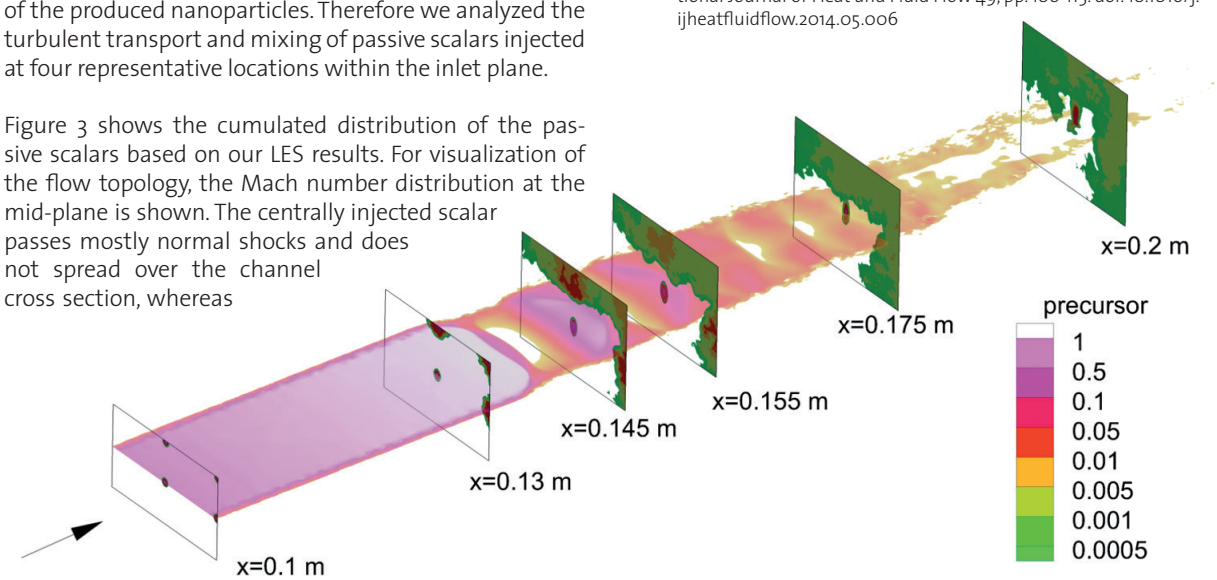

Figure 3: Mixing of passive scalars representing the precursor while passing the pseudo-shock system. For reference, the Mach number is shown at the mid-plane. At the inlet x=0.1 m, the boundary conditions for the four passive scalars is visible.

Shock Mach number influence in reacting shock-bubble interaction

Research Institution
Institute of Aerodynamics and Fluid Mechanics

Principal Investigator
Felix Diegelmann

Researchers
Nikolaus A. Adams, Stefan Hickel

Project Partners
Technische Universität München

SuperMUC Project ID: pr45wa

Introduction

The efficient mixing of fuel and oxidizer is essential in modern combustion engines. Especially in supersonic combustion the rapid mixing of fuel and oxidizer is of crucial importance as the detention time of the fuel-oxidizer mixture in the combustion chamber is only a few milliseconds [1]. The shock-induced Richtmyer-Meshkov instability (RMI) promotes mixing and thus has the potential to increase the burning efficiency of supersonic combustion engines [2]. However, the shock wave, necessary to induce RMI, causes a second effect in a reactive gas mixture. The compression and temperature increase over the shock front can ignite the gas mixture, followed by a subsonic deflagration or a supersonic detonation wave. The reaction wave in turn interacts with the RMI, which affects the flow field evolution und the mixing significantly.

Our LRZ project [3] is used for two-dimensional simulations of a reacting shock-bubble interaction (RSBI) to study the interaction between RMI and shock-induced reaction waves. A planar shock wave propagates through a gas bubble filled with a reactive gas mixture. The baroclinic vorticity generated at the interface causes the bubble to evolve into a vortex ring. Upon contact, the incident shock wave is partially reflected and partially transmitted. In our setup of a convergent geometry (a heavy gas bubble surrounded by a light ambient gas) the transmitted shock wave propagates at a lower velocity than the incident shock wave. Hence, the transmitted shock wave is deformed such that it is focused at the downstream pole of the bubble. Pressure and temperature increase at this shock focusing point, which can be sufficient to ignite the gas mixture.

Our setup contains a gas bubble filled with a stoichiometric composition of Hydrogen (H_2), Oxygen (O_2) and Xenon, surrounded by pure Nitrogen. The H_2-O_2 reaction is highly pressure sensitive. A previous study has already shown promising results by the variation of the initial pressure [4]. We use different shock Mach numbers to trigger a specific reaction wave type, either deflagration or detonation. The latter propagates up to 10E+08 times faster than a deflagration wave [5]. The influence of the reaction wave type on the mixing and the spatial evolution of the bubble is studied in detail.

Numerical Method

We use the parallelized numerical framework INCA [6] on the SuperMUC to solve the full set of compressible reacting multicomponent Navier-Stokes equations. The 2nd-order accurate Strang time splitting scheme is used to separate the stiff source term, containing the chemical reaction kinetics, from the Navier-Stokes equations, which results in a system of partial differential equations (PDE) and a system of stiff ordinary differential equations (ODE). The time integration for the PDE system is realized by the 3rd-order total variation diminishing Runge-Kutta scheme. The numerical fluxes at the cell faces are reconstructed from cell averages by the adaptive central-upwind 6th-order weighted essentially non-oscillatory (WENO-CU6) scheme. The 5th-order backward differentiation formula

Figure 1: Contour plots of SBI at Ma = 2.30. Upper parts show the reacting SBI, lower part the inert counterpart.

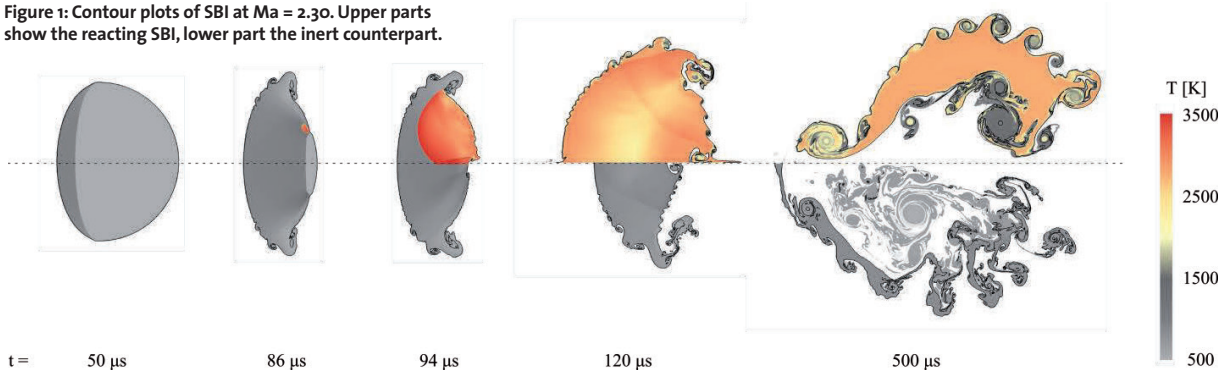

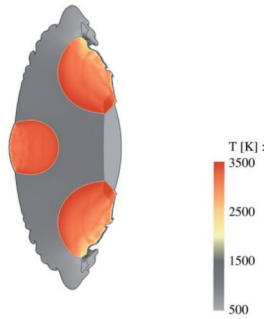

Figure 2: Multiple detonations at a shock Mach number of Ma = 2.50.

is applied to solve the stiff source term of the ODE, containing the chemical reaction kinetics. A complex H2-O2 reaction mechanism with eight species and 19 intermediate reactions is chosen to provide accurate results.

At the current status of our investigations we performed eight simulations of RSBI on a grid with 2.4 million cells. 1067 cores require approximately 200.000 CPUh for each simulation. The preceding grid study was conducted on grids with an even higher resolution. 4268 cores consumed approximately 1 million CPUh for a single simulation. Ten species and short timesteps require storage of about 1 TB per simulation.

Results and Methods

Simulations of RSBI at three different shock Mach numbers between Ma = 2.13 and Ma = 2.50 are investigated. The lowest shock Mach number induces a deflagration wave, higher shock strength leads to a detonation of the gas mixture.

Figure 1 shows the temporal evolution of a reacting SBI and its inert counterpart, exemplary at a shock Mach number of Ma = 2.30. The temperature contour plots of the gas bubble indicate ignition and propagation of the reaction wave. The first two contour plots outline the propagation of the initial shock wave through the bubble gas and the growth of instabilities at the outer interface. In the second timestep the ignition spot near the downstream pole of the bubble is visible. Thereafter the detonation wave propagates through the bubble gas within a few microseconds. The strong heat release leads to a rapid expansion. The comparison to the inert counter-part shows the different evolution of the bubble gas. The development of the primary vortex cores and the secondary instabilities is affected. The last timestep of Fig. 1 outlines the decelerated growth of the instabilities at the outer interface for the reacting setup. The inert simulation shows much finer structures and a higher degree in mixing.

An increase of the shock Mach number up to Ma = 2.50 shortens the induction time and leads to an earlier ignition. Furthermore, the complex flow field in a RSBI with regions of different compression levels causes two reaction zones with simultaneous ignition. Figure 2 shows a temperature contour plot shortly after ignition, outlining the detonation waves. An ignition near the upstream as well as the downstream pole of the bubble induces detonation waves propagating towards each other. The reactive gas mixture is consumed faster, which has a certain influence on the global properties of the RSBI.

The lowest shock Mach number of Ma = 2.13 induces a subsonic deflagration wave with a long ignition delay time. The low propagation velocity and the late ignition lead to a minor influence of the chemical reaction on the bubble evolution. The growth of the secondary instabilities is only affected in the long-term evolution, which reduces the influence of the deflagration wave on the mixing process distinctly.

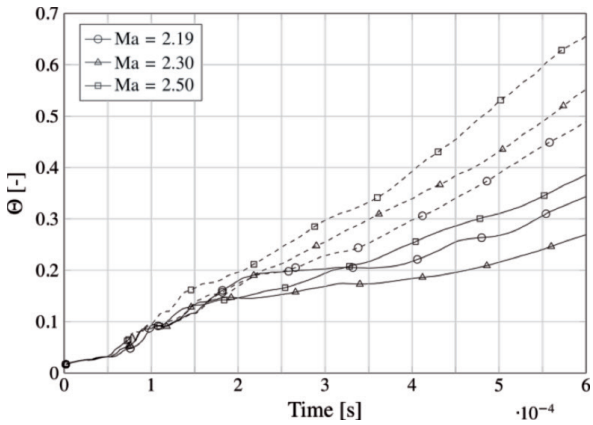

Figure 3: Molecular mixing fraction. ——: reaction; - - -: no reaction.

We use the molecular mixing fraction to determine the influence of the reaction waves on the mixing of the bubble gas with the surrounding gas. Figure 3 outline the temporal evolution of the mixing, which is significantly reduced by both reaction types. All reacting simulations show a decrease compared to their inert counterparts. The deflagration wave reduces mixing up to 30%. An even stronger effect is observed for detonation waves, leading to a decrease of up to 50%.

On-going Research / Outlook

The presented study will be extended to higher shock Mach numbers. We expect shorter induction times and more intense detonations, which in turn affects the global bubble evolution and mixing. Furthermore we are working on the extension to three-dimensional simulations of RSBI, which will require a certain amount of numerical optimization and performance enhancement. These computations will provide a deeper understanding of the three-dimensional effects during a RSBI and the influence on mixing processes.

References and Links

[1] Yang, J., Kubota, T., Zukoski, E., Applications of shock -induced mixing to supersonic combustion, AIAA J. 31 (1993) 854–862.
[2] Marble, F., Zukoski, E., Jacobs, J., Hendricks, G., Waitz, I., Shock enhancement and control of hypersonic mixing and combustion, Proceedings of AIAA 26th Joint Propulsion Conference, Orlando, 1990.
[3] https://www.lrz.de/projekte/hlrb-projects/0000000000F 43587.html
[4] Diegelmann, F., Tritschler, V. K., Hickel, S., Adams, N. A., On the pressure dependence of ignition and mixing in two-dimensional reactive shock-bubble interaction, Combustion and Flame 163 (1) (2016) 414–426.
[5] W. Fickett, W. C. Davis, Detonation: Theory and Experiment, Dover Publications, 2010.
[6] http://www.inca-cfd.com; http://www.aer.mw.tum.de

Large-eddy simulation of the coaxial injection of liquid nitrogen and gaseous hydrogen under supercritical pressures

Research Institution
Institute of Aerodynamics and Fluid Mechanics, Technische Universität München

Principal Investigator
Stefan Hickel

Researchers
Jan Matheis & Stefan Hickel

Project Partners
–

SuperMUC Project ID: pr47bu / Part 1

Introduction

Modern high performance Liquid Rocket Engines (LRE), like the Vulcain II engine, operate at combustion pressures (p>10 MPa) well above the critical pressures of the injected propellants. Furthermore, operating conditions of most main stage LREs are designed such that one or both propellants enter the thrust chamber at cryogenic temperatures. Hence, the injected propellants mix at a transcritical state, which is characterized by high liquid-like densities and viscosities. In this transcritical regime the physical properties, e.g., density, viscosity and specific heats, are strong non-linear functions of the local pressure and temperature. Therefore, trans- and supercritical mixing constitutes considerable challenges for numerical simulations with regard to physical modeling, numerical stability and computational efficiency.

For our studies we have developed a real gas thermodynamics framework within our in-house large-eddy simulation (LES) code INCA. As a highly demanding application example we confine our study to two selected operating conditions of a series of experiments of Oschwald *et al.* [1] in which quantitative density measurements in a coaxial LN2/GH2 jet at supercritical pressures (with respect to the critical pressure of the pure nitrogen) were obtained. This setup is highly challenging both numerically (density ratios between main nitrogen and annular hydrogen jet ranging from 43 to 166) and thermodynamically (real gas mixing effects). In general, shear coaxial injectors, as considered in this work, have been investigated in a number of experiments and numerical simulations, thus, there is a qualitative understanding of the underlying physics associated with trans- and supercritical injection and mixing. However, it is worth mentioning that especially quantitative comparisons between experimental data and high-fidelity numerical simulations are scarce.

Numerical Method

We solve the compressible multicomponent Navier-Stokes equations with our in-house LES code INCA (www.inca-cfd.org). In order to account for real gas effects, mass, momentum, and energy equation are closed by a cubic equation of state. The temperature is obtained from the equation of state using an iterative method, provided the density and energy are known. The governing equations are discretized by a conservative finite-volume scheme on a Cartesian grid.

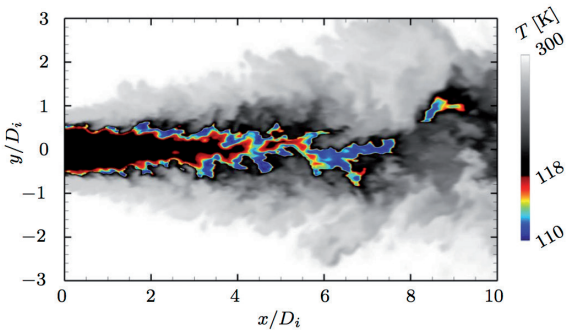

Figure 1: Instantaneous snapshots of the temperature distribution for the cases E4.

In order to avoid spurious oscillations at sharp density gradients, we use a 2nd order total-variation diminishing (TVD) flux function for the advection of mass and internal energy. Effects of unresolved subgrid scales (SGS) are modelled by the adaptive local deconvolution method (ALDM) of Hickel *et al.* [2]. The viscous flux is discretised using a 2nd order central difference scheme, and a 3rd order explicit Runge-Kutta scheme is used for time integration. Due to a high grid resolution, small time steps and a long integration time, high-performance computing power was necessary for all simulations. The computational grids had up to 116 million cells. The simulations were run on up to 8080 cores, and in total, approximately 16 million core hours were used. All simulations have been run on phase 2 of SuperMUC.

Results

Figure 1 depicts a contour plot of the instantaneous temperature distribution for the case E4. Contour levels are shown for 118 K < T < 300 K, from dark to light shades, superimposed by a second group of contour levels with

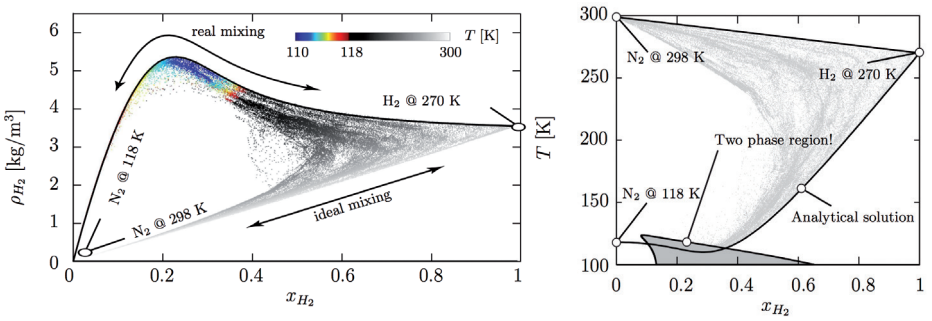

Figure 2: Left: Scatter plot of the instantaneous hydrogen density. Right: Scatter plot of the temperature together with the analytical solution of the adiabatic mixing temperature and prediction of the two-phase region.

110 K < T< 118 K, from blue to red shades. One can clearly identify the cryogenic nitrogen as 'dark core', surrounded by a co-flow of warm hydrogen. It is interesting to note that the temperature in the mixing layer drops below its inflow value of 118 K to approximately 110K. This observation can be attributed to real gas mixing effects. Figure 2 (right) is showing a binary phase diagram for from a vapor-liquid equilibrium (VLE) calculation with the Peng-Robinson EOS: In general, a binary system lying within the parameter space enclosed by the dew-point and bubble-point line (grey shaded region) may separate in a liquid and vapor phase.

The solid line labeled as 'Analytical solution', which partially intersect the two-phase region, depicts the adiabatic mixture temperature of a H_2/N_2 system at 4 MPa. From this analytical solution we found that the temperature drop in the shear layer is caused by real gas mixing effects (rather than by a heat transfer/ diffusion mechanism) [3]. Along with the temperature drop an increase in hydrogen density can be observed, c.f. Fig. 2 (left), exceeding its pure component value at the inflow. This phenomenon has also been recorded experimentally for the operating conditions of test case E4.

In the following we compare our numerical results with the experimental data of Oschwald *et al.* [1]. Figure 3 depict axial (centerline) nitrogen density profiles for test case E4 (nominal inflow). We observe significant differences in the potential core region ($x/D_i < 3$) with an experimental and numerical nitrogen density of 390.18 kg/m³ and 608.78 kg/m³, respectively. We note that the observed differences

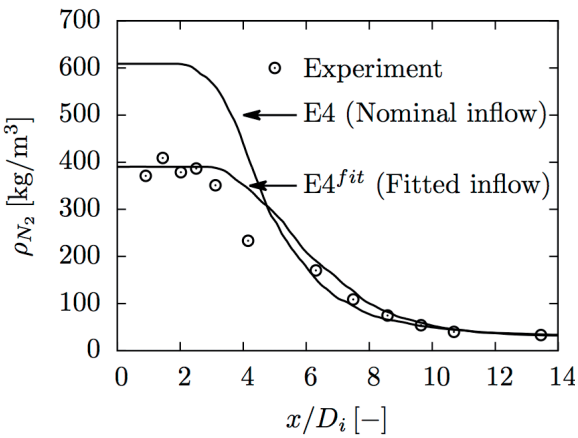

Figure 3: Axial (centerline) nitrogen density profiles for test case E4. Comparison between numerical simulation and experimental data of Oschwald *et al.* [1].

in density of approximately 200 kg/m3 cannot be attributed to an inaccurate equation of state. Measurements in a high-pressure low-temperature environment are very demanding and introduce a number of uncertainties. Oschwald *et al.* [1] report beam steering and reflection issues along with problems to determine the inlet temperature of the cryogenic nitrogen. Since the nitrogen density at injection is very sensitive to small changes in temperature, we believe that the observed discrepancies between experiment and simulation may be founded in a different nitrogen temperature prior to injection. Assuming a correct density and mass flow rate measurement (rather than temperature measurement), we define a test case $E4^{fit}$ with a fitted nitrogen bulk velocity u=7.49 m/s and inflow temperature T_{N2}=128.274K. Figure 3 depicts the corresponding centerline nitrogen density profiles. We now observe a very good agreement for the absolute value of the predicted nitrogen density, the potential core length (i.e. location of break-up) and the axial position for which a fully mixed state is obtained (x/D_i~10). These preliminary results suggest that the nitrogen temperature at injector exit could have been higher at the time of data collection.

Conclusions

The coaxial injection of cryogenic nitrogen and warm hydrogen into a supercritical nitrogen atmosphere was studied numerically by means of well-resolved large-eddy simulation. We were able to reproduce the increase in hydrogen density downstream of jet break-up that has been recorded experimentally. Based on the observation that the definition of inflow boundary conditions in terms of temperature and pressure must yield a density prior jet break-up much higher than what was measured experimentally, we defined a new test case for which bulk velocity and injection temperature are adjusted to match density and mass flow rate measurement. For this fitted boundary condition we observed a good agreement between experimental and numerical data.

References and Links

[1] Oschwald, M., Schik, A., Klar, M. and Mayer, W. (1999) Investigation of Coaxial LN2/GH2-Injection at Supercritical Pressure by Spontaneous Raman Scattering. In 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit.

[2] Hickel, S., Egerer, C. & Larsson, J. (2014) Subgrid-scale modeling for implicit large eddy simulation of compressible flows and shock-turbulence interaction. Phys. Fluids 26: 106101.

[3] Müller, H., Pfitzner, M., Matheis, J. and Hickel, S. (2015) Large-eddy simulation of coaxial LN2/GH2 injection at trans- and supercritical conditions. J. Propuls. Power. http://dx.doi.org/10.2514/1.B35827

www.aer.mw.tum.de/abteilungen/large-eddy-simulation
www.inca-cfd.org

Numerical Investigation of the Flow Field about the VFE-2 Delta Wing

Research Institution
Institute of Aerodynamics and Fluid Mechanics, Technische Universität München

Principal Investigator
Stefan Hickel

Researchers
Christian Zwerger & Stefan Hickel

Project Partners
–

SuperMUC Project ID: pr47bu / Part 2

Introduction

The industrial application of delta wings reaches from classical aerospace engineering, e.g., high-agility aircraft, aerodynamic devices or control surfaces, to unique environmental technologies, such as snow fences. Common to all these applications is the exploitation of leading edge vortices, see Figure 1.

Steadiness and stability of these leading edge vortices is of paramount importance, in particular for the controllability of high-agility aircraft. Given that vortices can undergo a sudden expansion often related to vortex breakdown, and considering that the occurrence of this phenomenon is critical for aircraft and not yet physically fully understood, a further investigation leading to a profound insight is required. A detailed understanding of vortex formation and breakdown necessitates a comprehensive insight into the entire unsteady flow, which can only be obtained from time-accurate high-fidelity simulations accompanied by experiments.

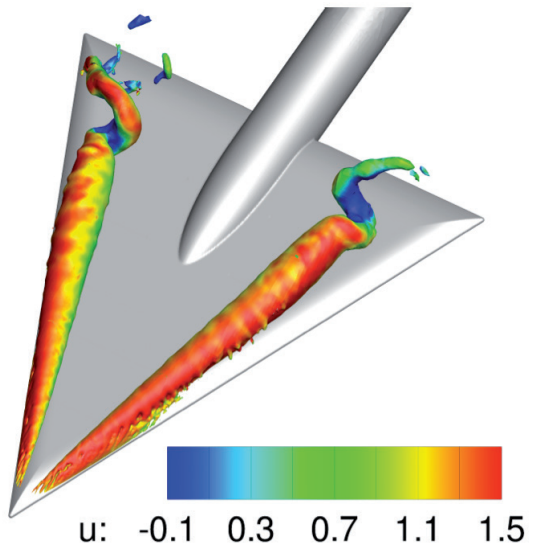

Figure 1: Isosurface of pressure coefficient colored by streamwise velocity for VFE-2 delta wing at AoA 23° and Re=2 million shows helical form of vortex breakdown [5].

RANS (Reynolds Averaged Navier-Stokes) simulations with appropriate statistical turbulence models are widely used in industry. However, this simplified approach often fails to accurately predict flow separation and reattachment. In the case of the VFE-2 delta wing, results obtained by RANS simulations poorly resemble existing experimental results.

One expects significantly better results from Large Eddy Simulation (LES). In LES, the large flow structures are resolved, while effects of small structures are modeled by subgrid-scale (SGS) turbulence models. In implicit LES (ILES), the truncation error of the discretization of the convective terms is carefully tailored to act as a SGS model. In this project the Adaptive Local Deconvolution Method (ALDM) for implicit LES is used. ALDM has been derived from spectral turbulence theory [3] and has shown considerable potential for the efficient representation of physically complex flows in generic configuration.

Numerical Method

The simulations have been performed with the flow solver INCA, which incorporates the ALDM [3] turbulence model for the compressible Navier-Stokes equations. The equations are discretized on a collocated Cartesian mesh and bounding surfaces of the flow that are not aligned with the grid are accounted for by a conservative cut-cell immersed boundary method [4]. For time advancement, an explicit third-order Runge-Kutta scheme is used. The diffusive terms are discretized by a second-order scheme. An efficiency improvement is achieved by modeling the turbulent boundary layer with a wall model [1] based on the Turbulent Boundary Layer Equations (TBLE) and by locally adapting the mesh resolution with local mesh refinement.

As a result of the high grid resolution and the requirement to resolve the unsteady behavior, which necessitates a large number of small time steps, high-performance computing power is required. The simulations have been mainly run on the fat node island of SuperMUC, which is equipped with Intel Xeon E7-4870 10C pro-

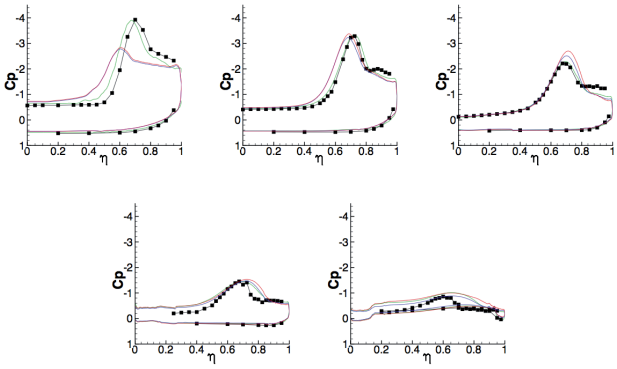

Figure 2: Results for the pressure coefficient C_p at cross sections x/c_r=0.2,0.4,0.6,0.8, and 0.95 (from left to right and from top to bottom) [5]. Red – LES without wall model, blue – LES with TBLE wall model, green – refined grid, black – experiment [2].

cessors and provides 6.4 GByte of memory per core. Tests have shown that our MPI parallelized Fortran code INCA has very good scaling properties on the fat node island. The employed computational grids had up to 74 million cells. The simulations were run on up to 2080 cores, and in total, approximately 57 million core hours were used.

Results

Our investigation focused on two aspects: (1) leading edge bluntness effects on the primary vortex separation and (2) vortex breakdown above the wing and its control.

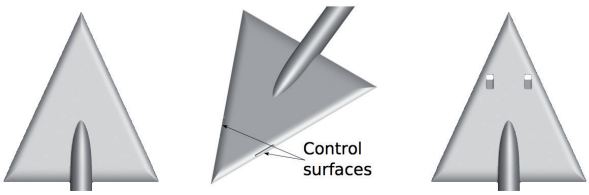

Figure 3: Configuration without flow control (left), with control surfaces in the front part of the wing (center), and with slots in the front part of the wing (right) [5].

With respect to aspect (1), a sharp leading edge (SLE) and a medium radius round leading edge (MRLE) were considered for angles of attack of 13°, 18°, and 23° leading to different overall flow characteristics. The large-eddy simulations correctly predicted the main flow phenomena and were quantitatively in reasonable to very good agreement with experimental measurements of steady and unsteady surface pressures (Figure 2) velocity distributions, and vortex breakdown position and frequency. With respect to aspect (2), flow control by oscillating control surfaces and flow control by a geometric modification leading to an injection of fluid from the pressure side were investigated, see Figure 3.

Considering the influence on vortex breakdown position, the numerical simulations confirmed experimental observations regarding oscillating control surfaces, and showed promising potential for flow control by the proposed geometric modification, see Figure 4.

Conclusions

Our simulations show that for the massively separated flow considered, which is dominated by large-scale fluctuations, LES outperforms RANS approaches, notably regarding the correct prediction of vortex breakdown, and is thus an expedient tool to investigate flow control mechanisms.

With regard to future numerical investigations of the flow around delta wings, our simulations suggest that further improvements can be obtained by using grids that are additionally refined in the near-wall region, and by resorting to more elaborate wall models. Oscillating control surfaces are less efficient for flow control than a geometric modification leading to bleeding of fluid from the pressure side to the suction side.

References and Links

[1] Chen, Z.L., Hickel, S., Devesa, A., Berland, J., Adams, N.A. (2014) Wall-modeling for implicit large-eddy simulation and immersed-interface methods. Theor. Comp. Fluid. Dyn. 28: 1-21.

[2] Furman, A. and Breitsamter, C. (2013) Turbulent and unsteady flow characteristics of delta wing vortex systems. Aerosp. Sci. Technol 24: 32-44.

[3] Hickel, S., Egerer, C. & Larsson, J. (2014) Subgrid-scale modeling for implicit large eddy simulation of compressible flows and shock-turbulence interaction. Phys. Fluids 26: 106101.

[4] Örley, F., Pasquariello, V., Hickel, S., Adams, N.A. (2015) Cut-element based immersed boundary method for moving geometries in compressible liquid flows with cavitation. J. Comput. Phys. 283: 1-22.

[5] Zwerger, C., Hickel, S., Breitsamter, C. & Adams, N.A. (2015) Wall-modeled large-eddy simulation of the VFE-2 delta wing. AIAA paper 2015-2572.

www.aer.mw.tum.de/abteilungen/large-eddy-simulation
www.inca-cfd.org

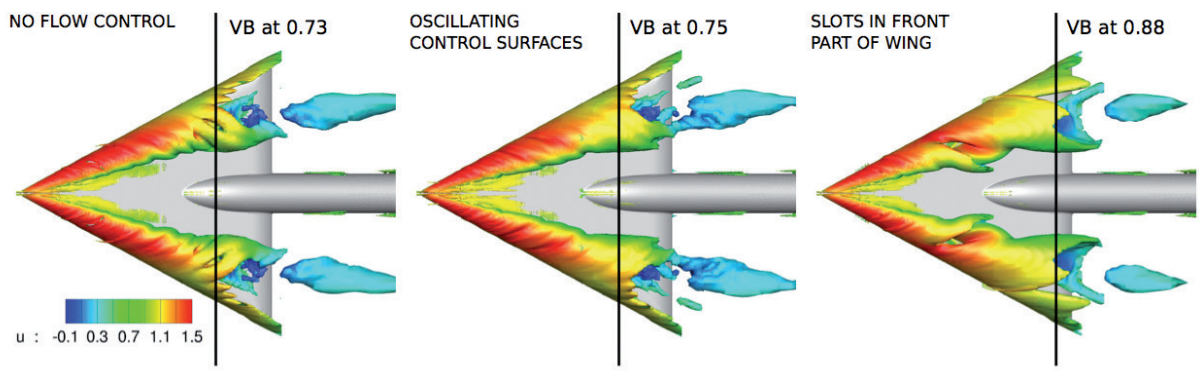

Figure 4: LES results for vortex breakdown (VB) position with different flow control configurations. Figures show isosurface of streamwise vorticity colored by streamwise velocity [5].

Optimal wall spacing for heat transport in thermal convection

Research Institution
Max Planck Institute for Dynamics and Self-Organization
Principal Investigator
Olga Shishkina
Researchers
Kai Leong Chong, Matthias Kaczorowski, Olga Shishkina
Project Partners
–

SuperMUC Project ID: pr47vi

Introduction

Thermal convection is a ubiquitous phenomenon that happens in many natural systems and engineering processes. It has been studied extensively by an idealized model, so-called Rayleigh-Bénard (RB) convection which simplifies the real situation to a fluid layer heated from below and cooled from above (for reviews, we refer to Ahlers, Grossmann & Lohse 2009; Lohse, Xia 2010).

In RB flow, the study of heat transport efficiency is one of the major concerns. Over the past decades, many efforts are devoted to studying how the heat transfer efficiency (in terms of Nusselt number *Nu*) depends on the Rayleigh number *Ra* (strength of thermal driving) and the Prandtl number *Pr* (intrinsic fluid property). The majority of studies have been conducted on the cell with the width-over-height aspect-ratio Γ about one or much larger than one for approaching the scenario in planetary convection. However, the small Γ situations are generally overlooked. Intuitively, one would expect less heat transport efficiency in a narrow container due to large effect of viscous drag from sidewalls. However, the recent experimental result by Huang *et al.* (2013) opposes this original thought. They have shown that geometrical confinement can lead to significant enhancement of global heat transport instead of reduction. This counter-intuitive result motivates more in-depth studies on small Γ Rayleigh-Bénard convection.

Previously, we have investigated the RB convection in highly confined cell by direct numerical simulations (DNS). Our results align the experimental findings by Huang *et al.* (2013) and further unveil that the enhanced heat transport is brought by the increased coherency in thermal plumes. In this study, we aim to study comprehensively the small Γ scenario by spanning a wider parameter range in Γ, which is difficult to accomplish in experiments. With this new parameter range, we further reveal the existence of the optimal spacing between sidewalls for the best heat transport efficiency. Our findings have achieved important advance in the passive thermal management and potentially benefit indoor ventilation.

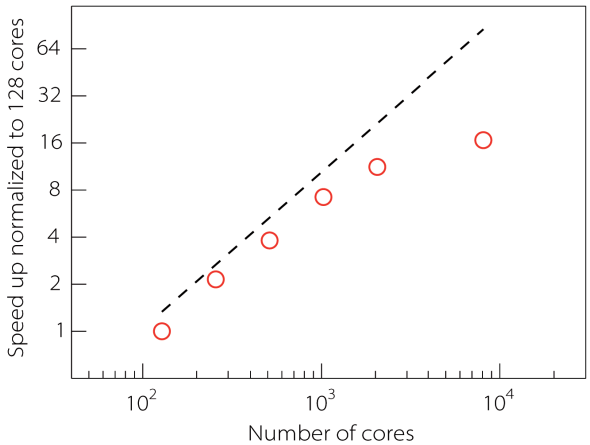

Figure 1: Scalability test of the code with 1024^3 grid points tested on Supermuc thin node island. The black dash line indicates the ideal speed-up.

Numerical Simulation

The approach to simulate thermal convection is to solve the incompressible Navier-Stokes equations and heat equation. The simulations are carried out by the well-tested fourth order finite volume code as described in [3]. The performance test had been done on Supermuc cluster thin node island up to 8192 cores with 1024^3 grid points and shown to have excellent speed up to 2048 cores as shown in figure 1.

In this study, *Ra* spans over three decades from 1×10^7 to 1×10^{10} and the aspect-ratio Γ spans from 1 to 1/64 with all simulations fixed at *Pr*=4.38 which correspond to water at 40°C. The temperature boundary conditions are adiabatic for all sidewalls and isothermal for top and bottom walls. For the velocity boundary conditions, all walls are no-slip. For simulations with *Pr*>1, the Batchelor length scale is the smallest length scale for the turbulent flow and thus the meshes are designed to resolve the global Batchelor length scale. Besides, much denser grid is needed near walls as suggested by Shishkina *et al.* (2010). Our meshes are designed based on the above requirements for simulating the reliable results.

Results

A major goal of this study is to compare the heat transport efficiency in a highly confined cell to that in cubic cell. Therefore, Nu is normalized with that in Γ=1 and $Nu/Nu(\Gamma=1)$ is plotted as a function of Γ as shown in figure 2. In our previous study, only cases of $Ra=1\times10^9$ and 1×10^{10} are simulated with Γ down to 1/32 and 1/16, respectively. The previous simulations solve the puzzle on the cause of enhanced heat transport efficiency under confinement by looking at the snapshots of temperature fields over the whole domain. Here, with wider parameter range both in *Ra* and Γ, it enables us to observe an optimal aspect-ratio of the convection cell for achieving the maximum heat transport efficiency compared to the cubic case at a given *Ra*. It is also found that the optimal Γ becomes smaller when Ra increased. This trend could be described by a power-law relation $\Gamma_{opt}=29.37Ra^{-0.31}$. Besides, the amount of enhancement is found to be larger at larger *Ra*. For instance, $Nu/Nu(\Gamma=1)$ can only reach less than 1.05 for $Ra=1\times10^7$ whereas $Nu/Nu(\Gamma=1)$ can reach more than 1.2 for $Ra=1\times10^{10}$.

The observation of the optimal aspect-ratio leads us to figure out the physical origin of this optimal point. In RB flow, the major heat carrier is the thermal plumes which transport heat between top and bottom plates. For the cubic case, only a portion of area near top or bottom plates is available for plume emission or impinging. The common feature found at Γ_{opt} is that half of area is available for plume impingement while the other half for plume emission. It, therefore, fully utilizes the region over plate for heat transport. All of the above results have been published in a peer-reviewed journal [6].

On-going Research / Outlook

The simulation of RB flow for *Ra* up to 1×10^{10} is computationally expensive in terms of computing power and hard disk storage. Thus, we gratefully acknowledge the computational resources supported by Leibniz-Rechenzentrum Munich. Compared to Γ=1 situation, a new physical picture of heat transport is identified here at Γ_{opt} for any explored Ra. Therefore, a detailed comparison between Γ=1 and $\Gamma=\Gamma_{opt}$ is valuable for our further research, for example, their vertical temperature and velocity profiles. Additionally, we plan to compare the fluid with different *Pr* under geometrical confinement, which are computationally expensive for the situations of $Pr<<1$ and $Pr>>1$.

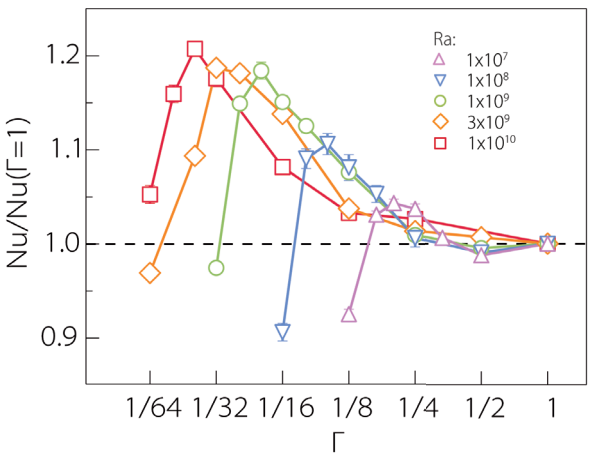

Figure 2: The relative global Nu as a function of Γ over three decades of *Ra*, figure taken from [6].

References and Links

[1] Ahlers, G., Grossmann, S., Lohse D. 2009. Heat transfer and large scale dynamics in turbulent Rayleigh-Bénard convection. Rev. Mod. Phys. 82: 503–537.

[2] Lohse, D, Xia, K.-Q. 2010. Small-scale properties of turbulent Rayleigh-Bénard convection. Annu. Rev. Fluid Mech. 42: 335–364.

[3] Huang, S.-D., Kaczorowski, M., Ni, R., Xia, K.-Q. 2013. Confinement-induced heat-transport enhancement in turbulent thermal convection. PRL. 111 (104501)

[4] Kaczorowski, M., Shishkin, A., Shishkina, O., Wagner, C. 2008. Development of a numerical procedure for direction simulations of turbulent convection in a closed rectangular cell. New Results in Numerical and Experimental Fluid Mechanics VI. 96:381–388.

[5] Shishkina, O., Stevens, R. J. A. M., Grossmann, S., Lohse, D. 2010. Boundary layer structure in turbulent thermal convection and consequences for the required numerical resolution. New J. Phys. 12 (075022)

[6] Chong, K. L., Huang, S.-D., Kaczorowski, M., Xia, K.-Q. 2015. Condensation of coherent structures in turbulent flows. Phys. Rev. Lett. 115 (264503)

Initiation and evolution of subaqueous patterns in channel flow

Research Institution
Institute for Hydromechanics

Principal Investigator
Markus Uhlmann

Researchers
Aman G. Kidanemariam

Project Partners
–

SuperMUC Project ID: pr84du

Introduction

Subaqueous sediment patterns, also commonly termed as dunes, ripples or simply bedforms, can be abundantly observed in natural streams or man-made canals. These sediment features, besides being simply fascinating, have important implications in many field of science and engineering. For instance, bedforms greatly influence the rate of sediment transport as well as the stability of hydraulic structures in a given river. Thus, fundamental understanding of the mechanisms which are behind their formation as well as predicting their characteristics is crucial. However, this task has been very challenging thanks to the complex interaction between the sediment particles and the driving turbulent flow.

In most previous theoretical work on pattern formation, the background turbulent flow is typically represented by a Reynolds averaged Navier-Stokes equations (RANS) model, while the sediment bed evolution is described by the Exner equation. The two problems are then coupled by an algebraic expression for the particle flux as a function of the local bed shear stress. Hydro-morphodynamic linear stability analysis is then performed to determine the stability of the sediment bed [2]. However, there is no clear consensus among the different predictions from these approaches, and, when compared to experimental observations, outcome of these models is unsatisfactory. This inadequacy can be linked to, among others, the poor predictive capability of the adopted algebraic expressions for the particle flux. Furthermore, these studies are unable to describe the subsequent nonlinear bedform evolution and the complex interaction between the turbulent flow over the evolving sediment bed.

Results and Methods

The simulations where performed employing a fluid solver which features an immersed boundary technique for the efficient and accurate treatment of the moving fluid-solid interfaces [3]. Moreover, in dense particulate flow problems, in addition to the fluid-particle interaction, the inter-particle contact dynamics is significant. To this end, we have developed a discrete element model (DEM), which is based on the soft-sphere approach, to treat inter-particle collisions. The DEM was successfully coupled with the DNS solver by taking into account the disparity in time scales between the Navier-Stokes dynamics and that of the grain contact [4]. The DNS-DEM coupling strategy was extensively validated with respect to available experimental data of the erosion and transport of a large number of spherical particles sheared by a laminar flow [4]. We have performed a large number of 'small-scale' simulations covering a wide range of values of the governing parameters. These simulations were typically running on less than 100 processor cores on SuperMUC. Beyond validation, these simulations have allowed us to tackle the problem of bedload transport by itself. In particular, we have addressed the scaling laws of particle flux, fluid and particle velocities inside and above the sediment bed, as well as the contribution of hydrodynamic and collision forces to the momentum transfer between the two phases.

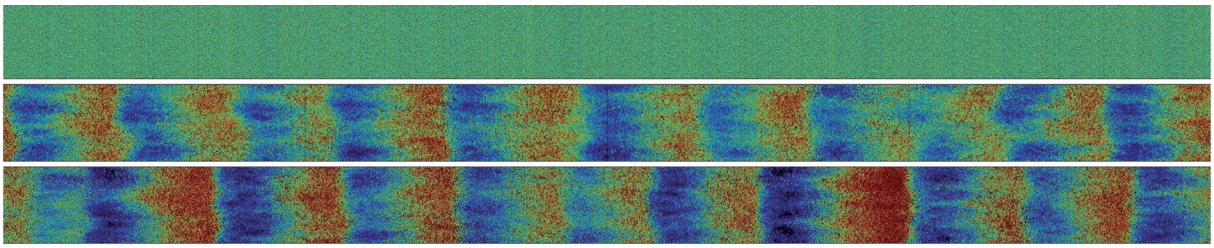

Figure 1 Top view of the particle positions at different times. Particle are colored according to their wall-normal location (increasing from blue to red). Flow direction is from left to right. Snapshots are at intervals of 200 bulk time units. Number of of particles considered: 1.1 million.

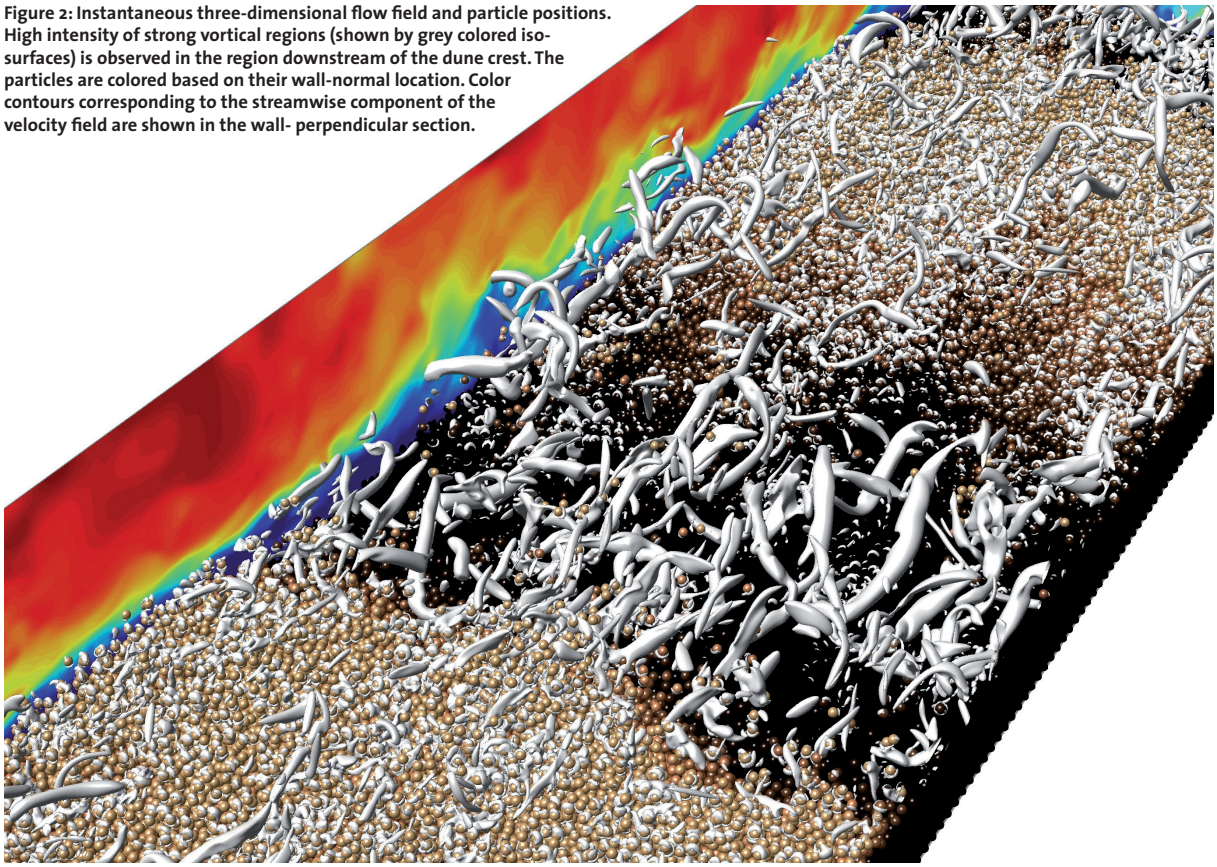

Figure 2: Instantaneous three-dimensional flow field and particle positions. High intensity of strong vortical regions (shown by grey colored iso-surfaces) is observed in the region downstream of the dune crest. The particles are colored based on their wall-normal location. Color contours corresponding to the streamwise component of the velocity field are shown in the wall- perpendicular section.

Subsequently, we have performed several medium to large-scale simulations of the formation of bedforms over an erodible bed, both in the laminar and turbulent regimes. In order to realistically capture the phenomenon, we have considered sufficiently large computational domains involving up to 3.6 billion grid nodes, while the mobile sediment bed is represented by up to 1.1 million freely moving spherical particles. These simulations were typically running on 144 to 2304 processor cores. It should be noted that our largest simulation is the first of its kind to pass the one-million fully resolved particle simulation milestone. As is shown in figure 1, sediment erosion-deposition process takes place over an initially flat bed, which ultimately leads to the emergence and evolution of the patterns. We have subsequently analyzed the temporal and spatial evolution of the sediment bed. The results of our simulations, with respect to the pattern wavelength, amplitude, asymmetric shape as well as their propagation velocity, are in excellent agreement with available experimental data [5]. Furthermore, by varying the computational box size, we were able to find a minimum box dimension which accommodates an unstable pattern wavelength, thus determining the cutoff length for pattern formation. Currently, we find ourselves at the stage of extensive data analysis. As is shown in figure 2, the interaction between the turbulent flow and the evolving sediment bed is intriguingly complex. For the purpose of characterizing such a flow, we are performing dune-conditioned statistical analysis, which takes into account the spatial and temporal variability of the sediment bed.

In performing the simulations, we have consumed a total of approximately 19 million CPU core hours.

On-going Research / Outlook

Although the domain sizes considered in our simulations are relatively large, they are still considered to be not sufficient to capture a number of important dune features, such as their three-dimensional patterning. The invaluable computing resources offered by LRZ will enable us to pursue a follow-up project to address the outstanding issues in the near future.

References and Links

[1] http://www.ifh.kit.edu/26_1222.php

[2] F. Charru, B. Andreotti, and P. Claudin. 2013. Sand Ripples and Dunes, Annu. Rev. Fluid Mech. 45(1), 469-493. DOI:http://dx.doi.org/10.1017/jfm.2011.528

[3] M. Uhlmann. 2005. An immersed boundary method with direct forcing for the simulation of particulate flows, J. Comput. Phys. 209(2), 448-476. DOI:http://dx.doi.org/10.1016/j.jcp.2005.03.017

[4] A. G. Kidanemariam and M. Uhlmann. 2014. Interface-resolved direct numerical simulation of the erosion of a sediment bed sheared by laminar channel flow, Int. J. Multiph. Flow. 67, 174-188. DOI:http://dx.doi.org/10.1016/j.ijmultiphaseflow.2014.08.008

[5] A. G. Kidanemariam and M. Uhlmann. 2014. Direct numerical simulation of pattern formation in subaqueous sediment, J. Fluid Mech. 750, R2. DOI:http://dx.doi.org/10.1017/jfm.2014.284

Large-eddy simulation of LOx/GCH4 injection at supercritical pressures

Research Institution
Institut für Thermodynamik, Universität der Bundeswehr München

Principal Investigator
Michael Pfitzner

Researchers
Hagen Müller & Michael Pfitzner

Project Partners
–

SuperMUC Project ID: pr63ze

Introduction

Today's main stage liquid propellant rocket engines (LRE) typically operate at supercritical pressures, i.e., at chamber pressures that exceed the critical pressure of the propellants, and at cryogenic injection temperatures. One or both propellants are thus injected at near-critical conditions and mixing, ignition and combustion are largely affected by non-ideal thermodynamic effects. In particular, the thermodynamic and transport properties, e.g., density, enthalpy, viscosity, are highly non-linear functions of temperature and pressure. Moreover, the surface tension between liquid and vapor is diminished at sufficiently high pressures and mixing is characterized by continuous-phase diffusion rather than by two-phase spray atomization. In these diffusion mixing layers, the fluid properties change drastically and the density may vary by two orders of magnitude within a few micrometers.

Using large-eddy simulations (LES) to perform unsteady simulations of such flows allows to generate an understanding, which is otherwise difficult to obtain, especially since experimental data are scarce. However, performing numerical simulations at LRE-typical conditions poses a serious challenge, mainly due to the non-linear behavior of the fluid properties near the critical point, and is subject of ongoing research. In this project, we aim at developing numerical tools that reliably predict trans- and supercritical injection and combustion. We extended the capabilities of the open-source CFD software OpenFOAM by real-gas thermodynamics and combustion models, which allow for an accurate prediction of such flows. The method has been applied to simulate the injection of transcritical nitrogen into a warm atmosphere [1] and the coaxial injection of LN_2 and GH_2 at supercritical pressures [2].

In the present report, we present results for the non-premixed combustion of liquid oxygen (LOx) and gaseous methane (GCH_4) at supercritical pressure. The present single-element configuration has been studied experimentally by Singla *et al.* [3], who measured the OH-radiation as a marker for the flame location and shape allowing for a qualitative comparison between experiment and simulation. They investigated several operating conditions, of which the most critical in terms of thermodynamics modelling is simulated herein. Methane is injected through an annulus of a coaxial injector element at ambient temperature (T_{CH4}=288 K), while oxygen is chilled (T_{O2}=85 K) and enters the combustion chamber through the center tube. The pressure in the chamber is supercritical (p_{ch}=56.1 bar).

Numerical Method

OpenFOAM solves for the compressible, multicomponent Navier-Stokes equations using a pressure-based al-

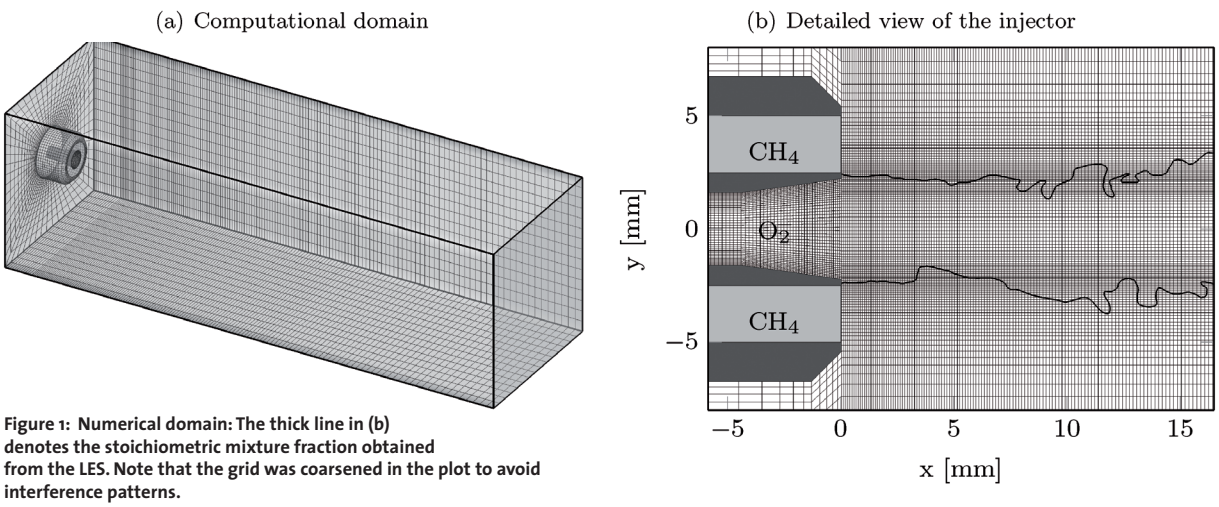

Figure 1: Numerical domain: The thick line in (b) denotes the stoichiometric mixture fraction obtained from the LES. Note that the grid was coarsened in the plot to avoid interference patterns.

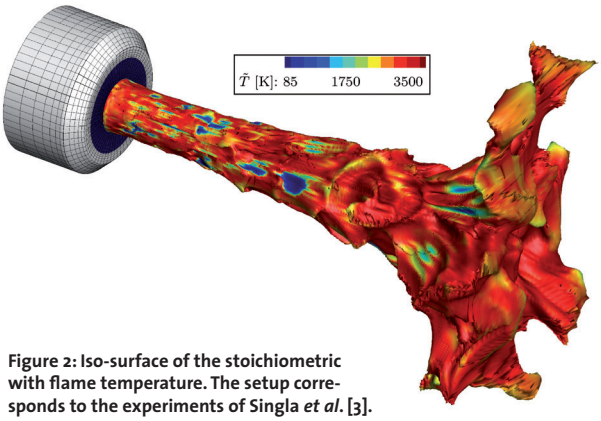

Figure 2: Iso-surface of the stoichiometric with flame temperature. The setup corresponds to the experiments of Singla *et al.* [3].

gorithm that is specifically adapted to real-gas flows [1]. The adaption is necessary, since the non-linearity of the fluid properties can cause inconsistencies in standard segregated solution approaches. In particular, the thermodynamic properties are modelled using a Peng-Robinson equation of state together with a novel volume translation method that corrects deviations at cold temperatures. The caloric properties are calculated with a departure function formalism [4].

To describe the chemical reactions, we use a modified Jones-Lindstedt mechanism (6 reactions / 9 species) that accounts for water and oxygen dissociation. Closure of the chemical source term on subgrid-scale level is obtained with the Eulerian stochastic field method. This is a numerically efficient method to describe the joint composition probability density function (PDF) in a purely Eulerian way. For spatial discretization we use a second-order central differences scheme with a van Leer limiter for momentum and scalar transport to avoid unphysical oscillations. A first-order implicit Euler scheme is used for temporal discretization. Figure 1 shows the computational domain that has been used for the present LES. The block-structured mesh comprises up to 16 million cells and is refined in regions of high density gradient. Due to the high grid resolution, the immense numerical effort to integrate the chemical reaction mechanism and due to real-gas thermodynamics modelling, high-performance computing power has been necessary for all simulations. Phase 2 of SuperMUC has been used.

Results

Figure 2 shows a typical LES realization and visualizes the characteristic flame shape. Shown is an iso-surface of the stoichiometric mixture fraction as a marker for the flame location and the surface is colored with the temperature. The flame shape is characterized by a comparably smooth region near the injector, which is followed by a strong radial expansion and an abrupt end of combustion further downstream. These specific features have been identified to be a result of the interplay between real-gas thermodynamics, combustion and the constraints imposed by the geometry, i.e., the interaction between flame and combustion chamber wall. In the first section, a strained diffusion flame forms between the warm methane stream and the cryogenic oxidizer. The high density gradient between oxidizer (ϱ_{O2}=1177 kg/m³) and flame, which is a characteristic feature at LRE-typical injection conditions, hinders the formation of vortices. Diffusion is the rate-controlling transport process and the heat release is rather small. Figure 1 also shows several locations of local extinction (blue spots) in this region of the flame.

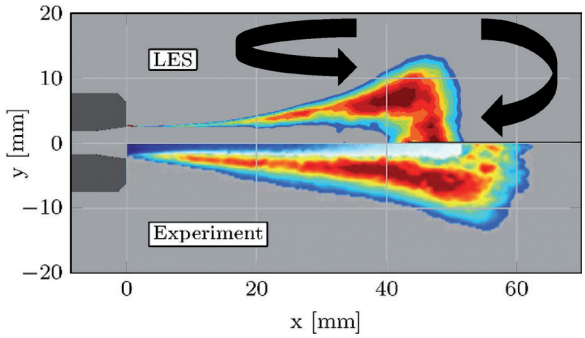

Figure 3: Averaged OH-radiation. Comparison between LES result and measurement of Singla *et al.* [3].

Further downstream, larger vortices evolve and accelerate combustion. The consumption of dense oxygen leads to a strong expansion and acceleration of hot gases in the radial direction. At the wall, the stream of hot gases is redirected and splits into two recirculation zones, one upstream and one downstream of the flame as indicated by the black arrows in Figure 3. Figure 3 also shows the averaged contour of OH-radiation. The upper half shows the LES result, the lower half the experimental result of Singla *et al.* [3]. The characteristic flame shape, i.e., thin diffusion flame, followed by strong expansion and abrupt end of combustion, is in qualitatively good agreement with the experiment, especially when considering the experimental scatter at these demanding conditions.

Conclusions

The non-premixed combustion of LOx/GCH4 at supercritical pressures was studied numerically by means of large-eddy simulation. The non-ideal behavior of the fluid has been modelled using a cubic equation of state and a novel volume correction method for accurate predictions at cryogenic temperatures. Turbulent combustion on unresolved scales has been modelled with a transported probability density function approach in a purely Eulerian way. The results demonstrate that the characteristic features of the flame are well predicted and a comparison with the available experimental data shows good agreement.

References and Links

[1] Müller H., Niedermeier C.A., Matheis J., Pfitzner M., and Hickel S. (2016) Large-eddy simulation of nitrogen injection at trans- and supercritical conditions. Phys. Fluids, vol. 28, 015102. doi:10.1063/1.4937948.

[2] Müller, H., Pfitzner, M., Matheis, J. and Hickel, S. (2015) Large-eddy simulation of coaxial LN2/GH2 injection at trans- and supercritical conditions. J. Propuls. Power. doi:10.2514/1.B35827.

[3] Singla G., Scouflaire P., Rolon C. and Candel S. (2005) Transcritical oxygen/transcritical or supercritical methane combustion. Proc. Combust. Institute, vol. 30, pp. 2921–2928. doi:10.1016/j.proci.2004.08.063.

[4] Matheis J., Müller H., Lenz C., Pfitzner M., and Hickel S. (2016) Volume translation methods for real-gas computational fluid dynamics simulations. J. Supercrit. Fluids, vol. 107, pp. 422–432. doi:10.1016/j.supflu.2015.10.004.

www.openfoam.com

Landing Aircraft and Enhanced Wake Vortex Decay

Research Institution
DLR-Institut für Physik der Atmosphäre; Oberpfaffenhofen, 82234 Wessling, Germany; Institute for Advanced Studies

Principal Investigator
Frank Holzäpfel

Researchers
Anton Stephan

Project Partners
–

SuperMUC Project ID: pr63zi

Introduction

As an unavoidable consequence of lift aircraft generate a pair of counter-rotating and long-lived wake vortices that pose a potential risk to following aircraft. The prescribed aircraft separations to avoid wake vortex hazards contribute significantly to capacity restrictions of large airports. Wake vortex behavior is largely controlled by the prevailing meteorological conditions like wind and the interaction with the ground.

The Deutsches Zentrum für Luft- und Raumfahrt (DLR) develops wake vortex advisory systems for airports to predict wake vortex behavior. Highly resolving large eddy simulations (LES) conducted on the SuperMUC supercomputer provide valuable insights in the physics of wake vortex behavior under various atmospheric conditions. These LES contribute indispensable guidance for the development of the real-time/fast-time wake vortex models [2].

A particular risk prevails during final approach, where the vortices cannot descend below the flight path, but tend to rebound due to the interaction with the ground. We could get several steps further in understanding the underlying physics with the help of numerical simulations.

Moreover, there is a strong appeal for artificial devices for destruction of wake vortices, or at least for weakening them to an uncritical level. A novel method to accelerate vortex decay, exploiting fundamental properties of vortex dynamics, was investigated by LES. Vortex decay can be initiated locally and accelerated in the vicinity with dedicated obstacles, so-called plate lines, installed at the ground, see Fig.1. It is found that a plate line may enforce the wake-vortex decay by more than 20% of the initial circulation. Based on simulations we developed a method how to optimize plate line design.

Results and Methods

A main achievement in that project was the development of a new Hybrid simulation method, coupling RANS (Reynolds averaged Navier-Stokes) and LES simulations [3]. Extending that method we could simulate the complete landing phase including final approach, flare, touchdown, and vortex decay [4], Fig.1.

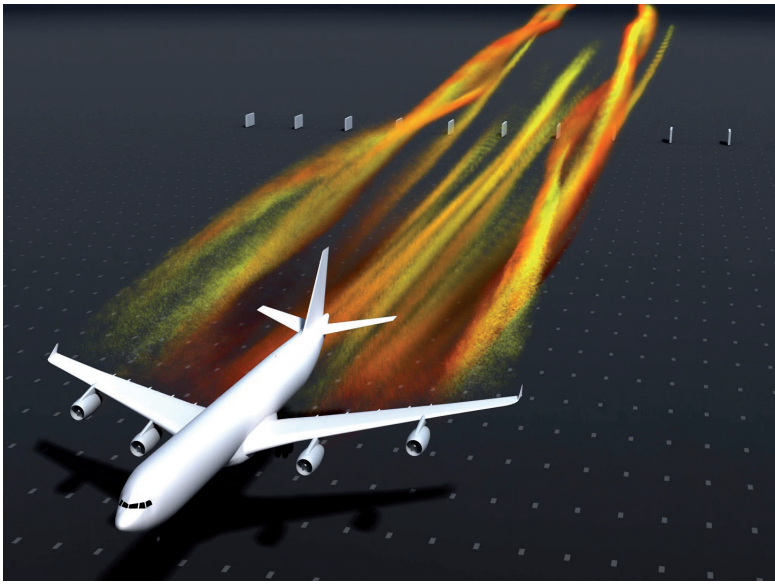

Figure 1: LES of wake vortex evolution during landing.

This allowed a detailed investigation of a landing aircraft and the simulation of the generated wake. Figure 1 depicts a snapshot of a high-end animation. Tracer particles were injected behind the aircraft. The complex vorticity distribution emerging directly behind the aircraft and the subsequent vortex roll-up to a single counter-rotating vortex pair can be observed. In the near wake flow the outer flap-tip vortices, the wing-tip vortices, the inner engine-nacelle vortices, and the vortices detaching from the wing-fuselage junction can be distinguished. Figure 2 depicts the wake evolution in the crosswind situation with a plate line during approach and after touchdown at the times -3.2 s, 0.4 s, 19.2 s and 60 s, where the time stamp 0 s is assigned to the instant of touchdown. The first effect of the plate line can already be identified at -3.2 s via a small gap of the vorticity layer generated at the ground. Already at

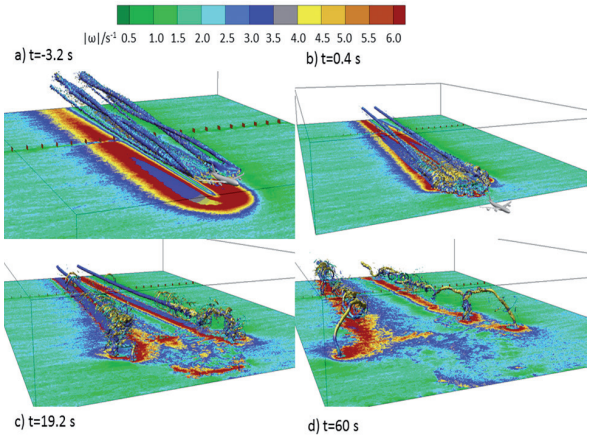

Figure 2: Simulation of wake vortex evolution of a landing aircraft at weak crosswind.

19 s Ω-shaped secondary vortices have detached from the plates and approach the primary vortices driven by self-induction.

The effect of the plate lines starts first at the downwind vortex and a few seconds later at the upwind vortex. A half vortex time scale later vigorous helical secondary vorticity structures have wrapped completely around the wake vortices and travel to either side of the plate line again driven by self-induction. The plate line concept exploits vortex dynamics by generating powerful secondary vortices that first actively approach the primary vortices and then actively propagate along the primary vortices finally leading to accelerated vortex decay. The simulations consumed approximately 3 million core hours per year. Typically the simulations run on 2048 processors. The amount of storage needed depends on the application. High end animations generated with *autodesk* consume about 5 TB of flow data.

Plate line optimization with Kriging
This study was performed in to find an optimal design of plate lines for operational use. It is dedicated to the numerical optimization of the above mentioned plate lines, see Fig. 1, in terms of plate distance as well as the aspect ratio of the plates. Varying both parameters, plate line distance and aspect ratio of the plates, we wish to find optimum values to maximize the effect. As we are dealing with a non-linear, unknown, cost-expensive black-box function, the global optimization is a difficult task. The number of possible simulations is strongly limited by cost and time. Here we employ a Gaussian process regression approach called Kriging that gives optimal values for functions in a multi-parameter space.

Figure 3 shows the mean value of the estimated plate line benefit in the entire parameter space. The A380 simulations reveal a optimum in the region of (plate aspect ratio = 0.25, plate separation 0.5).

On-going Research / Outlook

SuperMUC is our most important instrument for the study of fluid dynamics of wake vortices. Complex flows with a large spectrum of relevant scales, complex geometries with flow separation and high Reynolds numbers cannot be treated without smart algorithms on powerful supercomputers. One direction of research in fluid dynamics is the enhancement of the level of detail and the interaction of very small and very large scales. This leads to rapid increase of grid points and is strongly dependent on the development of supercomputers as the installation of SuperMUC Phase 2. However, many CFD codes originate from the early 80th where no one ever expected to use more than 100 processors. Appropriate scaling properties of memory and computation time where not a critical issue. In future we have to improve the scalability of our codes.

We already applied for a new project called *"Virtual Flight in a Realistic Environment"* to develop a two-way coupling of our RANS and LES codes. This project directly addresses the needs of the aircraft industry that classified the wake – vortex issue as "top level requirement". It will bring forward aircraft industry as well as air traffic management.

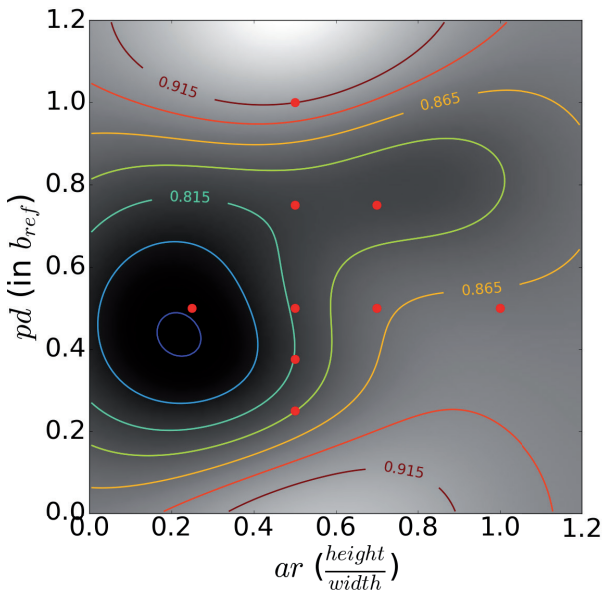

Figure 3: Kriging estimation of optimal plate design parameters for A380, simulation points (red).

References and Links

[1] http://www.pa.op.dlr.de/wirbelschleppe/
[2] Misaka et al.,(2012) Vortex Bursting and Tracer Transport of a Counter-Rotating Vortex Pair, Physics of Fluids
[3] Misaka, Takashi and Holzäpfel, Frank and Gerz, Thomas (2015) Large-Eddy Simulation of Aircraft Wake Evolution from Roll-Up until Vortex Decay. AIAA Journal, 53, pp. 2646-2670.
[4] Stephan A, Holzäpfel F, and Misaka T.(2014) Hybrid Simulation of Wake-Vortex Evolution During Landing on Flat Terrain and with Plate Line. Int. J. Heat & Fluid Flow, 49, pp.18-27.

Fully-resolved, finite-size particles in statistically stationary, homogeneous turbulence

Research Institution
Institute for Hydromechanics, Karlsruhe Institute of Technology

Principal Investigator
Markus Uhlmann

Researchers
Agathe Chouippe, Todor Doychev

Project Partners
–

SuperMUC Project ID: pr83la (Gauss Large Scale project), pr58cu

Introduction

Particle suspension in a turbulent flow is a phenomenon that occurs in many industrial and natural processes such as fluidized beds and combustion devices or pollutant dispersion in the atmosphere and raindrop formation in clouds. In such kind of systems several fundamental mechanisms interact making the whole picture highly complex.

In this project we aim at investigating with high-fidelity simulations the hydrodynamical interaction between a turbulent flow and a large number of finite-size particles (particles with diameter comparable or larger than the smallest flow structures) falling under the influence of gravity. We focus more specifically on homogeneous isotropic turbulence (therefore avoiding any wall turbulence influence), sustained at a statistically steady state. Particles are selected to be heavy, rigid and have a spherical shape [1].

Despite the academic aspect of such configuration the parameter space is very wide and a full description of the interaction mechanisms is still missing. More specifically, the need for high resolution with finite-size particles has made large systems almost impossible to accurately tackle with numerical simulation until the last few years. Our aim is here to improve the understanding of the fundamental processes involved with the aid of new fully-resolved data.

When gravity is accounted for, a relevant parameter for the particles is the Galileo number Ga that estimates the intensity of gravitational effects relative to viscous ones. Former works on single or many particle settling in initially ambient flow showed the complexity of particle falling even in the absence of turbulence. Indeed, a single particle exhibits different kinds of motion and wakes with increasing Ga [2], starting from steady axisymmetric regime, to steady oblique and then oblique oscillating regime, for finally becoming fully chaotic. In the work of [3] the influence of many particle settling has been investigated for two Galileo numbers corresponding either to a steady axisymmetric or a steady oblique regime, both in dilute systems (solid volume fraction equal to 0.005). They observed that in the second configuration particles tend to accumulate in vertical columns with the formation of large columnar flow structures. This clustering is then associated with a significant increase of the settling velocity. In the case of axisymmetric wake regime, no significant interaction mechanism has been highlighted.

The aim of the current project is now to investigate the influence of ambient turbulence on many-particle settling for Galileo number and solid (and mass) volume fraction similar to [3]. Cases with Ga=0 have also been explored with the aim of isolating the influence of pure turbulence on particle motion. For this latter case we carried two classes of simulations with two different particle sizes relative to Kolmogorov lengthscale (i.e. the characteristic lengthscale of the smallest structures of the flow).

Relevant parameters for the background flow are the Reynolds number and the characteristic length and time scales of the large and small structures. We consider here a Taylor micro-scale fluid with a Reynolds number of the order of 100-140 so that separation between large and small scales is effective, and particles with diameter ranging from 5 to 10 times the Kolmogorov lengthscale.

Results and Methods

An essential aspect of the study is the simulation with high fidelity of all the scales of turbulence as well as the hydrodynamics around each particle, while also considering a large number of inclusions. This implies on one side to use fine space resolution at particle scale but also relatively large simulation domains, leading to large simulation boxes. We employ the Immersed Boundary Method of [4] which enforces no slip boundary condition at particle surface and thus enables this full resolution of the flow (as visible on fig. 1). The method is MPI parallel with domain decomposition and optimized for application on large systems.

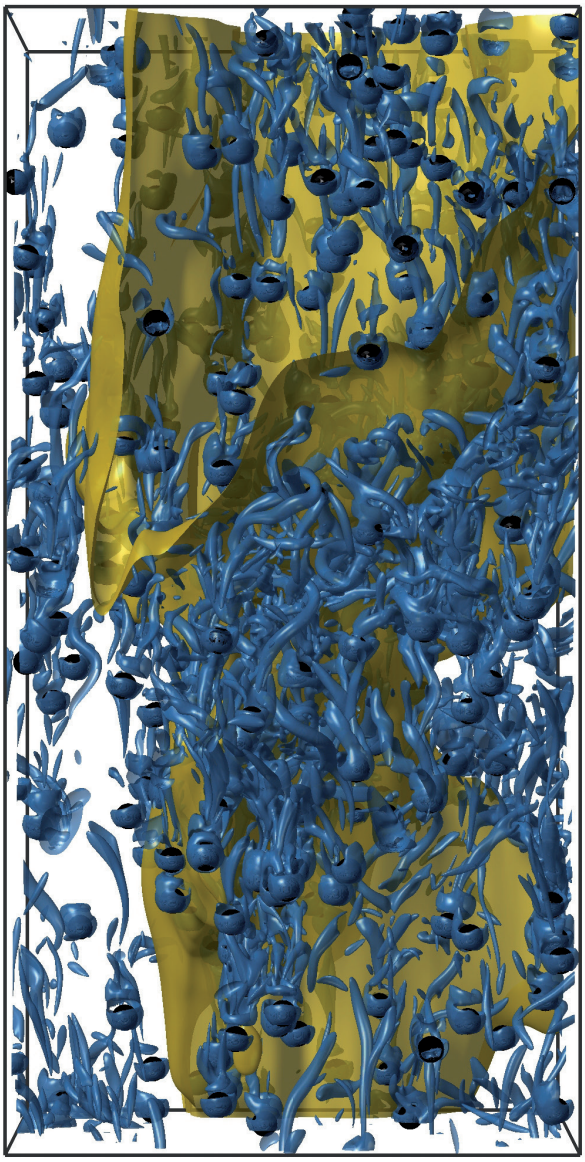

Figure 1: Flow visualization of a sub-volume of the domain (one fourth of each linear dimension is shown): Particles are represented in black, blue surfaces correspond to isosurfaces of the Q criterion and display the wakes of the particles. Isosurfaces of the vertical component of the velocity field are shown in yellow. For the velocity field, the flow has been box-filtered with a box width of approximately 5 particle diameters in order to extract the large scale structure.

Methods

The continuous phase is resolved with direct numerical simulation of the incompressible Navier-Stokes equations with standard fractional step. Temporal discretization is made with semi-implicit Crank-Nicholson method for the viscous term and a low storage three-step Runge-Kutta procedure for the non-linear part. The different spatial operators are evaluated by central finite-differences on a staggered grid with uniform and isotropic mesh. The temporal and spatial accuracy of the scheme are of second order. The translational and rotational motion of the particles is computed with the Runge-Kutta discretized Newton equations. Interpolation between Lagrangian end Eulerian quantities are done with regularized delta function yielding smooth variation of the hydrodynamic forces when particles move. Turbulence is generated with the large scale forcing of Eswaran and Pope [5], which has the advantage of sustaining a turbulent flow at statistically steady state that remains stable even in the presence of many finite-size particles [6].

We used periodic boundary conditions in the three directions and in the absence of gravity we used cubic simulation boxes with an extension of 64 to 128 particle diameters, and computational grid of 1024^3 to 2048^3. With the other cases we used domains elongated in the direction of gravity with typical extension of 85 and 170 particle diameters in the horizontal and vertical directions, respetcively, and typical computational grid of 2048x2048x4096. We used 1024 to 16384 processor cores, consumed 48 million core-hours and produced approximately 450TB of data.

Results

Our first results concern clustering events. We qualitatively observe that turbulence modifies accumulation scenario with the formation of clusters for both nonzero Galileo configurations, while it is observed only for steady oblique configurations with ambient flow. Voronoï tessellation analysis performed on the steady oblique regimes shows that turbulence also quantitatively decreases accumulation and tend to disturb particle columns as well as the columnar large scale structures of the flow. In the absence of gravity (Ga=0) particle accumulation is even smaller, and the strongest clustering is observed with the smallest particles.

As a possible consequence of the disturbance of the large flow columns by turbulence in the nonzero gravity case, the mean settling velocity of the particles is not enhanced anymore and tends to values of the same order as the settling velocity of one single particle in ambient flow.

On-going Research / Outlook

Data generated is still object of analysis, and we are now characterizing the impact of turbulence on the structure of the wake as well as the influence of the characteristic turbulence velocity relative to the settling velocity.

References and Links

[1] http://www.ifh.kit.edu/english/26_804.php

[2] M. Jenny, J. Dušek and G. Bouchet. 2004. Instabilities and transition of a sphere falling or ascending freely in a Newtonian fluid. J Fluid Mech. 508, 201-239. DOI:http://dx.doi.org/10.1017/S0022112004009164

[3] M. Uhlmann and T. Doychev. 2014. Sedimentation of a dilute suspension of rigid spheres at intermediate Galileo numbers: the effect of clustering upon the particle motion. J. Fluid Mech. 752, 310-348. DOI: http://dx.doi.org/10.1017/jfm.2014.330

[4] M. Uhlmann. 2005. An immersed boundary method with direct forcing for the simulation of particulate flows. J Comp. Phys. 209. 448-476. DOI:http://dx.doi.org/10.1016/j.jcp.2005.03.017

[5] V. Eswaran and S.B. Pope. 1988. An examination of forcing in direct numerical simulations of turbulence. Comp. Fluids. 16(3), 257-278. DOI:http://dx.doi.org/10.1016/0045-7930(88)90013-8

[6] A. Chouippe and M. Uhlmann. 2015. Forcing homogeneous turbulence in DNS of particulate flow with interface resolution and gravity, Phys. Fluids. 27, 123301. DOI: http://dx.doi.org/10.1063/1.4936274

High accuracy molecular dynamics simulation of fluids at interfaces

Research Institution
[1]Laboratory of Engineering Thermodynamics, University of Kaiserslautern

Principal Investigator
Martin Horsch[1]

Researchers
Stefan Becker[1], Katrin Stöbener[1], Stephan Werth[1], Stefan Eckelsbach[2], Wolfgang Eckhardt[3], Alexander Heinecke[3], Nikola Tchipev[3], Hans-Joachim Bungartz[3], Jadran Vrabec[2], Hans Hasse[1]

Project Partners
[2]Thermodynamics and Energy Technology, University of Paderborn
[3]Scientific Computing in Computer Science, Technische Universität München

SuperMUC Project ID: pr83ri (Gauss Large Scale project)

Introduction

At molecular resolution, the interface which separates coexisting fluid phases, e.g. between a liquid droplet and a surrounding vapour, is not sharp but continuous. Both the interface and the coexisting phases are subject to significant fluctuations which also contribute to the surface tension. Therefore, it is crucial to accurately reproduce the surface tension and analyse the properties of nanodroplets with a method that resolves the structure of fluid matter at the molecular level. Here, molecular dynamics (MD) simulation is applied to analyse interfacial phenomena qualitatively and capture interfacial properties quantitatively.

For MD codes, as for any software, there is a trade-off between generality and optimality for a single purpose, which no particular implementation can completely evade. The program *ls1 mardyn* (large systems 1: molecular dynamics) expands the temporal and spatial range of scales accessible to molecular simulation, with a focus on inhomogeneous systems (e.g. at interfaces) and non-equilibrium thermodynamics [1]. It is available as free software [2] and was applied and extended within the present project to accurately and efficiently simulate systems with vapour-liquid interfaces.

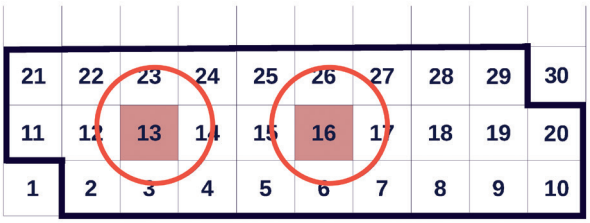

Figure 1. Hyperthreaded sliding window (represented here in two dimensions, instead of three as in the actual code), employed by the *ls1 mardyn* version optimized for the Intel Sandy Bridge EP architecture. At the present stage, the two highlighted cells 13 and 16 are concurrently processed by SMT threads. Thereby, cell 2 is considered as a neighbour cell for the last time, leaving the sliding window, and cell 27 enters the sliding window; cells 28 and 29 are prefetched.

By massively-parallel high performance computing, e.g. using the *ls1 mardyn* program, molecular simulation becomes an experiment *in silico*. Due to their heterogeneous structure, systems with a phase boundary have stronger long-range interactions, which also need to be taken into account [3]. Load balancing by recursive bisection of the simulation volume, using a linked-cell data structure, significantly improves the scalability of the code. In this way, supercomputers with a large number of cores can be used efficiently both for homogeneous and heterogeneous systems [1].

Results and Methods

Molecular simulations with system dimensions that far exceed the cut-off radius, beyond which a mean-field approach is employed for the intermolecular interactions, are most efficiently parallelized by space decomposition schemes. Thereby, the simulation volume is subdivided into smaller subvolumes (one for each process) that ideally carry the same load. In *ls1 mardyn*, an interface class for the domain decomposition scheme permits the generic implementation of different load balancing strategies operating on spatial subdomains with a linked-cell data structure [1].

To make use of hyperthreading, a lightweight shared-memory parallelization was implemented: By employing a sliding window, as sketched in Fig. 1, including prefetched cells, two threads can operate concurrently on independent cells. Thereby, it is avoided that threads work on directly neighbouring cells simultaneously. Therefore, a barrier, causing comparably little overhead on a hyperthreading core, is required after each thread has processed a cell. This allows the execution of one MPI rank per core with two OpenMP threads to create sufficient instruction level parallelism, leading to a significant performance improvement on SuperMUC.

Molecule data outside the sliding window are stored in the form corresponding to the object-oriented software

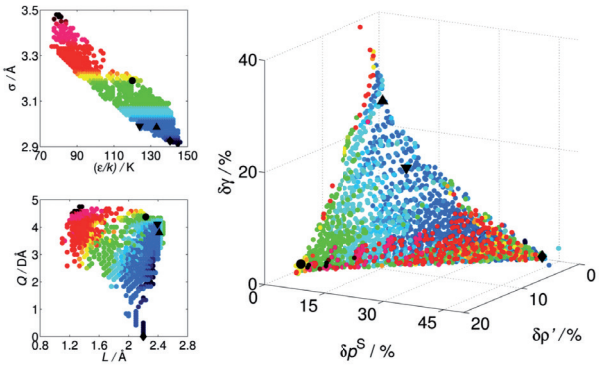

Figure 2. Pareto set within the parameter space (left) of the 2CLJQ class for modelling carbon dioxide as well as its representation in the objective space (right), taking three criteria into account, i.e. the average deviation of (correlated) model properties from (correlated) experimental data for the saturated liquid density, the saturated vapour pressure, and the vapour-liquid surface tension [3].

architecture, i.e. as a dynamic array of structures (AoS) containing Molecule objects. When the sliding window is shifted further and covers a new cell, the positions and velocities of the molecules in that cell are converted to a structure of arrays (SoA) form suitable for vectorization.
The two-centre Lennard-Jones plus point quadrupole (2CLJQ) class of molecular models provides a straightforward description of the intermolecular interactions for many low-molecular compounds [3]. Furthermore, vapour-liquid interfaces – for systems with large numbers of molecules – can for this model class be efficiently treated using a long-range correction that was developed within the present project, yielding a high accuracy for the computed surface tension. Accordingly, simulations of the vapour and liquid phases in direct coexistence were carried out along the whole vapour pressure curve, and the model parameters were varied systematically, covering the whole parameter range of 2CLJQ models for real fluids available from the literature.

Based on the simulation results for the surface tension, a precise correlation was obtained in dependence of the parameters, which is universally valid for the whole model class. In combination with the available knowledge of bulk fluid properties at vapour-liquid equilibrium conditions, this correlation can be used to adjust 2CLJQ models to multiple types of real fluid properties, e.g. by multi-criteria optimization based on an evaluation of the Pareto set, cf. Fig. 2, where this is illustrated for carbon dioxide using three criteria: Average agreement for the saturated liquid density, the vapour pressure, and the surface tension [3].

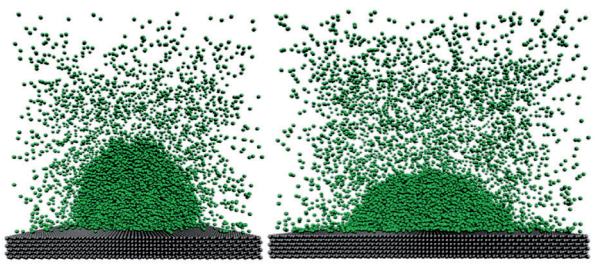

Figure 3. Snapshots from simulations of sessile droplets at the same temperature on walls of the same density and structure, differing only in the magnitude of the dispersive fluid-solid interaction, which is larger on the right side, yielding increased wetting, i.e. a smaller contact angle [4].

Furthermore, three-phase contact between the vapour-liquid interface of a sessile droplet and a solid surface was systematically investigated for the truncated-shifted Lennard-Jones potential (LJTS), considering the scenario with a perfectly planar solid substrate, cf. Fig. 3, as a reference case for future work on patterned and rough surfaces. Both the fluid and the solid were modelled by the LJTS potential (with different parameters), varying both the solid density and the dispersive fluid-solid interaction energy. From the present simulation results, a universal correlation was derived for the contact angle as a function of the well depth of the unlike fluid-solid interaction and further characteristic parameters, which was found to carry over to other similar systems with a good accuracy [4].

On-going Research / Outlook

The present project partners LTD (Kaiserslautern), SCCS (Munich), and ThEt (Paderborn) have applied for a follow-up computing project entitled *Scalable, Performant and Resilient Large-scale Applications of Molecular Process Engineering* (SPARLAMPE), project ID pr48te, which was approved and started in March 2016.

Based on the approaches for high-accuracy simulations established within the present computing project, a more direct orientation towards applications in mechanical process engineering has become feasible, based on the scalable and performant implementation of algorithms for heterogeneous sytems in *ls1 mardyn*. For validated and reliable molecular models, the SPARLAMPE project will further focus on linear transport coefficients in the bulk fluid and at interfaces as well as transport processes near and far from equilibrium.

These are necessarily HPC applications, since finite-size effects need to be assessed rigorously and the computation of transport properties requires a much more extensive sampling than other thermodynamic properties. In particular, as larger systems also require a longer relaxation time to reach equilibrium, fast equilibration techniques will be developed, implemented and employed for petascale simulations.

The project partners would like to thank Colin Glass, Christian Holm, George Jackson, Philipp Neumann, and Jayant Singh for fruitful discussions.

References and Links

[1] Niethammer, C., Becker, S., Bernreuther, M., Buchholz, M., Eckhardt, W., Heinecke, A., Werth, S., Bungartz, H.-J., Glass, C. W., Hasse, H., Vrabec, J., and Horsch, M., J. Chem. Theory Comput. 10(10): 4455-4464 (2014).
[2] http://www.ls1-mardyn.de/
[3] Werth, S., Stöbener, K., Klein, P., Küfer, K.-H., Horsch, M., Hasse, H., Chem. Eng. Sci. 121: 110-117 (2015).
[4] Becker, S., Urbassek, H. M., Horsch, M., and Hasse, H., Langmuir 30(45): 13606-13614 (2014).

Investigation of three-dimensional Dynamic Stall

Research Institution
Deutsches Zentrum für Luft- und Raumfahrt e.V. (DLR)

Principal Investigator
Anthony D. Gardner

Researchers
Kurt Kaufmann, Kai Richter

Project Partners
–

SuperMUC Project ID: pr83su

Introduction

Dynamic Stall is one of the most challenging flow phenomena existing in the field of helicopter aerodynamics. Under fast forward flight or maneuver, the retreating blade of a helicopter can temporarily stall (Dynamic stall), leading to a rapid change in pitching moment. Dynamic stall is an inherently unsteady phenomenon which is significantly affected by the interaction of the laminar/turbulent boundary layer transition, flow separation, vortex growth and propagation, and reattachment of the flow. During dynamic stall, large peaks in lift, pitching moment and drag appear, and these cause an undesirable increase in the mean drag. The flight envelope of a helicopter is limited to avoid the torsional impulse of dynamic stall, which can cause structural damage to the rotor by overloading the pitch links, and structural vibrations which reduce the comfort of passengers or damage the cabin.

Dynamic stall results in a completely different aerodynamic behavior compared to static stall and cannot be effectively approximated by static measurements or computations. The elasticity of helicopter rotor blades, and the complexity of the aerodynamics of fast forward flight mean that experimental data of sufficient quality to validate numerical models is only possible using simplified, structurally stiff wind tunnel models. These experiments create a complex, unsteady aerodynamic flow, which is nevertheless relatively simple when compared with the free flight of a complete helicopter. The resulting flow can be directly compared with high resolution numerical computations.

Historically, research into dynamic stall has strongly relied on prescribed pitching on two-dimensional (2D) models at constant Mach numbers. Wind tunnel experiments on finite span, nominally 2D models for these test cases, provide reference data for 2D and three-dimensional (3D) numerical simulations, and this method has led to significant advances in the numerical prediction of dynamic stall. To improve future rotorcraft, it is necessary to understand how 3D dynamic stall differs from 2D dynamic stall, so that the value of the existing 2D dynamic stall experimental data can be maximized.

Results and Methods

Unsteady three-dimensional Reynolds-Averaged Navier-Stokes (RANS) Simulations using the finite volume solver DLR-TAU were carried out on two finite wing geometries to fit the experiments carried out by ONERA in the F2 wind tunnel [1] and by the DLR in the Cross Wind Simulation Facility (SWG) [2], respectively. Each computation used 3,000 time steps per pitching period and 200-600 inner iterations leading to a consumption of 120,000 to 560,000 CPUh per dynamic stall computation.

Initial results were produced for the ONERA finite-wing configuration. This is a configuration for which wind tunnel results and comparative results with the elsA code of ONERA exist for comparison with the TAU results. The elsA code is a structured URANS code whereas the TAU code is an unstructured URANS code, but the approaches in each are similar enough that comparable results should be produced.

A test case for a freestream flow Mach number of M=0.16 with dynamic sinusoidal pitching of the airfoil at α=17+5sin(ωt) at a frequency of 5Hz was used. This

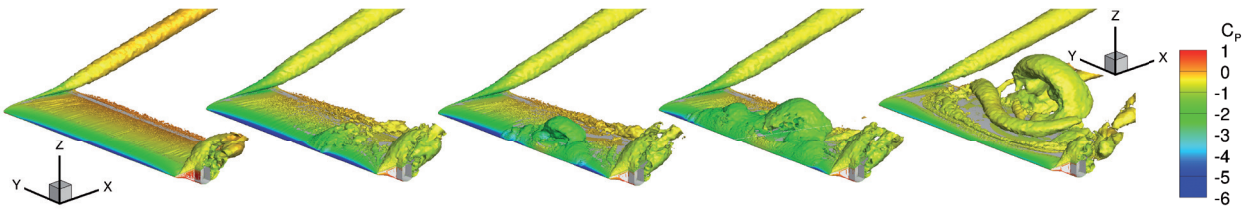

Figure 1: Dynamic Stall evolution around the ONERA finite wing. Visualized by means of isosurfaces of the λ2 criterion and Cp contour plots.

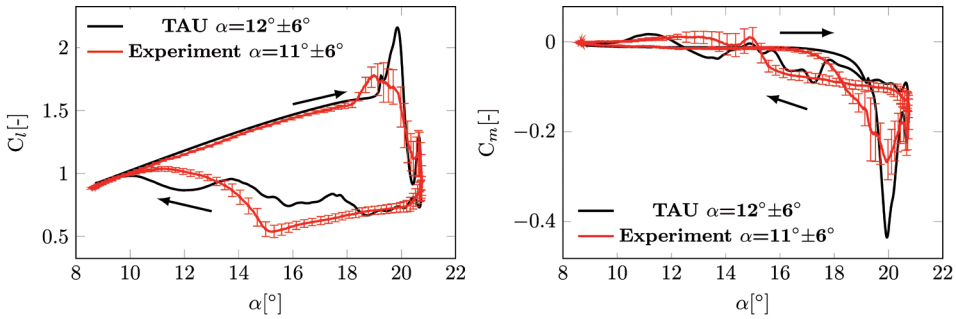

Figure 2: Comparison between the experiment α=11°±6° and the DLR-TAU computations α=12°±6° for the dynamic case at a single spanwise position at r/R=0.68. The numerical data is shifted by Δα=-1° to match the effective angle of attack.

caused dynamic stall at the start of pitch-down, as illustrated with contours of λ2 in Fig. 1 (Flow is from left to right). At the left, the attached flow is illustrated, and the tip vortex of the wing is visible, as well as flow separation due to the root geometry. With increasing time (Left to right), a dynamic separation vortex appears in the middle of the wing and propagates downstream, expanding in breadth until the whole wing is separated.

This grid used "best-estimate" values from static computations, and 2D dynamic stall computations, but the results are comparable with those from elsA and from the experiment. The results of these comparisons have been documented in a journal article [3] and a conference article [4].

To investigate the effects of the grid resolution on three-dimensional dynamic stall a convergence study was carried out on the DLR finite wing experiment [2] leading to a grid size of 21 Million nodes.

An additional investigation using the wind tunnel walls, shows that for the static case, the wall effects reduce the stall angle by approximately 1°. Therefore, in Figure 2 the experimental data is compared with numerical data using a 1° increased mean angle of attack. For clarity the numerical data was shifted by Δα=-1.0° to display approximately the same effective angle of attack. The lift overshoot occurs in the same region as in the experiment and the minimal pitching moment occurs at the same position. The beginning of the pitching moment drop is earlier in the experiment and the peak is lower. Nevertheless, a very good agreement with the experimental data is achieved.

Figure 3 (Flow is from right to left) shows the process of the dynamic stall event using isosurfaces of the λ2 criterion and color coded with Cp for several time instants. At α=10.1° on the upstroke (top left) the flow is attached and the wake as well as the blade tip vortex are visible. Separation starts to occur at the leading edge of the parabolic blade tip and the vortex formed wraps around the blade tip (top middle). Simultaneously the blade tip vortex has moved further inboard and trailing edge separation sets in along the entire span. At α=16.6° on the upstroke the dynamic stall vortex starts to evolve at the leading edge between r/R≈0.5 and r/R≈0.7 (top right). High spanwise propagation speeds are reached and at α=17.2° on the upstroke (bottom left) the dynamic stall vortex and the vortex of the parabolic tip vortex have jointed, building a separation line along the leading edge, from r/R≈0.2 up to the blade tip. At α=17.9° on the upstroke (bottom middle) the flow over the entire wing is separated and the first ring shaped structures are moving downstream. Eventually, at α=13.9° on the downstroke (bottom right) reattachment has set in, a suction peak along the span is visible and the remaining small vortex structures are propagating downstream.

On-going Research / Outlook

The investigations showed a good comparison between the experimental data and the numerical simulations and lead to a better understanding of the evolution of three-dimensional dynamic stall on pitching finite wings. In a further step, simulations investigating the influence of rotation using different speeds of rotation, pitching motions and advance ratios will be performed.

References and Links

[1] Le Pape, A., Pailhas, G., David, F., and Deluc, J.-M., Extensive Wind Tunnel Measurements of Dynamic Stall Phenomenon for the OA209 Airfoil Including Three-Dimensional Effects, 33rd European Rotorcraft Forum, Kazan, Russia, September 11-13, 2007.

[2] Merz, C.B., Wolf, C.C., Richter, K., Kaufmann, K., Mielke, A., Raffel, M., Spanwise Differences in Static and Dynamic Stall on a Pitching Rotor Blade Tip Model, 41st European Rotorcraft Forum, Munich, 1-4 Sept. 2015.

[3] Kaufmann, K., Costes, M., Richez, F., Gardner, A. D., Le Pape, A., Numerical Investigation of Three-Dimensional Static and Dynamic Stall on a Finite Wing, AHS Journal, Vol. 60, No. 3, 2015.

[4] Kaufmann, K., Gardner, A. D., Costes, M., Comparison between two-dimensional and three-dimensional dynamic stall, STAB 2014, Munich, 4-5 November, 2014.

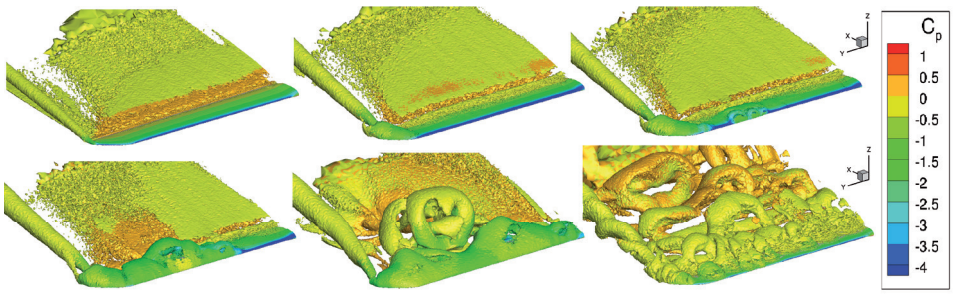

Figure 3: Dynamic Stall evolution around the DLR finite wing. Visualized by means of isosurfaces of the λ2 criterion and Cp contour plots.

In-situ Analysis of Flame Structures of Turbulent Non-Premixed Flames

Research Institution
NTFD, TU Bergakademie Freiberg
Principal Investigator
Christian Hasse
Researchers
Felix Dietzsch, Sebastian Popp, Danny Messig and Michael Gauding
Project Partners
ISUT, Otto von Guericke University, Magdeburg (Dominique Thévenin)

SuperMUC Project ID: pr83xa

Introduction

For non-premixed flames efficient mixing of fuel and oxidizer plays an important role. For turbulent diffusion flames understanding the interaction of mixing and chemical reactions is crucial. Usually these flames exhibit high mixing rates which enhance combustion efficiency due to increased reaction rates. However above a certain limit, the so called quench limit, excessive mixing rates lead to local extinction of the flame, c.f. figure 1.Extinction then leads to increased emissions or even to flame destabilization and blow out.

As a basis for our analysis we assume that chemical reactions are fast and only occur around thin layers of stoichiometric mixture. Furthermore, it is assumed that transport processes only occur along trajectories orthogonal to the scalar surface. This quasi one dimensional description of non-premixed combustion is known as the flamelet assumption. Since most technical combustion processes occur in the flamelet regime it is used in many simulations as a model to describe the turbulence-chemistry interaction. Previously, the flamelet approach was applied for flame structure analyzes of jet flames with few local extinction events, using scale resolving simulations [1]. In order to gain more insights into local extinction and get access to the temporal evolution of instantaneous flamelet structures, we conducted direct numerical simulations (DNS) of a turbulent temporally evolving jet flame exhibiting strong local quenching. The extraction of relevant quantities along flamelets is done in-situ.

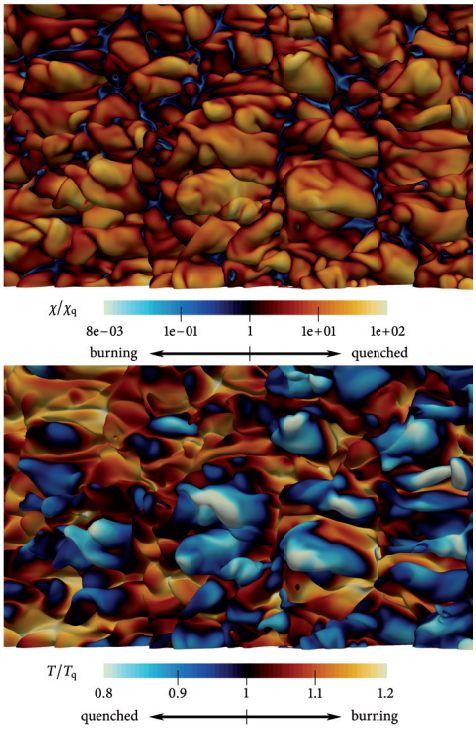

Figure 1: Snapshot of a simulation showing local extinction

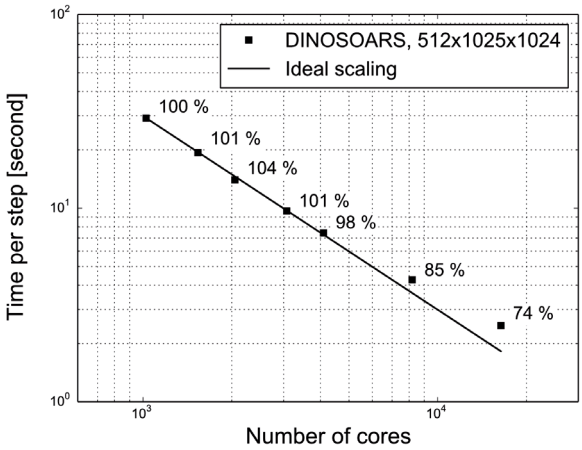

Figure 2: Strong scaling of DINOSOARS

Results and Methods

For our in-situ analysis we solved the three-dimensional reactive Navier-Stokes equations in a low Mach number formulation using our in-house solver DINOSOARS [2]. High order discretization is achieved using sixth order finite differences in space and a fourth order Runge-Kutta scheme for time integration. In order

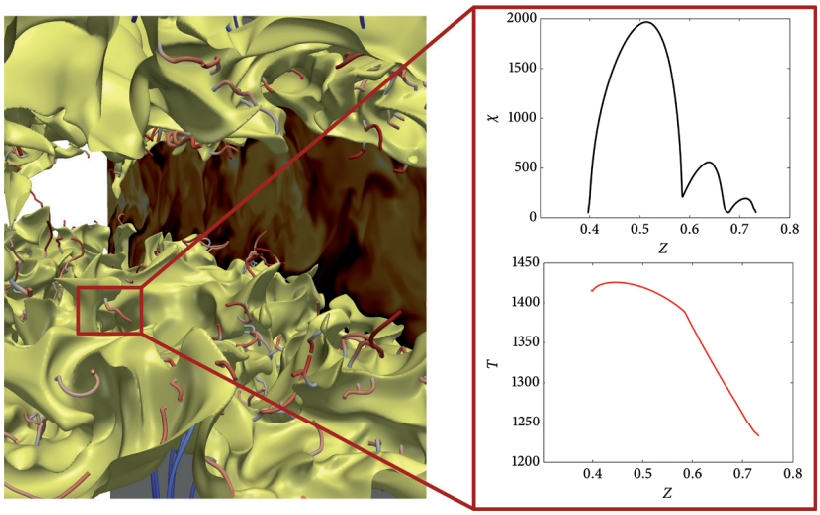

Figure 3: Sample snapshot of a simulation illustrating gradient trajectories interacting with the iso-scalar surface of stoichiometric mixture fraction (left). The right figure shows scalar dissipation rate and temperature over mixture fraction of the highlighted trajectory.

to preserve also high accuracy for the pressure solver DINOSOARS utilizes a spectral Poisson solver. Therefore, parallelization is done in two dimensions using the 2DECOMP&FFT library. During the first phase of our project we improved the scaling of our solver resulting in an excellent scaling behavior of up to 16384 cores as can be seen in figure 2.

The numerical setup we used for our analysis is based on a temporal evolving jet flame that was already extensively studied in the literature [3]. In addition to the standard flow equations we also solve transport equations for massless particles. For our analysis, we extended the procedure proposed in [4] following the concept of dissipation elements [5]. A key point of this concept is that the underlying scalar field is decomposed into local extremal points.These extremal points are connected by gradient trajectories which for our analysis can be understood as flamelets.

During the initialization of our simulation we distribute particles on the iso-scalar surface, describing the stoichiometric mixture fraction.Contrary to Lagrangian particles we consider particle transport due to both the convection of the flow field and the molecular diffusion of the iso-scalar surface. After particles have been advanced they are used as seed points for gradient trajectories. Using a second order Runge-Kutta scheme trajectories are integrated along the ascending and descending gradient until a local extremum is reached. All quantities needed during integration are interpolated by a tri-cubic interpolation procedure. The algorithm we use for tracking trajectories is based on a static allocation principle. Each processor holds a component of the mesh (block) and integrates trajectories until they leave the block owned by the processor. Since the path of a trajectory is not known a priori we use a buffer array that is communicated as soon as the trajectory leaves the block of the current processor. A globally communicated trajectory count is maintained so that all processors may monitor how many trajectories have yet to terminate. Once the count goes to zero, trajectory integration is finished and the code continues with the next time step. The left part of Figure 3 shows a snapshot of a simulation that was run on SuperMUC. It is clearly seen that a strong interaction of the yellow iso-scalar surface of stoichiometric mixture fraction with our extracted flamelets (thin tubes) exists. In the right part of figure 3 a plot of χ and T in flamelet space along one extracted trajectory is shown.

Having a jet Re=9000, the case required at least 512x1024x1024 grid points in order to resolve the smallest scales and was run on 4196 CPUs. A typical simulation of this size runs for about 4 days and generates nearly 15TB of data, depending on the number of particles considered. The storage requirements and the number of files usually generated were already optimized by implementing a fully parallel HDF5 based IO both for particles and trajectories.

On-going Research / Outlook

At the time of application SuperMUC had the largest amount of memory per core and one of the fastest interconnects which helped us in developing our code using different caching strategies. Our project also highly benefited from the high IO bandwidth and the large storage capacities of WORK.

We currently put considerable effort into the optimization of our in-situ algorithm to allow for a larger number of particles, and therefore also for a larger number of trajectories.n addition to that we wish to increase the resolution of our simulations since proper detection of extremal points is very sensitive to grid resolution. A large number of grid points, however, highly stresses the interconnect and also the available disk space, since we have to dump our data once every time step, to get proper temporal resolution.

References and Links

[1] Popp S., Hunger F., Hartl S., Messig D., Coriton B., Frank J.H., Fuest F., Hasse C., 2015 LES flamelet-progress variable modeling and measurements of a turbulent partially-premixed dimethyl ether jet flame. Combust. Flame 162, 3016-3029

[2] Abdelsamie A, Fru G, Oster T, Dietzsch F, Janiga G, Thévenin D. Towards direct numerical simulations of low-Mach number turbulent reacting and two-phase flows using immersed boundaries. Comput Fluids. 2016;131:123-141.

[3] Hawkes E.R., Sankaran R., Sutherland J.C., Chen J.H. 2007 Scalar mixing in direct numerical simulations of temporally evolving plane jet flames with skeletal CO/H2 kinetics. Proceedings of the Combustion Institute. 31, 1633-1640.

[4] Sripakagorn P., Mitarai S., Kosaly G., Pitsch H. 2004 Extinction and reignition in a diffusion flame: a direct numerical simulation study. J Fluid Mech. 518, 231-259.

[5] Wang L., Peters N. 2006 The length scale distribution function of the distance between extremal points in passive scalar turbulence. J Fluid Mechanics. 554, 457-475.

Large Eddy Simulation of turbulent flow interacting with complex structures

Research Institution
Fachgebiet Hydromechanik, Technische Universität München

Principal Investigator
Michael Manhart

Researchers
Wolfgang Schanderl

Project Partners
DFG

SuperMUC Project ID: pr84gi, h0022

Introduction

Turbulent flows are of intrinsic complexity due to the non-linear character of the governing equations. In a turbulent flow regime flow structures with a continuous spectrum of scales are interacting. With increasing Reynolds number, the range of the involved scales increases. Thus, typical technical or geophysical flows at large Reynolds numbers can only be predicted by the use of models reducing the degree of freedom. Computer simulations have emerged as a powerful tool to improve the understanding of such flows.

The presented project [1] aims at investigating the flow around a circular cylinder mounted on a flat plate. Even though the geometry is simple, the flow pattern around such a cylinder possesses several flow features interacting in a complex and highly dynamic way.

The main goals of the project are (i) to gain deeper understanding of the flow and its dynamics, and how these dynamics change with increasing Reynolds number and (ii) to perform simulations of such flows which can be used as reference for model development.

This project is currently being funded as a combined numerical/experimental study by the DFG. Parallel to the simulations performed at HLRB, experiments are performed at the Hydromechanics Laboratory at TUM. The experiments will be used for validation and to obtain complementary data to the numerical simulations, such as longtime records and high Reynolds number flow data. While the experiments have just started, first simulation results are already available and have been published [4].

For the large-eddy simulations within this project, the flow solver MGLET is being employed [2]. It uses a Finite Volume method to solve the incompressible Navier Stokes equations. A Cartesian grid with staggered arrangement of the variables enables an efficient formulation of the spatial approximations. An explicit third order low storage Runge-Kutta time step is used for time integration. Geometrically complex surfaces, arbitrarily curved, can be represented by an Immersed Boundary Method [3].

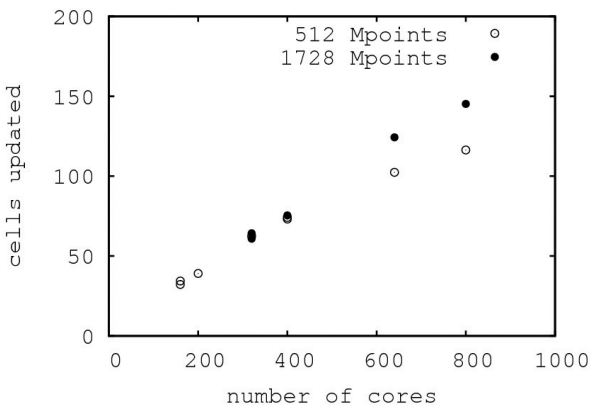

Figure 1: Total number of cells (in million) updated by one time step during one second as function of number of cores. Two problem sizes, 512×10^6 and $1{,}7 \times 10^9$ cells, are considered.

MGLET is parallelized by a domain decomposition method using MPI as a framework. For simple rectangular domains, the code has been tested so far for up to 8192 cores. As the purpose of the code is doing long runs, gathering statistics of the flow field, the applications are designed for maximum throughput on a number of cores as small as possible. Here the key number is the number of grid cells to be advanced by one time step within a CPU second. In Figure 1, this measure is plotted in terms of total number of grid points per second as a function of the number of cores on the SuperMUC. Figure 1 shows, that in one core-second, about 200000 grid points can be advanced in time by one time step. The problem sizes used for these benchmarks are 512 million and 1728 million cells,

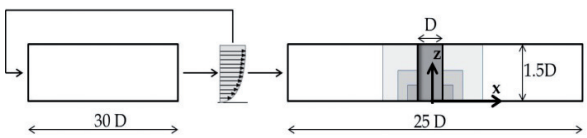

Figure 2: Side view of the setup. The grid around the cylinder is refined with three locally embedded grids [4].

respectively. The figure demonstrates that MGLET is well scalable for this number of cores and problem sizes. It has to be noted, however that this scalability is obtained in a regular one-block grid and with the basic second order solver. The scalability suffers, if multi-block grids are used, such as in the prediction of the flow around a cylinder on a flat plate which is described below.

Like other academic research codes, MGLET is being under constant development and improvement.

Results and Methods

We use MGLET to perform highly resolved LES of the flow around a wall-mounted cylinder at Reynolds numbers of Re=20000, Re=39000 and Re=78000 based on cylinder diameter and oncoming bulk velocity. The subgrid scale stresses are parameterized by the Wall-Adapting Local Eddy-Viscosity model (WALE). A free surface channel flow at a low Froude number is approximated by a free-slip condition at the upper wall. A fully turbulent open channel flow is set as inflow condition. It is generated by a so-called precursor simulation which is run in parallel to the main flow simulation, see Figure 2.

The region of interest around the cylinder/plate junction is resolved by three zonally embedded grids which give a total refinement by a factor of eight with respect to the global grid [4]. A grid study proofs convergence of the results over grid refinement. The configuration including the precursor simulation is documented in Figure 2.

For each Reynolds number, the grid has to be adapted according to the expected inertial stresses of the fluid. The finest configuration (Re=78000), using in total a number of 1.6 billion grid cells, can be run on 2048 cores with about 12 seconds per time step (SuperMUC). Compared to the Benchmark given in Figure 1, this is about a factor of 2-3 slower as expected for a configuration with a single block grid. Considering the complex communication patterns that arise from the zonally embedded grids, this can still be regarded as satisfying. After reaching a statistically stationary flow, the flow field has to be averaged over 1.2 million time steps to evaluate time-averaged values.

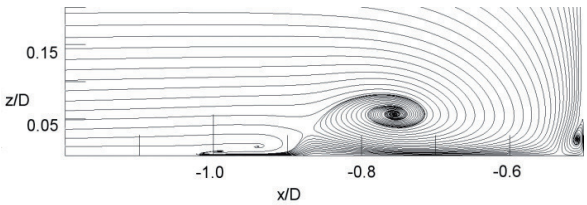

Figure 3: Time-averaged streamlines in the symmetry plane in front of the cylinder. Plotted is data at Re=78000.

The approaching boundary layer-type flow leads to a down-flow in front of the cylinder. This down-flow is forming a vortex when reaching the bottom plate, see Figure 3. Due to the main flow, this vortex is wrapped around the cylinder, forming the so-called horseshoe vortex system. The horseshoe vortex system can clearly be identified in the footprint of the instantaneous pressure field. In Figure 4, the instantaneous pressure distribution around the cylinder at a Reynolds number of Re=78000 is plotted in a wall parallel plane at z=0.06D above the bottom wall. Low pressure is marked with light colors, high pressure is marked in dark. One can distinguish the vortex shaped like a horseshoe upstream (left) of the cylinder, the blotchy shear layers where the flow separates from the cylinder and a region of low pressure downstream, where a von Karman-vortex is shedding from the cylinder.

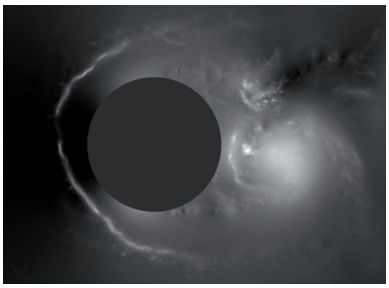

Figure 4: Instantaneous pressure distribution around the cylinder in a wall parallel plane at z=0.06D above the wall (Re=78000). Flow is from left.

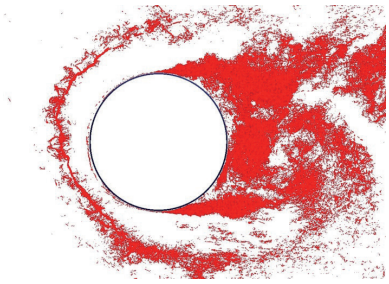

Figure 5: The Q-criterion visualizes coherent structures around the cylinder (Re=78000).

The same flow structures can also be observed in Figure 5, in which the second invariant of the velocity gradient tensor (Q-criterion) is evaluated. Figure 5 clearly illustrates the complex interaction of structures on a continuous spectrum of scales. One can identify structures having the size of the diameter of the cylinder (horseshoe vortex, von Karman-vortex) as well as events having the length scale of the grid resolution. Note that the grid has 440 grid points per diameter in directions parallel to the bottom plate and 1778 grid points per diameter in wall normal direction. This high resolution is necessary to obtain reliable and convergent results for this flow [4].

Outlook

As a next step, simulations of a cylinder placed in a scour hole will be conducted. It is expected that the required number of grid points will double with respect to the simulations with a flat plate. The goal of these new simulations is to investigate the change of the flow field and its dynamics with developing scour whole.

References

[1] www.hy.bgu.tum.de
[2] M. Manhart. A zonal grid algorithm for DNS of turbulent boundary layers. Computers and Fluids, 33(3):435–461, March 2004.
[3] N. Peller et al. High-order stable interpolations for immersed boundary methods. International Journal for Numerical Methods in Fluids, 52, 1175-1193, 2006.
[4] W. Schanderl and M. Manhart. Reliability of wall shear stress estimations of the flow around a wall-mounted cylinder. Computers and Fluids, "Computers and Fluids, 128:16-29, 2016.

Direct Numerical Simulation of an Adverse Pressure Gradient Turbulent Boundary Layer

Research Institution
Monash University, Universidad Polytécnica de Madrid, Istanbul Technical University

Principal Investigator
Vassili Kitsios

Researchers
Callum Atkinson, Juan A. Sillero, Guillem Borrell, Ayse G. Gungor, Javier Jiménez, Julio Soria

Project Partners
Leibniz Supercomputing Centre

SuperMUC Project ID: pr84lo (PRACE project)

Introduction

The efficient design and performance of many engineering systems rely on turbulent boundary layers (TBL) remaining attached to aerodynamic surfaces in regions of adverse pressure gradient (APG); for example wind turbine blades and aircraft wings. The boundary layer is the region where the mean flow velocity increases from zero at the wall to its maximum value at some distance away from the wall. If the near mean wall velocity points in the direction opposing the bulk motion, then the flow is deemed to have separated. Separation occurs in the presence of an APG, and can potentially result in catastrophic consequences (eg: aircraft losing lift) or at best sub-optimal performance (eg: wind turbines producing less energy).

There has been a long history of theoretical, experimental and numerical research into TBL. The vast majority of the research, however, has been centred on the zero pressure gradient (ZPG) case, while many aspects of turbulent structure and appropriate scaling of APG TBL remain largely unresolved. The study of APG TBL in an appropriate canonical form is, therefore, of utmost importance to understand the influence of local pressure gradient. Adverse pressure gradients typically arise due to the presence of convex curved surfaces, such as on a wind turbine blade. This type of configuration is difficult to systematically study, since the pressure gradient applied to the TBL is constantly changing. Here we study a self-similar APG TBL, in which the non-dimensional pressure gradient (β) is constant, and each of the terms in the governing equations have the same proportionality with position in the streamwise (bulk-flow) direction.

Specifically direct numerical simulation (DNS) is used to investigate the structure and dynamics of a self-similar TBL in a strong APG environment such that the flow is statistically maintained immediately prior to separation. The maximum momentum thickness based Reynolds number is $Re_{\delta 2}=10^4$. The results are also compared and contrasted to an associated ZPG TBL.

Results and Methods

The code uses the fractional-step method to solve the Navier-Stokes equations for the velocity and pressure fields. Fourier decomposition is used in the periodic spanwise (cross-flow) direction, with compact finite difference used in the aperiodic wall-normal and streamwise directions. The equations are stepped forward in time using a modified three sub-step Runge-Kutta scheme. The code utilises hybrid MPI / openMP parallelisation to decompose the domain [2]. All I/O is parallelised using the parallel HDF5 library and associated file format.

The DNS boundary conditions (BC) are such that the bottom surface is a flat plate with zero velocity. The spanwise boundaries are periodic. Due to the TBL growing in height as it develops in the streamwise direction, a streamwise normal plane prior to the outlet is copied, rescaled and then applied as the inlet boundary condition. The desired pressure gradient is applied to both the ZPG and APG cases through the application of an appropriate farfield suction velocity [1].

17 million CPU-hours were used in the course of this project. The code has been parallelised up to 32,768 cores, with typical jobs using 4096 cores. For a single instant in time, one set of restart files requires 168Gb of disk space. Statistics have been accumulated over millions

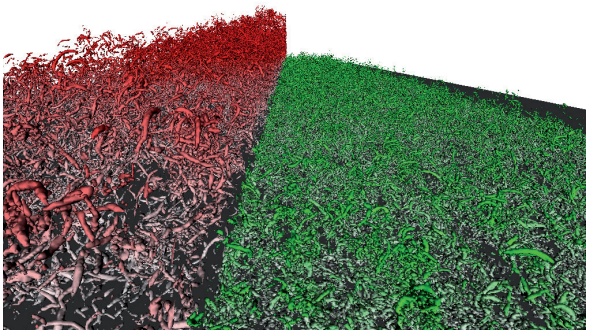

Figure 1: Instantaneous vortex structures of the APG TBL (red), and the ZPG TBL (green) illustrating the significant expansion of the APG TBL in comparison to the ZPG [1].

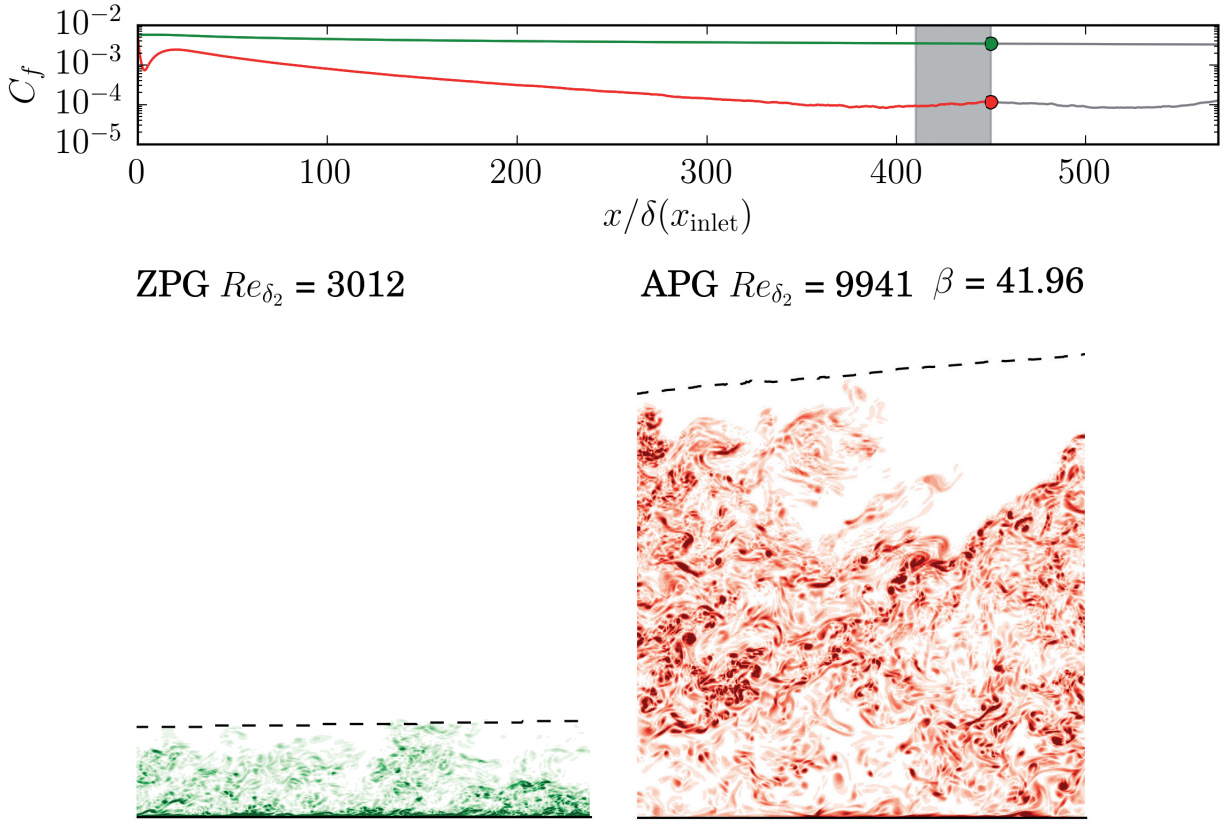

Figure 2: Instantaneous spanwise vorticity of the ZPG TBL (green) and APG TBL (red) within the domain illustrated by the grey box in the skin friction (C_f) plot [1].

of instances in time. Hundreds of restart files have been written to disk for further post-processing of the results.

Figure 1 compares instantaneous vortex structures from the APG TBL (red) and ZPG TBL (green), illustrating the vast complexity and size of such flows. The APG TBL also clearly expands more rapidly in streamwise direction due to the stronger farfield BC suction velocity.

Figure 2 further compares the two TBL on the basis of a slices of instantaneous spanwise vorticity of the ZPG TBL (green) and APG TBL (red) within the domain illustrated by the grey box in the top plot of skin friction (C_f) verses streamwise position (x). At the verge of separation C_f=0. It is again clear that the APG TBL is much taller than the ZPG TBL, from the slices of spanwise vorticity. The majority of the fluctuations in the ZPG case are located in the near wall region. This is due to the fact that the only source of shear in the ZPG TBL is the wall itself, and it is mean shear that gives rise to vortex structures. In the APG TBL, shear is imparted throughout the flow via the farfield BC. As can be seen there is less spanwise vorticity concentrated at the wall, with the vortex structures instead distributed throughout the wall normal domain. The APG TBL is deemed to be self-similar within the streamwise range where $C_f<10^{-4}$.

On-going Research & Outlook

One of the major difficulties with this project, was that we know the desired statistical properties of the flow, but not the boundary conditions that give us these statistical properties. This meant that in addition to initial theoretical work, some trial and error was also necessary. Access to the significant computational resources of the SuperMUC system enabled us undertake the massively parallel simulation of APG TBL, and also to test various boundary condition configurations, until the desired flow state was achieved.

Additional simulations have already been planned and are under development. The new simulations will allow the streamwise extension of the self-similar zone to better interrogate the physical mechanisms of aerodynamic separation. The long term objectives are to use this knowledge to revolutionise the design of current energy generation and transport platforms that operate in APG environments, leading to improved generation of energy and the efficiency of energy use under a wider range of operating conditions. Improvements in the performance of such systems will lead to more efficient and cleaner power generation and to a reduction in fuel consumption and minimization of CO_2 emissions.

References and Links

[1] Kitsios, V., Atkinson, C., Sillero, J.A., Borrell, G., Gungor, A. G., Jiménez, J. & Soria, J., Direct numerical simulation of an equilibrium adverse pressure gradient turbulent boundary layer at the verge of separation, 15th European Turbulence Conference, Delft, Netherlands, 25-28 August, 2015, 1pp.

[2] Sillero, J., Jiménez, J., and Moser, R., 2013, One-point statistics for turbulent wall-bounded flows at Reynolds numbers up to δ+ ≈ 2000. Phys. Fluids, 25:105102.

Large eddy simulation of pulverized coal and biomass combustion

Research Institution
University of Duisburg-Essen, Instituto Superior Técnico (University of Lisbon)

Principal Investigator
Andreas Kempf

Researchers
Miriam Rabacal, Martin Rieth

Project Partners
Imperial College London (UK), University of Sheffield (UK), TUB Freiberg (Germany), University of Stuttgart (Germany)

SuperMUC Project ID: pr84mu (Gauss Large Scale project), pr85qi (PRACE project)

Introduction

Pulverized coal and biomass combustion (PCBC) is currently among the major sources of energy supply and is expected to play an important role in future energy supply. However, coal combustion releases large amounts of carbon dioxide. Future power plants are required to be efficient and low-polluting, which could be achieved by carbon capture and storage or co-firing coal with biomass.

While experimental studies provide valuable and fundamental understanding of the processes of pulverized coal and biomass combustion, they cannot provide all information due to limited optical access and other issues related to the harsh combustion environment. Simulations, such as large eddy simulations (LES), can complement experimental findings by providing large data sets that can be analyzed in great detail. However, numerical methods and modeling approaches need to be developed further to facilitate a comprehensive investigation of the physics of PCBC. Our work is on developing such models and methods for PCBC. In particular, methods to treat particle conversion and the gas phase combustion are developed.

Results and Methods

The code used for all simulations is the in-house finite volume (FV) Fortran code PsiPhi. The code solves the implicitly filtered Navier-Stokes equations in the low Mach number limit. Continuity is enforced by a pressure-correction scheme and projection method using a Gauß-Seidel solver with successive over-relaxation. The code is parallelized by MPI with non-blocking communication and scaling is demonstrated for up to 128,000 cores. Cartesian, equidistant grids are used, which allow for an efficient usage of a large numbers of cells. A third order Runge-Kutta scheme is used for time-advancement.

Coal and biomass particles are treated as Lagrangian parcels. Their parallelization relies on the same domain decomposition as the gas phase treated by the FV method. Coupling between particles and gas phase is facilitated by tri-linear interpolation schemes. The discrete ordinates method is used to solve for the radiative heat transfer. The code is compiled with Intel Fortran and both IBM and Intel MPI is used. The code has been sped up by improving parallelization and algorithms by us and our collaborators during the projects on SuperMUC.

The overall CPU-hours used in project pr84mu were around 21 million and 10 million in project pr85qi. Typical, coarse grids consisted of ~500,000 cells and fine grids consisted of up to ~1,700,000,000 cells. Numerical particles were typically around 40,000,000. Coarse runs were conducted with ~1000 cores, whereas fine runs were conducted with ~15.000 cores. To reduce the issue of very

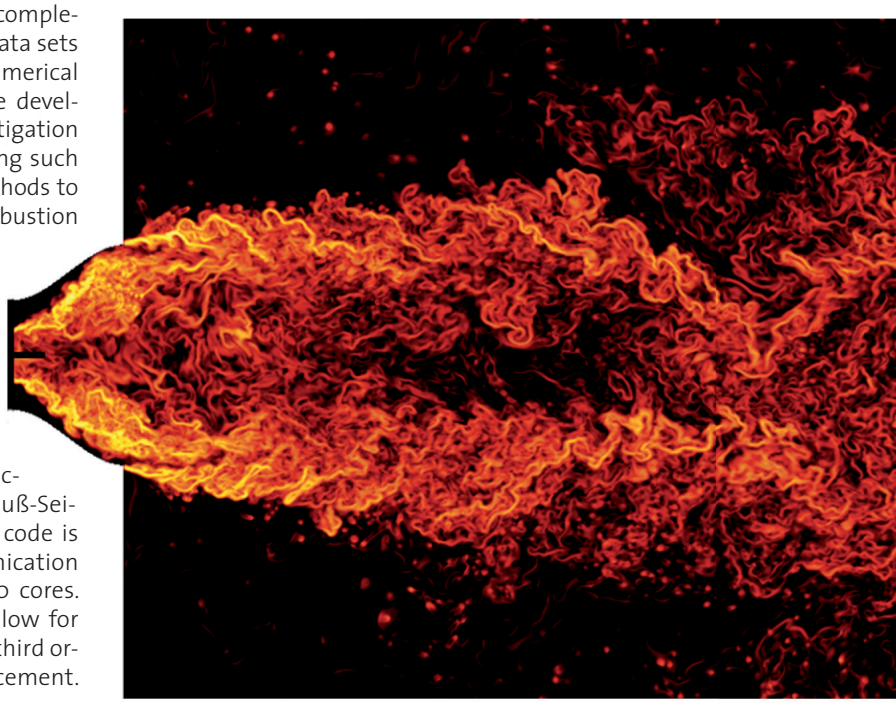

long initialization times, results from coarse grids were used to initialize runs on finer grids. The largest run conducted, 16384 cores, generated one restart file per core with the total size of four terabytes. Data relevant for post processing was combined to a hdf5-file of around one terabyte. The overall WORK storage required was 31 terabytes.

Large-scale coal and biomass flames in furnaces that have been studied in detail experimentally – the IST and the BYU furnace – were used as reference cases. The classical coal and biomass conversion models, originally developed in cooperation with Imperial College were tested, as well as improved models and strategies. The first step of particle conversion, pyrolysis, is too complex to be explicitly modeled in LES. In cooperation with TUB Freiberg, a pre-processing strategy was developed to optimize the parameters of a simple empirical model based on the predictions of advanced pyrolysis models [1]. The massively parallel simulations provided a good description of scalar and velocity fields, confirmed by the good agreement with experiments. Flame stabilization, flame structure and particle burnout are strongly affected by the fuel properties and the fluid dynamics, and LES is able to provide insights to the phenomena occurring in this type of application that are currently not available through experimental means. Single particles were tracked over time and instantaneous ensembles were collected to obtain a better understanding of the conditions that coal particles are subjected to [2]. The effect of conversion modeling, particularly the empirical devolatilization model and the mode of char combustion model, on the flame lift-off and flame length of a co-fired flame was also investigated [3].

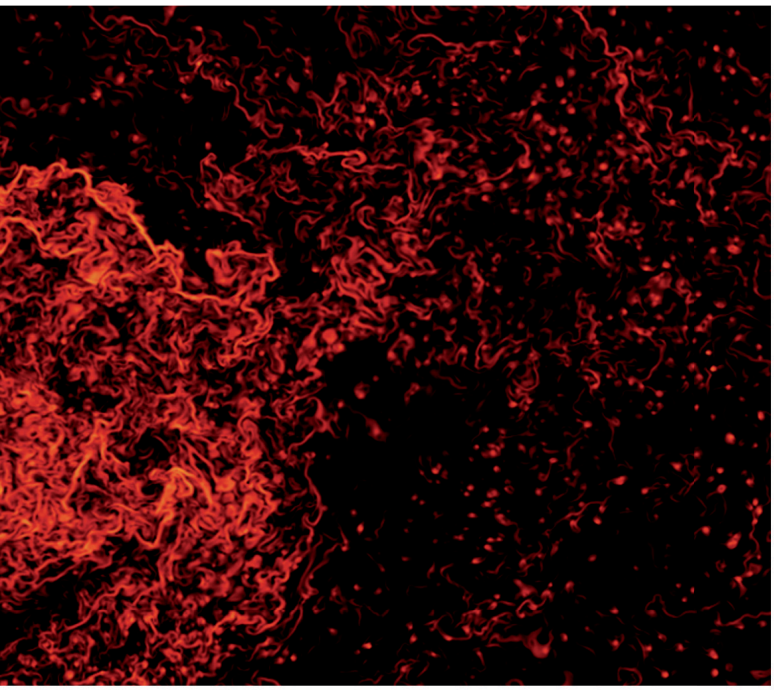

Figure 1: Logarithmic scalar dissipation rate of the sum of volatile and char-off gas mixture fractions in a section of around 1.5 x 3 m of the flamelet LES of the IFRF semi-industrial furnace.

A part of the project was to incorporate the flamelet model, which is particularly popular in gaseous turbulent combustion computations, to the LES of coal combustion. The test case used for the flamelet LES is a semi-industrial scale furnace with a thermal power of approximately 2.5 MW. The first flamelet table used to describe the chemical state of the reacting gas phase was based on two mixture fractions for volatile and char-off gases as well as on enthalpy and variance. The flamelet table was generated before the simulation, and chemical state variables looked up from the table based on the four parameters obtained during the simulation. First simulations provided good results and could demonstrate the suitability of the flamelet model for such simulations [4]. A further improvement was achieved by including scalar dissipation rate as a look-up variable [5]. Scalar dissipation rate is an important parameter in the flamelet model and can be understood as an inverse mixing time-scale. The large simulations show very good results compared to the experiment and reveal a wealth of information that is yet to be analyzed. A key feature of such furnace simulations is depicted in Figure 1, where the logarithmic scalar dissipation rate of the sum of the two mixture fractions is presented, showing regions of intense mixing in the volatile flame close to the inlet but also in the shear layers between flue gases and fresh combustion air at the edges of the quarl outlet.

On-going Research / Outlook

"SuperMUC Next Generation" will enable to simulate test cases even closer to industrial applications. But also to run laboratory scale configurations at higher resolution to resolve physics down to the scales of turbulence and particle-turbulence interaction, which provides valuable understanding of the physics of turbulent coal conversion.

References and Links

[1] M. Rabaçal, B. Franchetti, F. Cavallo Marincola, F. Proch, M. Costa, C. Hasse, A.M. Kempf. 2015. Large eddy simulation of coal combustion in a large- scale laboratory furnace. Proceedings of the Combustion Institute, 35, 3609-3617.

[2] M. Rabaçal, M. Costa, A.M. Kempf. 2016. Particle History from Massively Parallel Large Eddy Simulations of Coal Combustion in a Large-Scale Laboratory Furnace, Fuel, in revision.

[3] M. Rabaçal, M. Costa, M. Vascellari, C. Hasse, M. Rieth, A.M. Kempf. 2016. Large Eddy Simulation of Co-Firing Biomass and Coal in a Large-Scale Furnace, submitted to the Combustion Theory and Modelling.

[4] M. Rieth, F. Proch, M. Rabacal, B. Franchetti, F. Cavallo Marincola, A.M. Kempf. 2016. Flamelet LES of a semi-industrial scale furnace. Combust. Flame, under revision.

[5] M. Rieth, F. Proch, A.G. Clements, M. Rabacal, A.M. Kempf. 2016. Highly resolved flamelet LES of a semi-industrial scale furnace. Submitted to the Proceedings of the Combustion Institute.

Natural thermal convection at high Rayleigh numbers

Research Institution
Max Planck Institute for Dynamics and Self-Organization
Principal Investigator
Olga Shishkina
Researchers
Susanne Horn, Olga Shishkina, Sebastian Wagner
Project Partners
–

SuperMUC Project ID: pr84pu, pr94na, pr63ro

Introduction

Turbulent thermal convection is ubiquitous in nature; its investigation is needed for the better understanding of geophysical and astrophysical flows and for technological improvements in industrial applications. By means of Direct Numerical Simulations (DNS), we study various canonical setups of this problem, ranging from the classical example of Rayleigh-Bénard convection (RBC), over inclined and vertical convection, to horizontal convection [1].

The common parameters that describe these systems are the Rayleigh number (*Ra*), characterising how vigorous the flow is, the Prandtl number (*Pr*), specifying the fluid, and the geometrical aspect ratio of the convection cell. Our theoretical investigations are aimed on a better understanding of turbulent thermal convection in the above flow configurations and on the development of models to predict main mean flow characteristics such as mean momentum and heat transport, measured, respectively, by the Reynolds and the Nusselt number, and their scalings with the control parameters.

Results and Methods

For our simulations we use the finite volume code goldfish developed by the researchers of this project, Olga Shishkina, Sebastian Wagner and Susanne Horn. The code is written in fortran and parallelised with MPI; it solves the incompressible Navier-Stokes equations and the temperature equation. The input-output is also fully parallel using the netCDF-4/HDF5 format.

Since we perform almost exclusively DNS, i.e. our meshes have to be fine enough to resolve the smallest relevant turbulent scales, the number of cores used strongly depends on the studied configuration and parameters, especially on the Rayleigh number and to some extent also on the Prandtl number. Typically we run several simulations simultaneously with varying input parameters using between 128 and 1024 cores. To obtain reliable statistical convergence run times in the order of weeks up to months are required. For one setup we usually store 500 to 1000 instantaneous flow fields on the SCRATCH space, resulting in about 15 to 20 TB per project and user. These are necessary for many analysis methods, which would also otherwise slow down the simulations tremendously if conducted during run time. Furthermore, not all relevant quantities are known a priori, but can be recovered from the instantaneous fields. This enables us to re-use our data for our on-going research.

Inclined convection
In inclined convection a fluid layer is tilted with respect to the gravity direction, and thus, not only buoyancy, but also shear drives the flow in this case. The inclination angle $\beta=0$ corresponds to Rayleigh-Bénard convection, i.e. a fluid heated from below and cooled from above, and $\beta=\pi/2$ corresponds to vertical convection, where a fluid layer is confined between vertically aligned heating and cooling

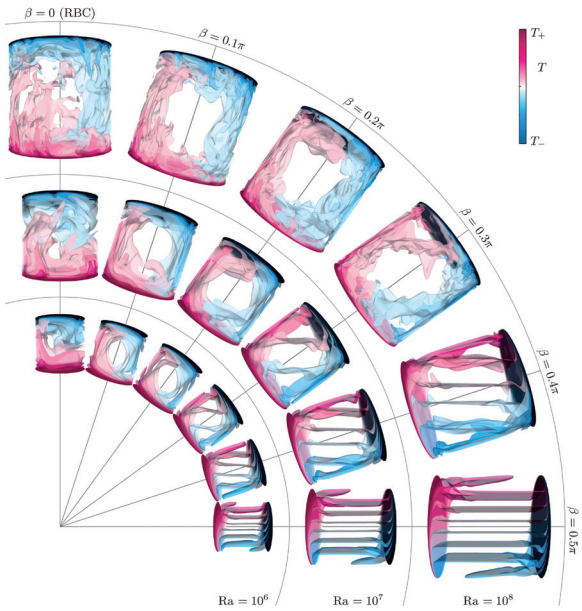

Figure 1: Isosurfaces of the instantaneous temperature fields in inclined convection in cylindrical containers, filled with a fluid of Prandtl number one, for varying Rayleigh numbers and different inclination angles ranging from β=0 (Rayleigh-Bénard convection) to β=π/2 (vertical convection). Adopted from [2].

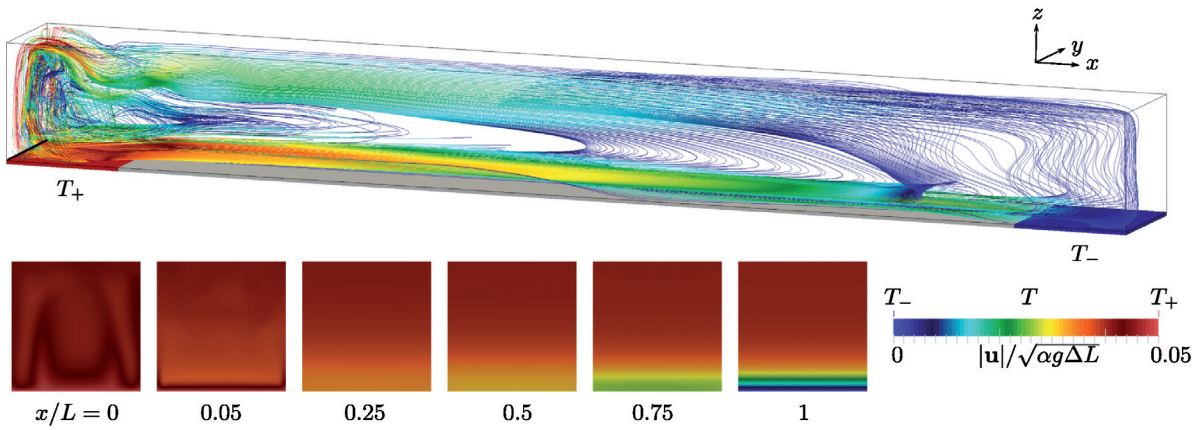

Figure 2: Scheme of the studied horizontal convection setup together with streamlines; the left one tenth of the bottom is heated (red), while the right one tenth of the bottom is cooled (blue). The rest of the bottom, as well as the top and sidewalls, are adiabatic. Cross sections of the temperature are given below. Adopted from [3].

plates, see Fig. 1. In Ref. [2] we have shown, that the heat flux dependence is not universal and is a complicated non-monotonic function of Ra and Pr. Thus, with a small inclination of the RBC cell, the Nusselt number can decrease or increase, compared to that in the RBC case, for large and small Pr, respectively. Close to $\beta=\pi/2$, the Nusselt number decreases with growing β in all considered cases. Moreover, a slight cell tilt may not only stabilise a large scale circulation but can also enforce one for cases where the preferred RBC state is a more complicated multiple roll state.

Boundary layer structures

Viscous and thermal boundary layers (BLs) play a critical role in the heat transfer. The classical Prandtl-Blasius-Pohlhausen BL theory cannot describe well the BLs in turbulent RBC as its inherent features like pressure gradients within the BLs, fluctuations and buoyancy, are assumed to be negligible in that approach. Therefore, based on our DNS, we advanced the BL theory for turbulent RBC by its extension to the case of a non-vanishing pressure gradient within the BLs, when a large-scale circulation approaches the heating/cooling plates not parallel to them. Further, we studied the effect of fluctuations within the BLs by considering the turbulent thermal diffusivity. The new developed BL equation [4] is analytically solvable and its solutions are found to be in excellent agreement with the DNS.

Scalings in turbulent horizontal convection

In horizontal convection, heating and cooling are applied to different parts of the same horizontal surface of a fluid layer, see Fig. 2. This is a paradigm setup relevant for many geophysical systems, in particular, in the large-scale ocean circulation, as heat is supplied to and removed from the ocean predominantly through its upper surface, where the ocean contacts the atmosphere. It is also important in process engineering. In Ref. [3] we studied the scalings in horizontal convection. The derived model is in perfect agreement with the numerical data.

Wall roughness in thermal convection

Surface roughness is known to have a large impact on the global heat transfer. Therefore we study numerically the effect of the roughness introduced by a set of distinct obstacles attached to the heating and cooling plates, for different *Ra*, *Pr* and configurations of the cell geometry [1]. Based on our DNS, we suggested a simple parameter-free model, which captures well the mean heat flux enhancement due to regular wall roughness, the height of which is larger or comparable with the thickness of the thermal BL in the case of smooth plates.

Rotating Rayleigh-Bénard convection

In rotating Rayleigh-Bénard convection, the flow does not only depend on *Ra* and *Pr*, but also very strongly on the applied rotation rate, often expressed by the inverse Rossby number, *1/Ro*. We conducted DNS for various *Ra*, *Ro* and *Pr*, and also taking into account the explicit temperature dependences of the fluid properties. One of the main objectives of the study is to investigate rotating RBC in the regime of geostrophic turbulence which is considered the most relevant for the description of geo- and astrophysical phenomena. It translates to very high *Ra* and *1/Ro*, i.e. turbulence in the presence of strong rotational constraints. Further, a novel method to identify the flow regimes was proposed [1], which is based on an analysis of the balance of the poloidal energy and toroidal energy.

On-going Research / Outlook

Numerical simulations of turbulent thermal convection are mainly limited by the required very fine resolution, in particular of the boundary layers, but also of the small structures, such as plumes in the bulk. These restrictions are imposed by the available computing power, the memory of the nodes and also the available storage. Hence, only very recently it has become possible to study, for example, geostrophic turbulence in rotating convection with DNS. However, we are still only in the border area, and every further advance into this utmost important regime will enrich our knowledge on geo- and astrophysical fluid dynamics. Similar applies to inclined and horizontal convection and turbulent thermal convection in general.

References and Links

[1] http://www.lfpn.ds.mpg.de/shishkina/projects.html
[2] Olga Shishkina and Susanne Horn. 2016. Thermal convection in inclined cylindrical containers. J. Fluid Mech. 790, R3.
[3] Olga Shishkina and Sebastian Wagner. 2016. Prandtl number dependence of heat transport in laminar horizontal convection. Phys. Rev. Lett., 116 (2016), 024302.
[4] Olga Shishkina, Sebastian Wagner, Susanne Horn and Emily S.C. Ching. 2015. Thermal boundary layer equation for turbulent Rayleigh-Bénard convection. Phys. Rev. Lett., 114, 114302.

Modulation of Turbulent Properties in a Spray Flame Burning n-Heptane: Direct Numerical Simulation

Research Institution
Lab of Fluid Dynamics and Technical Flows (LSS/ISUT), University of Magdeburg "Otto von Guericke"

Principal Investigator
Dominique Thévenin

Researchers
Abouelmagd Abdelsamie, Cheng Chi, Timo Oster

Project Partners
–

SuperMUC Project ID: pr84qo (Gauss Large Scale project)

Introduction

The current report shows the status of our HLRB project pr84qo [1], which is running on SuperMUC (GCS large scale project) since November 2014 and until December 2016. In total we were awarded 18 Million CPU-hours. The aim of this project is to investigate turbulent spray flames, quantifying possible modifications of the turbulent properties. Ignition is analyzed in homogenous isotropic turbulence and shear flows (temporal jet) by Direct Numerical Simulation (DNS), considering n-heptane liquid droplets. The droplets, being smaller than the grid resolution and Kolmogorov length scale, are modeled as point droplets, while the Navier-Stokes equations are solved in the low-Mach number regime. Detailed models are employed to describe chemical reactions and molecular transport in the gas phase. In the current DNS, the continuous (gas) phase is simulated in a standard manner (Eulerian frame) whereas the discontinuous (droplet) phase is tracked in a Lagrangian frame. Two-way coupling interaction between both phases is quantified via the exchange of mass, momentum and energy. The impact of different parameters is investigated, in particular: initial temperatures, initial pressure, equivalence ratio/droplet mass fraction, droplet size, turbulence level (with a Taylor Reynolds number up to 150) and mean shear effect. The in-house 3D DNS solver DINO [2] is used for all simulations.

Algorithms and Numerical Methods

The current project relies on an in-house DNS code. DINO is a new Fortran-2003 code, which has been developed in our group since the beginning of 2013 [2, 3]. DINO is a three-dimensional low Mach number DNS solver code with a 6th order finite-difference spatial discretization for reacting and multi-phase turbulent flows. The code is parallelized in two dimensions using the 2DECOMP&FFT library that acts on top of standard MPI and FFTW. The Poisson equation for pressure is solved by means of FFT, both for periodic and non-periodic boundary conditions (in the latter case with pre- and post-processing steps). Although in any low Mach number solver the time step restriction for acoustic wave is removed, the restriction of time step due to chemistry stiffness is still there. For that reason, a 3rd order semi-implicit Runge-Kutta scheme, or alternatively a combination (Radau5 for stiff terms, explicit 4th order Runge-Kutta for non-stiff terms) are used for time integration. Both techniques rely on standard linear algebra libraries, LAPACK and BLAS, for the implicit stiff part of the ODE. By default the chemical source terms are computed using the Cantera-1.8 library. The transport properties are computed either with the Cantera library or with the EGlib-3.4 library. The discontinuous phase in multi-phase flow simulations (droplets/spray) is tracked by using either a classical Lagrangian point force approach (for non-resolved particles) or with the Immersed Boundary Method (IBM) technique for fully resolved droplets (an approach that is not discussed further here). The initial turbulent field is generated by inverse Fourier transform with analytical energy spectrum (Passot-Pouquet or Von Karman-Pao). Input/output operations rely on MPI-I/O routines provided by the 2DECOMP&FFT library. These files are used for restarting the simulations while using parallel HDF5 saving for storing data used for postprocessing. The code is already under GIT version control, which helps all users to quickly and safely carry out changes or updates, if needed. As build environment DINO uses cmake and it can be compiled with both GNU and Intel Fortran compilers. Required resources for this project are summarized in Table 1.

Table 1: summary of required resources.

Total CPU-h	Overall storage	Typical #cores	#generated files
18 Mio	40 TB	2048	6000

Scientific Results

As mentioned in the previous section we are investigating the burning of n-heptane liquid droplets in two different configurations: Homogeneous isotropic turbulent (HIT) and temporally-evolving turbulent jet flows. Figure 1 depicts the time evolution of the burning process in

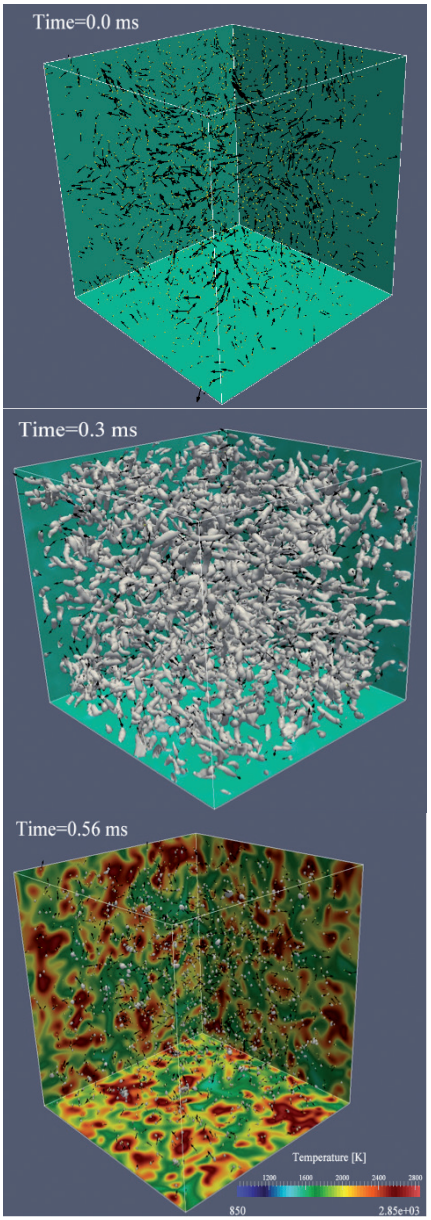

Figure 1: Evolution of n-heptane mass fraction iso-surface in the gas phase (gray) with slices of gas temperature (color) as well as droplet position and velocity vectors (black)

HIT: liquid distribution (time=0.0 ms), evaporation process (time=0.3 ms) and first ignition (time=0.56 ms). In this configuration, liquid droplets with a uniform temperature of 300 K are randomly distributed in a hot homogeneous isotropic turbulent environment (oxygen/nitrogen with temperature of 1500 K and atmospheric pressure). The corresponding results and more statistics had been presented during last Combustion Symposium [4]. Figure 2 shows the evaporation and ignition process in the temporal jet configuration. Here, the droplets are randomly distributed inside the central part of the domain (jet flow region) with initial temperature of 300 K and with the same velocity as the surrounding jet flow. The jet flow and co-flow mixture consist of oxidizer with a uniform temperature of 1700 K and pressure of 5 bar. In Fig. 2 the yellow spheres represent the size and location

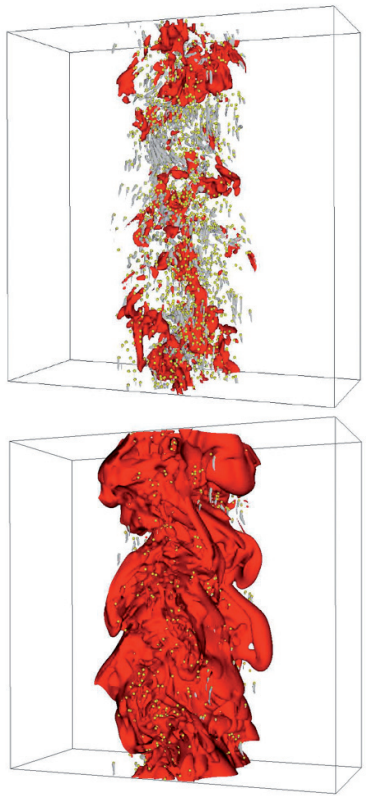

Figure 2: Time evolution of evaporation and ignition process of gaseous mixture. Red iso-surfaces correspond to a temperature of 1800 K, gray iso-surfaces represent n-heptane in the gas phase (after evaporation), yellow spheres represent the liquid droplets. Top: at time t=67 µs. Bottom: at time t= 75 µs.

of the droplets (the size is multiplied by a factor of 4 for easier visualization), while the gray and red iso-surfaces represent the evaporated heptane vapor in the surrounding mixture, and the temperature of 1800 K, respectively. Part of the result corresponding to this configuration have been already reported in Refs. [2, 5]. The output statistics and published data of the current project can ultimately become a reference data-set for many practical and academic works dealing with turbulent spray combustion. The complete 3D output data generated in the project will thus be gathered in a database accessible for other researchers who would like to validate their own spray evaporation or ignition model by analysis and comparison with DNS data.

References and Links

[1] https://www.lrz.de/projekte/hlrbprojects/0000000000F4 3804.html

[2] Abdelsamie A., Fru G., Oster T., Dietzsch F., Janiga G., Thévenin D. Towards direct numerical simulations of low-Mach number turbulent reacting and two-phase flows using immersed boundaries. Comput. Fluids 131: 123-141, 2016.

[3] Thévenin D., DNS and LES of transitional and two-phase flows, in Direct and Large-Eddy Simulation X, Limassol, Cyprus, Keynote Lecture, 20, 2015.

[4] Abdelsamie A., Thévenin D. Modulation of turbulent properties in a spray flame burning n-Heptane using Direct Numerical Simulation. In 35th Symposium (International) on Combustion, San Francisco, Poster presentation, 2014.

[5] Abdelsamie A., Thévenin D. DNS of burning n-heptane droplets: auto-ignition and turbulence modulation mechanisms. In Proceeding of Direct and Large-Eddy Simulation X, Limassol, Cyprus, 152, 2015.

Coupled multi-physics simulations on the massively parallel system SuperMUC

Research Institution
Simulation Techniques and Scientific Computing, University of Siegen

Principal Investigator
Sabine Roller

Researchers
Verena Krupp, Harald Klimach, Jens Zudrop

Project Partners
Benjamin Uekermann, Institute for Advanced Study, Technical University Munich

SuperMUC Project ID: pr84xu

Introduction

Due to the increasing computing power available on modern supercomputers it becomes feasible to include more and more physical effects and scales in a single simulation. Fluid-structure-acoustic interaction is a typical application for such multi-physics, as well as multi-scale simulations. Within the DFG funded SPPEXA project ExaFSA [1] we are aiming for such a simulation of sound generation in turbulent flows, together with its propagation across large distances as well as the interaction with structural mechanics. In this context the transition from turbulent flow to acoustic wave propagation is an essential task that includes many computational challenges. Solving this problem with a monolithic approach is still too expensive and might not be feasible from a numerical point of view. However, these different effects typically appear in spatially separable areas (Figure 1) and a coupled approach can be used to address each domain with the best-suited numerical approximation. The interaction between the domains is realized via the exchange of information at the joint interfaces.

To enable individual discretization, different governing equations as well as adapted numerical methods in each domain (Structure, Fluid, Acoustic Far Field), we exploit a partitioned coupling approach: for each domain we apply the appropriate solver, which allows us to reuse existing software packages with their individually high scalability and efficiency characteristics. We then 'glue' them together by using a dedicated coupling library.

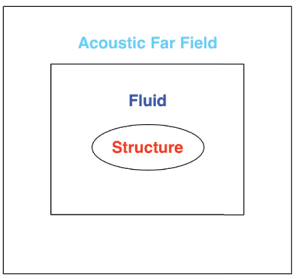

Figure 1: Sketch of a typical fluid-structure-acoustic set up.

Typically, the acoustic far field has the largest extent (e.g. for the noise generation of wind turbine it is interesting how much noise is reaching the houses in the neighborhood) and needs to be solved efficiently. Additionally, the fluid-acoustic interface is computationally intense in terms of communication as well as data mapping.

Results and Methods

Acoustic simulations require the transport of waves over long distances without damping or phase shifts in the waves. On the other hand, the solution in this part of the domain is smooth. Thus, high order methods are well suited from a numerical quality point of view as well as under a computational efficiency focus. We thus apply the high order Discontinuous Galerkin solver Ateles for the acoustic far field. Ateles is part of the end-to-end parallel framework APES [2], which provides the required tools for pre- and post-processing on the basis of the common mesh library TreElM [3].

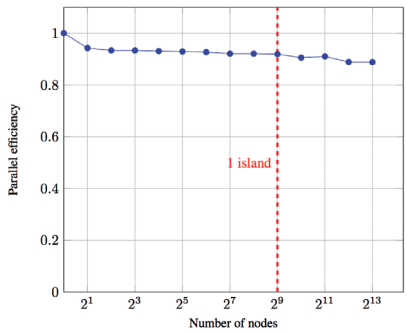

Figure 2: Weak scaling of the 64th order spatial scheme with 16 elements per process on up to 8192 nodes of SuperMUC with hybrid parallelization.

Ateles is designed for large-scale simulations on massively parallel systems and shows good scaling behavior on systems such as the SuperMUC (Figure 2).

In addition to the Linearized Euler Equations, Ateles also features a non-linear Euler and Navier-Stokes kernel. These kernels are applied for the simulation of the turbu-

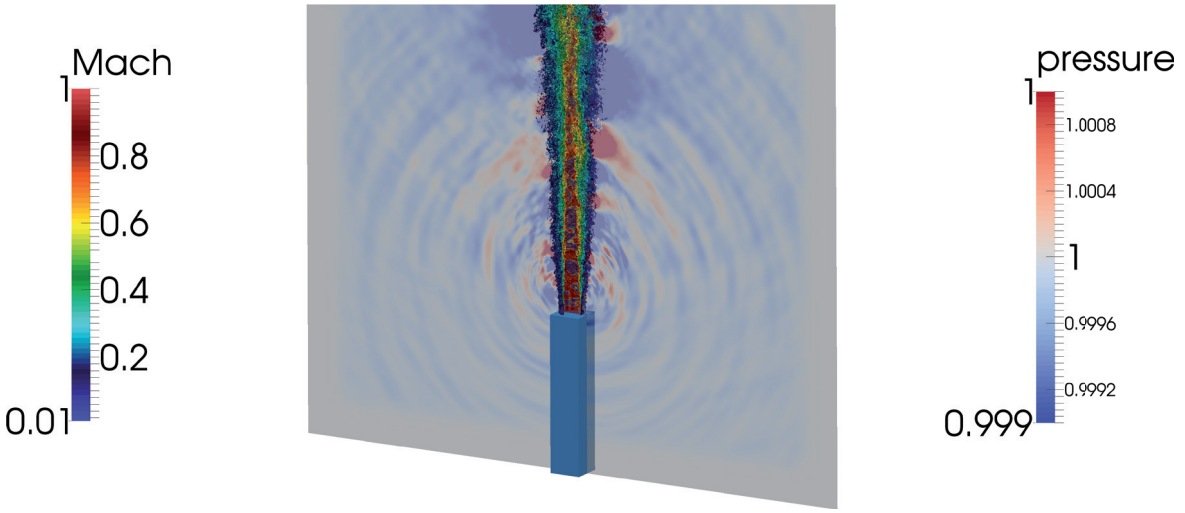

Figure 3: Iso-surfaces of vorticity magnitude colored by Mach number and acoustic pressure on a slice around the core of the free-stream jet. Multi-level mesh with 139,424 elements, O(16), yielding more than 2.85billion degrees of freedoms. The simulation was run on 2048 nodes, i.e. 32,768 cores, of Super-MUC [5].

lent flows that generate the sound waves and transport them through the acoustic near field before the waves are handed over to the acoustic far field solver. With the high order solver Ateles we are aiming for direct numerical simulations (DNS). Figure 3 shows a flow field of a supersonic jet with more than 2.85 billion degrees of freedom that resolve the smallest scales in the flow.

For the partitioned coupling approach, the coupling library preCICE [4] is used which takes care of the data mapping, the data exchange and the steering of the simulation time steps. An explicit parallel coupling method with point-to-point communication is used in preCICE and, therefore, both domains can be efficiently solved in parallel. At each time step, data is exchanged at the interface. For the partitioned coupling approach, a static load balancing based on heuristics is exploited. Therefore, idling processes due to different work of the non-linear flow domain in contrast to the linear acoustic domain can be minimized according to the usecase. For e.g. the coupled simulation of a subsonic jet, shown in Figure 4, each domain spends respectively 20% and 14% of the whole execution time in the coupling library, which represents the time of communication and data mapping with nearly no waiting time observed.

On-going Research / Outlook

The just finalized phase 1 of the ExaFSA project focused on setting up the framework and gaining experience with the general quality of the coupling algorithms, especially on the data mapping between the processes. Future work focuses on the simulation of real world fluid-structure-acoustic interaction to bring new insights into such applications and realize computational optimization of e.g. the sound design of an aircraft or wind energy plants.

The DFG funded project is extended in the 2nd phase where we will focus on challenges regarding stability at the coupling interface when coupling different meshes and time integration schemes, tailoring the discretization to the physics as well as to the machine.

With the help of compute resources on SuperMUC, we were able to explore and overcome the bottlenecks in the optimization of data mapping within the coupling tool. In the on-going research efforts, we are also optimizing the methods of static load balancing between the domains that are parts in the coupled simulation. Efforts also involve improvements in the sustained performance of the Ateles framework.

References and Links

[1] http://ipvs.informatik.uni-stuttgart.de/SGS/EXAFSA/
[2] H. Klimach, K. Jain, and S. Roller. End-to-end parallel simulations with Apes. In Michael Bader, Arndt Bode, Hans-Joachim Bungartz, Michael Gerndt, Gerhard R. Joubert, and Frans Peters (Ed.). 2014. Parallel Computing: Accelerating Computational Science and Engineering (CSE), Vol.25. Munich, Germany. IOS Press.
[3] https://bitbucket.org/apesteam/treelm
[4] H.-J. Bungartz, F. Lindner, B. Gatzhammer, M. Mehl, K. Scheufele, A. Shukaev and B. Uekermann. preCICE -- A Fully Parallel Library for Multi-Physics Surface Coupling. In Computers and Fluids. Elsevier, 2015. accepted.
[5] Jens Zudrop. 2015. Efficient Numerical Methods for Fluid- and Electro-dynamics on Massively Parallel Systems, Ph.D. Dissertation, RWTH Aachen, Germany

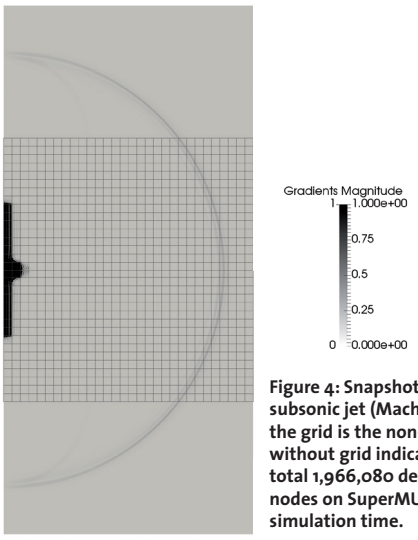

Figure 4: Snapshot of a coupled 2d simulation of a subsonic jet (Mach 0.1), where the domain showing the grid is the non-linear flow domain and the one without grid indicates the linear acoustic domain. In total 1,966,080 degree of freedoms, running on 66 nodes on SuperMUC, require 50,000 CPU-h for 250 s simulation time.

Dynamics of Shedding Partial Cavities and Erosion Aggressiveness of Cavitating Ship Propeller Flow

Research Institution
Institute of Aerodynamics and Fluid Mechanics

Principal Investigator
Bernd Budich

Researchers
Bernd Budich, Steffen J. Schmidt, Nikolaus A. Adams

Project Partners
Office of Naval Research

SuperMUC Project ID: pr85ki

Introduction

Hydraulic machinery, such as water turbines, turbo pumps and marine propulsion systems, are frequently affected by *cavitation*, which occurs when the local static pressure in the liquid drops below the vapor pressure. This causes evaporation and the generation of vapor pockets (*cavities*). Being advected again into regions of increased pressure, a sudden implosion-like re-condensation takes place. During these collapses, the pressure level can reach magnitudes of several hundreds to thousands of atmospheres and strong shock waves are created. The aggressiveness can be severe enough to damage even metal surfaces. When exposed to cavitation over a sustained time, this may eventually lead to the failure of the affected components.

In this project, we focus on the numerical investigation of cavitating flow in the context of ship propellers. While material erosion reduces the lifespan of propeller blades and the rudder, cavitation also causes a reduction in thrust and efficiency. Propeller performance is degraded, and fuel consumption increased. Simultaneously, cavitation induces noise and structural vibrations. To avoid these highly undesirable consequences, an in-depth understanding of the relevant mechanisms is necessary. A key aspect of this project is to develop the ability to assess local flow aggressiveness and to quantify the potential of material erosion.

Computational Fluid Dynamics (CFD) is advantageous, as it enables the analysis of flow structures and dynamics at both a temporal and spatial accuracy that is difficult to obtain with experimental studies. The formation and propagation of complex shock wave systems is crucial for understanding cavitation and its impact on materials. However, the time-scales between the convective flow and wave dynamics typically span several orders of magnitudes. To resolve these effects, time steps smaller than nanoseconds are required and thus millions of iterations are necessary to cover characteristic time intervals. This leads to a substantial numerical effort, requiring massively-parallel, high-performance computing resources.

Numerical Method

In order to capture collapse-induced pressure peaks and associated shock wave dynamics, the numerical approach developed at the TUM Institute of Aerodynamics and Fluid Mechanics [1], takes into account the two-phase compressibility of the water-vapor-mixture. The density-based, 3D finite volume method is based on a homogeneous mixture model [2]. Spatial reconstruction utilizes a 2nd-order TVD scheme on body-fitted, structured grids. An explicit, 4-stage Runge-Kutta method is used for time integration. In this study, we focus on inertia-dominated flow physics and thus neglect viscosity. Furthermore, the flow is assumed barotropic and the effects of gas content are neglected.

Results

Shedding Partial Cavity

For a validation of the numerical approach, the canonical configuration of shedding partial cavities developing over a prismatic test body was studied [3]. The simulations exhibit complex flow patterns that are also observed experimentally. Figure 1 shows an instantaneous view of the flow field. Cavities, visualized by iso-surfaces of the 10% volume fraction α in Fig. 1(a) are compared to iso-surfaces of vorticity magnitude $|\omega|$, i.e. vortical structures, in Fig. 1(b). The simulations show cavitating horse-shoe vortices and cavitation of streamwise-oriented hairpin vortices, which match the experimental observations. Propagating

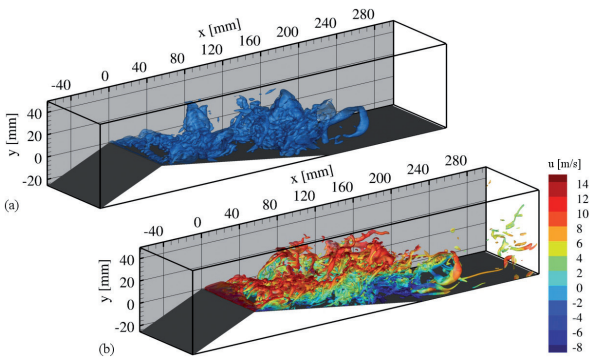

Figure 1: Comparison of instantaneous vapor structures (iso-surfaces α=0.1, top) and vortical structures (iso-surfaces $|\omega|$=4000 s^{-1}, colored by axial velocity u, bottom).

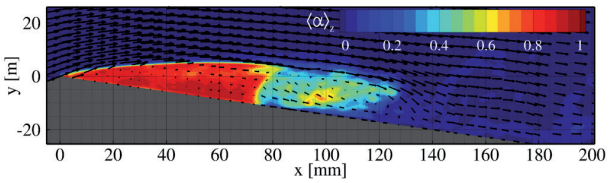

Figure 2: Spanwise-averaged instantaneous void fraction, exhibiting a condensation shock at x≈70 mm . Velocity vectors show the spanwise-averaged local flow velocity.

shock wave structures, also known as crescent-shaped regions that occur for the localized collapses of vapor clouds, are also correctly reproduced in the simulation.

The partial cavity experiences periodic shedding with a characteristic Strouhal-number of *St*=0.28, which is in close agreement with the experimental findings. At the chosen operating point, the shedding is dominated by an abrupt condensation of the water-vapor mixture, originating at the cavity aft portion when the attached sheet reaches its maximum spatial extent. Observed by experiments only recently, it is correctly reproduced using the numerical model, as depicted in Fig. 2. It was demonstrated that this phenomenon represents a compressible shock wave.

Cavitating Propeller

Subsequently, the developed method is utilized for the computation of cavitating flow around the model propeller *VP1304* [4,5]. Reproducing the *Potsdam Propeller Test Case 2011*, three operating points exhibiting cavitating flow were simulated. The obtained results show good agreement for the computed propeller efficiency, blade pressure distributions, and cavitation patterns. Location and extent of cavitating regions is in agreement with the experiments and previous numerical studies, supporting the choice of our numerical approach. A snapshot of the computed flow field is shown in Fig. 3, visualizing vapor structures in Fig. 3(a), and vortical structures in Fig. 3(b). The flow aggressiveness for this configuration is analyzed by two criteria: first, pressure maxima are monitored for each grid cell. Second, using a collapse detection algorithm, collapse events of isolated cavities within the flow field are recorded. The visualization in Fig. 4 compares the maximum detected pressure on the blade surface

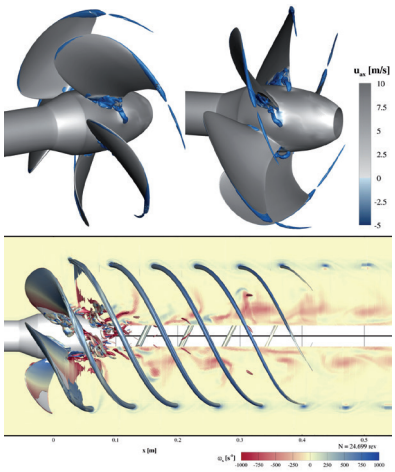

Figure 3: Flow field around VP1304: vapor structures (iso-surfaces α=0.1, top) and vortical structures (iso-surfaces |ω|=1000 s^{-1}, colored by axial vorticity ω_x, bottom).

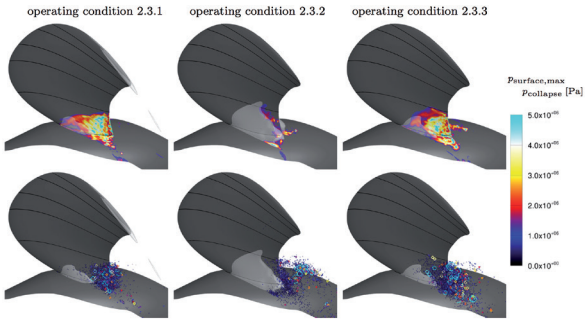

Figure 4: Assessment of flow aggressiveness (view on the suction side of a single propeller blade). Comparison between maximum surface pressures p_{max} (top) and scatter plot of collapse events (bottom). The size and color of the symbols correspond to the magnitude of the pressure pulse upon cavity collapse $p_{collapse}$.

with collapse events for the three operating conditions. By performing a statistical analysis of recorded collapse locations, rates and intensity, this enables the computation of collapse spectra. The results indicate that the cavitating flow is least aggressive at the design point of the propeller.

HPC resources are a necessity for the presented studies due to the disparity between the extremely small time scales of compressible wave dynamics (~$1 \cdot 10^{-9}$ s) and time intervals of interest, e.g. the revolution of the propeller (~$1 \cdot 10^{-2}$ s). Utilizing a maximum of up to 2028 cores for the parallel computations, the investigations required a total of approximately $20 \cdot 10^{6}$ CPU-hours.

The presented analyses allowed us to validate the chosen numerical approach. They helped us to better understand the flow dynamics of partial shedding cavities experiencing sheet-to-cloud transition. Furthermore, it was demonstrated that the approach can be used to assess local flow aggressiveness on ship propellers.

On-going Research and Outlook

Current research topics focus on the mutual interaction between boundary layers, fluid turbulence and cavitation dynamics. Furthermore, the flow solver is extended to enable a closer modelling of the actual operating environments of ship propellers. This includes the presence of a ship wake or oblique inflow, as well as to investigate the effect of control surfaces, e.g. struts or rudders. The extremely small time steps required by the model currently demand high-performance computing resources. Therefore, with the aim to reduce the computational effort, several acceleration techniques, such as implicit time stepping methods, will also be investigated.

References and Links

[1] http://www.aer.mw.tum.de/abteilungen/gasdynamik
[2] Schmidt, S.J., Sezal, I. H., Schnerr, G. H. and Thalhamer, M. (2008). "Riemann techniques for the simulation of compressible liquid flows with phase-transition at all Mach numbers - shock and wave dynamics in cavitating 3D micro and macro systems," 46th AIAA Aerospace Sciences Meeting and Exhibit, Reno
[3] Budich, B., Neuner, S., Schmidt, S. J. and N. A. Adams (2015). "Numerical investigation of shedding partial cavities over a sharp wedge," 9th International Symposium on Cavitation, Lausanne
[4] Budich, B., Schmidt, S. J. and N. A. Adams (2015). "Numerical Investigation of a Cavitating Model Propeller Including Compressible Shock Wave Dynamics," 4th International Symposium on Marine Propulsors, Austin
[5] Budich, B., Schmidt, S. J. and N. A. Adams (2015). "Numerical Simulation of Cavitating Ship Propeller Flow and Assessment of Erosion Aggressiveness," 6th International Conference on Computational Methods in Marine Engineering, Rome

Specific turbulent kinetic energy transfer in very anisothermal flow

Research Institution
CEA Grenoble DEN/DANS/DM2S

Principal Investigator
Benoit Mathieu

Researchers
Frederic Aulery, Adrien Toutant, Gauthier Fauchet, Françoise Bataille

Project Partners
CNRS-PROMES laboratory, Rambla de la thermodynamique, Perpignan France

SuperMUC Project ID: pr86bo (PRACE project)

Introduction

At the PROMES laboratory, we study the effect of very strong temperature gradients on the turbulence of wall bounded flows. This study is motivated by the flow characteristics inside solar receiver of concentrated solar power tower plants. As an example, PEGASE (Production of Electricity from GAs turbine and Solar Energy) is a technology of solar plant that uses pressurized air at very high temperature.

In the situation of high temperature gradients, the interaction between energy and momentum equations are strong and classical models are not valid anymore. The temperature gradient can be considered as a strong external agency that modifies the turbulence properties. In the case of low speed flow, only very few studies are dedicated to this coupling. In particular, there is no reference data that concerns subsonic flow without low Reynolds number effect and with dilatational effect due to strong thermal gradient.

TrioCFD is a fluid mechanic c++ code of the TRUST platform designed for HPC with a natural MPI parallelization that allows to run over 10 000 cores. With a specific module developed for this project, the operator and pressure solver use SEE vectorization, IJK structured grid and a cache optimized multigrid solver that give us a speed up of 3 from the previous version of TrioCFD.

Using TrioCFD code, we run a large Direct Numerical Simulations (DNS). With this DNS, we will be able to study the thermal boundary layer in the case of dilatational flows. Furthermore, it can help to better understand the coupling between the dynamic and the thermal part in low Mach flows.

Results

In solar receiver air flows, the velocity dilatational effect is negligible and the thermal dilatation has strong effect. Consequently, we use the low Mach flow assumption (Ma<0.3). This implies that continuity, momentum and energy conservation equations are coupled. For the numerical point of view, the equations are coupled using a fully conservative algorithm that ensures local and global conservation mass and energy independently of the mesh.

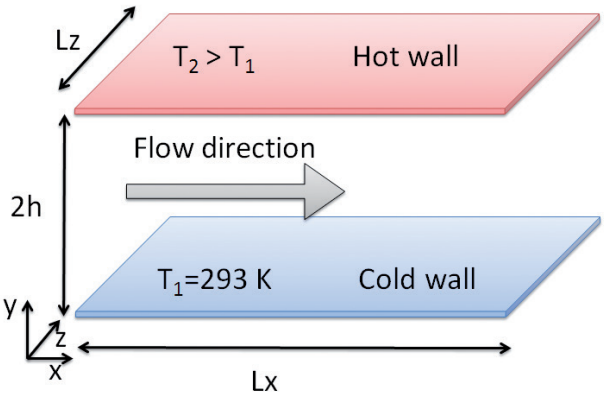

Figure 1: Channel flow configuration, thermal gradient is applied between the two walls.

Simulation set up
The time integration is done using a Runge-Kutta third-order scheme. The convective operators are centered fourth order for velocity and third order QUICK for mass. We use a suitable geometry for the study of wall-fluid interaction: the bi-periodic channel flow (figure 1). The lengths of the channel are Lx=4πh, Ly=2h, and Lz=2πh. We apply a ratio between the two walls Tr=T2/T1=2. The turbulent Reynolds number (based on friction velocity) is 395.

The mesh is regular in periodic directions (1536 for x and 1536 for y) and hyperbolic in wall-normal direction (896 cells). The dimensionless first cell size is y^+=0.125.

In total, the 2.1 billion cells are divided between 8192 cores (one whole island). This simulation has run over 500 k time step for a total time of 6 000 000 cpuhours. The multigrid solver uses 6 grid levels and a mixed precision method (3 simple precision iterations followed by

a double precision iteration) consequently the domain is splited only on eight parts in wall normal direction (respectively 32 and 16 in x and y directions). The amount of data generated is around 8 To.

Simulation profiles

We have compared the results from the DNS with a fine LES (100 times less cells and this study algorithm) computed on JADE cluster. The LES has the same root mean square profiles (temperature and velocity) than the new DNS with a magnitude error near to 5% (not shown here). This good result allows using the LES for the study of the turbulent kinetic energy transfers in the spectral space.

The turbulent kinetic energy evolution can be decomposed into three main mechanisms which are the turbulent production, the total transfer and finally the viscous part. These mechanisms can be split into a part that corresponds to the incompressible limit and another part that corresponds to the dilatational effect. This last part is specific to very anisothermal flow and is equal to zero in the incompressible limit. In the following, this part is called *thermal part*.

The figure 2 shows the thermal part of the total transfer at the cold and hot sides of the channel in function of the wave number (horizontal axis) and the dimensionless distance to the wall (vertical axis). This transfer corresponds to a destruction of turbulent kinetic energy at the hot side and to a production of turbulent kinetic energy at the cold side.

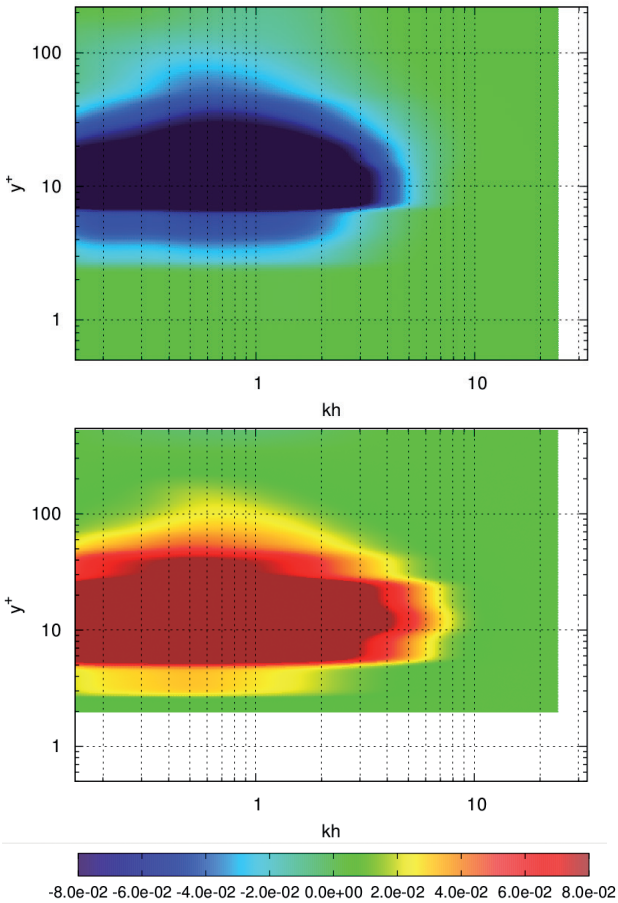

Figure 2: Thermal part of the total transfer in function of the wave number (horizontal axis) and the dimensionless distance to the wall (vertical axis). Hot side of the channel at the top and cold side at the bottom. The magnitude of the thermal part is rescaled by the maximum magnitude of production.

The maximum magnitude of the transfer thermal part represents ten percent of the maximum magnitude of the production. Obviously this transfer does not correspond to a Reynolds effect. It shows a complex coupling between temperature and velocity that creates a dissymmetry of turbulent statistics at the hot and cold sides. Concerning solar power tower plant, this transfer has to be avoided. Indeed, we want a high intensity of turbulent kinetic energy at the hot side in order to increase the heat exchange between the wall and the working fluid. We have shown that the transfer thermal part is due to the velocity normal to the wall that is created by the temperature gradient. In order to reduce this transfer and to increase the heat exchange between the wall and the working fluid, we propose to use vortex generators. These vortexes remove the velocity normal to the wall and decrease the thickness of the viscous sublayer. These two effects increase the heat transfer between the wall and the working fluid without increasing too much pressure drop.

On-going Research / Outlook

This very fine simulation constitutes a very interesting database. We will compare the DNS results with other LES and DNS. In the case of very anisothermal flows, the variations of fluid properties are important. Concerning turbulence modelling, they create specific subgrid terms. It can be interesting to compute these terms usually neglected in subgrid models (like viscosity-velocity gradient correlation). With these terms, we should be able to improve thermal subgrid models for LES. Furthermore, we can study the turbulent kinetic energy equation in physical and spectral space. In addition, by comparisons to other setup (thermal rate, friction number, boundary conditions, geometry...) we will be able to understand the effects of the interaction between turbulence and temperature and propose some geometry modifications for solar receivers.

Finally, this academic work has for goal to improve the understanding on the coupling between turbulence and thermal gradient in low-Mach flows using numerical simulations. Following, this work is used to propose new innovative heat exchanger designs. Implemented in solar tower power plant, these exchangers will decrease the cost of electricity production by a better efficiency.

References and Links

[1] A. Toutant, F. Bataille, "Turbulence statistics in a fully developed channel flow submitted to a high temperature gradient", International Journal of Thermal Sciences, 74, 0, 104-118, December 2013.

[2] Aulery F., Toutant A., Bataille F. and Zhou Y., Energy transfer process of anisothermal wall-bounded flows, Physics Letters A, vol. 379, p. 1520-1526, 2015.

TrioCFD project web site: http://www-trio-u.cea.fr/
Promes Laboratory web site: www.promes.cnrs.fr

CFD based low-order modeling of the nonlinear flame dynamics and of swirl fluctuations

Research Institution
Professur für Thermofluiddynamik

Principal Investigator
Wolfgang Polifke

Researchers
Stefan Jaensch, Alp Albayrak

Project Partners
Siemens, Alstom and Rolls Royce

SuperMUC Project ID: pr86fa

Introduction

Despite the development of power plants using renewable energy sources, gas turbines will play an important role in future energy production. They offer operational flexibility and at the same time low emission of greenhouse gases. These are important properties to serve as backup solution in the age of renewable energy sources, which is essential to maintain net stability and a reliable energy supply. Thermoacoustic oscillations limit the development of gas turbines aimed to lower emission of pollutant and a higher operational flexibility. The basic mechanism behind thermoacoustic oscillations is as follows: Small initial fluctuation of the velocity, say, yield fluctuation of the global heat release rate of the flame. This unsteady heat release rate acts as a volume source, which in term creates acoustic waves. These waves are reflected at the boundaries of the burner back to the flame and perturb the flame, again. This feedback can get unstable and yield very large oscillations. If the machine is not turned down, these oscillations can cause significant damage.

The occurrence of those instabilities in a gas turbine depend on the interaction of all parts of the engine. However, due to limited computational power it is by no means possible to simulate a whole gas turbine within a single LES (Large Eddy Simulation). Low order network models are the state of art approach for estimating these instabilities. First, a low-order model of each component in a gas turbine is determined. Then the low order models are interconnected in order to predict the global heat release rate of the flame. Entirely acoustic elements without reactive flows can be modeled with a linearized version of the Navier-Stokes equations. However, a low-order model for the flame cannot be found on this way.

Therefore, in the present project the so called CFD/SI approach is investigated [3]: The flame is simulated with a LES. These simulations are expensive and necessitate the use of SuperMUC. In order to deduce low-order models efficiently from the LES, the LES is perturbed with a broadband excitation signal. The resulting fluctuation of a reference velocity and of the global heat release rate are measured. The time series collected are post-processed with system identification methods in order to determine the low-order models.

This method has already been proofed to be both accurate and efficient. In the scope of the present project two new aspects are investigated: (1) It is investigated how non-linear low order models can be deduced. (2) The impact of swirl waves on the flame dynamics is investigated. For this purpose, two different swirl burners are investigated, one having axial swirler (BRS burner) and the other radial swirler (FVV burner) shown in Fig 1. and 2., respectively.

Results and Methods

The object-oriented C++ Software package OpenFOAM [4] is used to perform the LES simulations. OpenFOAM employs an implicit finite volume

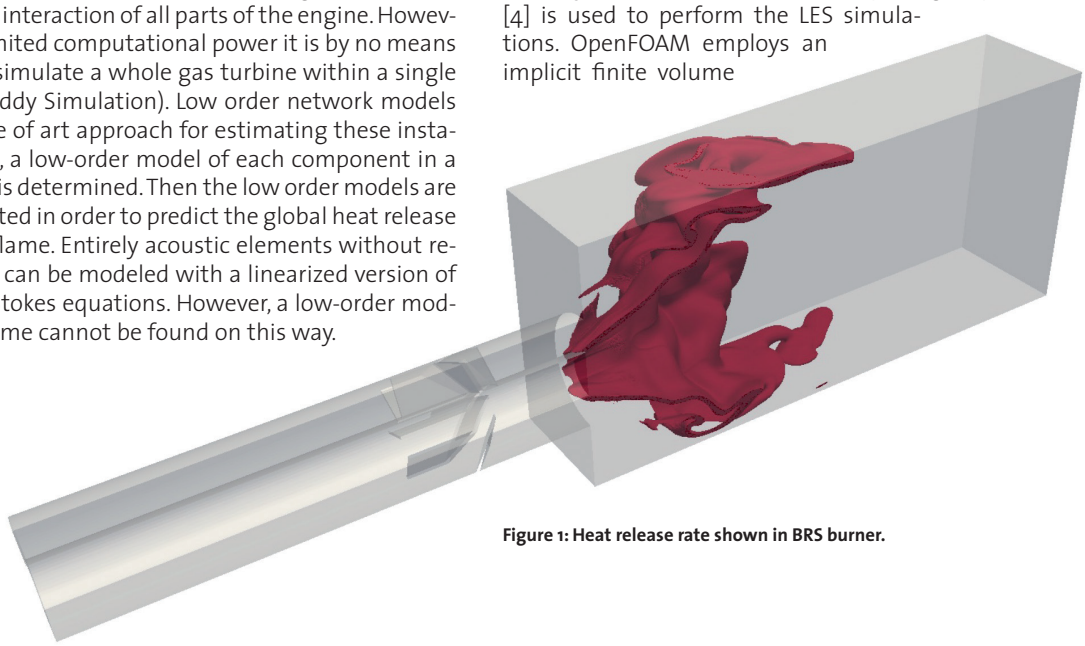

Figure 1: Heat release rate shown in BRS burner.

Figure 2: Streamlines shown in FVV burner.

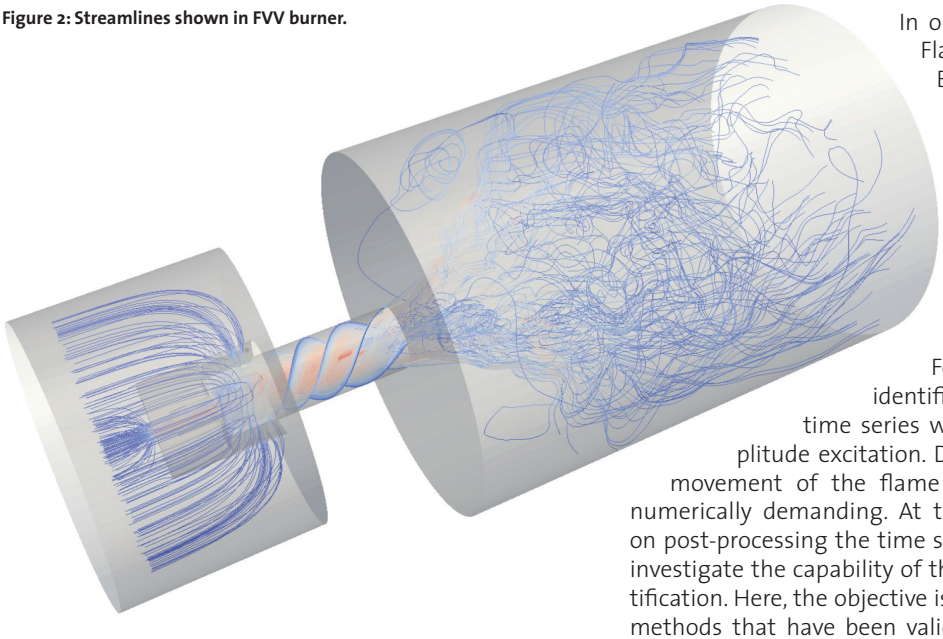

scheme with well-known PISO and SIMPLE algorithms. The solver is based on the standard solver reactingFOAM.

The turbulence is modeled with Smagorinsky sub-grid scale LES model. Global 2-step chemistry is used to model Methane-Air combustion. The Thickened Flame Model is implemented for decreasing the mesh resolution requirement by artificially thickening the flame. The low Mach number assumption is used in order avoid acoustic wave reflection at boundaries. Adaptive Mesh Refinement capability is added to the solver for refining the mesh only in flame region and therefore saving computational time.

The BRS burner rig is used for code benchmarking. Experiments and numerical results using AVBP [3] are available in literature. In Fig. 3, the Flame Transfer Function calculated by means of experiment is compared against numerical approach (LES-SI) by different solvers AVBP and OpenFOAM. A good agreement with experiment is achieved with OpenFOAM simulations.

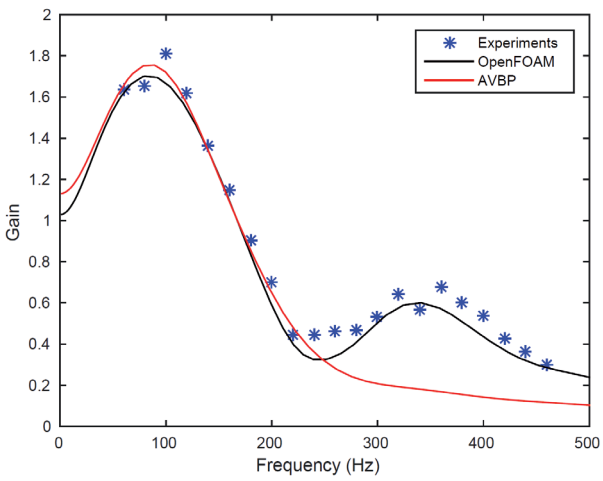

Figure 3: Comparison of Flame Transfer Functions from experiment and numerical simulations by OpenFOAM and AVBP.

In order to compute a single Flame Transfer Function for BRS burner, time series around 0.35 seconds are required. This can be achieved in 120000 CPU hours with 840 processors.

On-going Research / Outlook

For the nonlinear system identification we created long time series with broadband, high amplitude excitation. Due to the resulting large movement of the flame these simulations were numerically demanding. At the moment the focus is on post-processing the time series obtained in order to investigate the capability of the nonlinear system identification. Here, the objective is to investigates how well methods that have been validated for laminar flames [6] can be used to model thermoacoustic oscillations of turbulent flames. As the focus of this work is now on post-processing the CFD data the SuperMuc is not required anymore for this part of the work.

With the SuperMuc our current focus is on the investigation of the swirl flames. Here, a model for the influence has been developed by the authors of the present project [7]. This model is to be validated against a turbulent swirl flame. Extensive numerical studies with the SuperMuc will be necessary.

One proposal continuing the work done within this project has just been submitted. The idea is to use the CFD/SI approach to deduce low-order models for the combustion noise. This is of high industrial interest, as such model combined with a thermoacoustic network models allows to predict the noise emitted by the engine.

In future it is planned to investigated the uncertainty of the prediction in more detail. Here, physical parameters as the wall temperature and the turbulence model as well as numerical parameters as as the discretization scheme or the mesh are to be investigated. These study will require are huge computational and thus, necessiate the use of SuperMUC

References and Links

[1] http://www.tfd.mw.tum.de/index.php?id=99
[2] http://www.tfd.mw.tum.de/index.php?id=105
[3] W. Polifke, "Black-box system identification for reduced order model construction," Annals of Nuclear Energy, vol. 67C, pp. 109–128, May 2014.
[4] www.openfoam.org
[5] http://www.cerfacs.fr/avbp7x/
[6] S. Jaensch, M. Merk, E. Gopalakrishnan, S. Bomberg, T. Emmert, R. I. Sujith, and W. Polifke, "Hybrid CFD/ low-order modeling of nonlinear thermoacoustic oscillations," in submitted to the 36th Symposium of the Combustion Institute, Seoul, Korea, 2016.
[7] A. Albayrak, W. Polifke, "On the propagation velocity of Swirl Waves in Annular Flows",Proceedings of the 21st International Congress on Sound and Vibration, Beijing, China, 2014.

Aerodynamic Investigations of Vortex-Dominated Wing Configurations with Active and Passive Flow Control

Research Institution
Institute of Aerodynamics and Fluid Mechanics, Technische Universität München

Principal Investigator
Christian Breitsamter

Researchers
Andreas Hövelmann, Jae-Hun You, Stefan Pfnür, Andrei Buzica

Project Partners
German Research Foundation (DFG), Airbus Group, NATO Science and Technology Organization (STO, AVT-183), Group for Aeronautical Research and Technology in Europe (GARTEUR, AG-49)

SuperMUC Project ID: pr86fi

Introduction

The flow field characteristics of swept low-aspect-ratio wing configurations with delta or diamond wing shape are typically dominated by large-scale vortex structures, which originate from flow separations at the wing leading edges. With increasing angle of attack, the vortices grow in size and strength, and unsteady effects commonly become relevant. The aerodynamic characteristics then may change considerably, and large impacts on stability and control of the wing configurations are observed. Nonlinear flow interactions occur over a wide range of analyzed scales and consequently, complex flow fields are often present. By use of active and passive flow control mechanisms such as dynamic blowing at the wing surface or rather varying leading-edge contours (active leading-edge flaps and contour modifications), the vortex characteristics are manipulated and controlled in some extent. Stability and control of respective low-aspect-ratio wing configurations can therefore mainly be improved.

Results and Methods

Scientific Results
Extensive analyses have been conducted in the present SuperMUC project on leading-edge contour modifications of a diamond wing configuration [1,2]. Depending on sharp and rounded leading-edge segments in spanwise direction, the flow separation and the formation of leading-edge vortices is thus controlled. For sharp leading edges, the flow separation is fixed geometrically, but for rounded leading-edge contours, the flow separation occurs as a smooth surface separation on the upper wing surface. Consequently, additional parameters such as leading-edge radius or Reynolds number become more relevant. Figure 1 highlights corresponding results of the analyzed diamond wing configuration with spanwise-varying leading-edge contours. At the considered angle of attack of $\alpha = 24°$, in total three dominant leading-edge vortices occur, which arise in the respective wing segments from both sharp and rounded leading edges. They interact with each other and thus influence the overall aerodynamic characteristics as well as the stability behavior of the diamond wing configuration. Several combinations of varying leading-edge contours have been considered in the CFD computations, and flow separation onset characteristics and vortex-vortex interactions have mostly been understood. The results thereby contribute to an improved knowledge on relevant flow phenomena in the respective field of research.

Moreover, the effect of active flow control on leading-edge vortex systems has been analyzed in this SuperMUC project for a delta wing configuration. The occurring large-scale vortical structure is influenced by pulsed blowing jets, which operate from several slots located close to the leading edge. The stability of the leading-edge vortex system is thereby controlled. Figure 2 presents resulting differences between the baseline

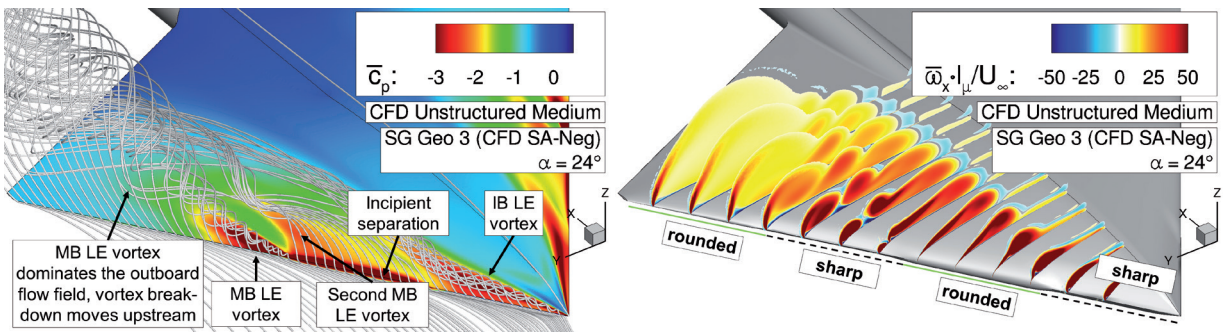

a) $\alpha = 24°$, surface pressures and field streamlines. **b) $\alpha = 24°$, axial vorticity contours.**
Figure 1: Leading-edge vortex formation due to spanwise-varying leading-edge contours, SAGITTA configuration.

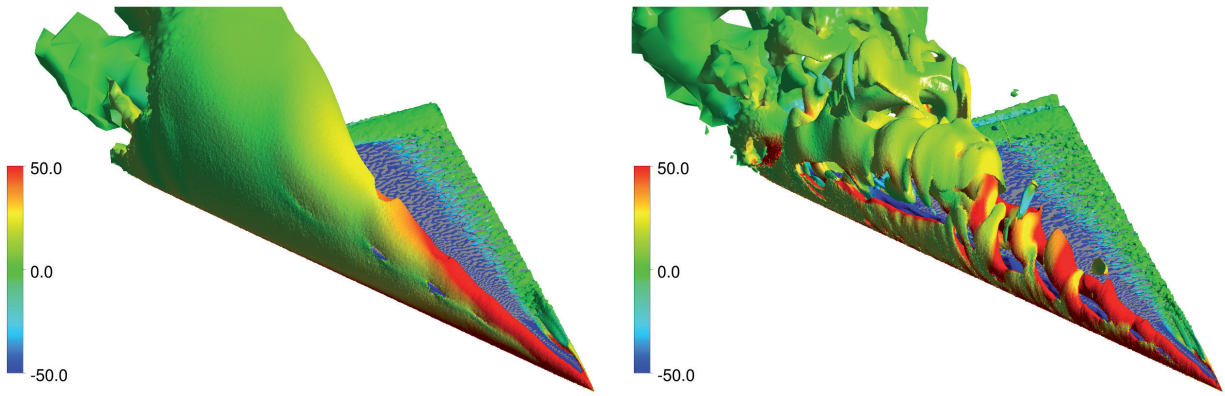

a) α = 23°, baseline case. **b) α = 23°, actuated case, fpulse = 65 Hz.**
Figure 2: Iso-surface of Q-criterion (Q=50 1/s²) colored by contour levels of normalized axial vorticity (ω_x * l_{ref} / U_∞), VFE-2 configuration (half wing).

and the actuated case at α = 23°. One can notice that the flow structures originating from the blowing jets affect the overall flow structure of the main leading-edge vortex. Thereby, unsteady flow phenomena such as vortex bursting are delayed, which leads to more stable aerodynamic characteristics. For even higher angles of attack, related CFD computations indicate that the existence of the leading-edge vortex can be maintained up to α = 45° by the applied active flow control. This consequently improves the stability and control of the wing configuration, and the reachable flight envelope of practical applications can thus be extended.

Numerical Methods
The flow computations in this SuperMUC project have mostly been performed with the Reynolds-Averaged Navier-Stokes (RANS) equations. They are considered to be the most complete description of flows in continuum mechanics and describe the exchange of momentum in the fluid considering friction. Since turbulent fluctuations within the flow field are not resolved, but modeled, turbulence models must be applied to closure the set of partial differential equations. In case of vortex-dominated wing configurations, turbulence modeling is particularly important, as a correct detection of the flow separations and the large-scale vortex structures is decisive for the quality of the results. Both steady-state RANS computations and time-accurate (URANS) flow simulations have been considered. Unsteady flow phenomena have been accounted for by small time steps, which overall leads to large computer resource requirements. For an even better spatial and temporal resolution, hybrid algorithms such as Detached Eddy Simulation (DES) have also been applied. Compared to the (U)RANS equations, high-energy turbulent scales are not modeled anymore in DES/LES computations, but resolved. The computational effort increases drastically, but the results become more accurate both in space and in time.

Most of the numerical investigations have been computed with the TAU-Code, a CFD solver developed at the DLR (German Aerospace Center) Institute of Aerodynamics and Flow Technology [3]. It is developed for optimized parallel efficiency on high-performance computers and uses several features such as the multigrid technique for convergence acceleration. The DLR TAU-Code is based on a finite volume discretization, and several upwind and central schemes can be applied. If turbulent flows are considered, a variety of one- and two-equation models as well as different kinds of Reynolds stress models (RSM) can be chosen. The commercial flow solver ANSYS CFX has been used as well for the numerical investigations. The widespread flow simulation software combines an advanced solver with powerful pre- and postprocessing capabilities. It runs in parallel mode and uses an element-based finite volume method. Various turbulence models are available, and for instance, the SAS model can resolve spectral contents of unstable flow even in URANS computations. Consequently, it can provide a LES-like resolution of turbulent eddies in the unsteady flow region.

Typical computations of the present investigations have been run with up to 840 cores per job, and about 10 Mio core hours have been spent so far. Approximately 20 TB data has already been accumulated, and the scheduled computations currently require increased resources. The analyses are characterized by various parameter studies, which range from changed flight conditions to different flow control parameters and wing shape variations.

On-going Research and Outlook

The on-going research currently concentrates in particular on the flow separation characteristics of low-aspect-ratio wing configurations with rounded leading-edge contour. Moreover, the flow control aspects on suchlike configurations are still in the scope of the analyses. Consequently, future work will focus in more detail on the computation of cases with pulsed blowing and leading-edge contour actuation to efficiently control the vortex structures. So far, the GCS Supercomputer SuperMUC has mainly contributed to the success of the research investigations. With respect to the upcoming CFD investigations, it will be essential for the progress in flight physics understanding.

References and Links

[1] Hövelmann, A. and Breitsamter, C.; "Leading-Edge Geometry Effects on the Vortex Formation of a Diamond-Wing Configuration"; Journal of Aircraft, Vol. 52, No. 5, 2015, pp. 1596–1610.
[2] Hövelmann, A., Pfnür, S., und Breitsamter, C., "Flap Efficiency Analysis for the SAGITTA Diamond Wing Demonstrator Configuration", CEAS Aeronautical Journal, Vol. 6, Iss. 4, 2015, pp. 497–514.
[3] Gerhold, T.; "Overview of the Hybrid RANS Code TAU"; MEGAFLOW – Numerical Flow Simulation for Aircraft Design, Vol. 89 of Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Springer Verlag, 2005, pp. 81–92.

waLBerla – A massively parallel framework for fluid simulations

Research Institution
Chair for System Simulation (Informatik 10), Friedrich-Alexander-Universität Erlangen-Nürnberg

Principal Investigator
Harald Köstler

Researchers
R. Ammer, D. Bartuschat, M. Bauer, C. Godenschwager, K. Pickl, F. Schornbaum

Project Partners
Katarina Gustavsson (School of Computer Science and Communication, Royal Institute of Technology (KTH), Stockholm), Ana-Sunčana Smith (Cluster of Excellence "Engineering of Advanced Materials", FAU Erlangen-Nürnberg)

SuperMUC Project ID: pr86ma

Introduction

About a decade ago, the development of the waLBerla software framework [1] (widely applicable lattice Boltzmann solver from Erlangen) has started as a joint effort of several researchers at the University of Erlangen-Nürnberg. This framework provides a common basis for fluid simulation codes using the lattice Boltzmann method. Over the years, the capabilities and the number of contributors also from other universities and institutes have grown. Recent applications of the software are:

- electrokinetic flows,
- fluid-particle interaction phenomena of non-spherical particles,
- blood flow in complex geometries and local refinement,
- electron beam melting,
- microscopic swimmers.

In order to simulate real-world scenarios, waLBerla relies on using an immense compute power only available on modern supercomputers. This is reflected in the architecture of waLBerla that has been designed systematically in every step to run efficiently on massively parallel architectures.

Results

In the scope of this SuperMUC project, we performed scaling experiments and production runs for the previously listed applications realized withing waLBerla. In the following, we will present the principal results for electrokinetic flows and fluid-particle interactions of elongated particles. For the simulations presented below, in total 15 million core hours were used on SuperMUC.

Electrokinetic flows
Electrokinetics is the mechanism of choice for microfluidic manipulation and actuation in lab-on-a-chip (LoC) systems. The increasing importance of LoC devices is attributable to the fact that they can be used as portable biological analysis devices for point-of-care diagnostics. At the small scales of LoC systems, measurements of the flow are very difficult. Therefore, simulations are required for the design and optimization of such systems. In order to accurately capture the multiple physical effects at small scale, a very fine discretization and small time steps are necessary, resulting in the need for a large amount of computational resources.

To facilitate particle-laden electrokinetic flow simulations, waLBerla was augmented by a module of efficient parallel solvers for sparse linear systems arising from finite volume discretizations of electric potential equations. Moreover, to simulate the fluid-particle interaction, modeled with the momentum exchange method, waLBerla is additionally coupled to the physics engine pe.

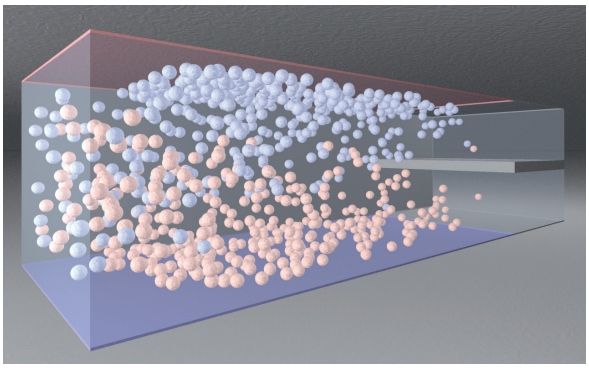

Figure 1: Multiphysics simulation of bisecting micro-channel: separation of oppositely charged spherical particles with a radius of 60µm in fluid flow [2].

The implemented algorithm for the electric migration of charged particles in non-electrolyte solutions was validated against analytical solutions for a single spherical particle. In the simulations that were running for 3.1 to 14.5 hours on 1815 to 8192 cores of a thin-node island, 34000 to 79000 time steps were performed. By means of the simulations, the impact of different simulation parameters on physical accuracy was examined.

Moreover, as example scenario, the separation of interacting charged macromolecules in fluid flows inside

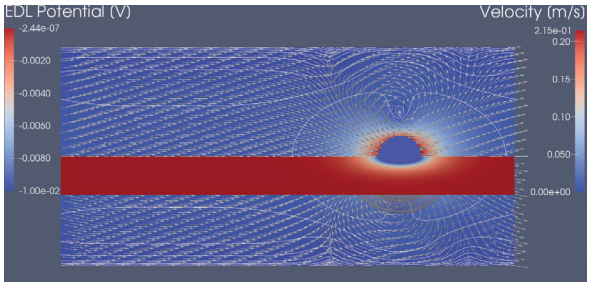

Figure 2: Electrophoresis of spherical particle with a radius of 120nm in micro-channel with insulating no-slip walls. Visualization of flow field in vertical plane, EDL potential in horizontal plane, and ion charge distribution around charged particle as semi-transparent hemispherical shells (after 30001 time steps) [3].

channels of dimensions relevant for LoC was simulated. The result of such a multiphysics simulation on 144 cores after 108301 time steps is visualized in Figure 1.

Furthermore, weak and strong scaling experiments were conducted on SuperMUC, showing good parallel scaling on up to 32768 cores (4 thin-node islands) with more than seven million fully resolved charged particles [2]. In this project, also the electrophoretic motion of charged particles in electrolyte solutions was simulated. These simulations additionally include the effect of ions in the fluid on the particle motion. Simulation results for the electrophoresis of a charged sphere in a micro-channel are illustrated in Figure 2.

To validate the electrophoresis simulations, the electrophoretic velocity of a charged sphere due to an applied electric field was simulated on SuperMUC for various simulation parameters, mostly on 8192 cores (1 thin-node island). The results obtained within 37 to 48 hours runtime were compared against analytical solutions. In these simulations, 70000 to 95000 time steps were performed to ensure that the particles had reached their terminal velocities. The validations show that the implemented algorithms can correctly simulate the electrophoretic motion of charged particles [3].

Tumbling spherocylinders

To verify the correctness of the implemented LBM method coupled to the rigid body dynamics method for the fluid-particle interaction of non-spherical objects in the flow, the translational and rotational motions of single spherocylinders in Stokes flow were compared against analytical solutions in [4].

From these simulations, the influence of shape, inertial, and wall effects on the motion of elongated particles were systematically examined. Additionally, visualizations of the flow field around sedimenting and tumbling spherocylinders were presented in [4].

The large-scale LBM simulations for the translational and rotational motion with 37000 to 88800 time steps were performed on SuperMUC on up to 8192 cores (1 thin-node island) within 8 to 19 hours. For the tumbling motion simulations, 605000 time steps were conducted on 768 cores of SuperMUC within 48 hours.

On-going Research / Outlook

The great parallel performance and the efficient scheduling of SuperMUC helped performing the computationally challenging multiphysics simulations within a short time. SuperMUC Next Generation may facilitate even larger simulations and more detailed models.

In the future, we plan to perform simulations of the electrophoretic motion of large numbers of charged particles in electrolyte solutions for real-world LoC scenarios. Moreover, we are going to perform simulations of complex geometry flow problems in human vascular structures. Especially we want to explore the flow in the coronary arteries and the heart muscle on SuperMUC Phase 2.

References and Links

[1] C. Godenschwager et al. (2013). A framework for hybrid parallel flow simulations with a trillion cells in complex geometries. In Proc. of SC13: International Conference for High Performance Computing, Networking, Storage and Analysis. doi:10.1145/2503210.2503273

[2] D. Bartuschat and U. Rüde (2015). Parallel Multiphysics Simulations of Charged Particles in Microfluidic Flows. Journal of Computational Science, 8(0):1-19. doi:10.1016/j.jocs.2015.02.006

[3] D. Bartuschat (submitted 2015). Direct Numerical Simulation of Particle-Laden Electrokinetic Flows on High Performance Computers. PhD thesis, University of Erlangen-Nürnberg.

[4] D. Bartuschat, E. Fischermeier, K. Gustavsson, U. Rüde (2016). Two computational models for simulating the tumbling motion of elongated particles in fluids. Computers & Fluids, 127(0):17-35. doi:10.1016/j.compfluid.2015.12.010

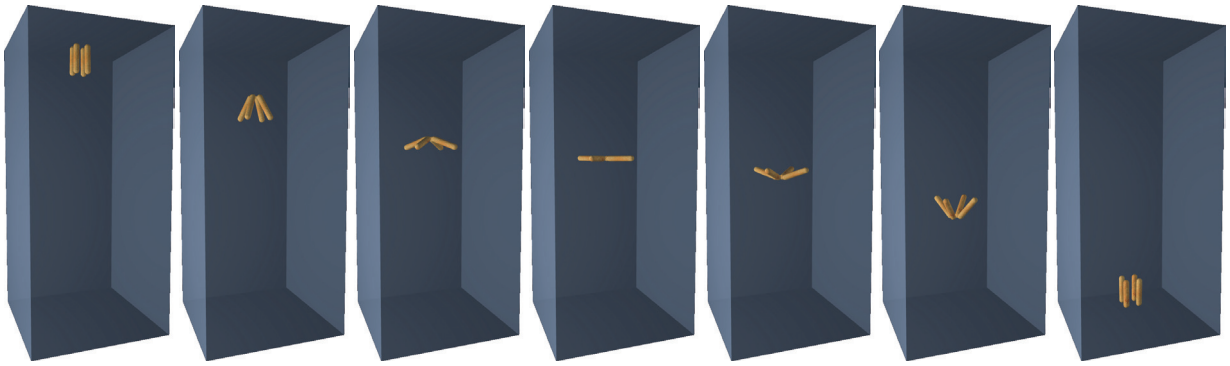

Figure 3: Tumbling spherocylinders sedimenting in Stokes flow [4].

Computational acoustics of supersonic turbulent free round jets

Research Institution
Institute for Fluid Mechanics and Engineering Acoustics

Principal Investigator
Jörn Sesterhenn

Researchers
Juan José Peña Fernández

Project Partners
Institute for Fluid Mechanics and Engineering Acoustics, TU Berlin

SuperMUC Project ID: pr86po (Gauss Large Scale project)

Introduction

Noise prediction is one of the most relevant topics for Computational Fluid Dynamics due to the fact that noise optimization, energy saving and pollutant emission minimization complement each other. In this project headed by Professor Jörn Sesterhenn of the Technische Universität Berlin, numerical simulations of a supersonic jet were performed on HPC system SuperMUC of LRZ, focusing on the research of the acoustic field.

The most accurate method to simulate this problem is based on Direct Numerical Simulation. However, this kind of simulation requires a huge amount of computational resources which to date are not available to most of the researchers or the industry. To overcome these limitations, a feasible approach to said problem is based on Large Eddy Simulations (LES). The downside of LES, however, is that only the large scales of the problem are solved whilst the small structures created by the turbulence are modeled with a mathematical algorithm. This is due to the fact that the small structures show universal behavior for high Reynolds numbers in which the large scales are decoupled of the small ones.

Of special interest in this project was the so called 'near to *intermediate field*' in which the acoustic properties can be directly related to the fluid structures that generate these acoustic waves.

The main objective of this project is to describe the noise generation mechanisms that take place in the supersonic jet case. The so called '*turbulent mixing*' noise is created in both subsonic and supersonic jets. Turbulent mechanisms in the shear layer of the jet are the main reason for noise generation. This kind of noise is mainly radiated downstream of the jet. If a supersonic jet is not in adapted condition, a quasi-periodic structure with shock-waves is generated. These shock-cell structures generate some additional noise, which shows two different characteristics depending on the generating mechanism. '*Broadband shock-associated*' noise radiates mainly in upstream direction and is generated by the present shock-cell structure. The other mechanism is the one that is being referred to as *"screech"*. This is a single frequency tone which is the result of a feedback loop at the nozzle exit. The generated noise and the eddies which are detached from the shear layer generate a single frequency tone. Depending on the relative velocity of the jet to the speed of sound, different large coherent structures are present in the flow field, and the noise mechanism change. Examples of such structures are single or double helical modes and toroidal modes.

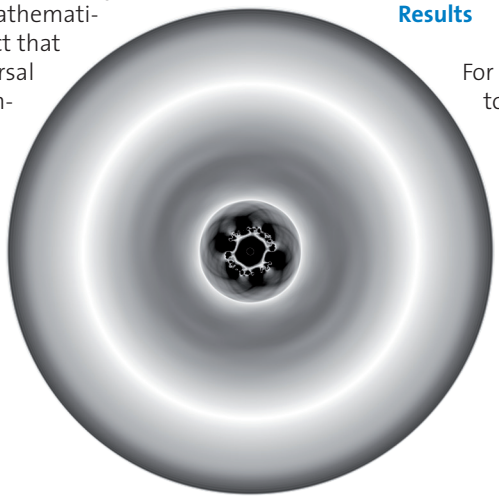

Figure 1: In the first stage of the jet, the first acoustic wave is shown over a plane normal to jet axis. In addition, the shear layer can be seen and the turbulent structures that in the shear layer are generated. The region in which the acoustic phenomena are radiated can be also perceived. Direct Numerical Simulation of a free round jet with Reynolds number 10.000.

Results

For free round jets focusing the near to intermediate field, a simulation should be performed in a box of dimensions 25D x 15D x 15D (while D is the diameter of the nozzle exit). In our case, a grid of 2048 x 1024 x 1024 = 2.1 billion cells was necessary. To perform the needed statistics, we had to run the simulation for a minimum of 30.000 time steps. This means that 2.1 billion points and the governing equations had to be solved at every individual time step. Taking all these factors into account, this simulation required at least 22.9 million CPU hours which means this project would have taken more than 2614 years to compute if executed on a single CPU sys-

tem. Thanks to now available petascale HPC systems like SuperMUC of the LRZ, though, projects of this magnitude can be tackled. Due to the fact that on SuperMUC more than 8000 CPUs simultaneously computed the numerical solution, we were able to execute this project within the time frame of just one year.

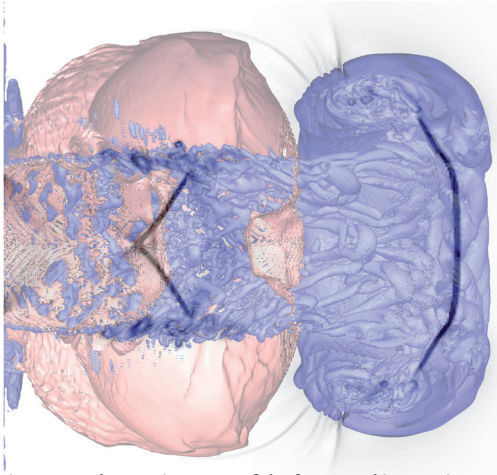

Figure 2: In the starting stage of the free round jet, Q-criterion and pressure isosurfaces are plotted in blue and red, respectively. In black and white the density gradient is presented as a pseudo-schlieren over the mean plane. Here the interaction between flow and acoustic phenomena can be observed.

The simulations were performed with a self developed code, written in FORTRAN language, that solves the compressible Navier-Stokes equations in characteristic formulation with finite differences. For this case is especially important to have a high order numerical code in order not to damp the acoustic waves that we try to investigate. Our code is sixth order and compact in space, and through a Runge-Kutta time integration routine we have fourth order in time.

In this project, both Direct Numerical Simulations (DNS) and Large Eddy Simulations (LES) are performed. The DNS are essentially the reference of the problem in which we are interested. With them we infer the physical phenomena that are occurred and we get more information about the noise generation process. The Large Eddy Simulations are the feasible solution for the current industry problems. Today DNS are not affordable for the industry, and therefore only the academia is able to perform this kind of simulations. The problems that the industry have to face, the academia is not able to solve, but in this context the Large Eddy simulations are the wire that connect both of them. In order to run enough accurate LES, the mathematical models behind have to be developed and validated, and this is exactly the aim of this part of the project, the validation of two LES models.

The most important finding of this project is the difference between the structures that are generated in the shear layer in a turbulent jet in comparison with a laminar one. The spatio-temporal evolution of these structures, and more specifically the effect that these structures have in the radiated noise is under detailed investigation.

In figure 1 the shear layer has a hexagonal shape in which different noise sources can be identified. Around the shear layer a ring can be identified. This ring is the region in which the noise generated has been radiated, but due to the early stage of the fluid flow, the noise has not yet reached the whole domain. The big circle shown in figure 1 is the first acoustic wave that is generated when the jet starts.

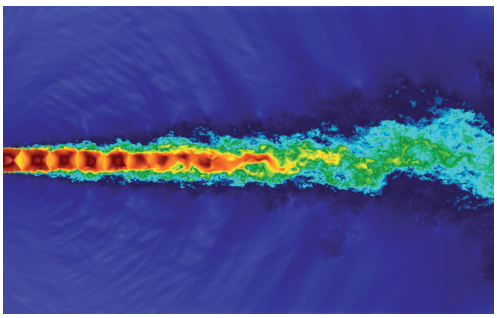

Figure 3: Continuous stage of the free round jet. Contour of velocities in logarithmic color scale. The flow phenomena as well as the acoustic ones can be clearly seen in this dual image.

Of special interest is also the first vortex ring that is generated because of the expansion when the first stage of the jet passes through the nozzle exit. In figure 2 an isosurface of the Q-criterion has been plotted in blue. In this figure a lot of structures that belong to the vortex ring can be identified.

Another relevant finding is that strong pressure perturbations are attached to the rear part of the vortex ring and this is the reason why some slots are formed in this region of the vortex ring. In figure 2 one can see how the center of the pressure perturbations, plotted in this figure as red isosurface, have their center on the shear layer, and more specifically where the shock-waves are present. In a supersonic jet, both the shear layer and the shock-cell structure contribute to the noise generation process. In this project we expect to identify more accurate the process of noise generation by investigating in detail the structures of the shear layer and the acoustic field.

On-going Research

The most important limitation of this project has been the WORK disk space. In order to investigate the statistically steady regime of the jet, more than 30000 time steps have to be taken into consideration. This problem was overcome by doing the statistics during the simulation take place.

In the close future, an extensive post-processing of all the generated data will be done. Special focus on the present structures of the shear layer will be made and the global objective of the post-processing stage is the deep understanding of the involved phenomena of the noise generation process.

References

[1] Sesterhenn, J. 2000. A characteristic-type formulation of the Navier-Stokes equations for high order upwind schemes. Computers and Fluids.

[2] Schulze, J and Sesterhenn, J. 2011. Adjoint based noise minimization of a round supersonic jet. Journal of Physics: Conference Series

Cavitation Erosion in Injection Systems

Research Institution
Institute of Aerodynamics and Fluid Mechanics

Principal Investigator
Christian Egerer

Researchers
Felix Örley, Theresa Trummler, Stefan Hickel, Steffen Schmidt

Project Partners
–

SuperMUC Project ID: pr86ta

Introduction

Developments in direct Diesel injection systems increase rail pressures to more than 2500 bar. This trend aims at enhancing jet break-up and mixing to improve combustion and reduce emissions. Higher flow acceleration, however, implies thermo-hydrodynamic effects, such as cavitation, which occur when the liquid evaporates locally. The collapse of such vapor structures causes strong shock waves. When bubbles collapse near a solid wall, high-velocity liquid jets directed towards material surfaces are created. Imposed structure loads can lead to material erosion, which may be so strong that the performance degrades severely or devices may fail. On the other hand, these loads are used to clean nozzle holes and throttles from surface deposits, and can promote jet break-up. Furthermore, two-phase flows can be used to maintain choked nozzle conditions and a constant mass flow rate.

Understanding the flow phenomena inside an injection system is necessary to quantify the effects of turbulence and cavitation, and their influence on jet and spray characteristics. Small dimensions, high operating pressures and short timescales make the instrumentation of fuel nozzles with experimental equipment challenging. Computational Fluid Dynamics (CFD) can provide time-resolved information on flow structures in arbitrary small geometries. Numerical simulations thus have become an important tool in the design process of injection systems.

The present research project focuses on the prediction of cavitation erosion in fuel injection systems using a CFD approach. This includes wave dynamics, interaction of cavitation and turbulence as well as flow transients due to moving geometries. In our project, we use Large Eddy-simulation (LES) to understand the flow dynamics. Our simulations run on SuperMUC in the massively parallelized numerical framework INCA [1].

Numerical Method

With LES, the smallest turbulent flow scales are not resolved on the computational grid. Effects of these scales thus must be modelled. We employ an implicit LES approach based on the Adaptive Local Deconvolution (ALDM) method [2]. To consider two-phase effects, we apply the homogenous-mixture cavitation model. The actual vapor-liquid interface of cavitation structures is not reconstructed in this barotropic model. Surface tension thus is neglected. Recently, we have extended the single-fluid two-phase model by a component of non-condensable gas. Complex (moving) bodies are considered by a conservative cut-cell method [3].

Results

Interaction of turbulence and cavitation
We have performed wall-resolved LES of the flow through a generic throttle to validate our model in the context of turbulent, cavitating nozzle flows [4]. At pressure difference 300 to 115 bar, we observe periodic formation of vapor in the detached shear layer at the throttle inlet (left column of Fig. 1). When the pressure difference is increased (300 to 55 bar, right column of Fig. 1), a stable vapor sheet develops at the throttle inlet. In the throttle center, cavitation in large, stable vortices is observed.

We find that while turbulence and vortex dynamics play a dominant role at low pressure differences, the formation of a stable sheet cavity and cavitating vortices suppress turbulence at high pressure differences.

Break-up of cavitating liquid jets
Recently, we have studied the break-up of cavitating liquid jets in a free gas phase [5]. The setup resembles a generic, scaled-up automotive fuel injector and consists of a cavitating water jet emanating from a rectangular nozzle injected into air. We investigated several operating conditions that lead to different cavitation characteristics.

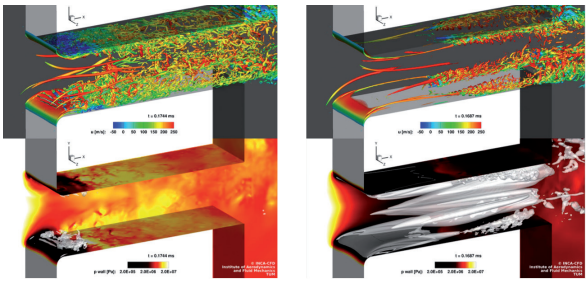

Figure 1: Instantaneous turbulence (top) and cavitation structures (bottom) in a generic throttle valve. Pressure difference 300 to 115 bar (left) and 300 to 55 bar (right).

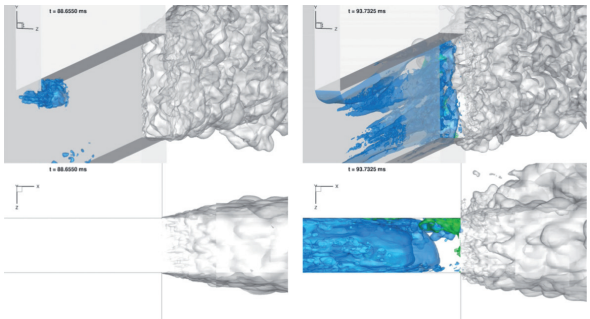

Figure 2: Cavitation structures (blue) and jet surface (grey) for developing cavitation (left) and supercavitation (right): perspective view (top) and top view (bottom).

In a supercavitating state, see right column of Fig. 2, we observe a stronger break-up of the liquid jet than in case of developing cavitation, see left column of Fig. 2.

From an analysis of the transient data we have identified three main mechanisms that lead to distortions of the jet surface and, ultimately, to a widening and break-up of the jet. First, turbulent fluctuations, which are induced by collapse events in the proximity of the exit plane of the nozzle, add to the momentum in wall-normal direction. Second, low pressure vapor regions near the nozzle exit and the gas filled plenum form a pressure gradient, which enables entrainment of gas from the outlet region into the nozzle. When the gas is being ejected again, the water is

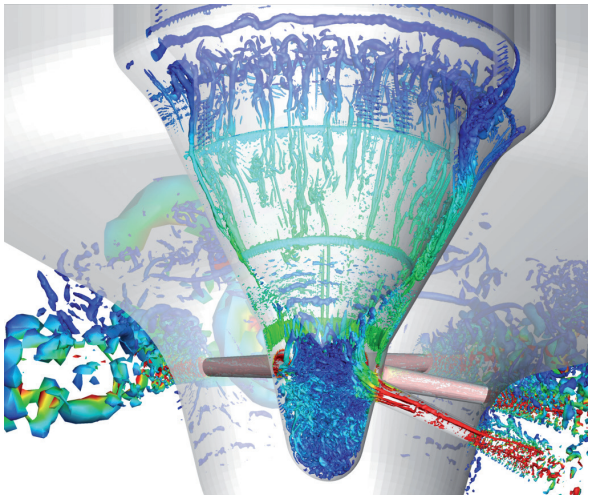

Figure 3: Coherent vortical structures during the main injection phase colored by velocity in the half-domain.

accelerated towards the side walls and creates large scale bulges of liquid. Third, collapse events of cavitation structures inside the jet near the liquid-gas interface induce high velocity liquid jets directed towards the interface.

9-hole Diesel injector with moving needle

To study our models in realistic environments, we investigate the turbulent multiphase flow inside a 9-hole common rail Diesel injector during a full injection cycle of ISO 4113 Diesel fuel at a pressure of 1500 bar into air. Our simulation includes a prescribed needle movement. The nozzle holes have a mean diameter of 150μm.

The analysis of the turbulent flow field reveals that the opening and closing phase are dominated by small-scale turbulence, while in the main injection phase large vortical structures are formed in the volume upstream of the needle seat, and reach into the nozzle holes, see Fig. 3. In each hole, several of these structures are present at the same time. During and after the closing phase, cavitation structures are detected in the nozzle holes and in the sac hole region, see Fig. 4, and cause violent collapse events. Subsequently, the collapse of the sac hole cavity and

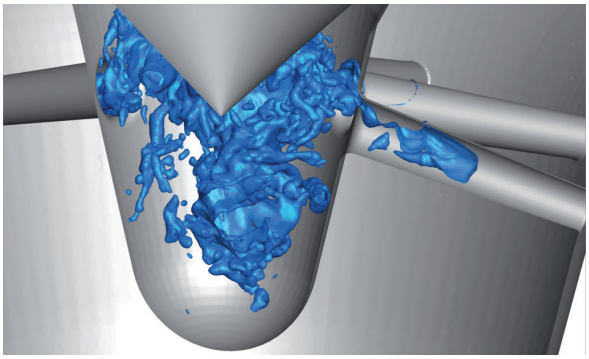

Figure 4: Cavitation structures inside the nozzle holes and sac hole shortly after full closing of the injector needle.

rebound effects cause a large number of strong events near the lowest point of the sac hole. These events during this phase thus are considered to be most likely to cause surface erosion inside the device during operation.

On-going Research and Outlook

Our studies helped us to better understand the dynamics of cavitating, turbulent fluid flows in injection systems. High-performance computing is a necessary tool to address the requirements that the investigations of turbulent, cavitating flows impose on spatial and temporal resolution. Due to the high speed of sound, the time step size usually is on the order of less than a nano-second, while near-wall turbulence requires a high grid resolution and thus causes a large number of cells. Our simulations usually run on up to 5600 cores to compute physical time scales on the order of micro-seconds. Current research topics include the development of improved gas models with degassing and solution of gas in liquid, the analysis of fluid structure interaction in the context of cavitation, the modeling of quantitative erosion prediction in cavitating flow environments, and acceleration methods for CFD codes for the simulation of cavitating flows.

References and Links

[1] http://www.inca-cfd.com, http://www.aer.mw.tum.de

[2] Hickel, S., Egerer, C. P., & Larsson, J. (2014). Subgrid-scale modeling for implicit large eddy simulation of compressible flows and shock-turbulence interaction. Physics of Fluids, 26(10), 106101. doi:10.1063/1.4898641

[3] Örley, F., Pasquariello, V., Hickel, S., & Adams, N. A. (2015). Cut-element based immersed boundary method for moving geometries in compressible liquid flows with cavitation. Journal of Computational Physics, 283(C),1–22. doi:10.1016/j.jcp.2014.11.028

[4] Egerer, C. P., Hickel, S., Schmidt, S. J., Adams, N. A., 2013. Large-eddy simulation of turbulent cavitating flow in a micro channel. Physics of Fluids, 26(8), 085102. doi:10.1063/1.4891325.2

[5] Örley, F., Trummler, T., Hickel, S., Mihatsch, M. S., Schmidt, S. J., Adams, N. A., (2015). Large-eddy simulation of cavitating nozzle flow and primary jet break-up. Physics of Fluids, 27(8), 086101. doi:10.1063/1.4928701

Direct Numerical and Large-Eddy Simulation of Wall-Bounded Flows for Aeroacoustic Source Term Identification

Research Institution
Institute of Process Machinery and Systems Engineering, University Erlangen-Nuremberg

Principal Investigator
Stefan Becker

Researchers
Christoph Scheit, Katrin Nusser

Project Partners
Georg Hager, Gerhard Wellein, RRZE Erlangen

SuperMUC Project ID: pr86xe

Introduction

In order to gain a deeper understanding of the aero-dynamic noise generation mechanisms and transmission for automotive applications, researchers from the University Erlangen-Nürnberg leveraged the HPC system SuperMUC to develop a hybrid aeroacoustic method. The turbulent flow over a forward-facing step served as a test case for the final validation of a hybrid scheme for the computation of broadband noise, as typically caused by turbulent flows. The scheme was applied to the flow around a simplified car model, which was studied experimentally in the research project FORLärm [1].

Results and Methods

In order to make DNS/LES simulations of the flow over a forward-facing step and of the flow around a simplified car model (SAE body) feasible, SuperMUC was used. The typical job sizes were around half an island, although scaling measurements showed good parallel efficiency for more than one island. For the DNS, a pure MPI based, asynchronous communication was employed, while for the SAE body a hybrid MPI/OpenMP approach was used for higher flexibility. Scaling measurements for a setup similar to the one used for the DNS are shown in Fig. 1, comparing different communication strategies.

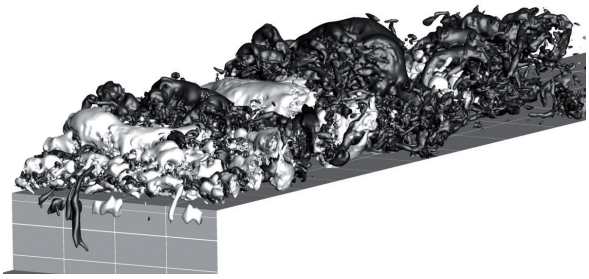

Figure 2: Isosurfaces of turbulent pressure fluctuations p', bright color corresponds to –9 Pa, dark to 9 Pa.

For the simulation of the forward-facing step, 400 million CVs were used, so that an even better efficiency as for the test case could be achieved here.

The software employed, Fastest-3D, is a block-structured finite volume solver with explicit and implicit time stepping algorithms, several linear and non-linear flux discretization schemes and LES turbulence models. The linear equation solver is based on an incomplete LU factorization with a geometric multigrid acceleration. Acoustic

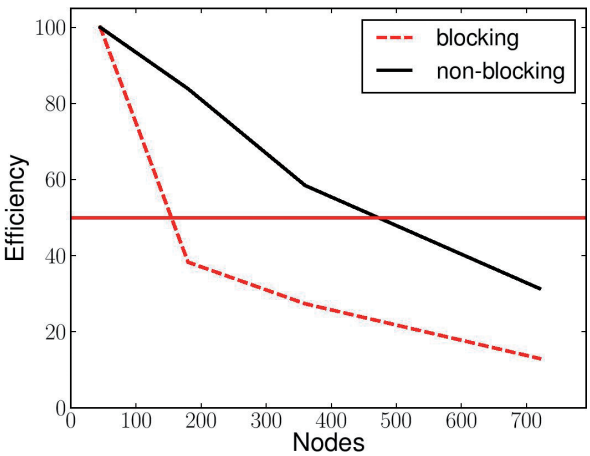

Figure 1: Scaling measurements; on SuperMUC, optimized (non-blocking) and non-optimized (blocking communication calls) code versions.

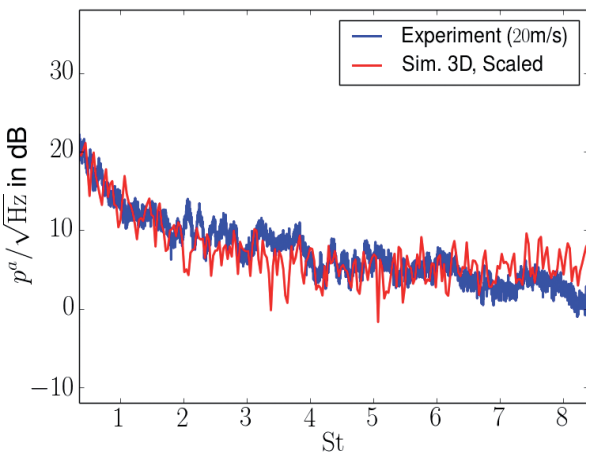

Figure 3: Comparison of measured and calculated acoustic spectrum in 1m distance above the step edge.

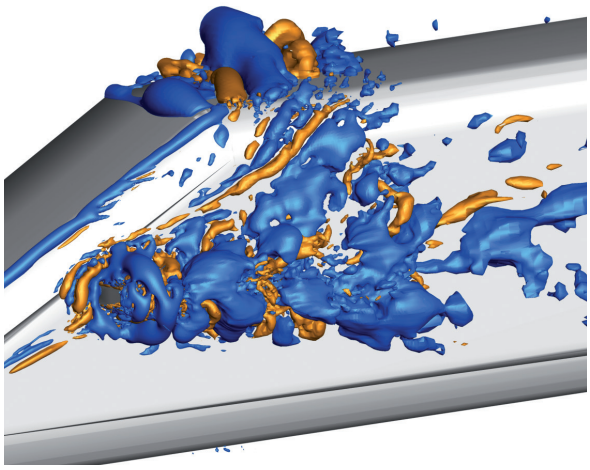

Figure 4: Isosurfaces of the turbulent pressure fluctuations for p' = ± 20 Pa, generated by the flow around the side view mirror.

source terms are based on Lighthill's acoustic analogy, vortex sound, or the acoustic perturbation equations (APE). Acoustic evaluation of data is done in an acoustic postprocessing step using CFS++ [4]. About 10 million CPU hours were spent for the DNS, including a grid independency study. For the LES of the flow around the SAE body around 7 million CPU hours were needed. In order to be able to store the acoustic source term data, a conservative data projection to a coarser grid level was implemented, such that source term data in the range of only 2 – 3 TB had to be stored.

As a main result it was possible to validate the implemented hybrid acoustic method using the DNS data from the step flow. In addition, deeper insight into the flow structures developing from the step was achieved. Fig. 2 shows isosurfaces of the turbulent pressure fluctuations. Small structures are generated in the vicinity of the step and are convected downstream of the step where they merge to larger structures again. Clearly, a highly turbulent flow field is caused by the step flow. Fig. 3 shows the acoustic spectrum obtained by the hybrid simulation approach compared to measurements. As can be seen, predicted and measured acoustic pressure are in good agreement. This also validates the projection of the acoustic source terms to a two-times coarser grid level.

The second part of the project was concerned with the application of the developed hybrid acoustic scheme to the flow around a simplified car model. The turbulent flow around the model was computed using a Large Eddy Simulation. To illustrate the instantaneous flow and the developing structures behind the a-pillar and the generic side view mirror, isosurfaces of the turbulent pressure fluctuations are visualized in Figure 5. A great variety of sizes of the vortex structures can be observed and also very small structures have been resolved by the simulation. On various points on the side window, local pressure values were gathered and compared in frequency and amplitude to experimental data. Figure 4 shows the spectrum of these wall pressure fluctuations at a point in the wake behind the side view mirror. A good agreement between simulation and experimental data can be observed up to a frequency of 1200 Hz.

On-going Research / Outlook

Altogether it was possible to compute highly-resolved turbulent flow structures using the computational power of SuperMUC. Besides the large number of compute nodes available, it turned out that the parallel file system and storage capacities were of major importance for the success of the project.

Research on the SAE body and especially on the interior acoustics is still on-going. As a next step, a coupling with the exterior flow field will be implemented.

Follow-up projects could be used to identify the influence of wall-models in LES and also to scale up simulations to higher Mach and Reynolds numbers. This would be of interest especially for the simulation of the flow around the SAE body.

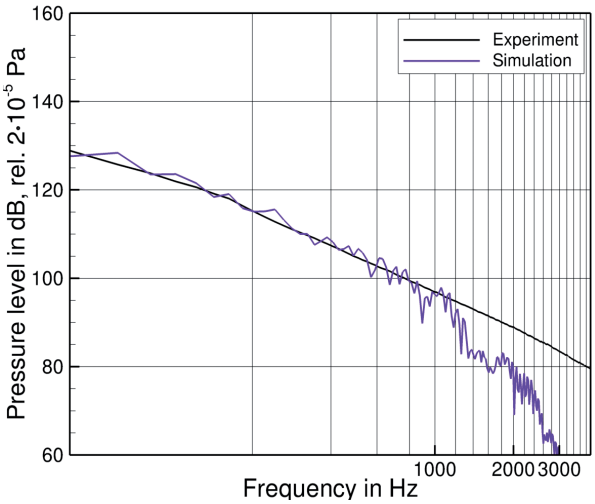

Figure 5: Comparison between measured and simulated frequency spectrum of wall-pressure fluctuations.

References and Links

[1] FORLaerm: www.bayfor.org/de/geschaeftsbereiche / forschungsverbuende/welt-der-materie/forlaerm.html

[2] C. Scheit, G. Hager, J. Treibig, S. Becker and G. Wellein. Optimization of Fastest-3D for modern multicore systems. http://arxiv.org/abs/1303.4538

[3] K. Nusser, S. Müller, C. Scheit, M. Oswald and S. Becker: "Large Eddy Simulation of the Flow around a Simplified Car Model." In: Direct and Large-Eddy Simulation X, 2015, Limassol. To appear in: Direct and Large-Eddy Simulation X, Springer

[4] M.Kaltenbacher, "Numerical Simulation of Mechatronic Sensors and Actuators." Springer, 3rd edition, 2015

Coupling 4 Molecular Dynamics Codes in a Massively Parallel Molecular-Continuum Fluid Dynamics Framework

Research Institution
TU München, Department of Informatics
Principal Investigator
Hans-Joachim Bungartz, Philipp Neumann
Researchers
Nikola Tchipev, Wolfgang Eckhardt, Piet Jarmatz
Project Partners
–

SuperMUC Project ID: pr87cu

Introduction

Many problems from micro- and nanofluidics require the capturing of molecular effects. However, respective simulations are often limited by the computational intensity of molecular dynamics simulations and – as a consequence – by length and time scale restrictions. Molecular-continuum simulations address this issue by decomposing the overall computational problem into molecular dynamics and coarse-grained flow simulations. For example, only small, but important sub-regions of interest are resolved by molecular dynamics whereas the rest of the computational domain can be treated by continuum or mesoscopic flow solvers.

Over the past years, we have developed the macro-micro-coupling tool *MaMiCo* [1] which facilitates the development of molecular-continuum flow simulations in the sense that it encapsulates and thus hides coupling components and coupling algorithms from the coarse-grained and the molecular dynamics solvers. In this project, we have demonstrated the flexibility and scalability of our software with regard to coupling various pieces of simulation software and executing the arising coupled simulation on high-performance compute systems. We successfully coupled four Lattice Boltzmann simulations (*PeanoLB, waLBerla, OpenLB, Palabos*) and four molecular dynamics packages (*SimpleMD, ls1 mardyn, ESPResSo, LAMMPS*) using *MaMiCo* [2].

Results and Methods

Computational Setup
We considered three channel flow experiments named MD-30, MD-60, and MD-120; going from MD-30 to MD-60 to MD-120 doubles the size of the overall computational domain in each dimension, cf. Tab. 1. In each scenario, a spatially adaptive Lattice Boltzmann (LB) simulation covers the whole channel, cf. Fig. 1. In the middle of the refined LB region, molecular refinement is activated and a molecular dynamics (MD) simulation is executed. For MD, we considered short-range, single-centered Lennard-Jones interactions. The cut-off radius for these interactions was chosen as $r_c=2.5\sigma$ where $\sigma=0.34$nm is the molecule diameter (corresponding to argon).

Coupling Method
Both LB and MD solvers are coupled via velocity exchange in a Schwarz-like sense to simulate quasi-steady state flow scenarios. Each coupling cycle consists of two phases. In the LB phase, we run the LB simulation until steady state and impose space-time averaged flow velocities from MD in all LB cells which coincide with parts of the inner MD domain, that is in the whole MD domain except for its boundary strip which has a thickness of two LB cells. In the subsequent MD phase, the steady-state LB velocities are imposed via a relaxation method onto the molecules in the boundary strip of the MD system. The MD system is then equilibrated, averaged velocities are sampled from MD, and the simulation continues with the next coupling cycle.

Table 1: Number of grid cells and molecules for scenarios MD-30, MD-60, MD-120. No. molecules: number of molecules in MD simulation. No. LB cells (coarsest/finest): number of LB grid cells on the coarsest and the finest grid level of refinement. The grid cells on the finest level have a cell size dx:=r_c .

Name	No. molecules	No. LB cells (coarsest)	No. LB cells (finest)
MD-30	16,250	48x24x24	24x24x24
MD-60	130,050	96x48x48	48x48x48
MD-120	1,040,502	192x96x96	96x96x96

Coupling LB: waLBerla
waLBerla [3] is a spatially adaptive, massively parallel Lattice Boltzmann framework developed at the chair of system simulation of the FAU Erlangen-Nümberg. To couple *MaMiCo* and *waLBerla*, we implemented a new sweep which carries out a typical LB time step on all LB grid cells and locally incorporates the averaged MD velocities via an additional forcing term.

Coupling MD
To couple an MD simulation and *MaMiCo*, each MD simulation needs to provide three interface implemen-

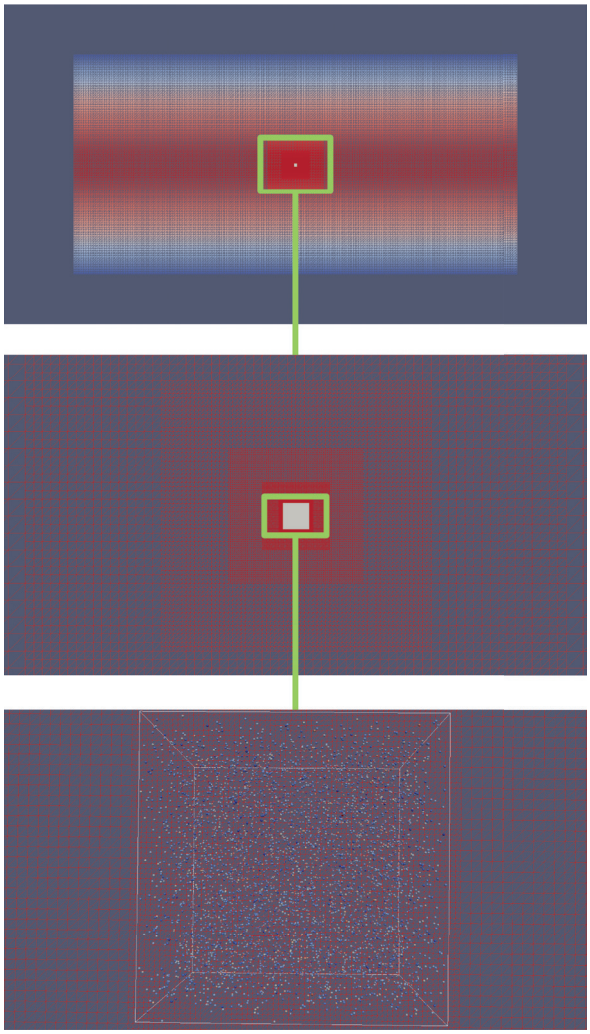

Figure 1: Zoom into channel flow for a modified 6-level MD-120 setup. The small white-framed box corresponds to the MD simulation domain.

tations: a *MoleculeInterface* to access the properties (position, velocity, ...) of a single molecule, a MoleculeIterator to traverse molecules on a cellwise basis, and a MDSolverInterface which provides access to simulation-wide MD parameters such as the MD domain size, or Lennard-Jones parameters.

SimpleMD [1], *ESPResSo* [4], *ls1 mardyn* [5], and *LAMMPS* [6] were successfully coupled to *waLBerla* via *MaMiCo*. Only minor code modifications/extensions on MD side were required for *ESPResSo* and ls1 mardyn. *LAMMPS* required additional sorting routines to map molecules onto the grid cell structure used by *MaMiCo*.

Strong Scaling

We evaluated the parallel performance of all simulation codes in various node-level and strong scaling experiments. Good scaling was achieved on up to 512 compute cores. The most time-consuming part in the coupling is the MD phase. We found that the coupling comes at an overall relative overhead of ca. 15%, compared to a standalone MD simulation. Bigger overheads only occurred in the *waLBerla-ls1* mardyn coupling due to *ls1 mardyn's* highly vectorized kernel routines. This high level in performance can currently not be sustained by *MaMiCo* due to some modularity and encapsulation aspects.

We conducted further studies using a modified MD-120 and MD-240 scenario with up to 6 levels of grid refinement – cf. Fig. 1 – on up to 4096/16384 cores (256/1024 nodes of SuperMUC), cf. Fig. 2 for strong scaling results of the MD-120 variant. For MD-240, the channel diameter corresponds to 6.96µm; our simulation could thus resolve a domain from continuum down to single molecules which is as thick as a strand of spider silk and which is basically visible to the naked eye. A single coupling cycle (LB phase + MD phase) of these large-scale experiments requires up to 8000 compute cores for 5-20 hours.

On-going Research / Outlook

A bottleneck consists in the strong scalability limits of the MD simulation. Although the short-range MD simulations are known to scale well, the molecular-continuum methods cannot fully exploit this feature since we require a huge number of time steps on moderately to rather small-sized MD domains. Amongst others, the huge number of time steps is required to obtain high-quality sampling data. We currently investigate other coupling schemes based on embarrassingly parallel ensemble averaging, i.e. coupling many MD simulations to a single Lattice Boltzmann simulation. This allows to sample over independent MD simulations and thus shifts the computational load from the execution of many time steps to the execution of many simultaneous MD simulations. First implementations based on *LAMMPS* and *MaMiCo* have been successfully established for this purpose.

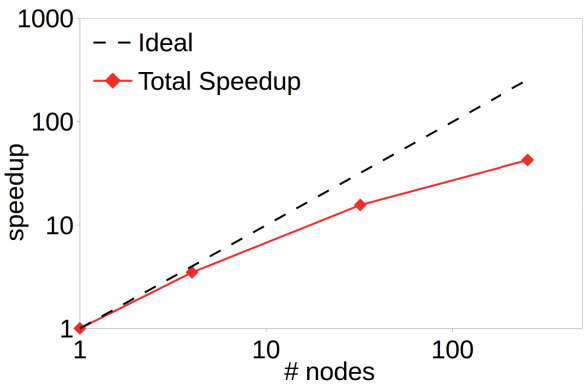

Figure 2: *waLBerla-SimpleMD*, strong scaling (1-256 nodes), modified MD-120 scenario.

References and Links

[1] www5.in.tum.de/mamico (for software download) and www5.in.tum.de/wiki/index.php/Software_Developments
[2] P. Neumann, H. Flohr, R. Arora, P. Jarmatz, N. Tchipev, H.-J. Bungartz. MaMiCo: Software design for parallel molecular-continuum flow simulations. Computer Physics Communications 200: 324-335, 2016
[3] www.walberla.net
[4] espressomd.org
[5] www.ls1-mardyn.de
[6] lammps.sandia.gov

Direct Numerical Simulation of the interaction of a wall-mounted cube with a turbulent boundary layer

Research Institution
Imperial College London

Principal Investigator
Christos Vassilicos

Researchers
Sylvain Laizet, Carlos Diaz Daniel

Project Partners
EU project Multisolve [1]

SuperMUC Project ID: pr87gu (PRACE project)

Introduction

The velocity deficit region of a moving fluid over any solid surface is called a boundary layer and it has a fundamental role in the overall aerodynamic drag and noise generation of an object. The theoretical study of this flow is a research topic of great interest in fluid mechanics and high-fidelity simulations can provide very valuable information to unravel the secret of turbulent boundary layers.

Protuberances located at the wall of an object can modify the acoustic signature of the flow and can create a tonal noise which can be undesirable in some engineering applications. The current project investigates boundary layer physics and the interaction with solid obstacles from a fundamental point of view. A generic protuberance is modelled as a single wall-attached cube, immersed in the turbulent boundary layer with dimensions similar to the boundary layer thickness.

Numerical Methods

The Navier-Stokes equations govern the dynamics of the flow, assumed incompressible in this study. They are solved using the high-order flow solver Incompact3D [2, 3], based on 6th order finite-difference spatial schemes and a spectral treatment for the pressure equation. The viscous terms are discretized with a semi-implicit temporal scheme so that large time-steps can be used for the fractional step method used for the time advancement of the simulations. Incompact3D is parallelized thanks to a powerful 2D domain decomposition strategy based on the open-source 2D DECOMP & FFT library [4]. Incompact3d can scale with up to one million computational cores [5].

The present investigation starts with a Direct Numerical Simulation (DNS) of a fully turbulent zero-pressure gradient boundary layer with Reynolds numbers up to 2000 (based on the momentum thickness and the free-stream velocity). Afterwards, a wall-mounted cube is included in the computational domain using a customized immersed boundary method. The idea is to add an extra forcing term in the Navier-Stokes equations in order to force the velocity to zero at the wall of the cube. The cube of size L is located at a streamwise position corresponding to Re=750. For these simulations, the numerical domain is discretized with 4097x513x256 mesh nodes for a size of 420δx40δx15δ (where δ is the height of the initial boundary layer). The simulations were carried out using 8,192 computational cores. Over 15 TB of data were collected in form of 2D/3D statistics, instantaneous flow field snapshots and time-probes. A visualization of the flow is presented in Figure 1 where the transition from a laminar boundary layer to a turbulent boundary layer can be observed via the vorticity field.

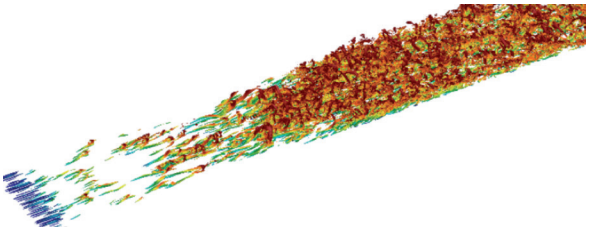

Figure 1: Turbulence structures using vortex visualization contours.

Results

Boundary layer simulations have provided results of theoretical relevance alongside comparisons with previously published data. The baseline simulation without the cube was validated with reference data and the agreement for a wide range of statistics was excellent. One of the main studies focused on the shear stress vector at the wall. The influence of large-scale structures on the fluctuating shear stress vector is responsible for a slow increase of the dissipation at the wall with increasing Reynolds number as seen in Figure 2. An empirical correlation for wall dissipation was proposed based on the current results, in good agreement with previous simulations from other authors.

Other results obtained from our simulations are related to a theoretical model to relate skin friction and velocity

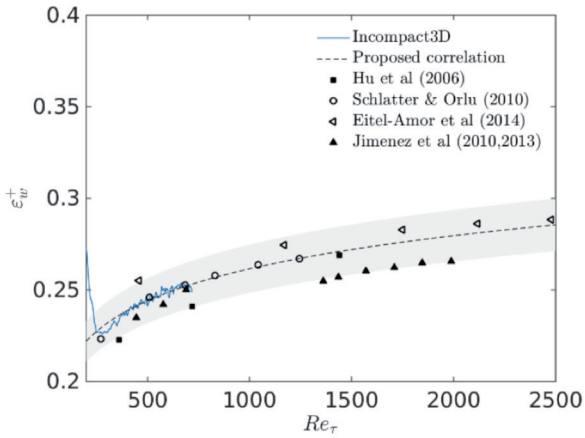

Figure 2: Empirical correlation for wall dissipation in boundary layers.

fluctuations. Motivated by the ideas of Kolmogorov [6] and following a more recent study [7], this theoretical model links statistics of the wall shear stress fluctuations to the structure function of streamwise turbulent velocity fluctuations at a distance from the wall inside the turbulent boundary layer. This relationship relies on and broadens the Townsend-Perry theory of wall-normal turbulence profiles, in which velocity and wall shear stress fluctuations are connected by wall-attached eddies. The Townsend-Perry theory is the only one we have for predictions of the critically important turbulence profiles in a turbulent boundary layer.

From a fundamental point of view, the vortices formed behind a three-dimensional bluff body immersed in a fully turbulent boundary layer have such a complex nature that they pose difficult and interesting problems. First, the finite height of the bluff body (the size of the boundary layer) implies non-negligible end effects, enhancing the three dimensionality of the flow. Secondly, the wall boundary layer may interact with the quasi-periodic flow around the bluff body, resulting in a highly complex flow structure. In the neighbourhood of the base of the three-dimensional bluff body an adverse pressure gradient is produced as a result of the deflection of the flow as seen in Figure 3.

Therefore, parts of the boundary layer are be forced to separate from the wall and vortices are induced, being stretched around in the shape of a horseshoe. Then the flow rolls up into a number of continuously generated horseshoe-shaped vortices, wrapping up the base of the bluff body and trailing downstream subsequently in multiple vortex pairs with their axes parallel to the direction of the main flow. It is expected that those vortices will persist far downstream of the bluff body with a clear and distinct acoustic signature. The peak in the energy spectra in Figure 4 is most probably the signature of the near-periodic vortex street emanating from the wall-mounted cube. This well-defined low-frequency peak corresponding to a Strouhal number of around 0.05 is in agreement with previous experimental studies.

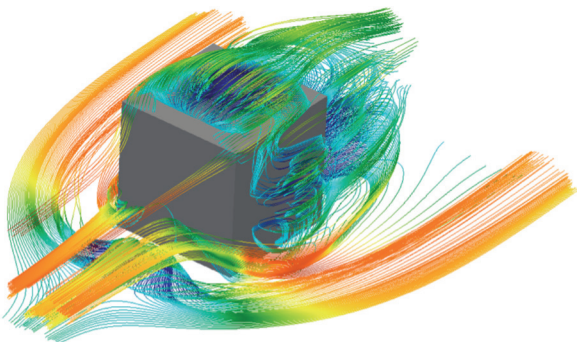

Figure 3: Time-averaged streamlines illustrating the flow structures around the wall-mounted cube.

On-going Research / Outlook

Future work will focus on high-fidelity simulations with different heights for the cube. At the moment, the height of the cube is smaller than the boundary layer thickness and investigations will be carried out for a cube that is up to 10 times higher than the boundary layer thickness. Future works will also involve simulations at realistic Reynolds numbers using a hybrid RANS-LES strategy. Finally, we will investigate strategies to mask, deplete and/or scramble the acoustic signature of the cube.

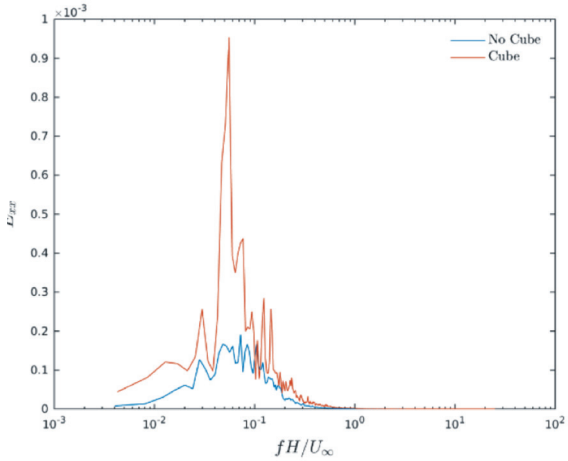

Figure 4: Streamwise velocity spectra for the simulations with and without the cube at a streamwise distance of 45L after the cube (where L is the size of the cube) and at a vertical height of 4.7L.

Our published papers can be found on our website http://www.imperial.ac.uk/tmfc and updates on our research can be followed on Twitter @ImperialTMFC.

References and Links

[1] www.multisolve.eu
[2] www.incompact3d.com
[3] Laizet S. & Lamballais E., High-order compact schemes for incompressible flows: a simple and efficient method with the quasi-spectral accuracy , J. Comp. Phys., 228(15): 5989-6015
[4] http://www.2decomp.org/
[5] Laizet S. & Li N., Incompact3d, a powerful tool to tackle turbulence problems with up to o(10^5) computational cores, Int. J. of Numerical Methods in Fluids, 67(11): 1735-1757
[6] Kolmogorov A.N., A refinement of previous hypotheses concerning the local structure of turbulence in a viscous incompressible fluid at high Reynolds number, Journal of Fluid Mechanics, 13(01):8285, 1962
[7] Vassilicos, J.C., Laval, J., Foucaut, J., and Stanislas, M., The streamwise turbulence intensity in the intermediate layer of turbulent pipe flow, Journal of Fluid Mechanics, 774:324-341, 2015

Fractal flame structure due to the Darrieus-Landau instability

Research Institution
Div. Fluid Mechanics, Dept. of Energy Sciences, Lund University, 22100 Lund, Sweden

Principal Investigator
Rixin Yu

Researchers
X. Bai, V. Bychkov

Project Partners
Dept. of Physics, Umea University, SE 901 87, Umea, Sweden

SuperMUC Project ID: pr87lu (PRACE project)

Introduction

The hydrodynamic Darrieus-Landau (DL) instability [1] of deflagration fame fronts is a fascinating and important hydrodynamic phenomena emerging in combustion, inertial confinement fusion and thermonuclear Supernovae. The DL instability develops because of the density drop (expansion of the burning gas) at a flame front described by the ratio of the unburned fuel/oxidizer mixture to the burnt gas density. Due to the gas expansion, a propagating flame modifies the gas flow, which leads to unstable spontaneous wrinkling of an originally smooth front. Recently, instabilities with similar properties have been also obtained for transformation fronts in advanced materials – organic polymer semiconductors and crystals of nanomagnets.

The nonlinear outcome of the DL instability has been heavily debated. It was first supposed the DL instability makes burning intrinsically turbulent. Later an opposite scenario has been favored that strong linear (thermal) and nonlinear (Huygens) stabilization of DL instability leads to smooth cellular flame structures, with minor increase of flame propagation speed [2], also see Fig 1.a. However experimental data demonstrated self-accelerative growth of expanding flame ball radius at sufficient large length scales [3]. The acceleration had been interpreted as emergence of fractal flame structure with cascade of flame cells of different sizes imposed on another. The fractal excess *d* (i.e. the excess of the fractal dimension over the embedding dimension) was assumed to be a universal constant and was found to be around 1/3 [3]. However later refined experiments report considerable spreading of *d* with different properties in different fuel mixture. Since accurate experimental measurement of *d* are complicated by influence of gravity, confinement, flame-acoustic interaction, etc. Thus the question about the DL fractal flame parameter *d* persists; it was not even clear if this value is a university constant or it depends on flame properties.

Results and Methods

In this work we perform heavy numerical simulations to study the development of DL instability within periodic channel of different widths, the simulations covers realistic density ratio range of 5-10. At very large channel width the fractal cascading structure of the flame front was clearly obtained (shown in Fig 1.c): one large cell of size of the whole channel, 4-6 intermediate size cells imposed on the large one, and about 4-7 minor cells arising on the intermediate cells. Formation of unburned pockets at the DL fractal flame fronts is clearly observed in Fig 1.c. A mean fractal excess *d* was found be around 0.21-0.32 and also depends on density ratio, which is estimated by fitting a power law for the scaled statistically stationary flame propagation speed against the channel width, with *d* as the power exponent. Another fractal excess *d'* was computed using boxing counting of the fractal flame front shapes, it was found *d'* is significant smaller (around 0.07-0.12) than the mean excess *d*, however it also depends on density ratio.

This work is published [4].

Computational issues

For this research, the first author writes a reacting flow C++ solver for the Naiver-Stokes equations [5] on a moving observing coordinate which moves to keep the fractal flame staying inside the computational domain. The solver is implemented for general space dimension (which can handle 1D, 2D and 3D simulations). The solve is parallelized by MPI with the simple domain decomposition strategy, a variable-coefficient Poisson equation for pressure correction is solved with a multigrid method[]. Small time step is required to track the development of DL instability, which takes quite long time to evolve into statistical stationary state, after which more time is need for computing statistics. A single heavy simulation for very large channel width case takes around 1728 CPUs for 8 days. Before simulations, the first author implemented run-time data-processing function to extract key results, this greatly eases the demand on data storage, only very

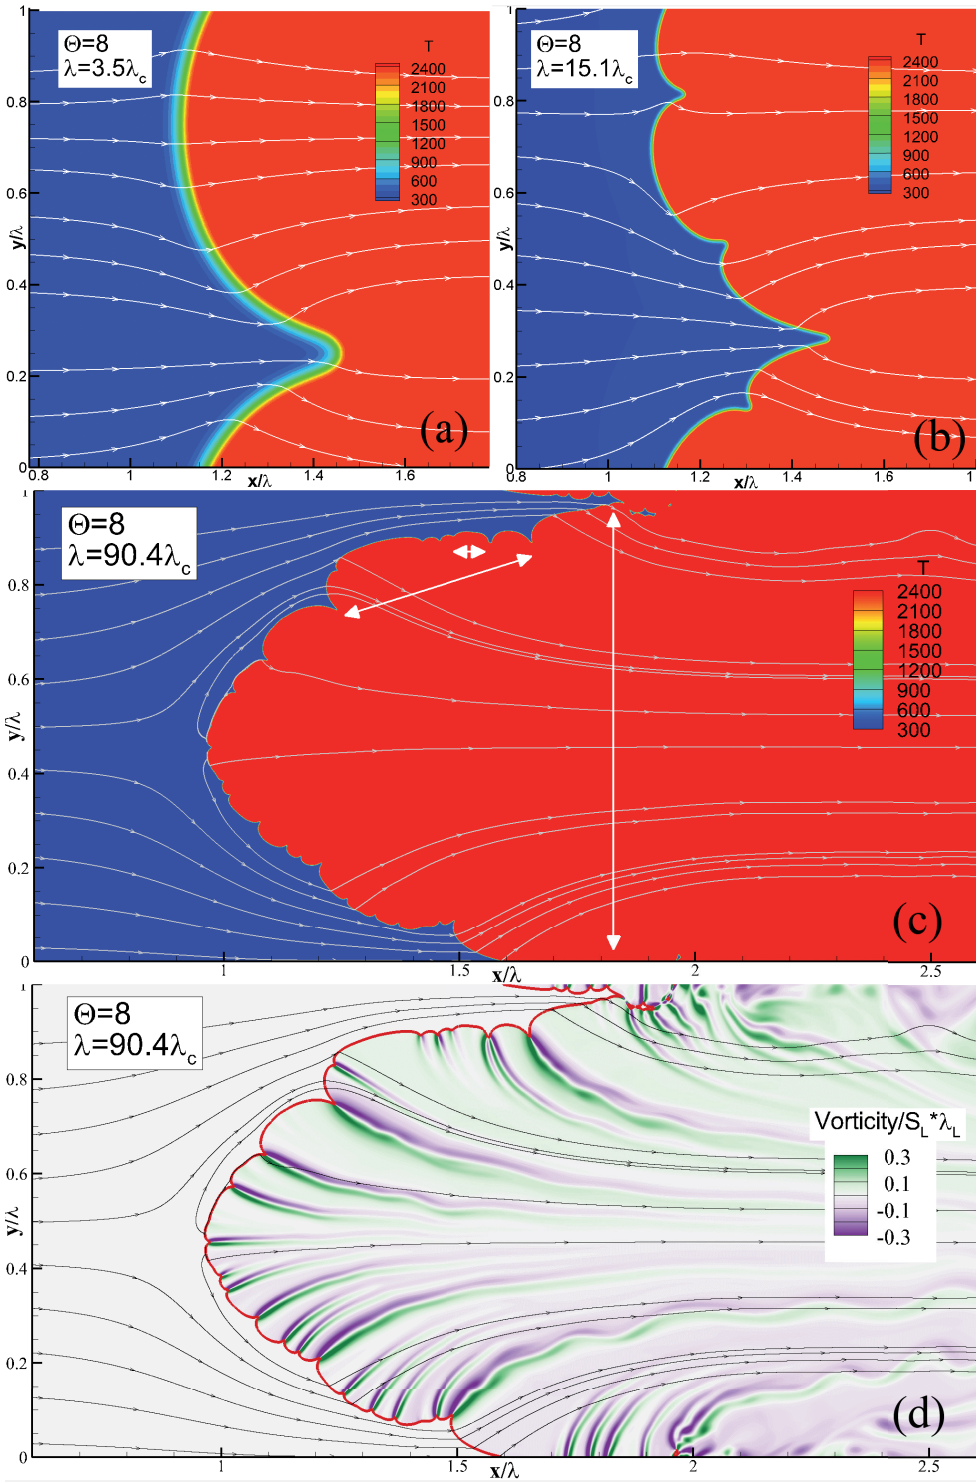

Figure 1: Characteristic snapshots of temperature distribution for a flame front with density ratio 8 in a periodic channel of small width $\lambda=3.5\lambda_c$ (a); medium width $\lambda=15.1\lambda_c$, (b); large width $\lambda=90.4\lambda_c$ (c), the white arrows indicate characteristic steps in the fractal cascade; vorticity distribution for the large width $\lambda=90.4\lambda_c$, (d). In above, λ refers to the channel width, λ_c is the cutoff length below which thermal conduction suppresses the DL instability.

little is left for off-line data post analysis. To avoid outputting too many files flooding the file system, the MPI-IO functions and a newly implemented general subroutine for gathering data scattered across different processors are all used for outputting action such as storing data field, plotting, finally one simulations cases only generated less than 30 files.

References and Links

[1] L.D. Landau, E.M. Lifshitz, Fluid Mechanics, ergamon Press, Oxford, 1989.

[2] V. Bychkov, M. Liberman, Phys. Rep. 325, 115 (2000).

[3] Y. Gostintsev, A. Istratov, Y. Shulenin, Combust. Expl. Shock Waves 24, 63 (1988).

[4] R. Yu. XS Bai, V. Bychkov. Fractal flame structure due to the hydrodynamic Darrieus-Landau instability. Physical Review E 92.6 (2015): 063028.

[5] R. Yu, J. Yu, XS Bai, J. Comput.Phys., 231,5504-5521,(2012)

Relativistic effects on the Richtmyer-Meshkov instability

Research Institution
Department of Civil, Environmental and Mechanical Engineering, University of Trento, Italy

Principal Investigator
Michael Dumbser

Researchers
Olindo Zanotti, Michael Dumbser

Project Partners
–

SuperMUC Project ID: pr87ni, pr94zu (both PRACE projects)

Introduction

The Richtmyer-Meshkov (RM) instability develops when a shock wave impacts through a contact discontinuity within a fluid. This instability has been studied predominantly under terrestrial physical conditions, and especially in experiments of inertial confinement fusion. However, the RM instability is also present in astrophysical systems, such as in the formation of supernova remnants, where it is supposed to be responsible of the mixing of the gas. When the velocity of the gas becomes very large, i.e. when it approaches the speed of light, relativistic effects must be taken into account, but it is not at all obvious how they may affect the development of the RM instability. The main objective of this project has been precisely to investigate the extent to which relativistic effects may alter the dynamics of the RM instability [1].

Results and Methods

Our innovative numerical scheme [2, 3] is based on the following ingredients

1. A WENO finite volume scheme;
2. A local space-time Galerkin predictor allowing for the construction of a single-step time-update ADER scheme;
3. Adaptive Mesh Refinement (AMR) implemented according to an element-by-element refinement criterion.

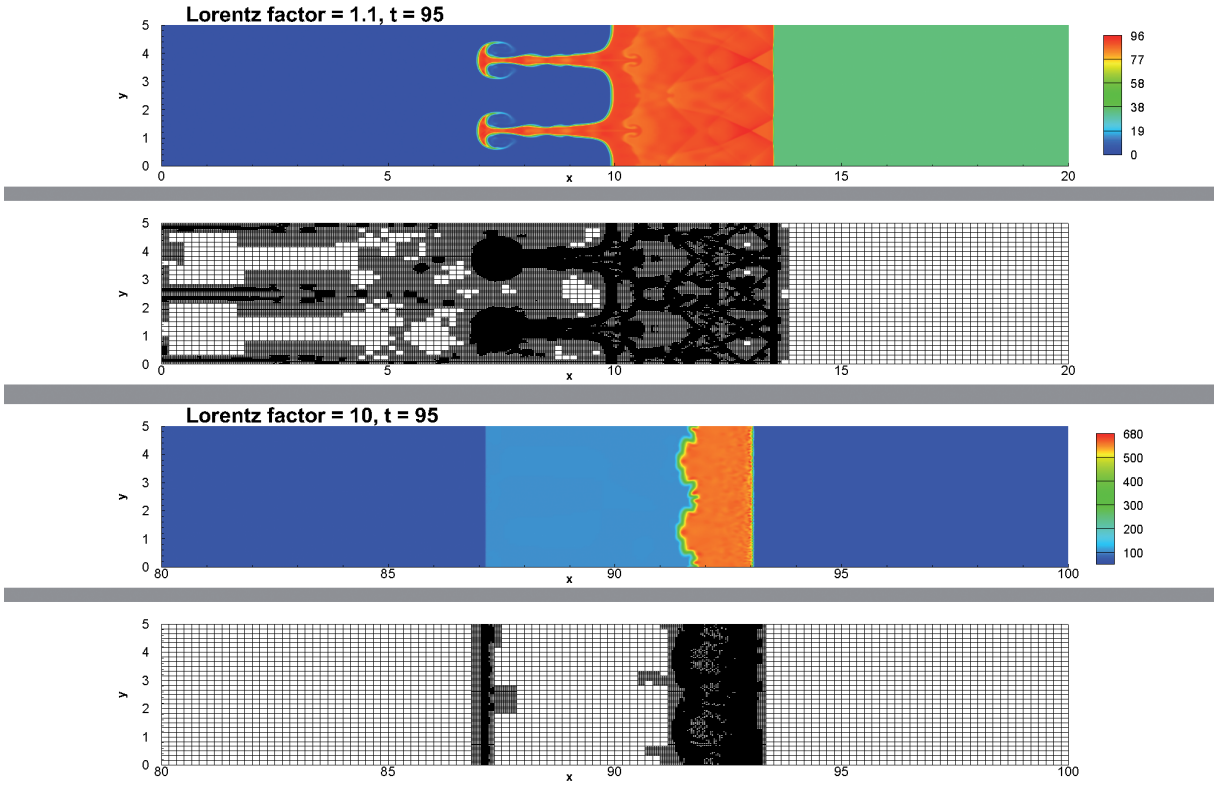

Figure 1: Suppression effect of the RM instability when the Lorentz factor is increased

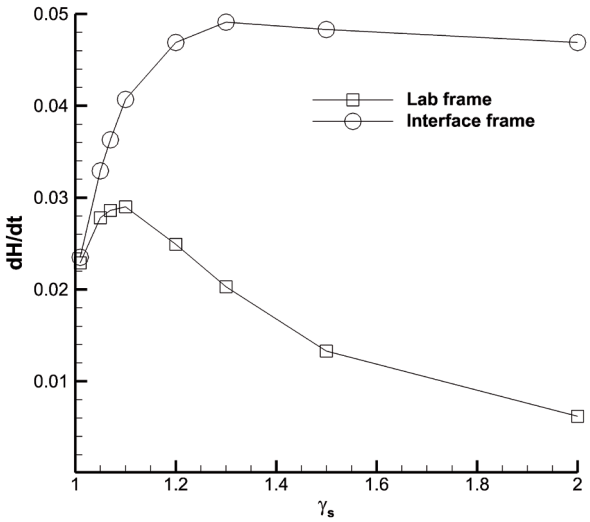

Figure 2: Linear growth rate of the instability as the Lorentz factor of the incident shock wave is increased

We have performed two and three dimensional third order (in space and time) numerical simulations of the relativistic RM instability spanning a wide parameter space and with an unprecedented accuracy. We have considered both the case in which a light fluid penetrates into a higher density one (Atwood number A>0), and the case in which a heavy fluid penetrates into a lower density one (Atwood number A<0). For large Lorentz factors of the incident shock wave, the relativistic RM instability is substantially weakened and ultimately suppressed. The growth rate of the RM instability in the linear phase has a local maximum which occurs at a critical value of the Lorentz factor ~ [1.2, 2]. We have also revealed a genuinely relativistic effect, absent in Newtonian hydrodynamics, which arises in three dimensional configurations with a non-zero velocity component tangent to the incident shock front. Namely, in A>0 models the tangential velocity has a net magnification effect, while in A<0 models the tangential velocity has a net suppression effect.

A limited sample of these results, fully available in reference [1] is reported in Figure 1, Figure 2 and Figure 3. An approximate number of ~50 models have been evolved. Typically, 1024 cores per job have been used for the 2D models, and 2048 cores for the 3D models. Each model generates ~200 files on average.

Our code has been implemented in Fortran 2003 and parallelized through standard MPI. At the heart of our DG scheme there is a local predictor step, which allows to perform the time update through a single time-step. This is particularly advantageous in terms of MPI communications, and it is also very rewarding when AMR is also activated. All these features have been fully exploited on SuperMuc. In three space dimensions, and when AMR is strongly activated, the memory usage can become an issue, and we have found convenient to resort to the fat-node. The CPU distribution of the mesh is managed through the free library Metis and Parmetis [4], which is available as a pre-compiled module on SuperMuc.

On-going Research / Outlook

SuperMUC computational resources have been fundamental for the development of this innovative project, providing an ideal infrastructure both in terms of computational efficiency and data storage.

For the near future, we plan to fully exploit our ADER-AMR numerical scheme to study high energy astrophysical systems, with a particular focus on the solution of the Einstein equations, for which the Phase 2 configuration would be extremely beneficial.

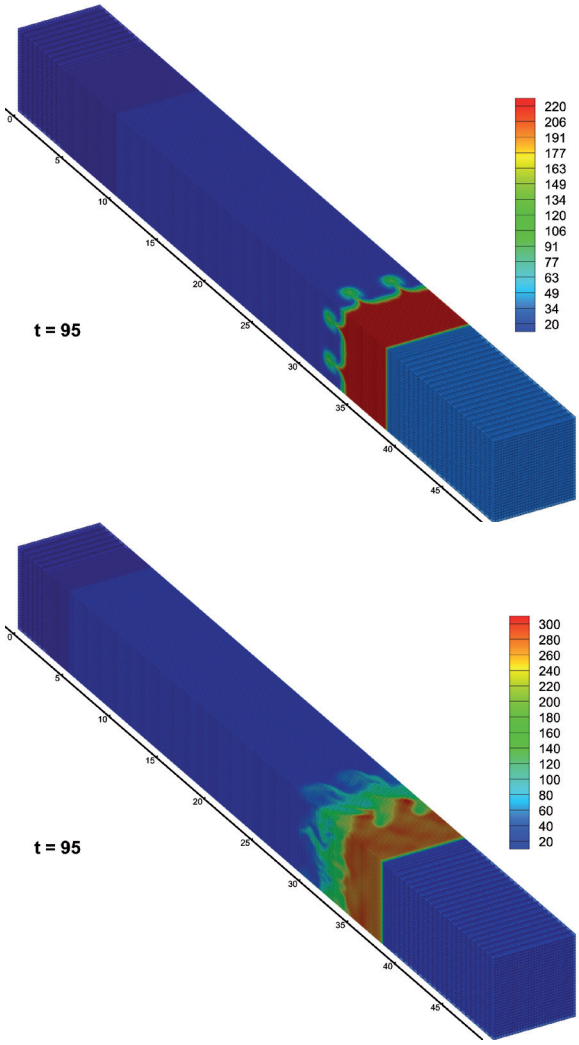

Figure 3: Top (Bottom) RM instability without (with) velocity component tangential to the shock front

References and Links

[1] Olindo Zanotti and Michael Dumbser. High order numerical simulations of the Richtmyer-Meshkov instability in a relativistic fluid, Phys. of Fluids 27, id.074105 (2015)

[2] Michael Dumbser, Olindo Zanotti, Arturo Hidalgo, Dinshaw Balsara. ADER-WENO finite volume schemes with space-time adaptive mesh refinement, JCP 248, 257-286 (2013)

[3] Olindo Zanotti and Michael Dumbser. A high order special relativistic hydrodynamic and magnetohydrodynamic code with space-time adaptive mesh refinement, Comp. Phys. Comm., 188, 110-127 (2015)

[4] http://glaros.dtc.umn.edu/gkhome/metis/parmetis/overview

Direct Numerical Simulation of Open-Channel Flow at Fully-Rough Regime

Research Institution
Institute for Hydromechanics (IfH), Karlsruhe Institute of Technolgy

Principal Investigator
Markus Uhlmann

Researchers
Marco Mazzuoli

Project Partners
–

SuperMUC Project ID: pr87yo

Introduction

The reliability in quantifying the transport of sediments in fluvial and estuarine environments is limited by our capability to predict the flow resistance at the river bed. Even in the simplified case in which regular monosize roughness-elements are considered, it is still unclear how vortex structures form and interact with the roughness and how far from the bottom the fluid-roughness interaction affects the flow structure. In particular, in the case of shallow open-channel flows at transitional and fully-rough regimes, a detailed description of the flow in the vicinity of individual roughness-elements is necessary to understand the overall flow-resistance effect. Although a wide set of laboratory experiments were carried out in channels and open-channels at moderate and high Reynolds numbers (e.g. the recent contribution of Amir et al. 2014), it is hard, if not impossible, with experimental facilities to obtain accurate and high-resolved measurements of velocity and pressure in the vicinity of the bottom, in particular in the interstitial spaces between the roughness elements. On the other hand, so far the extremely high computational cost has limited the use of direct numerical simulations of the phenomenon to cases in which the transitional regime was attained. The results obtained with the present project contribute to reduce the distance between laboratory experiments and the capability to reproduce them numerically. Indeed, with regard to the previous work of Chan-Braun et al. (2011), who investigated the transitional open-channel flow over an array of wall-mounted spheres in square arrangement, in one of the present runs, referred to as S1, the bulk Reynolds number was increased up to 7000 and the fully-rough regime was attained. Consequently, also the Reynolds number based on the friction velocity increased while the ratio of the open-channel height to the diameter of the spheres (roughness elements) was kept constant equal to 5.5. Moreover, the effect of the bottom geometry was investigated by means of a second simulation (run S2) which differed from the previous one only in the arrangement of the spheres on the plane bottom. In the latter simulation the spheres were arranged randomly taking care of maintaining the solid area fraction, namely the ratio of the area occupied by the spheres to the bottom area as a function of the distance from the wall, similar to that of run S1 where spheres were arranged in a square pattern.

Results and Methods

The two runs S1 and S2 were performed on both SuperMUC Phase 1 and Phase 2. The Navier-Stokes and continuity equations were approximated by a second-order finite-difference scheme and solved numerically with a fractional-step method. The velocity was forced to vanish at the surface of the roughness elements by means of the direct-forcing immersed-boundary method proposed by Uhlmann (2005). A uniform and equispaced computational grid was used with grid-spacing of about one wall unit. Table 1 shows the details of the domain size in the streamwise (x), wall-normal (y) and spanwise (z) directions, where H denotes the open-channel height and N the number of grid points in each direction.

Table 1: Grid size and parameters of present simulations.

runs	(L_x, L_y, L_z)	(N_x, N_y, N_z)	k^+	Re_τ	Re_b
S1	(12, 1, 3)H	(6912, 576, 1728)	120	560	7000
S2	(12, 1, 3)H	(6912, 576, 1728)	137	630	7000

For the run S1, 1024 spheres were mounted on a smooth wall in square arrangement while for the run S2, the spheres were initially "shaken" and left free to roll on the wall until their relative position was nearly random (see Figure 1a,b). Then, they were crystallized on the wall and simulation S2 was performed.

From the comparison of run S1 with the simulation F50 performed by Chan-Braun et al. (2011) in the transitional regime, we found that (i) the logarithmic profile scaled by the friction velocity is shifted of about -6.5 with respect to the case where the roughness was absent and of -1.5 with respect to the case F50; (ii) a significantly larger recirculation region forms downstream of the roughness elements; (iii) in the fully-rough regime strong turbulent fluctuations are present deep into the riblets-like

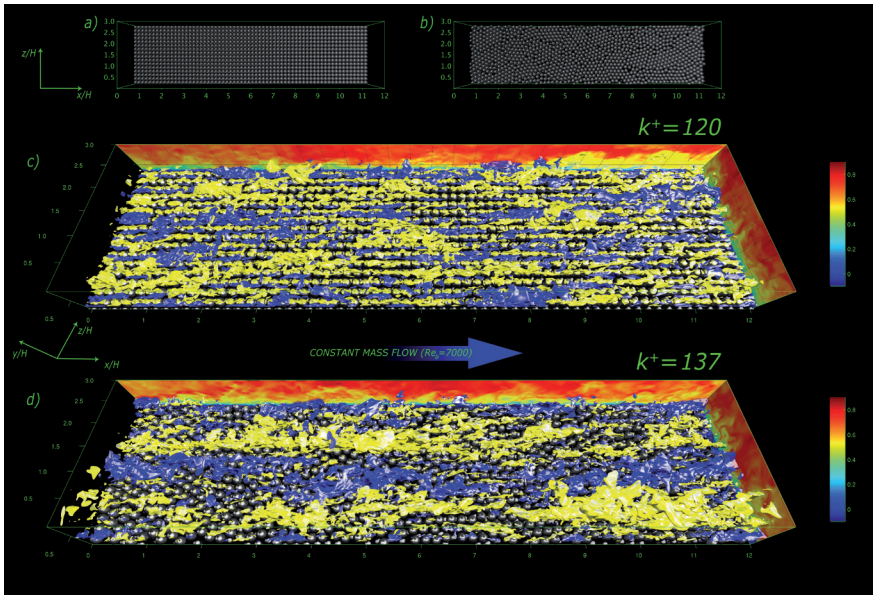

Figure 1: Top view of the bottom geometry of runs a) S1 and b) S2. Panels c) and d) show two instant snapshots of runs S1 and S2, respectively. Yellow and blue isocontours indicate low- and high-speed streaks while lateral panels are coloured by the streamwise component of the velocity.

grooves between the spheres promoting the formation of diameter-periodic low/high-speed regions (see Figure 1c,d); (iv) lift force acting on the individual roughness elements increases.

Simulation S2, which is still running on SuperMUC, showed a surprising increase of the form drag with respect to the run S1 (~30%), apparently because the surface of the spheres is more exposed to the flow. An interpretation of this purely geometrical issue arose from the comparison of the stress distribution on the roughness elements in the simulations S1 and S2 (Mazzuoli & Uhlmann, in prep.).

Nearly 9.7M CPU hours were allocated on SuperMUC for the present project and were equally assigned to runs S1 and S2. Each run was initiated on a coarser grid (with spacing of about two wall units) until turbulence was well developed and a statistically steady flow was attained. Preliminary runs used 729 cores while final runs (on the fine grid) were preformed over 6561 cores of SuperMUC Phase 1 (411 nodes) and Phase 2 (235 nodes). Each time step of the final runs was computed in ~24 s, even though the best performance (15 s per time-step) was observed on Phase 2 by using 6561 cores over 411 nodes.

Output files were generated which consisted in particle-related files, runtime turbulence-statistics computations and instantaneous snapshots of the flow field. Each snapshot consisted of a number of binary files equal to the number of cores and required ~230GB of space. About 60TB of data were generated on WORK and archived step-by-step by the tape-archiving system. The post-processing was performed on SuperMUC with Octave using home-made MPI-like-parallel scripts while graphical object were subsequently created on the facilities of the Steinbuch Center of Computation (SCC, Karlsruhe).

Finally, a time-resolved series of flow-field snapshots was generated during the run S2, which should allow us to describe the evolution of vortex structures shed by the roughness elements. The entire series, consisting of more than 400 snapshots (about 100TB), was generated directly on the SuperMUC SCRATCH file-system and transferred for the post-processing to the HPSS storage system at the SCC by means of gtransfer (rate ~1GB/s).

On-going Research

The present project showed that an overlap in parameter space between laboratory and numerical experiments is possible and new challenges are now open. By maintaining the domain size and roughness configuration of the run S2, further simulations should be performed to explore the parameter space at different flow rates. Nonetheless, a significant extension of the computational domain in the streamwise and spanwise directions may allow us also to describe the dynamics of turbulence super-structures in the fully-rough regime. Since purely MPI parallelization was adopted, the advantage of having 28 cores per node in SuperMUC Phase 2 was not significantly observed in our tests. The possibility to hybridize the code with an OpenMP parallelization to enhance the code efficiency may be considered if the domain were extended.

We acknowledge that the present work was funded by the DFG project UH 242/4-2.

References and Links

[1] Mohammad Amir, Vladimir I. Nikora, and Mark T. Stewart.. 2014. J FLUID MECH 757, (Oct 2014), 458-497. DOI: http://dx.doi.org/10.1017/jfm.2014.498

[2] Clemens Chan-Braun, Manuel García-Villalba, and Markus Uhlmann. 2011. Force and torque acting on particles in a transitionally rough open-channel flow. J FLUID MECH 684 (Oct 2011) 441-474. DOI: http://dx.doi.org/10.1017/jfm.2011.311

[3] Markus Uhlmann. 2005. An immersed boundary method with direct forcing for the simulation of particulate flows.J COMPUT PHYS 209, 2 (Nov 2005) 448-476. doi:10.1016/j.jcp.2005.03.017

Towards Large-Eddy/Filtered-Density Function Simulations of Turbulent Sooting Flames

Research Institution
Deutsches Zentrum für Luft- und Raumfahrt e.V. (DLR)

Principal Investigator
P. Gerlinger

Researchers
C. Eberle, A. Fiolitakis

Project Partners
–

SuperMUC Project ID: pr87zi

Introduction

Combustion is one of the oldest heat and power generation technologies and continues to play an important role in covering the energy demand of the world. The intense use of combustion, however, has led to several environmental problems. One of them are soot emissions. Soot and soot precursors are suspected to be carcinogenic. Furthermore, soot particle emmisions from aircraft engines influence the formation of cirrus clouds at high altitudes and have thus an impact on the climate. From a technical point of view, soot indicates incomplete and hence less efficient combustion. By its high radiative emissivity, soot contributes to locally elevated heat loads on combustion chamber walls. Therefore, continuous efforts are made to improve combustion systems and to reduce their soot emissions.

Due to the increasing availability of computational resources, CFD (Computational Fluid Dynamics) has become an important tool in the design process of combustion systems. CFD provides detailed, time-resolved information about the three dimensional, reactive flow field and thereby complements experimental investigations which are in many cases limited to exhaust gas analysis, since optical access to the flame is not realizable. Soot predictions in technical combustion are particularly challenging. Firstly, the chemical and physical processes of soot evolution are highly complex and involve countless of intermediate steps. Also some aspects of soot evolution are not yet fully understood and are a topic of ongoing research. Secondly, technical combustion occurs predominantly in turbulent flows, which are characterized by a highly unsteady, fluctuating velocity distribution. Due to the high disparity of length scales (large geometries versus the smallest turbulent structures) the direct numerical simulation of such flows is computationally very expensive and therefore limited to a small range of problems. It is therefore common practice to derive transport equations for statistical quantities (termed "Reynolds Averaged Navier Stokes", RANS) or to remove the smallest turbulent scales by spatial filtering operations ("Large Eddy Simulation", LES). In either case the governing equations contain unclosed correlations which require modelling. Especially challenging in this context are the highly nonlinear, unclosed terms which originate from averaging or filtering the chemical source term (turbulence-chemistry-interaction).

Due the high computational cost, it is at present not possible to use highly elaborate modelling approaches for each of the cited problems at the same time (combustion, soot evolution, turbulence, and turbulence-chemistry-interaction). Thus, two modelling approaches are followed in this project, where a simple approach is deliberately chosen for one of the cited fields in order to develop highly elaborate models at realistic simulation times. One approach uses a computationally efficient RANS turbulence modelling approach, in combination with a Transported Probability Density Function (TPDF) method, which was developed in previous work [1] to describe turbulence-chemistry-interaction. A big advantage of TPDF is that the chemical source term appears in closed form. In the second approach, the focus is on a detailed and time-resolved description of turbulent structures by LES. A computationally efficient model is used for sub-grid scale turbulence chemistry interaction, namely a **F**inite-**R**ate-**C**hemistry (**FRC**) model with an **A**ssumed **PDF** (**APDF**) closure [2], where (in addition to the species transport equations) only two more equations for second order moments are required.

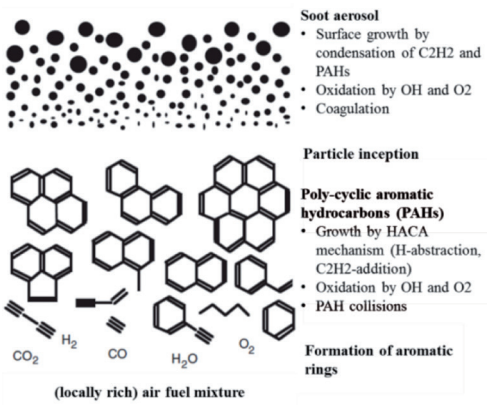

Figure 1: Phenomenology of soot formation, based on [4].

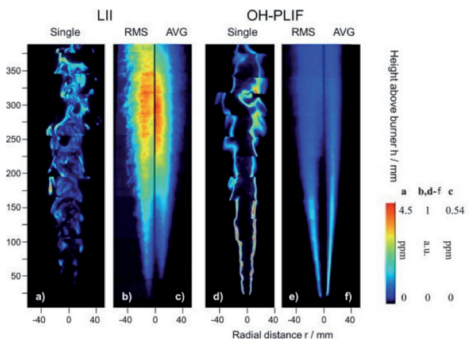

Figure 2: Experimental data for OH and soot [5].

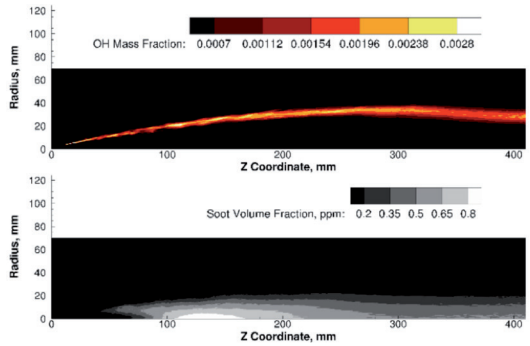

Figure 3: RANS/TPDF results for OH and soot.

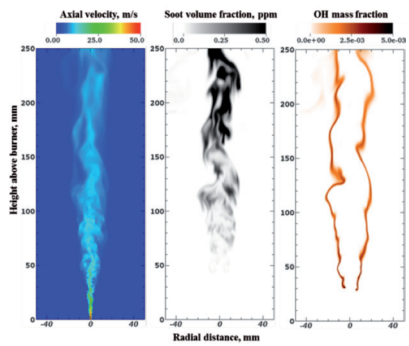

Figure 4: Instantaneous LES results for axial velocity, soot volume fraction, and OH mass fraction.

Both, the RANS/TPDF and the LES/APDF approach use the same models for gas-phase-chemistry, PAHs, and soot. Gas-phase-chemistry is modeled by a detailed reaction mechanism, which describes the formation of small aromatics. PAHs and soot are treated by sectional approaches, where the particle size distribution is discretized by sections with averaged chemical and physical properties. Commonly accepted PAH and soot surface chemistries are considered as summarized in Fig. 1 (for details refer to [3]). The chemical model consists of 72 species and 803 reactions. By developing LES and TPDF methods separately, this work represents a first step towards an improved LES/Filtered-Density-Function for the simulation of soot formation.

Results and Methods

For the simulations in this project the DLR in-house code THETA (Turbulent Heat-Relase Extension of the Tau Code) is used. THETA is an unstructured finite-volume solver which has been optimized for low Mach number combustion problems and features efficient, matrix-free linear solvers. Parallelization is achieved via domain decomposition. State of the art turbulence models and special numerical strategies for the solution of a stiff system of coupled partial differential equations are implemented. To solve the transport equation of the joint thermochemical PDF for the TPDF model, a Monte-Carlo method is available. Stochastic particles are simulated, whose evolution mimics the behavior of the PDF of the thermochemical system. The Monte Carlo solver features a hybrid MPI/OpenMP parallelization paradigm and is two way coupled to the finite-volume-solver in THETA.

The flame investigated in this project is a turbulent, lifted, sooting, ethylene-air jet flame [5]. Experimental data for temperature, soot volume fraction, and velocity are available. Examples are provided in Fig. 2.

First computational results of the combined RANS/TPDF method are shown in Fig. 3. Averaged OH mass fraction and soot volume fraction are illustrated here. The "noise" which is visible in these images is due to the stochastic error in the Monte Carlo solver for the TPDF method. Comparing first the OH mass fraction to the averages of OH-PLIF in Fig. 2 a qualitative good agreement is found. In comparison to that the qualitative agreement in average soot volume fraction shows some discrepancies. The soot formation starts too early in the simulation and the peak value is reached further upstream compared to the experiment. However, the maximum value of about 0.5 ppm agrees well to the experiment. This is an excellent first result. The early onset of soot formation might be due to short comings in the turbulence modelling ("round jet anomaly") and requires still further investigation.

Figure 4 shows representative, instantaneous LES results. The axial velocity distribution reveals a high disparity of turbulent time scales which results from the interaction of the high-velocity fuel jet with the co-flowing air. This is particularly challenging for time resolved simulations, since small time steps are required for reasons of numerical stability and long simulation times are required for the acquisition of converged statistics. The shape of the wrinkled flame front which is indicated by the OH mass fraction compares well to measurements. In agreement to experimental data shown in Fig. 2, soot formation starts in fuel rich regions at about 50 mm above the burner. The soot structures are enveloped by the wrinkled flame front and overlapping between soot and OH is minor due to the high oxidative potential of OH.

Outlook

The simulations will be continued to obtain converged statistical quantities for model validation. The numerical results will be analyzed and the capability of the different modelling approaches (RANS/TPDF and LES/APDF) for soot predictions in turbulent will be assessed. The conclusions drawn within this project will help to continue the work towards a filtered density function approach (LES/TPDF) for soot prediction.

References

[1] A. Fiolitakis, P. Ess, P. Gerlinger, M. Aigner. Combustion and Flame (2014) 161:2107-2119
[2] P. Gerlinger. Combustion Science and Technology (2003) 175:841-872
[3] T. Blacha, M. Di Domenico, P. Gerlinger, M. Aigner. Combustion and Flame (2012) 159:181-193
[4] H. Bockhorn: Soot Formation in Combustion: Mechanisms and Models. Springer Series in Chemical Physics, Volume 59, 1994
[5] M. Köhler, K.P. Geigle, T. Blacha, P. Gerlinger, W. Meier. Combustion and Flame (2012) 159:2620-2635

Shock-wave/turbulent boundary-layer interaction over a flexible panel

Research Institution
Institute of Aerodynamics and Fluid Mechanics

Principal Investigator
Vito Pasquariello

Researchers
Vito Pasquariello, Stefan Hickel, Nikolaus A. Adams

Project Partners
German Aerospace Center, Institute of Aerodynamics and Flow Technology

SuperMUC Project ID: pr94we

Introduction

In the context of launch vehicles, shock-wave/boundary-layer interactions (SWBLI) are common flow features that may generate high-magnitude transient side loads. During the start-up of liquid propellant fueled rocket engines, the rocket nozzle operates in an overexpanded condition, which consequently leads to unsteady internal flow separation. The interaction can critically affect the rocket nozzle performance in case of shock-induced boundary-layer separation and it is a main source of maximum mean and fluctuating pressure loads that the underlying structure is exposed to. These high-magnitude transient loads can be severe enough to fail interfacing components as well as the complete nozzle in the rocket engine. With future rocket technologies focusing on optimal weight systems, fluid-structure interactions (FSI) become significant and must be taken into account in the design process in order to ensure structural integrity. Multi-disciplinary numerical tools are necessary for a correct prediction of the complex flow physics influenced by structural deformations. For our studies we developed a Finite Volume – Finite Element coupling approach for the solution of compressible FSI problems based on a staggered Dirichlet-Neumann partitioning. In this study we will perform high-fidelity large-eddy simulations (LES) of SWBLI in the context of overexpanded rocket nozzles for both rigid and elastic wall-boundaries at high Reynolds numbers.

Numerical Method

The governing equations for the fluid domain are the compressible Navier-Stokes equations, which are solved with our in-house Finite Volume LES code INCA [1]. Within the LES framework, the smallest turbulent flow scales are not resolved on the computational grid, but must be modeled. The Adaptive Local Deconvolution Method [2] is used which implicitly provides subgrid-scale effects. INCA operates on Cartesian grids, allowing for an efficient blocking-strategy and thus high parallel performance. The structural field is solved with the Finite Element method. We recently developed a coupling framework based on a conservative cut-cell method [3]. The communication between both solvers is performed through the Message Passing Interface within a Multiple Program Multiple Data environment.

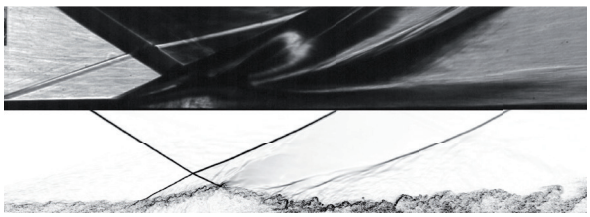

Figure 1: Schlieren comparison between experiment (top) and LES (bottom) for the baseline SWBLI.

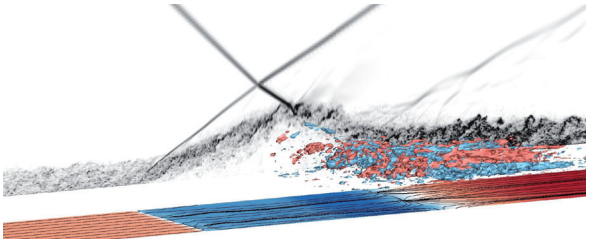

Figure 2: Three-dimensional flow visualization of the baseline SWBLI investigated in this study.

Results

Wall-resolved LES and the high Reynolds number for this SWBLI study require the use of a large amount of computational cells. A total number of 364×10^6 cells has been used. Exploiting the very good scaling properties of our flow solver INCA, we run the simulations on 13860 cores on SuperMUC (Phase 1). SWBLI inherently cover multiple time scales, requiring long integration times for statistically reliable data at low-frequencies. The high grid resolution near the wall (cell size on the order of micro-meters) and the high-speed flow result in a physical time-step size on the order of nano-seconds, thus requiring millions of time-steps. The baseline SWBLI presented in the following ran for 5×10^6 iterations and consumed a total amount of approximately 15×10^6 core-hours, indicating the necessity of high-performance computing (HPC) in the context of SWBLI at high Reynolds numbers.

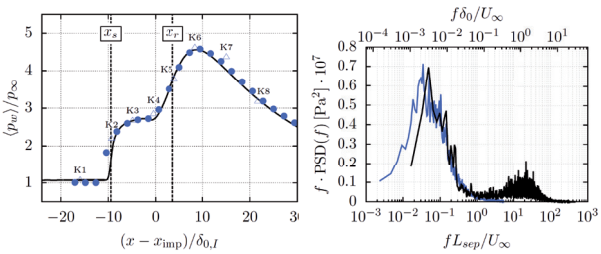

Figure 3: Mean wall-pressure evolution (left) and power spectral density near the mean separation location (right) for experiment (blue) and LES (black).

Baseline SWBLI

The topology studied in this work is an oblique shock-wave impinging on a flat plate turbulent boundary-layer (TBL). Experiments for this case were conducted at the German Aerospace Center [4]. A qualitative comparison between experiment and simulation is given in Fig. 1 by means of instantaneous Schlieren pictures. The adverse pressure gradient imposed by the incident shock is large enough to cause boundary-layer separation and consequently forms a separation shock. Figure 2 gives a three-dimensional impression of the SWBLI. Isosurfaces of time-averaged streamwise vorticity indicate the presence of pairs of counter-rotating Görtler-like vortices. Surface-streamlines show distinct saddle and nodal points along the reattachment line, with footprints of streamwise vortices further downstream. The mean wall-pressure evolution is shown in Fig. 3 (left) and validated against experimental data. The pressure evolution is in good agreement. The massively separated flow exhibits a low-frequent shock motion, which is investigated in Fig. 3 (right) by means of power spectral densities evaluated close to the mean separation location. Both spectra agree well and predict a low-frequency unsteadiness in the range $St_L = 0.03 \ldots 0.05$, which is consistent with previous findings. To highlight the flow-structure at this characteristic low-frequency, we conducted a dynamic mode decomposition (DMD). The pressure and velocity mode, see Fig. 4, clearly highlight a breathing motion of the separation bubble accompanied by an oscillatory movement of the separation shock.

Coupled SWBLI

The coupled SWBLI considers the experimental configuration with a pitching shock generator and an elastic wall [4]. Preliminary numerical results are summarized in [5]. Figure 5 shows a space-time diagram of recorded pressure for the undeflected panel position.

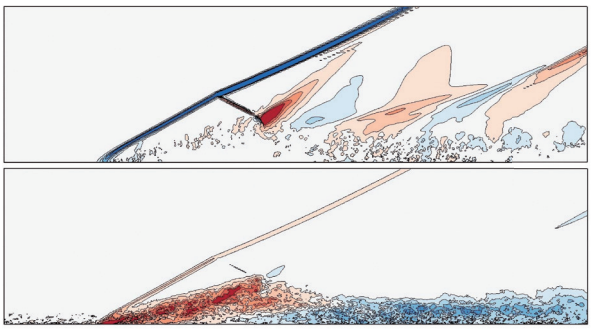

Figure 4: Typical low-frequency pressure and velocity mode (top / bottom) for the baseline SWBLI.

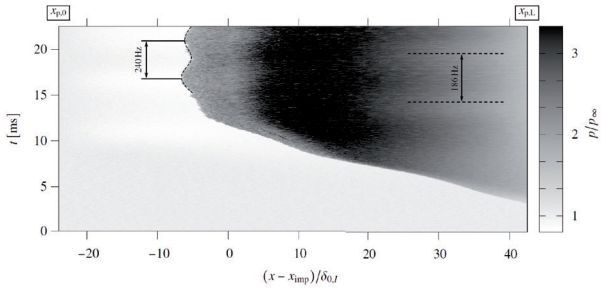

Figure 5: Space-time diagram of recorded pressure.

Compression and expansion waves emanating from the panel oscillation can be clearly seen in the rear part of the panel, together with a large-scale separation shock movement. Figure 6 gives a direct comparison between the coupled and baseline configuration in terms of resolved Reynolds shear stress. Compared to the uncoupled case, the separation, reflected and reattachment shock exhibit a low-frequent shock motion with greater spatial extent, probably enhanced by the panel motion. Moreover, the negative static displacement of the panel contributes to an overall weaker interaction.

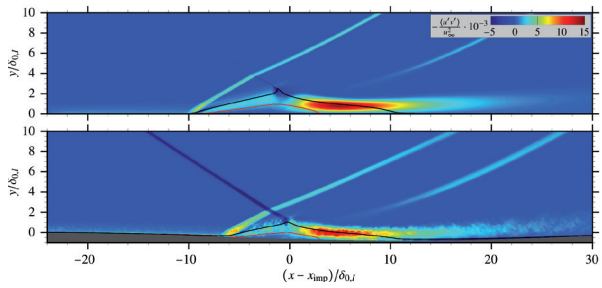

Figure 6: Mean resolved Reynolds shear stress for the baseline (top) and the coupled (bottom) SWBLI.

On-going Research and Outlook

We are currently working on a detailed post-processing of the baseline SWBLI by means of spectral analysis and DMD, incorporating thousands of three-dimensional snapshots of the flow with approximately 80 TB of storage. Further numerical studies with flexible panels will be carried out with a focus on the aeroelastic interaction of the massively separated flow-region with structural modes. The preliminary results show a good agreement between experiment and numerics for both mean and unsteady aspects of the SWBLI. Thus, we are confident that our LES is accurate and will enable us to gain deeper insight into the low-frequency unsteadiness once advanced post-processing techniques are applied.

References and Links

[1] http://www.aer.mw.tum.de/ http://www.sfbtr40.de/ http://inca-cfd.org/

[2] Hickel, S., Egerer, C.P., & Larsson, J. (2014). Subgrid scale modeling for implicit large eddy simulation of compressible flows and shock-turbulence interaction. Physics of Fluids, doi:10.1063/1.4898641

[3] Pasquariello, V., Hammerl, G., Örley, F., Hickel, S., Danowski, C., Popp, A., Wall, W.A., & Adams, N.A. (2016). A cut-cell finite volume – finite element coupling approach for fluid-structure interaction in compressible flow. Journal of Computational Physics, doi:10.1016/j.jcp.2015.12.013

[4] Daub, D., Willems, S., & Gülhan, A. (2015). Experiments on the interaction of a fast-moving shock with an elastic panel. AIAA Journal, doi: 10.2514/1.J054233

[5] Pasquariello, V., Hickel, S., Adams, N.A., Hammerl, G., Wall, W.A., Daub, D., Willems, S., & Gülhan, A. (2015). Coupled simulation of shock-wave/turbulent boundary-layer interaction over a flexible panel. 6Th European Conference for Aeronautics and Space Sciences, Krakau (Poland)

Chemistry and Material Sciences

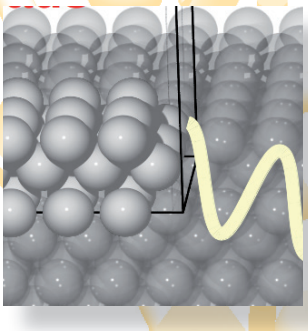

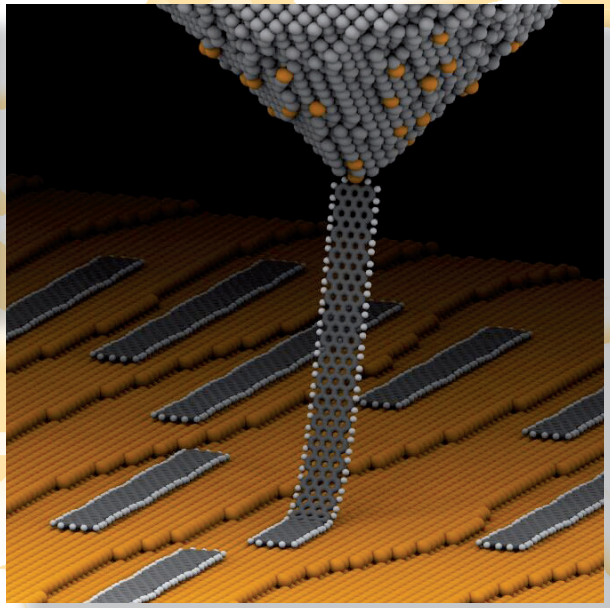

Quantum Monte Carlo and Exact Diagonalization Studies of Correlated Electron Systems

Research Institution
Institut für theoretische Physik und Astrophysik, Universität Würzburg

Principal Investigator
Werner Hanke, Fakher F. Assaad

Researchers
Florian Goth, Ewelina Hankiewicz, Gang Li, Manuel Laubach, Andrzej Fleszar, Rolf Reinthaler, Johannes Hofmann, Dietrich Rothe, Martin Bercx

Project Partners
–

SuperMUC Project ID: h014z

Introduction

Correlation dominated phenomena occur generically in narrow-band systems and in low dimensions. Here, one can cite the canonical example of the Luttinger liquid fixpoint characteristic of the fractionalization of the electron in the one-dimensional limit. More surprisingly, correlation effects can occur at the interface between two semi-conductors even when the semi-conductors themselves can be well described at the single electron level [1]. This example characterizes the overall motivation of the research carried out in this proposal. The key is the understanding of the interplay of interface/surface physics where Rashba spin-orbit coupling (RSOC) is necessarily present, correlation effects and topology. As a step towards this overall goal, the first part on dynamic quantities of 1D Hubbard chains with RSOC summarizes accomplished work leading to an understanding of the interplay between RSOC and correlation effects in one dimension [3]. The following part on Nontrivial edge states of 2D topological insulators on semiconductor surfaces summarizes numerical calculations aiming at generating topological states with adatoms on semiconductor surfaces before we will give a small lookout.

Access to the LRZ supercomputing center was imperative during the grant period to do the relevant simulations that resulted in numerous papers [2,3] in a wide range of topics on correlated electrons. In all cases access to supercomputing facilities allows to carry out simulations on larger and larger system sizes so as to be able to extrapolate to the thermodynamic limit relevant for the understanding of experiments and collective phenomena.

Dynamic quantities of 1D Hubbard chains with Rashba-type spin-orbit interaction
The physics of interacting 1D electrons has attracted considerable interest since the inherently present strong confinement marks a breakdown of the usual nearly free electron picture that is valid in higher dimensions. Due to this strong confinement no individual movement is possible. Therefore any excitation is a collective mode of the system. The effective canonical model applied for 1D systems is the Luttinger liquid that predicts the emergence of collective modes carrying only spin that move at a different speed from excitations carrying charge. The microscopic model we consider is the Hubbard-model in 1D with an added RSOC $H(\mu,\lambda,U) = H_t+H_\lambda+H_U$ where H_t models hopping of the electrons between neighbouring sites and controls the filling. H_λ is the RSOC with coupling strength λ and H_U is the Hubbard interaction of electrons on the same site with strength U. Guided by Quantum Monte Carlo simulations that were performed on SuperMUC we were able to analytically verify an identity that connects simulations performed at different parameter values including the point λ=0 without RSOC, $H(\lambda,U) = f(\lambda)^{-1}H(0,Uf(\lambda))$ with a certain function f. This relation enables us to reinterpret insights gained from simulations of the Hubbard model in terms of a model with RSOC. We exemplify this relation and its use with the help of the single particle spectral function $A(k,\omega)$, which is a measure for the propability of finding an electron of momentum k and energy ω, in Figure 1. The spectral function depicted in (b) is connected to the one in (a) via the relation $A^s(k,\omega,\lambda) = f(\lambda)\, A(k+s \arctan(\lambda), f(\lambda)\,\omega, 0)$ valid between spectral functions.

Nontrivial edge states of 2D topological Insulators on semiconductor surfaces
Topological states of matter differ from conventional materials in that a topological system is insulating in the interior of the bulk but metallic at the surface/edge. Since the discovery of the first two topological states of matter, i.e. two-dimensional (2D) integer and fractional quantum Hall (IQH and FQH) states, topological order became one

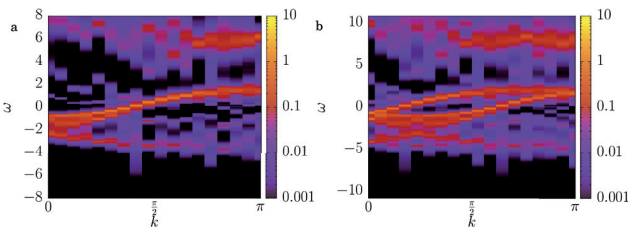

Figure 1: a shows a single particle spectrum for $\lambda = 0$, $U = 6$ and b shows a spectrum for $\lambda = 0.91$, $U = 8.11$ whichcan be connected to a via the relation above.

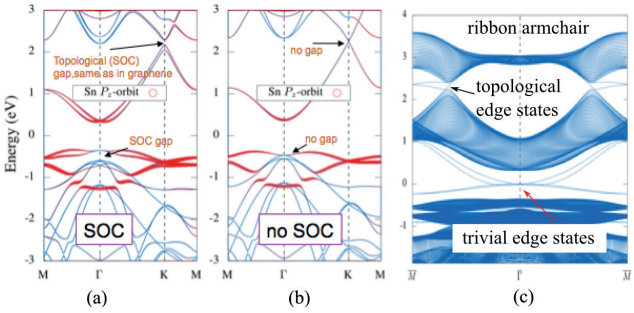

Figure 2: Electronic structure of the bulk stanene/SiC(0001) (a) with and (b) without SOC. The ribbon edge states with an armchair shape are shown in (c).

of the most appealing topics to study in condensed matter physics [4]. In IQH and FQH systems, the bulk Fermi level is located in the middle of two Landau levels and there exist metallic edge states. Their ground states cannot be understood or classified by the conventional symmetry breaking theory, which features their topological nontrivial natures. The first discovered symmetry-protected topological order is the quantum spin Hall (QSH) state, which can be viewed as two copies of a quantum Hall system with opposite magnetic field. The topological property of the QSH state is protected by time-reversal symmetry, and it was experimentally realized in the interfaces of ultra-clean mercury telluride quantum wells.

Another intriguing possibility of realizing unconventional topology in 2D is to grow a topological insulator (TI) on semiconductor surfaces. As a 2D TI, the freestanding Bi(111) bilayer has been shown to exhibit nontrivial topological edge states. The band inversion of this system is sensitive to the buckled honeycomb structure. Similarly, tin bilayer, which is known as stanene, also shows quantum spin Hall states. Chemical functionalized stanene has been shown to contain a strong band inversion between the Sn-s and Sn-p_x+p_y bands, which leads to the appearance of edge states at the Fermi level with a Dirac linear dispersion. Due to the fact that freestanding/decorated stanene and Bi(111) bilayer have not been experimentally reported, it is very important to know, for both academic studies and spintronics applications, if stanene and Bi(111) bilayer preserve a non-trivial topological nature when they are grown on realistic substrates.

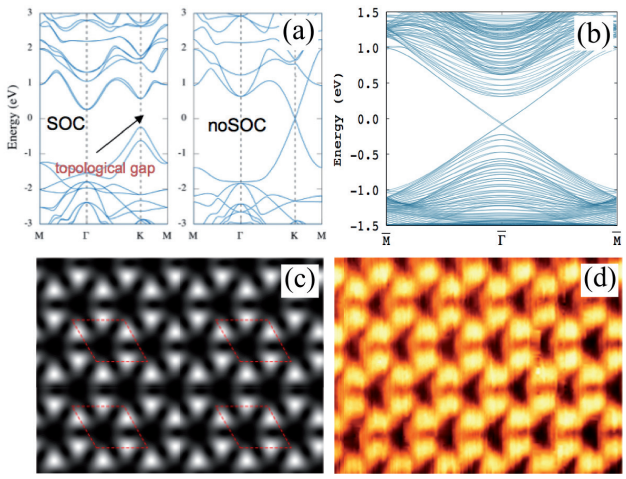

Figure 3: The electronic structure of (a) the bulk and (b) the ribbon with an armchair edge of the Bi(111)/SiC(0001).The simulated STM images (c) show an excellent agreement with experimental ones displayed in (d).

The 6H-SiC(0001) √3 x √3 surface, as one of the candidate substrates with a matchable lattice constant, is a large gap insulator. The growth of stanene and Bi(111) bilayer on this substrate is supposed to strongly reduce the interference of the adatom states around the Fermi level with the states from the substrate. Thus, the topological nature of stanene and Bi(111) bilayer should be unaffected. By using Density-Functional Theory (DFT) with the generalized gradient approximations (GGA), we studied the electronic structure of the stanene/SiC(0001), in which tin adatoms form a honeycomb ring on the SiC(0001) surface. Due to the strain from the substrate, the stanene becomes completely flat with a 2/3-ML coverage. We found that the SOC has a critical influence on the electronic structure of stanene/SiC(0001), as shown in Figure 2 (a) and (b) [5]. The SOC opens a small gap at the Γ-point below the Fermi level and at the K-point at 2 eV above the Fermi level. The latter is a topological nontrivial gap, similar to that in graphene. When the SOC is absent, at the K-point, the electronic structure of stanene/SiC(0001) displays a linear Dirac-type dispersion, see Figure 2 (b). The SOC breaks the band degeneracy here and opens a small gap which supports the non-trivial edge states, as can be seen in the ribbon calculation with an armchair edge of Figure 2 (c).

Inside the energy gap at 2 eV above the Fermi level, there exist edge states due to the strong SOC effect and the nontrivial band topology. In contrast, the charge gap at the Fermi level is induced by the chemical bonding and, thus, is irrelevant to the SOC. The edge states inside this gap are, thus, trivial ones.

To realize the topological insulating phase of adatom systems on semiconductor surfaces based on stanene/SiC(0001), one needs to bring the Fermi level up to the energy gap at 2 eV. One efficient way is to replace tin with bismuth, which effectively dopes electrons to the system. As shown in Figure3, Bi(111)/SiC(0001) displays especially around the K-point a very similar electronic structure as that of stanene/SiC(0001). At the K-point, the SOC opens a gap which is now positioned right at the Fermi level (Figure3 a). From a ribbon calculation with an armchair edge, we confirmed that this is a topological gap that supports nontrivial edge states, see Figure3(b). This energy gap is as large as 1 eV, thus, Bi(111)/SiC(0001) can be a very promising candidate for a large gap 2D QSH insulator. We also examined the experimental feasibility of growing the honeycomb ring of bismuth on SiC(0001) surface. From the preliminary results of the scanning tunnelling microscopy (STM) images, we found a nice agreement between theory (Figure3 (c)) and experiment (Figure3 (d)), confirming the possibility of realizing a large gap 2D QSH insulator in reality.

References and Links

[1] F. F. Assaad, "Interface superconductivity: Get it strained", Nat. Phys. advance online publication (Nov. 2014).
[2] F. F. Assaad, T. C. Lang, and F. Parisen Toldin, "Entanglement spectra of interacting fermions in quantum Monte Carlo simulations", Phys. Rev. B 89 (Mar. 2014) 125121.
[3] F. Goth and F. F. Assaad, "Equivalence of Rashba-Hubbard and Hubbard chains", Phys. Rev. B 90 (Nov. 2014) 195103.
[4] M. Z. Hasan and C. L. Kane, "Colloquium," Rev. Mod. Phys. 82, (Nov. 2010), 3045-3067
[5] Gang Li et al., submitted to PRL.

Hybrid Density Functional Modeling of Transition Metal Clusters

Research Institution
Fachgebiet Theoretische Chemie, Department Chemie, Technische Universität München

Principal Investigator
N. Rösch

Researchers
T. Soini, A. Nikodem, A. Matveev, A. Genest

Project Partners
–

SuperMUC Project ID: h0351

Introduction

Catalysis plays a decisive role for a variety of processes in chemical industry. Catalysts facilitate or steer chemical reactions which otherwise would be slow or inefficient due to a large fraction of undesired byproducts. Most of the catalysts applied today in industry are heterogeneous catalysts, especially small particles of transition metals, supported on oxides or zeolites. This motivates a considerable scientific effort exploring transition metal particles and reactions on their surfaces. In this project we applied density functional methods to larger clusters of late transition metals comprising more than 100 atoms. The goal of this project is to examine the performance of various density functionals and to study the properties of transition metal clusters.

In the fundamental equation of density functional theory, the Kohn-Sham (KS) equation, electronic orbitals and their energies are calculated as eigenvectors and eigenvalues of the density dependent KS operator. These orbitals yield the electronic density. As the KS operator depends in turn on the density, the KS equation has to be solved iteratively (self-consistent field procedure, SCF). While the KS method is exact in principle, the quantum mechanical part of the electron-electron interaction, the density dependence of the so-called exchange-correlation (XC) term, is not explicitly known. Various approximate forms have been suggested. In the simplest formulations, the local density approximation (LDA) and the generalized gradient approximation (GGA), the XC functionals depend on the electronic density and, in addition, its gradients, respectively. More accurate functionals are constructed by also invoking the kinetic energy density of the electrons as variable (meta-GGA) or by including "exact exchange" contributions (hybrid functionals). The largest data sets to be calculated and treated in the Kohn-Sham approach using a local representation are analytic six-dimensional two-electron integrals to construct the Coulomb term of the electron repulsion. For the corresponding "classical" repulsion of electrons these so-called four-center integrals can be simplified by representing the electron density via a set of auxiliary functions. For the exact exchange term in hybrid XC functionals, four-center integrals have to be calculated explicitly. For our calculations we used the parallel density functional program ParaGauss [1], developed in our group. This software uses Gaussian functions to represent electronic orbitals. Recently ParaGauss has been improved by introducing a new algorithm for the high-performance parallel treatment of the more demanding tasks of the KS approach. A work stealing algorithm without a central steering unit has been implemented, applying MPI-2 [2]. The new algorithm avoids a bottleneck at a steering master process for a large number of workers and is applicable to any type of task proratable in independent work packages. Initially, these work packages are distributed evenly to the cores used. To equilibrate the computational load, a process that has finished its tasks steals untreated tasks from another process. Thus, each of the parallel processes stays occupied until all tasks have been treated. The resulting load balancing is more

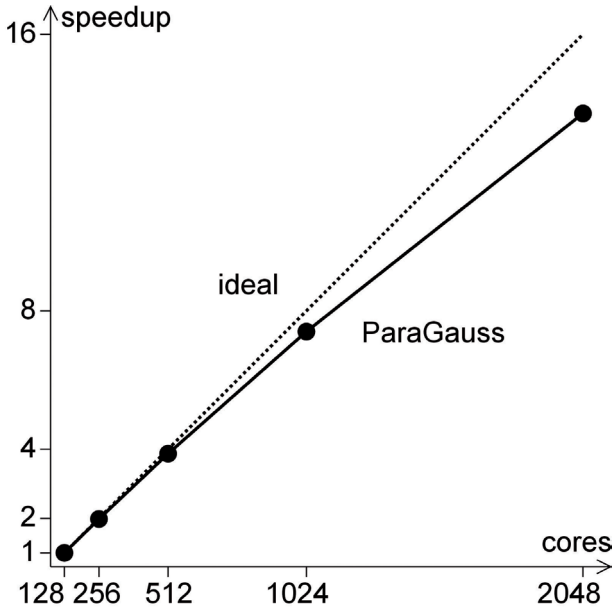

Figure 1: Scaling of a density functional electronic structure calculation of $Pt_{140}(CO)_8$.

favorable than that due to a static distribution [2]. This algorithm has been used in ParaGauss to parallelize the calculation of all types of analytic integrals and the numerical evaluation of the XC functionals.

Results

The new parallelization library has been validated by calculations on metal clusters applying hybrid XC functionals as the analytic four-center integrals needed for the exact exchange part of these functionals are the most demanding task of such density functional calculations. A series of all-electron test calculations of the four-center integrals of medium-sized copper clusters Cu_n, n = 10–22, without exploiting structural symmetry, showed that the new algorithm yields efficiencies above 93% for calculations on 512 cores for all these systems, increasing to 100% for larger systems. Work stealing performed always slightly better than a static distribution of tasks. For Cu_{22} with work stealing an efficiency above 95% is achieved even on 2048 cores, while a static task distribution yields only 70 %. Models of metal nano particles can efficiently be calculated by choosing spatially symmetric structures and exploiting this symmetry. Full electronic structure all-electron test calculations applying the TPSSh hybrid XC functional including energy derivatives for geometry optimization of Cu_{79} (O_h symmetry) were carried on 512 to 2048 cores with an efficiency of 92% on 1024 cores and 72% on 2048 cores. This drop in efficiency for 2048 cores for this medium sized calculation is due to a limited scaling of various small tasks of the SCF procedure while the demanding calculation of four-center integrals still yields an efficiency of 96%, demonstrating that the implementation is capable of dealing with larger systems. As a pertinent example of a larger system, treated again with the TPSSh functional, we mention electronic structure calculations using pseudopotentials of the cluster Pt_{140} (D_{4h} symmetry) covered by 8 CO molecules. Test calculations on this system, on 128 to 2048 cores, showed very favorable scaling (Figure 1), with an overall efficiency of 86% on 2048 cores.

Our new efficiently parallel implementation of four-center integrals allowed for the first time a comparison of various hybrid XC functionals for larger transition metal clusters which are interesting models of heterogeneous catalysts. The hybrid functionals PBE0, M06 and TPSSh, have been compared to their semi-local congeners PBE (GGA), M06L, and TPSS (both meta-GGA). Cube-octahedral particles of Ni, Pd, and Pt (in O_h symmetry) with 13, 38, 55, 79, and 116 atoms were chosen as test set [3]. Calculated average metal-metal bond lengths and cohesive energies (binding energies per atom) were extrapolated to infinite cluster size to compare the results to the corresponding experimentally known bulk properties. Hybrid functionals underestimate the cohesive energy, in line with their overestimation of bond lengths, while meta-GGA functionals tend to overestimate the cohesive energy. The known preference of hybrid functionals for high-spin states was confirmed also for the larger transition metal clusters inspected. Overall, the hybrid functional TPSSh yields the best results and thus is suggested when a hybrid functional is to be applied in studies on transition metal catalysis [3].

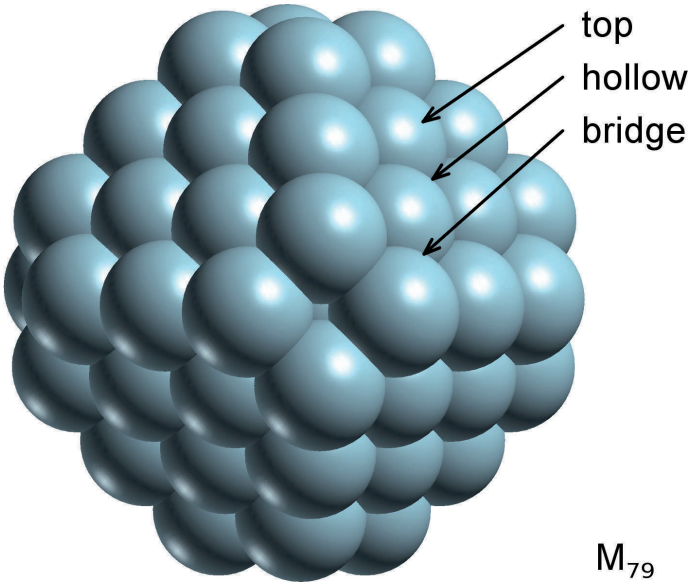

Figure 2: Various types of adsorption sites for CO on the truncated octahedral metal cluster M_{79} as example.

We also explored the preferred adsorption site of the CO molecule on Pt, as common density functional methods predict preferred adsorption at hollow sites (Figure 2) in contrast to the experimental finding of top sites on the close-packed Pt(111) surface. This failure of the KS approach is interpreted as a self-energy artifact because the lowest unoccupied molecular orbital of CO, $2\pi^*$, is calculated too low in energy. This failure may be resolved by hybrid DFT methods. Therefore, we calculated CO adsorption at the (111) facets of the cluster model Pt_{79} (D_{4h} symmetry), resulting in the species $Pt_{79}(CO)_8$. In agreement with other density functional results for surface models, we also find for all GGA, meta-GGA, and hybrid functionals considered that CO adsorption is more stable at hollow sites than at top sites on Pt_{79} (Figure 2). The most preferred site for all methods is the bridge site. The smallest energy difference between hollow and top sites is provided by the TPSSh functional. Calculations with the larger cluster Pt_{140} yield hollow sites as most preferable, showing that the preferred bridge site of Pt_{79} seems to be a feature of smaller clusters. Interestingly, while current hybrid methods seem to fail in predicting the correct CO adsorption site, we were able to demonstrate that the DFT+U approach, applying an on-site correction, is able to resolve this problem for CO at Pt_{140} at the price of slightly underestimating the CO binding energy [3].

Exploiting the improved parallel efficiency of ParaGauss with the PBE functional, we pursued the question how the preference for an hcp packing of ruthenium develops with increasing metal cluster size [4]. For this question we picked ruthenium particles with 55 to 250 atoms. The Ru bulk has a hcp packing, but very small Ru particles have been shown to prefer a fcc lattice. We picked two series of fcc and hcp type clusters with nuclearities in a comparable size range and applied an extrapolation

strategy to infinite cluster size to compare the accuracy of our calculated results with experimental data for the bulk. Interestingly, one hardly notices a difference in the extrapolated nearest neighbor distances: 271.6 pm for fcc clusters and 271.8 pm for hcp clusters [4]. One would have expected that the hcp lattice leads to smaller values but seemingly the variations in the chosen set of clusters are too small to express this in the extrapolated values. We also noticed typical contractions of bond lengths in the outermost layer of a cluster, by up to 10 pm for both structure types. The cohesive energy of the bulk is by ~8% underestimated. However, the hcp clusters are correctly predicted to be more stable than their fcc congeners. Regarding the evolution of the hcp lattice we determined that already 55 atom clusters show a mild preference of a few kJ mol^{-1} per atom for hcp, hence the structure transition tends to occur at cluster sizes in that range or slightly below.

On-going Research / Outlook

Ongoing improvements of the parallel performance of ParaGauss aim at routine applications to transition metal and oxide particles that comprise a few hundred atoms.

References and Links

[1] Belling, T.; Grauschopf, T.; Krüger, S.; Mayer, M.; Nörtemann, F.; Staufer, M.; Zenger, C.; Rösch N. in: High performance scientific and engineering computing, Bungartz, H.-J., Durst, F., Zenger, C., Eds., Lecture Notes in Computational Science and Engineering, Vol. 8., Springer, Heidelberg, 1999, 439-453.
[2] Nikodem, A., Matveev, A., Soini, T., Rösch, N., Int. J. Quantum Chem., 114, 813-822.
[3] Soini, T., Rösch, N., Phys. Chem. Chem. Phys., 17, 2015, 28463-28483.
[4] Soini, T., Ma, X., Aktürk, O., Suthirakun, S., Genest, A., Rösch, N. Surf. Sci., 643, 2016, 156-163.

Uranyl Adsorption at Clay Mineral Surfaces

Research Institution
Fachgebiet Theoretische Chemie, Department Chemie, Technische Universität München

Principal Investigator
N. Rösch

Researchers
A. Kremleva, S. Krüger

Project Partners
–

SuperMUC Project ID: h0351

Introduction

The aqueous and sorption chemistry of actinides is of special concern for the assessment of environmental safety with regard to treating and storing radioactive waste of nuclear power plants and other facilities of nuclear industry. In this project the adsorption of actinide ions on clay minerals is studied computationally. Clay minerals are ubiquitous in soils and various sedimentary rocks. Clays are used as technical barriers to prevent actinide migration and clay rocks represent a possible host formation for a final repository of highly radioactive waste. Cation adsorption and precipitation is an important mechanism to prevent the transport of actinide ions in the environment. Therefore detailed knowledge of the adsorption chemistry of actinides on clay minerals is of key importance for predicting their behavior and controlling these hazardous elements in the environment.

We employed the well-established Vienna Ab Initio Simulation Package (VASP) [1] for calculating the electronic structure of periodic bulk and surface models (Figure 1). Surfaces are modeled as repeated slabs, separated by a "vacuum" layer, to achieve three-dimensional periodicity that allows one to apply the plane-wave band-structure approach to electronic structure. This software uses MPI for a parallel treatment by distributing bands and coefficients of the wave function representation. For typical calculations with periodic unit cells of 100–200 atoms 64–128 cores can be used efficiently. Larger systems allow a higher degree of parallelization. Test calculations for uranyl adsorbed on montmorillonite showed efficiencies (referenced to 32 cores) of 78 % on 128 cores and 50 % on 256 cores for a small model with 200 atoms. For large models with 800 atoms an efficiency of more than 90 % is obtained with 256 cores and of 76 % for 512 cores. These results demonstrate a very favorable scaling with increasing systems size. Dynamic simulations of solvated mineral surfaces present a particular computational challenge because for realistic simulation times of about 10–20 ps in the order of 50–100 thousand electronic structure calculations have to be carried out. In optimizations of solvated mineral surfaces, surface solvation is approximately modeled by adsorption of 1–3 layers of water molecules (Figure 1), but the results suffer from the arbitrary initial structure of the solvation layer [2]. To avoid a costly fully dynamic treatment, we suggested a simulated-annealing scheme at low temperature [3]. Equilibration for 3–5 ps at 200 K and lowering of the temperature to 0 K during 1 ps, followed by an optimization step, leads to more trustworthy energies and structures. This procedure provides an essentially equilibrated solvation layer and allows one to compare various, also metastable surface complexes, while a fully dynamic treatment yields only the most stable species.

Results

Clay minerals show a layered structure at the atomic level, composed of interconnected sheets of silica tetrahedra SiO_4 and alumina octahedra AlO_6. In this project the very common class of 2:1 clay minerals has been inspected. This class of clay minerals shows the common structure of a central alumina sheet sandwiched between two silica sheets. This basic structure corresponds to the mineral pyrophyllite, exhibiting a neutral unit cell. Various other 2:1 clay minerals are derived from this structure by cation substitutions, inducing permanent negative charges of the mineral layer that are balanced by counter ions between the layers, e.g., Na^+ and K^+.

Clay minerals typically form platelets of micro- to millimeter size. These platelets show predominantly basal surfaces parallel to the mineral layers, exposing chemically saturated silica tetrahedra that show a low reactivity. With respect to metal ion adsorption, edge surfaces oriented (essentially) perpendicular to the minerals sheets

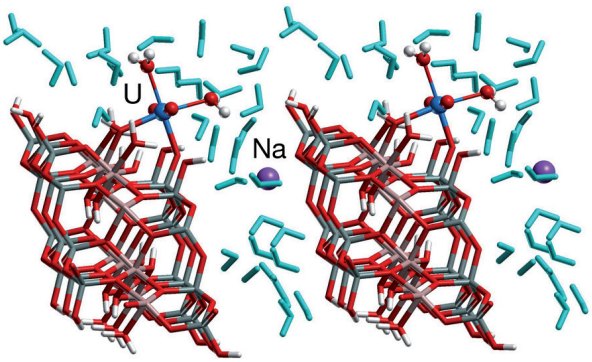

Figure 1: Uranyl(VI) adsorbed at the solvated (110) edge surface of a smecite mineral.

(Figure 1) are more reactive as they expose a variety of aluminol, silanol, and mixed surface groups for various orientations and terminations. In addition to the negative charge of the layers resulting from substitutions, edge surfaces exhibit variable surface charges, depending on the pH of the surrounding aqueous solution, due to adsorption (positive charge, low pH) or release (negative charge, high pH) of protons. In this project, surface structures of pyrophyllite of various edge surface orientations have been examined. Dynamic simulations showed, in agreement with earlier studies, that some of these surfaces have a dynamic character as protons are easily exchanged between $AlOH_2$ and AlOH surface groups.

In a large survey we studied the adsorption of uranyl(-VI), UO_2^{2+}, on the more important edge surfaces (010) and (110) (Figure 1) of pyrophyllite, montmorillonite, and beidellite [2]. The charged 2:1 smectite clay minerals montmorillonite and beidellite differ from pyrophyllite by substitutions of Al^{3+} by Mg^{2+} in the octahedral sheet and of Si^{4+} by Al^{3+} in the tetrahedral sheet, respectively. The permanent layer charges have been balanced by solvated Na^+ counter ions in the interlayer space. We inspected surface models where the cation substitutions are on the surface or in subsurface positions. This large set of surface models, together with the presence of various adsorption sites on each of them, resulted in more than 100 adsorption complexes to be optimized. Surprisingly, essentially no changes were found in the structure of the mineral surfaces due to substitutions, even for substitutions at the surface. Exceptions, besides rare proton exchanges between surface groups, are aqua ligands at surface Al^{3+} centers of the (010) edge surface. These ligands are loosely bound and may be exchanged with the solution or lost in the course of close-by uranyl adsorption. As the surface structures of the minerals compared are rather similar, also the same types of uranyl adsorption complexes have been determined. While the structure of complexes at the same site of different surfaces and minerals is rather similar, the complex formation energies vary. The energy differences between various adsorption complexes tend to be wider distributed on pyrophyllite than on the substituted minerals. For the example of the (110) surface we were able to show by means of the simulated annealing procedure, that the charged smectites bind uranyl stronger than pyrophyllite and that binding at cation-exchanged surface sites is preferred (Figure 2). Independent of the mineral considered, the most stable adsorption complex is uranyl bridging in bidentate fashion an AlOH and a SiO surface group (Figure 2) [3]. Besides the typically fivefold coordinated uranyl ion, we also found uranyl monohydroxide to appear as adsorbate. This change of the adsorbed species takes place when an aqua ligand of uranyl close to a surface AlOH group deprotonates and the proton moves to the surface, forming an $AlOH_2$ group [2]. In addition to our finding of uranyl monohydroxide as a novel adsorbed species, yet to be confirmed experimentally, we calculated for some complexes a reduced equatorial coordination number of 4, instead of the commonly found one 5. Notable geometric differences of uranyl adsorption complexes between the minerals are only obtained when adsorption takes place at sites involving a substituted cation directly at the surface. From these results we conclude that the type of surface chemical groups forming various adsorption sites determines the adsorption complexes present. In contrast, surface orientations as well as the specific mineral play only a secondary role. This result agrees with available X-ray absorption fine structure (EXAFS) experiments which provide similar geometric parameters for uranyl adsorption at various clay and related minerals.

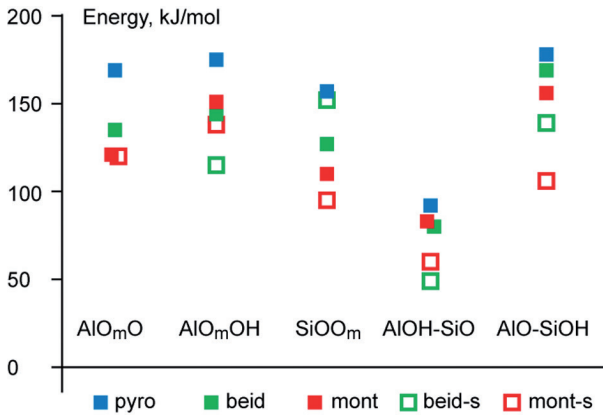

Figure 2: Formation energies of adsorbed uranyl(VI) complexes at the (110) edge surfaces of pyrophyllite, beidellite, and montmorillonite. Open symbols represent sites at substituted surfaces (s).

EXAFS results for uranyl adsorption on clay minerals showed a slight elongation of the uranyl U-O bond, U-Al/Si distances of about 310 pm and 330 pm and 5–6 U-O equatorial bonds to surface O centers as well as to aqua ligands, which in EXAFS are accessible only in an averaged form. In some experiments this equatorial shell of bonds was resolved into 2–3 shorter and some longer bonds. The short ones are interpreted as bonds to surface O centers while the longer ones have been attributed to bonds to aqua ligands. We extended this interpretation, as we calculated short as well as longer bonds to surface oxygen centers. In addition, short equatorial U-O bonds also appear as bonds of hydroxide to uranyl for uranyl monohydroxide as adsorbate.

By comparing various adsorption complexes we traced the variation of lengths of equatorial U-O bonds to the empirical charge of the oxygen centers involved [4]. Thus, a differentiation of bonds to ligands or to the surface by their length only is not possible. Overall we showed that uranyl complexes binding to the same surface groups of edge surfaces of various minerals are rather similar; several sites may simultaneously be occupied. In addition, the coordination number may be lowered to 4 and hydrolysis on the surface may occur, leading to new adsorbed species. Comparison of optimized geometries and improved energetic considerations allowed us to suggest probable sites. These results will be helpful for the construction of thermodynamic models of actinide adsorption at clay mineral surfaces and the interpretation of spectroscopic experiments.

On-going Research / Outlook

This first systematic survey of actinide adsorption at complex clay mineral surfaces, which provided new insights at the atomic level, is currently being extended to neptunyl NpO_2^+ and more complex minerals, like iron-substituted phyllosilicates. In this way we will examine if the concepts developed so far can be applied more generally to support the interpretation of pertinent experiments. A further facet of these studies will be to account also for the dynamic nature of the mineral/water interface by means of exemplary dynamic simulations.

This work has been supported by Bundesministerium für Wirtschaft und Energie, grant No. 02E11001.

References and Links

[1] Kresse, G., Hafner, J. Phys. Rev. B 47 (1993) 558-561; Blöchl, P. E., Phys. Rev. B 50 (1994) 17953-17979.
[2] Kremleva, A., Krüger, S., Rösch N., Phys. Chem. Chem. Phys. 17 (2015) 13757-13768.
[3] Kremleva, A., Krüger, S., Rösch N., Phys. Chem. C 120 (2016) 324–335.
[4] Kremleva, A., Krüger, S., Rösch, N., Surf. Sci. 615 (2013) 21-25.

Simulation of Electron Transfer and Electron Transport Processes in Molecular Systems at Surfaces

Research Institution
Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

Principal Investigator
Michael Thoss

Researchers
Michel Bockstedte, Pedro B. Coto, Susanne Leitherer, Veronika Prucker, Dominik Weckbecker

Project Partners
–

SuperMUC Project ID: pr28lo

Introduction

The transfer or transport of electrons (ET) is at the heart of many processes in physics, chemistry, technology and biology. Efficient top-down and bottom-up fabrication techniques have made possible the development of increasingly miniaturized components of electronic devices. This has motivated a renewed interest over the last decade in the research of ET in self-assembled monolayers of molecules adsorbed at surfaces and voltage driven charge transport in molecular junctions (i.e. molecules which are chemically bound to metal or carbon-based electrodes).

From a theoretical point of view, modelling electron transfer and electron transport at molecule-substrate interfaces requires a methodology that can describe simultaneously systems with a discrete energy spectrum such as molecules, and extended systems with a (quasi-)continuous energy spectrum, such as semiconductor or metal surfaces or electrodes. Furthermore, it must be able to account for quantum and non-equilibrium effects, which for the type of systems investigated usually play an important role in both heterogeneous ET reactions and voltage driven electron transport processes. Therefore the methodology, in addition to the detailed characterization of the electronic structure of the systems investigated, needs to account for the dynamics of the process using quantum dynamical methods.

In this contribution, we report recent results on the simulation of heterogeneous ET processes. Specifically, we analyze the dynamics of electron injection in a series of nitrile-substituted phenylthiolate self-assembled monolayers (SAM) adsorbed at the Au(111) surface [1] and investigate voltage-driven charge transport processes in molecular junctions that use graphene as electrode material [2].

Results

In the simulation of the ET processes, we have used a two-step approach that combines electronic structure calculations with quantum dynamical simulations. In the first step, we use ab initio electronic structure methods to characterize the electronic structure of the system of interest obtaining the relevant parameters needed for the simulation of the dynamics. Due to the extended nature of the systems investigated we have used periodic density functional theory (DFT) methods for the characterization of their electronic structures. Specifically, we have solved the Kohn-Sham equations using the projec-

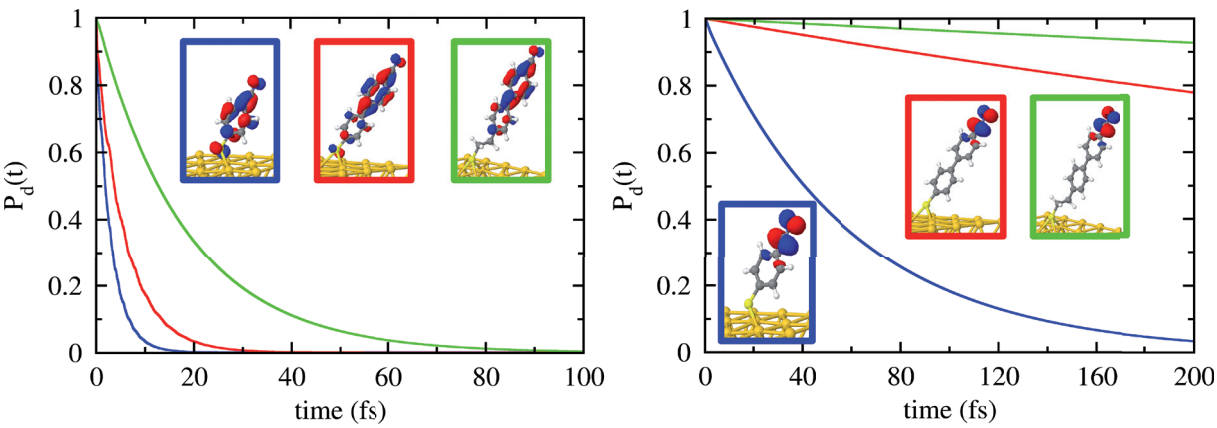

Figure 1: Population dynamics of π_1^* (left) and π_2^* (right) donor states for A (blue), B (red), and C (green). The orbital densities of the donor states for the different systems investigated are also shown.

tor augmented wave method and a plane wave basis as implemented in the Vienna Ab-initio Simulation Package (VASP) [3]. These calculations strongly benefit from the use of a massively parallel computing environment like SuperMUC. Information obtained in this step includes the energies of donor and acceptor electronic states, donor-acceptor electronic couplings and vibrational normal modes of the systems investigated. This information is used in the second step of the approach to simulate the dynamics of electron injection in heterogeneous ET processes or current-voltage characteristics of molecular junctions employing model Hamiltonians [4] and non-equilibrium Green's function techniques [2].

Electron injection dynamics in nitrile-substituted phenylthiolate SAMs adsorbed at the Au(111) surface

We have investigated the electron injection dynamics in a series of nitrile-substituted (poly)phenylthiolate SAM models characterized by having spacers exhibiting a different number of conductor (phenylene) and hybrid conductor-insulator (phenylene-methanediyl) groups ($-C_6H_4-$ (A), $-(C_6H_4)_2-$, (B) and $-(C_6H_4)_2-(CH_2)_2-$ (C), see Fig. 1) adsorbed at the Au(111) surface. Specifically, we have simulated the injection dynamics from the π_1^* and π_2^* donor states corresponding to the two π^* resonances of the -CN moiety (see Fig. 1). For all systems, we found equilibrium geometries where the adsorbed molecule is tilted against the surface normal in a fcc bridge-like configuration, in agreement with the experimental information. The results obtained show that, in contrast to what was found in previous work on ET in nitrile substituted alkanethiolates SAMs [4], the π_1^* and π_2^* donor states are not quasi-degenerate, as a consequence of their different coupling with the phenylene spacer. Specifically, π_1^* shows a large delocalization over the (poly) phenylene spacer in all systems investigated, featuring also significant density at the sulfur atom involved bonded to the surface. This is reflected in the electron injection dynamics (see Fig. 1). The results obtained show that the π_1^* donor states exhibit ultrafast (femtosecond) electron injection times that increase with the length of the spacer chain. In detail, A, B, and C exhibit electron injection times of $\tau_{1/2}$=1.8 fs, $\tau_{1/2}$=3.7 fs and $\tau_{1/2}$ = 12.7 fs, respectively. In the case of π_2^*, the electron injection times are significantly longer ($\tau_{1/2}$ = 40.5 fs, $\tau_{1/2}$ = 553.8 fs and $\tau_{1/2}$ > 1 ps, respectively). These results can be rationalized in terms of the distance dependence of the donor-acceptor couplings that depends on the length of the (poly)phenyl spacer and the degree of delocalization of the donor state. Addition of a small chain of methanediyl groups to the biphenyl spacer (system C) does not qualitatively modify the delocalized nature of the π_1^* donor state, which shows density over the biphenyl spacer. However, there is negligible density at the methanediyl groups or the sulfur atom (see Fig. 1, left). As a consequence, the electron injection rate in this system takes longer due to the reduction in the donor-acceptor couplings caused by the increase in the donor-acceptor distance.

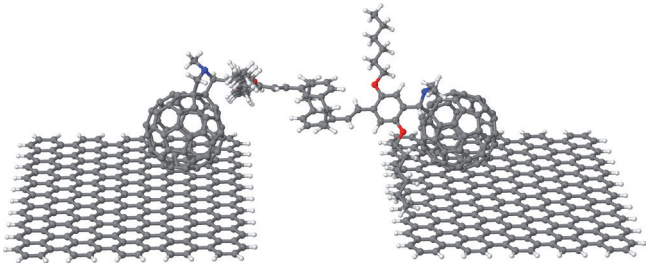

Figure 2: Single-molecule junction investigated with a fullerene end-capped molecule comprising two styrene units covalently connected through a [2,2']paracyclophane as bridge and zigzag terminated graphene electrodes.

Charge transport in molecular junctions with graphene electrodes

The development of new fabrication techniques has enabled the use of graphene as material for electrodes in molecular junctions. This synthetic carbon allotrope is endowed with different structural (rigidity and mechanical stability) and electronic (high electron mobility) properties that make it a promising material for nanoelectronic applications. In a first step towards the development of new devices using graphene-based single-molecule junctions the conductance properties of these systems have to be analyzed. As a model system we have investigated the nanojunction depicted in Fig. 2, which features fullerene anchor groups physisorbed at the graphene electrodes. Recent experiments [2] revealed asymmetric current-voltage characteristics as well as switching behavior in this junction. Our theoretical investigation shows that the asymmetry of the current voltage-characteristics can be rationalized on the basis of the intrinsic asymmetry of the molecular bridge. Furthermore, we have also found that edge states in graphene electrodes with zigzag termination play a significant role, strongly affecting the conductance at small bias voltages.

On-going Research / Outlook

We plan to continue and extend the study of heterogeneous ET in molecular systems at surfaces with the goal to achieve a comprehensive understanding of the dynamics of these processes. Specific processes to be investigated include electron transfer in organic molecules adsorbed at surfaces, focusing on the role of vibronic effects and charge transport in molecule-graphene nanojunctions, addressing in particular the role of vibronic effects, spin-sensitive transport, and light-driven transport phenomena. In addition, we are developing new theoretical methods for the characterization of these processes.

References and Links

[1] Prucker, V., Bockstedte, M., Wang, H., Coto, P. B., Thoss, M. 2016. Manuscript in preparation.

[2] Ullmann, K., Coto, P. B., Leitherer, S., Molina-Ontoria, A., Martín, N., Thoss, M., Weber, H. B. 2015. Nano Lett. 15, 3512.

[3] Viena Ab-initio Simulation Package (VASP), V5.3.5, 2014. https://www.vasp.at//

[4] Prucker, V., Rubio-Pons, O., Bockstedte, M., Wang, H., Coto, P. B., Thoss, M. 2013. J. Phys. Chem. C. 117, 25334.

http://www.thcp.nat.uni-erlangen.de

Coupling kMC and CFD in Heterogeneous Catalysis

Research Institution
Department Chemie – Technische Universitaet Muenchen

Principal Investigator
Karsten Reuter

Researchers
Matteo Maestri, Sebastian Matera, Stefano Rebughini

Project Partners
–

SuperMUC Project ID: pr47ma

Introduction

The worldwide rapidly growing demand for more efficient and sustainable exploitation of energy and material resources puts catalysis higher and higher up on the research agenda. New and improved catalysts are needed to deliver security of supply and industrial competitiveness. The required extreme targets on activity, selectivity and stability under very demanding operating conditions represent a formidable challenge and require the ability to nano-design catalytic materials well beyond our present capabilities. In this respect, it is pivotal to distinguish between the intrinsic catalytic properties of the catalyst material and its macroscopic observed functionality. The observable (and exploitable) functionality of heterogeneous catalysts results from the interplay of two quite distinct aspects. On the one hand, there are the intrinsic catalytic properties of the active catalyst material when interacting with molecules of the surrounding gas phase. They derive from the material's property to make and break chemical bonds at its surface(s) and are ultimately determined at the electronic structure level. On the other hand, there is the macroscopic heat and mass flow in the actual reactor geometry and at the actually employed operation conditions, which determines the local gas-phase concentrations and temperature at the catalyst surface. Both aspects are generally intricately coupled. In order to reach a rational understanding of the mechanisms underlying the observed functionality, an increased level of scientific understanding is required at each characteristic scale involved in the catalytic process (as depicted in Figure 1).

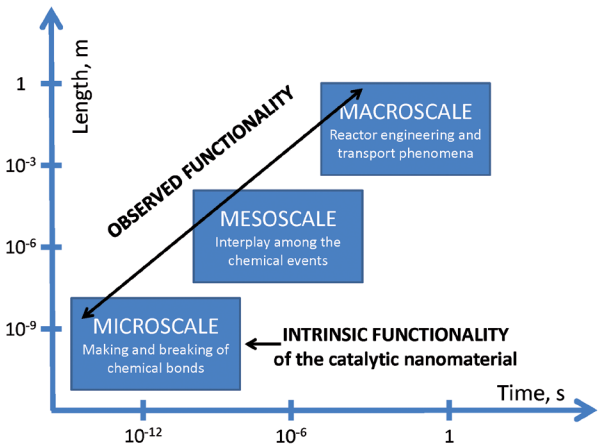

Figure 1: Time and length scales in heterogeneous catalysis.

With reference to Fig. 1:

- At the microscale, ab-initio - i.e. first-principles - electronic structure theory calculations are used to unravel the making and breaking of chemical bonds at the surface, by explicitly treating the electronic degrees of freedom;
- At the mesoscale, statistical simulations explicitly account for the interplay among all the chemical events, largely improving upon effective formulations that inherently rely on uncontrollable averagings;
- At the macroscale, energy, mass and momentum conservation equations are employed to model transport phenomena in the reactor and in the catalyst support, fully integrating the first-principles based reactive surface chemistry unraveled at the lower scales.

In the first two-years of the project pr47ma (2010-2012), we have successfully demonstrated how first-principles calculations and semi-empirical models can be used for the estimation of kinetic constants for complex reaction networks [1,2]. In the extension period (2012-2014) we have focused our attention on the coupling of the microscale and meso- scales) with the macro-scale (Computational Fluid Dynamics, CFD). In the period 2014-2016 we have put our attention in particular on the testing of our proposed algorithm for coupling kinetic-Monte-Carlo simulations (kMC) and CFD.

Thanks to the granted computational time, we proved that the resulting methodology allows to efficiently integrate the microscopically correct account of the spatial arrangements and interactions of the adsorbed chemicals with the fundamental description of fluid-dynamics and transport at the catalyst interface, in an effective seamless flow from the atoms to the reactor. *As a proof of the impact of the work, the resulting paper was included in the "ACS Editors' choice" list and published as "open-access" [6].*

Results

The integration of the kMC simulations into the fluid dynamic framework has been achieved using the interpolated kMC method [3]. In particular, Matera and Reuter employed an instantaneous steady-state approximation to generically present any irregularly gridded steady-state reactivity data of any complex (first-principles) kinetic model in the form of an interpolated data field for the boundary condition. For the computational solution of the fluid-dynamics we have used the recently developed solver catalyticFOAM, that allows for the solution of Navier–Stokes equations for complex and general geometries for reacting flows at surfaces, based on a detailed microkinetic description of the surface reactivity [4,5]. Noteworthy, the CatalyticFOAM solver exploits the operator-splitting technique to separate transport and reaction terms, and therewith allows already for the solution of the Navier-Stokes equations for complex and general flows at reactive surfaces described on the level of mean-field microkinetic rate equations. The CatalyticFOAM solver is intrinsically dynamic, and the steady state is reached by time-dependent simulation. In this respect the separation between reaction and flow-time scales is essential. This feature turned out to be extremely important when multiple steady-state solutions are possible depending on the initial conditions.

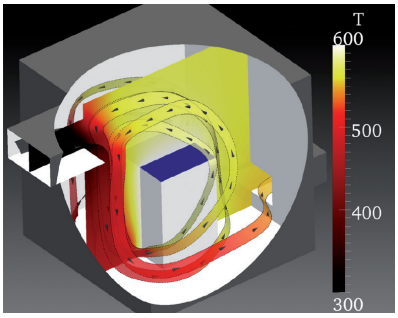

Figure 2: Illustration of the complex stationary flow pattern arising during CO oxidation at Pd(100) inside the reactor (Adapted from Ref. [6]).

The *project on superMUC* consisted the testing and assessing of the interface between the interpolated kMC method and catalyticFOAM solver. In particular, we have tested the capability of the code in managing complex computational domains. Both random and structured geometries were employed for the tests.

Our tests on superMUC clearly showed that the combination of the domain decomposition technique and operator splitting algorithm enable the simulation of very complex geometries along with a microkinetic description of the surface reactivity even when kMC derived turn-over-frequencies are employed. Moreover, contrary to what observed for the direct calculations of the reaction rates under the mean-field approach (see report for the period 2012-2014), the possibility of importing TOF calculated beforehand *strongly improved the efficiency of the simulation in terms of computational time*. In fact, under the assumption that the intrinsic catalytic activity at every resolved surface mesh point of the macroscopic fields is solely dependent on the local temperature and partial pressures at this point. The activity at neighboring mesh points only enters indirectly through its effect on these local gas-phase conditions, not directly through some coupling mediated by the surface population itself. In particular, for macroscopically homogeneous catalysts, it then suffices to determine one interpolated data field for the intrinsic catalytic activity as a continuous function of the temperature and all partial pressures. This data field can then be used as a CFD boundary condition independently at every CFD surface mesh point. In this way, the micro/meso-scales are fully decoupled in time from the macroscale. Thus, by the combination of the operator splitting algorithms with the domain decomposition technique, each run was easily parallelized on superMUC, with performances very similar to non-reactive CFD simulation with openFOAM. Also, the possibility of using TOF calculated beforehand allowed us to avoid the dependence of the efficiency of the scalability on the distribution of the catalytic cells in the domains, which was one of the main drawbacks reported in our report for the period 2012-2014. In fact, since the calculation of the source term is fully decoupled from the CFD simulation, there is no an unbalanced workload associated to the catalytic cells (i.e., where surface reactions take place), and thus the domain decomposition techniques turned out to be a very suitable technique for the parallelization of the simulation. Very good scaling properties were found for all the investigated cases. For instance, for a random packed bed of spheres, which consists of about 2,000,000 cells, was found to scale almost linearly up to 1024 processors (2000 cells/core), in full analogy with our previous mean-field simulations.

We demonstrated the capabilities of the established approach by applying it to a series of problems of increasing flow complexity that are particularly geared toward in situ experiments at flat-faced model catalysts. In all cases we observed intricate couplings of reactive surface chemistry and flow that significantly modify the observable catalytic function. As an example, Fig. 2 illustrates the complex flow and temperature fields arising during CO oxidation at Pd(100) inside the spectroscopic cell.

This strongly underlined the necessity of reliable accounts of both aspects in integrated multiscale catalysis simulations when aiming to establish an atomic-scale mechanistic understanding of this function in technologically relevant environments.

On the whole, through the granted computational time, we were able to prove that the proposed numerical algorithm makes possible the simulation of multidimensional systems with complex and detailed kinetic mechanisms, overcoming the unfeasible computational effort (especially in terms of memory) that would be required by fully-coupled algorithms.

References and Links

[1] M. Maestri, K. Reuter, Chem. Eng. Science, 74 (2012) 296
[2] M. Maestri, K. Reuter, Angew. Chemie Int. Ed. 50 (2011) 1194
[3] S. Matera, K. Reuter, Catal. Letters, 133 (2009) 156
[4] M. Maestri, A. Cuoci, Chem. Eng. Science, 96 (2013) 106
[5] www.catalyticfoam.polimi.it
[6] S. Matera, M. Maestri, A. Cuoci, K. Reuter, ACS Catalysis, 4 (2014) 4081-4092

Theoretical investigation of photo-catalytic water splitting

Research Institution
Chair for Theoretical Chemistry, TU Munich

Principal Investigator
Karsten Reuter

Researchers
Harald Oberhofer, Daniel Berger, Christoph Schober

Project Partners
–

SuperMUC Project ID: pr58ba

Introduction

Oxide surfaces – such as TiO_2 or ZnO – have long been known to possess the ability to split water if irradiated by light. Unfortunately, the yields are small and most of the reactive surfaces only work with highly energetic UV-light.

In order to make the process of water splitting viable for large scale application one therefore has to either find new materials or surfaces or introduce co-catalysts, like e.g. small clusters of coinage metals which have shown much promise in experiments.

Our project was mainly concerned with the formulation of methods to investigate the mechanisms involved in the water splitting process and finding promising new catalyst materials. Due to the size of the involved systems and the need to accurately describe not only structures but also chemical reactions we exclusively used methods based on Density Functional Theory. Apart from data analysis and simulation control scripts using the python ASE[1] frameworks all simulations were either performed with the FHI-aims[2] or CASTEP[3] DFT packages.

Results

Computational screening of co-catalysts
We formulated a theoretical approach for the screening of viable catalyst/co-catalyst combinations [4]. At the same time we were able to show that DFT beyond the GGA level is necessary to achieve even qualitatively correct results. Looking at Au clusters up to 55 atoms (see below) this means the fully non-local treatment of up to 4345 electrons in a sufficiently large basis. For this we employed the numeric atomic orbital all-electron code FHI-aims [2], which shows excellent parallelisation and near linear scaling with system size. It is very well suited for use on architectures such as SuperMUC due to a customised scalapack implementation which at its heart has the highly efficient ELPA [5] eigenvalue solver.

Non-local exchange calculations for such a large number of heavy atoms is not only expensive in terms of CPU time but also with regards to memory consumption. Additional optimisation of the code was necessary to reduce the memory necessary per core below the 1.5 Gb available on SuperMUC thin nodes.

Additionally, we implemented a solid state embedding scheme based on the QM/MM approach which allows us to even go beyond hybrid DFT functional accuracy [6]. With our embedding approach we can also study surface and solvent effects acting on the Au-clusters. As this approach effectively eliminates the need for periodic boundary conditions, QM/MM surface calculations show the same parallel efficiency as vacuum runs with only a minimal overhead due to the polarisation of the MM region. A first

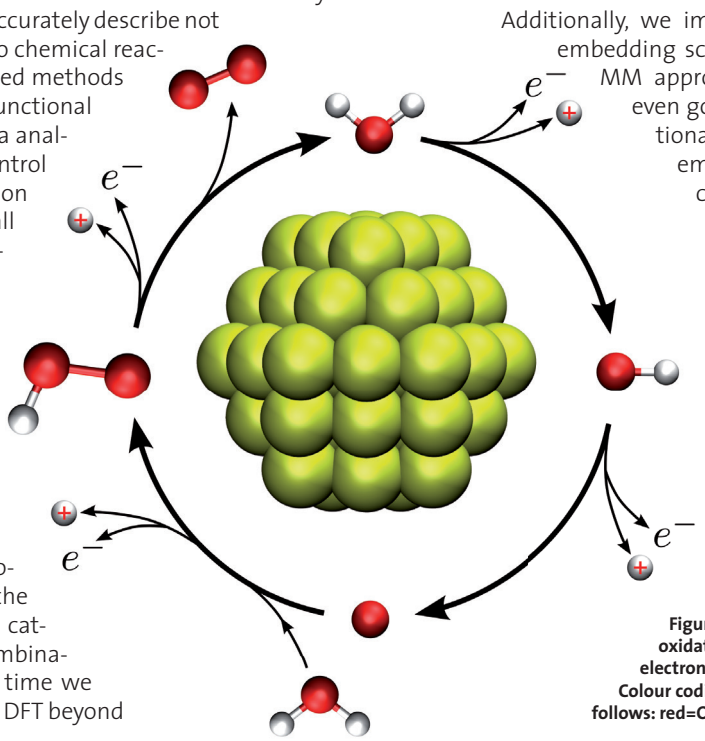

Figure 1: Schematic of a water oxidation reaction via proton coupled electron transfer on a gold nano-particle. Colour coding of the depicted atoms is as follows: red=O, white=H, yellow=Au

application of this method, successfully predicted charge state and stability of point defects on the catalytically highly important TiO2 (110) surface [7].

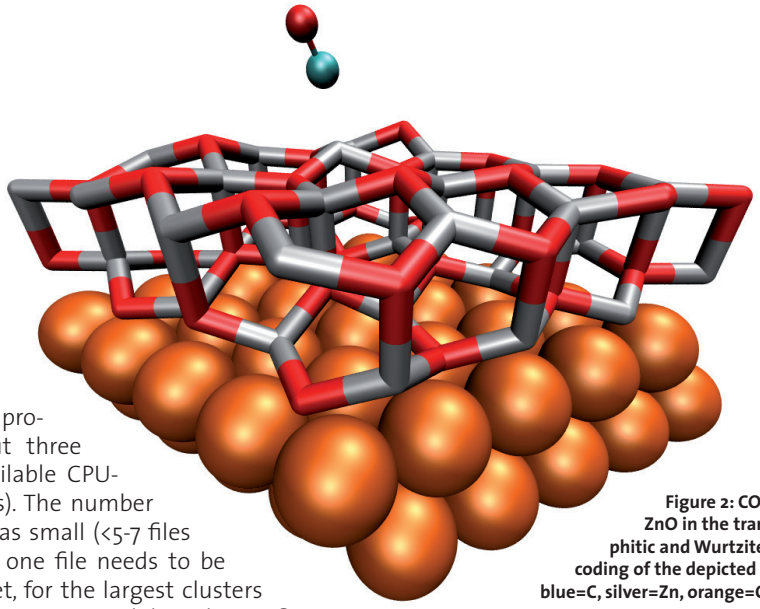

Figure 2: CO adsorbed on double layered ZnO in the transitionary phase between graphitic and Wurtzite on a Copper substrate. Colour coding of the depicted atoms is as follows: red=O, blue=C, silver=Zn, orange=Cu

For this part of the project we used about three quarters of our available CPU-time (3M CPU hours). The number of files generated was small (<5-7 files per run where only one file needs to be stored as output). Yet, for the largest clusters checkpointing files on SCRATCH did reach significant sizes of up to ~100 Gb. As these are only relevant in case of a crash or wall-time overrun no permanent storage of checkpoint files is necessary. Number of CPUs per job depended strongly on the size of the system at hand. Small jobs were typically conducted on up to 256 cores, while for the large ones we used up to 4096 cores. Average job size was around 512 cores.

A new thin film phase of ZnO

Together with the experimental group of Prof. Wöll at the Karlsruhe Institute of Technology we studied the formation of ZnO mono-, di- and tri-layers on Copper surfaces. Such systems are relevant for chemical applications, as they allow direct coupling of a chemically active oxide with a metal, and can be produced relatively easily by oxidising brass surfaces. Using CO as a probe molecule we were able to show the existence of a previously unknown thin film phase of ZnO [8], which is structurally situated halfway between the already known graphitic and Wurtzite phases. An example of the new phase is depicted in Figure 2.

In order to rule out any polarisation and finite size effects calculations had to be performed on a 5x5 Cu surface unit cell mirrored around the centre of the simulation box. As calculations were performed with the pseudopotential based plane-wave code CASTEP[3] such large box sizes were rather costly. In order to compare with experiment we calculated the harmonic vibration frequency spectrum and the transition dipole for the adsorbed CO molecule, on every unique adsorption site. Vibrational analysis was performed via the method of finite differences with the ASE[1] python framework. CASTEP parallelises well on SuperMUC architecture especially regarding reciprocal space sampling over k-points.

In total we used approximately 1M CPU hours for this project. Due to the nature of our implementation of the Tkatchenko Scheffler van der Waals correction a large number (~300 per run) of temporary files were generated on scratch. Yet again, only the actual output files containing energies and frequencies were retained for later use. Here, memory limitations were not an issue. On average we again used around 512 cores per run. Calculation of the transition dipole moments was performed in a post-processing step on a workstation.

Outlook

Calculations for projects such as this are extremely costly due to the large size of the systems involved and the level of theory necessary to gain meaningful results. Only with a supercomputer such as SuperMUC were these calculations even possible. During the two years run-time of our project, progress was hindered by frequent breakdowns and longer outages of services at SuperMUC. These, which in part can be attributed to the newness of the system, added up to a total of not much less than 10% downtime over the course of our project. Additionally, SuperMUC policy to completely eliminate access to project data at the very day the project ends seems curiously strict, at least compared to other large supercomputing centres, as it essentially prohibits users from running jobs during the last days of their allotted project time.

Nevertheless, the results outlined above represent only the beginning of our research into more effective water splitting catalysts. Follow up proposals for a more in-depth study of surface and solvent effects, as well as dynamical simulations of water splitting are currently in preparation.

References and Links

[1] https://wiki.fysik.dtu.dk/ase/
[2] https://aimsclub.fhi-berlin.mpg.de/
[3] http://www.castep.org/
[4] H. Oberhofer and K. Reuter; J. Chem. Phys. 139 044710 (2013)
[5] T. Auckenthaler, V. Blum, H.-J. Bungartz, T. Huckle, R. Johanni, L. Krämer, B. Lang, H. Lederer, and P.R. Willems; Parallel Comput. 37 783 (2012)
[6] D. Berger, A.J. Logsdail, H. Oberhofer, M.R. Farrow, C.R A. Catlow, P. Sherwood, A.A. Sokol, V. Blum, and K. Reuter; J. Chem. Phys. 141, 024105 (2014)
[7] D. Berger, H. Oberhofer, K. Reuter; Phys. Rev. B 92, 075308 (2015)
[8] V. Schott, H. Oberhofer, A. Birkner, M. Xu, Y. Wang, M. Muhler, K. Reuter, C. Wöll; Angew. Chem. Int. ed. 52 11925 (2013)

Heterogeneous catalysis from advanced molecular dynamics

Research Institution
Lehrstuhl für Theoretische Chemie, Ruhr-Universität Bochum

Principal Investigator
Dominik Marx

Researchers
Johannes Frenzel, Daniel Muñoz-Santiburcio, Niklas Siemer, Luis Martínez-Suárez

Project Partners
–

SuperMUC Project ID: pr63ce

Introduction

Heterogeneous catalysis is a pivotal technology in chemical industry. Improvement or even development new of catalysts in terms of selectivity, stability, and energy efficiency is most desirable. Key to rational design strategies is the detailed understanding of the structure– activity relationship such as knowledge of active sites and reaction mechanism(s). Yet, harsh reaction conditions of may largely hamper such insight via experiments.

One example is the the synthesis of the bulk chemical methanol. Nowadays it is produced from syngas (a mixture of CO_2, CO and H_2) over the Cu/ZnO catalyst (ICI process). The efficiency of this process results from a complex scenario of surface chemical reactions at high temperatures and pressures. In response to these physical and chemical conditions of the gas phase the catalyst undergoes dynamical morphological changes. These transformations give rise to a huge space of possibilities in terms of structural and chemical configurations of both, adsorbates and continuously altering Cu/ZnO surface. Here we show that the underlying reaction mechanism(s), active sites and chemical species can be automatically "synthesized" in silico using our molecular dynamics approach to computational heterogeneous catalysis [1,2,3] while keeping any a priory input on species and mechanisms at a minimum.

Results and Methods

Accelerated finite temperature *ab initio* molecular dynamics in conjunction with a thermodynamically optimized Cu_8/ZnO nanocatalyst model [4,5] was used to sample a reduced but multidimensional free energy landscape of the chemical transformations of molecular carbon species from CO_2 toward methanol, see Figure 1 and Ref [6]. This landscape discloses an overwhelmingly rich network of parallel, competing and reverse reaction channels over a *lively* catalyst surface, see Figure 1. Over twenty carbon species have been generated in addition to side species and adspecies decorations on Cu. Highly reactive formal-

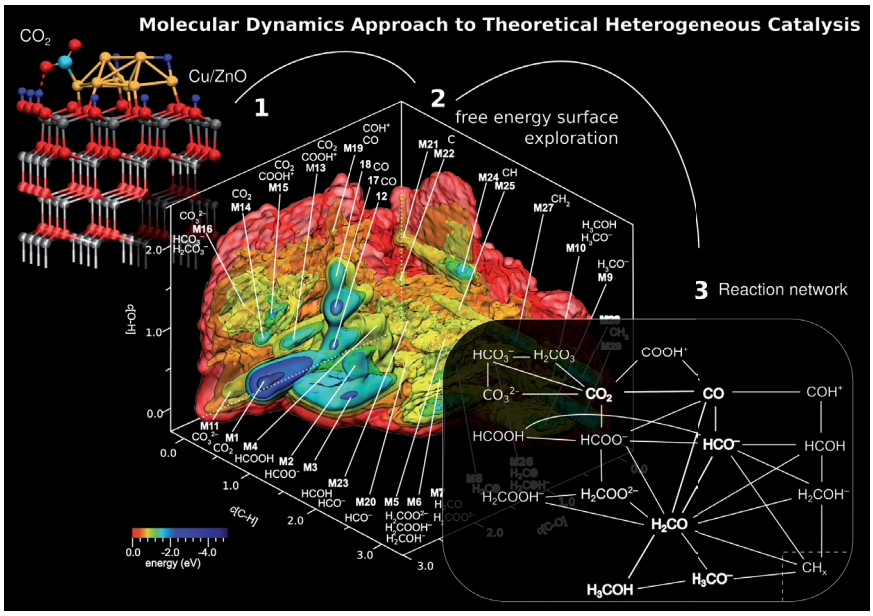

Figure 1 Process of generating a global reaction network of chemical species and elementary reactions from pure computational means.
1 Thermodynamic optimization: Activated CO_2 over the Cu_8/ZnO catalyst model.
2 Metadynamics sampling: gross 3D free energy surface with the "synthesized" chemical species which were obtained from analysis of the ab initio molecular dynamics trajectory underlying the metadynamics.
3 Schematic reaction network of methanol synthesis from CO_2 over the Cu_8/ZnO nanocatalyst model including parallel reactions, prominent side reactions such as methanation, reverse and forward water gas shift reaction, and reactions leading to catalyst deactivation.

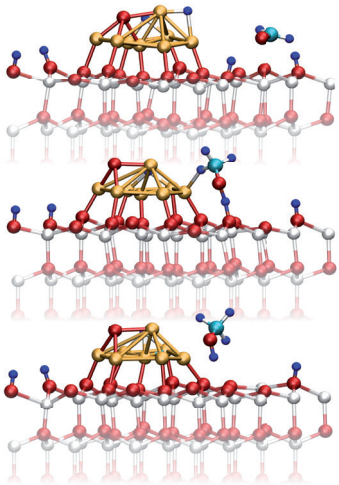

Figure 2: Potential reaction pathway of methanol formation from formaldehyde over the Cu_8/ZnO nanocatalyst model via the Eley–Rideal type surface chemical reaction. Representative snapshots are taken from the ab initio molecular dynamics trajectory underlying the sampling of the free energy surface as shown in Figure 1.

dehyde species could be identified to be the key intermediate toward methanol. Analysis of the trajectories underlying the ab initio metadynamics sampling of the free energy landscape results several well known side reactions such as reverse and forward water gas shift reaction, and methanation but also those reactions leading to catalyst deactivation such as coking. Further analysis of individual reaction steps within the full reaction network identified three prototypical types of surface chemical reactions, namely Eley–Rideal, Langmuir– Hinshelwood, and Mars–van Krevelen. We stress that the former and the latter mechanisms can only be realized when the active reaction space systematically included in all simulations, namely the near surface at the interface between catalyst and gas phase, for an example see Figure 2, and dynamical transformations of catalysts sites, respectively.

Performance and scaling

The AIMD simulations were carried out using CPMD [7,8] which is a DFT-based molecular dynamics code. The electronic structure of the nanocatalyst model was described by a plane wave basis set in combination with ultrasoft pseudopotentials. The Kohn-Sham equations of DFT were propagated in an extended Lagrangian scheme. Computationally this is the most demanding part of the calculation which require a high band width/ low latency interconnect. We assigned these tasks to processor groups (PGr). Within a PGr parallelization is realized via MPI and either MPI or OpenMP/Vector for inter-node and intra-node communication processing, respectively. At this level of parallelization the allocation of 160 cores allowed for the highest speedup for our Cu/ ZnO nanocatalyst model. Herein the 3D-FFT of the electronic wavefunctions is split up most efficiently into one 2D-FFT per processor core. Within one 2D-FFT the CPMD code allows more levels of parallelization, e.g. via DFT orbitals. Utilizing these levels for our system resulted in lower speed-up. Production runs were carried out with 2D-FFT parallelization because it gave the best utilization of the granted cpu time and time to solution.

On SuperMUC CPMD propagation of a single Cu/ZnO nanocatalyst model required 8.9 core·min per MD step on fat nodes, 6.1 core·min per MD step on Phase 1 thin nodes and 3.5 core·min per MD step on Phase 2 nodes, that is full utilization of 4 nodes, 10 nodes, and 6 nodes, respectively. There exists another level of parallelization which is the multiple walker extension to the metadynamics simulation. Walkers interact in building up the biasing potential and this interaction requires negligible communication effort. Such, the algorithm parallelizes intrinsically with respect to the number of replicas (walkers) of the total system allowing for some tenths of walkers. On SuperMUC partitions best performance between queue waiting vs. production was achieved using up to ten walkers, which was 1600 cores. In addition to the MD sampling electronic structure analysis was performed using 2D-FFT parallelization, too. On SuperMUC this project has used approximately ten million core hours. A typical job used 80 to 1600 cores depending on the number of walkers to explore the FES of Figure 1. Besides bookkeeping of the trajectory data at every MD step larger disk I/O demands during such jobs were complete restart files which were written every two hours and were 0.5 GiB in size. The maximum storage needed in SCRATCH, WORK and PROJECT was 1 TiB and 5 TiB and 100 GiB, respectively. We note that typical node utilization is best with the number of cores being powers of two, as is the thin nodes of Phase 1, whereas Phase 2 with 28 cores is not optimal for complete filling of nodes by our typical jobs.

On-going Research / Outlook

With the installation of SuperMUC Phase 2 computations with our Cu/ZnO catalyst model were speed up by about 70% compared to Phase 1. This allows for larger system sizes and ab initio molecular dynamic simulations with computationally more demanding electronic structure methods for calculation of atomic forces. Currently we are investigating mild and selective methanol oxidation over Titania supported gold nanocatalyst both, in gas phase and in aqueous solution, by employing advanced MD techniques similar to the present project. These simulations require twice to three times of computational resources in the setup compared to Cu/ZnO. However, accurate calculation of the electronic structure for these Titania gold systems will be critical in terms of available resources. We estimate an increase in computational cost by 100 to 500 times which would be required and hopefully accessible by "SuperMUC Next Generation".

References and Links

[1] J. Kiss, J. Frenzel, N. N. Nair, B. Meyer, and D. Marx, J. Chem. Phys. 134, 064710 (2011).
[2] J. Frenzel, J. Kiss, N. N. Nair, B. Meyer, and D. Marx, Phys. Status Solidi B 250, 1174 (2013).
[3] J. Frenzel, D. Marx, J. Chem. Phys. 2 141, 124710 (2014)
[4] L. Martínez-Suárez, J. Frenzel, D. Marx, and B. Meyer, Phys. Rev. Lett. 110, 086108 (2013).
[5] L. Martínez-Suárez, J. Frenzel, D. Marx, Phys. Chem. Chem. Phys. 16, 26119 (2014)
[6] L. Martínez-Suárez, N.Siemer, J. Frenzel, D. Marx, ACS Catal. 5, 4201 (2015)
[7] D. Marx and J. Hutter, "Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods", Cambridge University Press (2009)
[8] CPMD, http://www.cpmd.org/, Copyright IBM Corp 1990-2015, Copyright MPI für Festkörperforschung Stuttgart 1997-2001

http://www.theochem.rub.de/research/marx
http://www.ruhr-uni-bochum.de/solvation

Sub-Kelvin cooling with Magnetic Molecules

Research Institution
Universität Bielefeld, Fakultät für Physik

Principal Investigator
Jürgen Schnack

Researchers
–

Project Partners
–

SuperMUC Project ID: pr63fa

Introduction

That magnetic materials can change their temperature is known since the pioneering work of Emil Warburg and other scientists in the nineteenth century. He discovered that plain iron changes its temperature when an applied magnetic field is removed. This magnetocaloric effect (MCE) can be used for a variety of cooling applications, for instance in special room-temperature refrigerators that work without refrigerant fluids. The magnetocaloric effect is also employed in order to achieve the lowest possible temperatures in thermodynamic cycles such as the Carnot or Ericsson cycles, which work with paramagnetic substances. Sub-Kelvin temperatures were experimentally obtained already in 1933 by W.F. Giauque who received the Nobel prize for his achievement in 1949. In contrast to paramagnets, magnetic molecules should offer more flexibility concerning the cooling process, but until now this remained a hypothetical option.

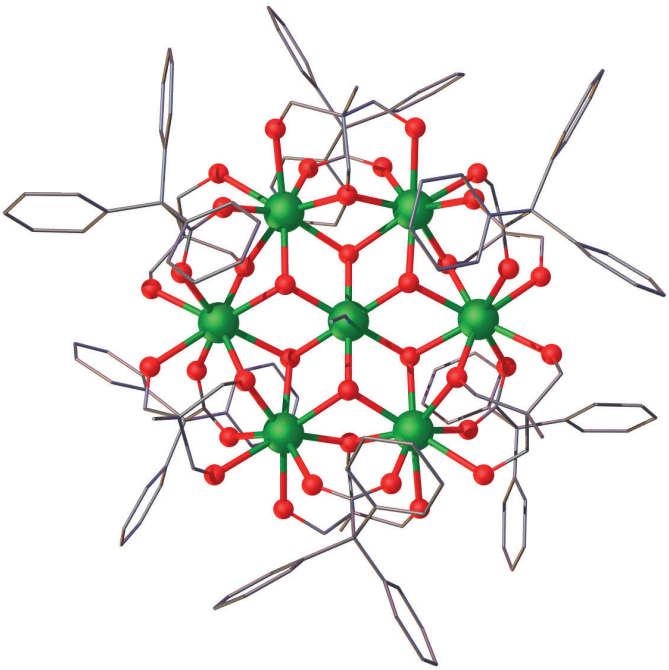

Figure 1: Structure of $[Gd_7(OH)_6(thmeH_2)_5\ (thmeH)(tpa)_6(MeCN)_2](NO_3)_2$ ["Gd_7"; H_3thme = tris(hydroxymethyl)ethane; Htpa = triphenylacetic acid].

Our group [1] investigates small quantum spin systems such as for instance magnetic molecules for their use as low-temperature refrigerants. Since those substances constitute correlated many-body systems powerful approximations to the Schrödinger equation are needed in order to understand their magnetic properties properly. We employ powerful Krylov space methods in order to evluate thermodynamic observables such as magnetization or specific heat. Our experience is that our method of choice – the Finite-Temperature Lanczos Method (FTLM) [2] – delivers very accurate quantum results, which enable us to study thermodynamic functions such as the entropy in great detail [3].

Results and Methods

Our partners from Chemistry at the University of Manchester in the UK synthesized a molecular cluster containing seven gadolinium ions, with six forming a hexagon and one sitting in the center, see Figure 1. Each gadolinium ion possesses a spin of s=7/2. For Gd one expects a weak interaction of these magnetic ions, therefore a large MCE should be possible.

MCE experiments with samples of the crystals were carried out in Zaragossa (Spain). In an adiabatic process the temperature fell to around 0.2K. Besides the fact that this investigation constitutes the first sub-Kelvin cooling experiment with magnetic molecules, it also exemplifies the richness of adiabatic processes in interacting magnetic quantum systems.

For a paramagnet the curves of constant entropy, on which the important adiabatic processes run up or down, are straight lines heading towards the origin of the T-B plane, compare black lines in Figure 2. Their slope, which is the cooling rate, is always T/B for each pair of temperature T and field B, from where one wants to cool down, for instance. Magnetic molecules allow one to achieve very different cooling rates in certain parts of the T-B plane, especially close to the B axis. Such a behavior could be theoretically predicted for Gd_7 and successfully measured. The bumpy structure of the isentropes of Gd_7 shown in Figure 2 as colored curves reflects the unusual structure of magnetic energy levels present in Gd_7 and

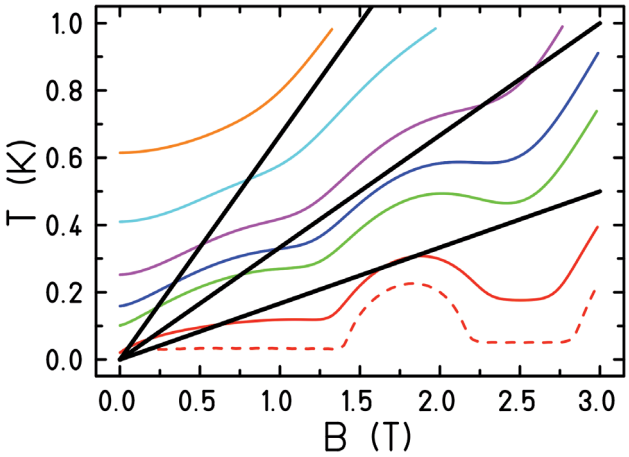

Figure 2: Theoretical curves of constant entropy, i.e. isentropes of Gd_7 (coloured) compared to those of a paramagnet (straight black lines).

simultaneously demonstrates that an interacting quantum magnet may produce cooling rates that are much larger in certain T-B regions compared to a paramagnet, a phenomenon that is termed enhanced magnetocaloric effect. In addition processes such as heating upon decreasing the field are possible, too, which never happen with paramagnets.

The unusual and pronounced bumpy structure of the isentropes of Gd_7 (Fig. 2), which could be followed in the cooling experiments, is an outcome of the frustrated nature of the antiferromagnetic interactions in Gd_7. Since the centered hexagon consists of triangles, the antiferromagnetic interaction is unable to constitute a magnetic ground state of pairwise "happy" combinations of up and down spins. Such a situation is termed "frustrated". Frustration often leads to an unusual bunching of low-lying energy levels, which as a function of applied magnetic field may vary strongly. It is this strong variation of the density of low-lying energy levels that produces the beautiful bumpy entropy landscape. The hope is thus that in the future we would be able to design isentropes according to our needs by means of rational design of magnetic molecules.

For the theoretical modeling we used a combination of FTLM and complete numerical diagonalization. FTLM was employed to search in parameter space for good fits to the static magnetization as well as susceptibility. Here openMP schemes were used to parallelize matrix-vector operations on the nodes. MPI was employed to start simulations with various parameter sets. In a second step complete numerical diagonalization was performed using a decomposition of the Hilbert space according to the irreducible representations of the underlying symmetry groups.

On-going Research / Outlook

Magnetocalorics remains a very active field of research, both for sub-Kelvin as well as for room temperature applications. One aim is to replace the very rare and thus expensive ^{3}He that is used in many cryostats for low-temperature cooling. A thorough understanding of the magnetocaloric properties of quantum spin systems rests on an accurate evaluation of their thermodynamic functions such as the entropy.

Even for moderately larger spin systems such functions can only be evaluated on HPC systems since matrix-vector operations in gigantic Hilbert spaces are involved. But thanks to the progress made in development of methods (FTLM, Density Matrix Renormalization Group, Quantum Monte Carlo) and to the technical progress of HPC systems the theoretical investigation of quantum spin systems has greatly advanced, and so has our understanding of the magnetocaloric properties.

We thus like to close with a cartoon of Gd_7 and a snowflake, Figure 3, in order to express our hope, that more and better magnetic refrigerants will be discovered in the future.

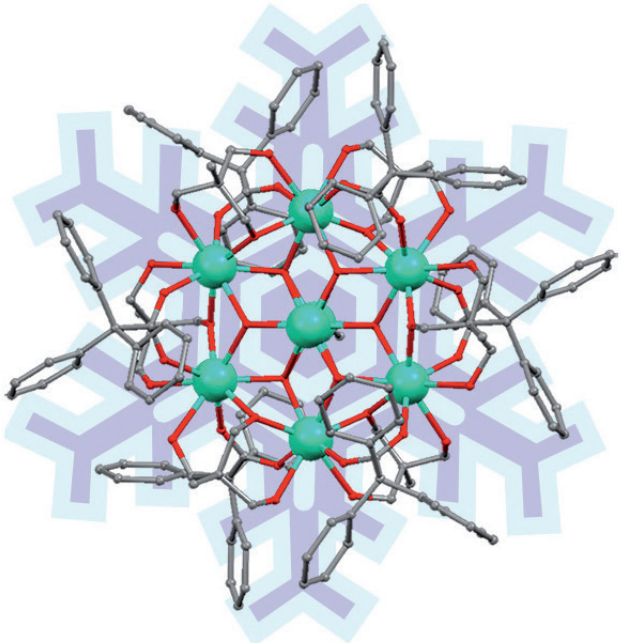

Figure 3: Graphical combination of the molecule Gd_7 with a snow flake signaling its magnetocaloric properties.

References and Links

[1] http://obelix.physik.uni-bielefeld.de/~schnack
[2] J. Jaklič and P. Prelovsěk. 1994. Phys. Rev. B. 49, 5065–5068.
[3] (a) T.N. Hooper, J. Schnack, S. Piligkos, M. Evangelisti, E.K. Brechin. 2012. Angew. Chem. Int. Ed. 51, 4633-4636; (b) Yong Zheng, Qian-Chong Zhang, La-Sheng Long, Rong-Bin Huang, A. Müller, J. Schnack, Lan-Sun Zheng and Zhiping Zheng. 2013. Chem. Commun. 49, 36-38; (c) J. Schnack and C. Heesing. 2013. Eur. Phys. J. B 86, 46; (d) E. Garlatti, S. Carretta, J. Schnack, G. Amoretti and P. Santini. 2013. Appl. Phys. Lett. 103, 202410.
[4] J.W. Sharples, D. Collison, E.J.L. McInnes, J. Schnack, E. Palacios, M. Evangelisti. 2014. Nature Communications 5, 5321.

Numerical Renormalization Group studies of deposited Magnetic Molecules

Research Institution
Universität Bielefeld, Fakultät für Physik

Principal Investigator
Jürgen Schnack

Researchers
Martin Höck, Henning-Timm Langwald, Felix Kaiser, Jürgen Schnack

Project Partners
–

SuperMUC Project ID: pr63fa

Introduction

Magnetic molecules offer the prospect of encoding and storing information in their magnetic state. This applies, in particular, to bistable molecules such as single molecule magnets (SMMs). The possibility to store, e.g., one bit of information in the state of a single molecule would constitute an enormous miniaturization and could lead to data storage technologies with significantly increased areal density. However, to make this technological application feasible, the molecules need to be individually addressable so that their magnetic state can be probed and manipulated on a molecule-by-molecule basis. In the last years, there has been an increasing interest in the question whether this functionality can be achieved by a controlled deposition of magnetic molecules on suitable substrates. This approach can introduce new complications due to interactions between the molecules and the surface. Therefore, even if the magnetic response of the isolated molecule is well understood, its magnetic properties in contact with the surface have to be reinvestigated.

Kondo impurity models can be used to theoretically describe magnetic molecules that are deposited on a me-

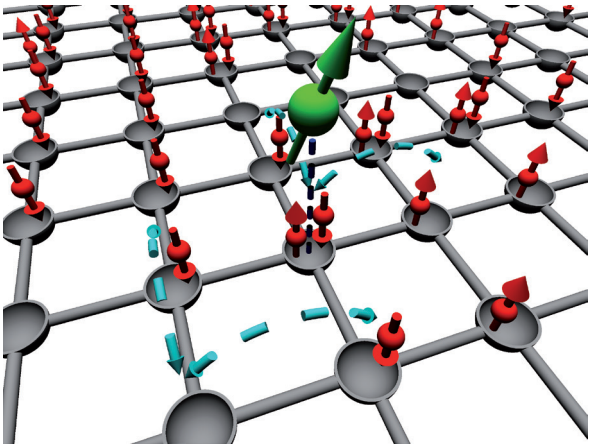

Figure 1: Schematic structure of a single spin deposited on a non-magnetic metallic substrate. Red spheres with arrows represent the conduction electrons.

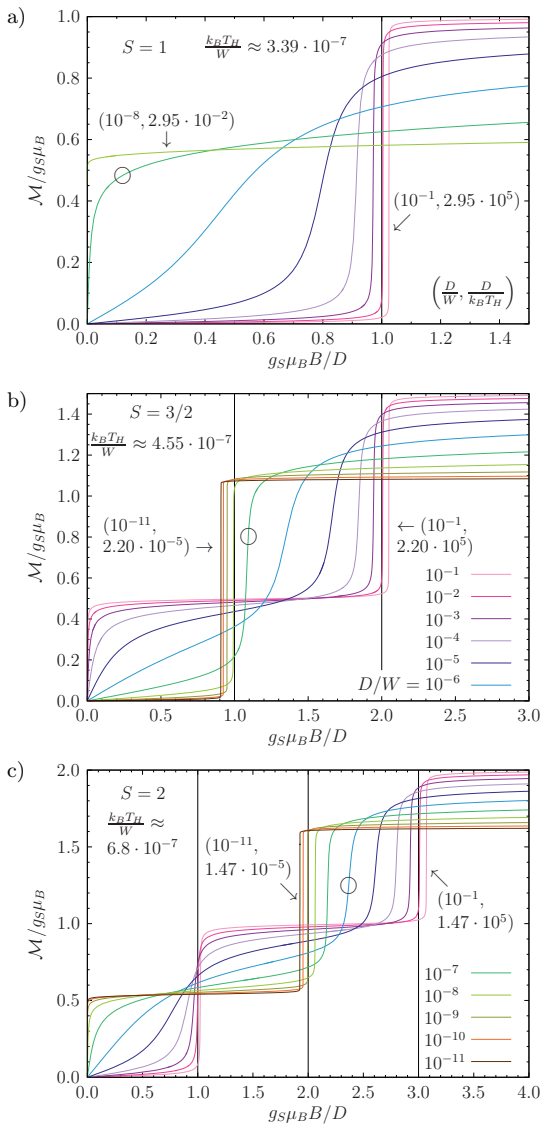

Figure 2: Impurity magnetization M for varying hard axis anisotropy D>0 as function of magnetic field for vanishing temperature, and impurity spin a) S=1, b) S=3/2, and c) S=2. The value of D increases from bottom to top or from left to right, respectively. For details see [3].

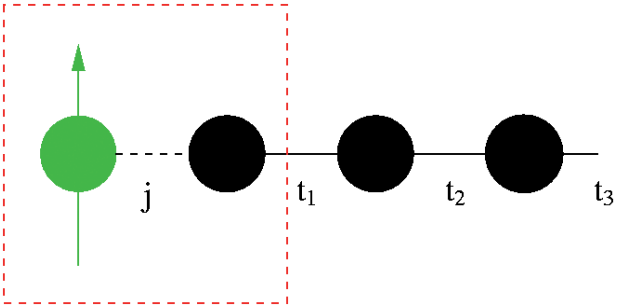

Figure 3: Sketch of the Wilson chain that is used for the evaluation of observables by means of NRG [4].

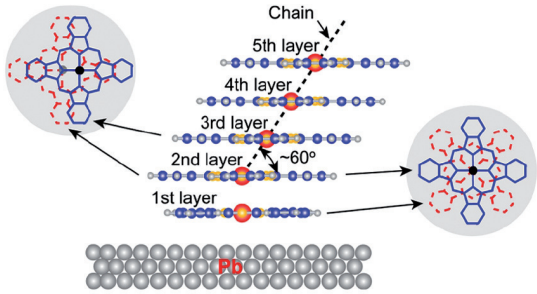

Figure 4: Schematic arrangement of stacked cobalt- phthalocyanin molecules on a lead surface. The cobalt ions of a stack form an effective spin chain. For details see [4].

tallic substrate. In our group [1] we study such models with non-zero magnetic field and additional uniaxial anisotropy for the impurity spin in the framework of the Numerical Renormalization Group (NRG) invented by Wilson [2].

Results and Methods

A first project dealt with the magnetic properties of deposited single spins with uniaxial anisotropy. The model contains the conduction electrons of a non-magnetic metallic substrate, which interact with a deposited spin, see Figure 1.

The antiferromagnetic interaction of spin and conduction electrons leads to screening effects, i.e. when the local magnetization of the spin is measured or calculated it is reduced compared to the non-interacting case. For spins larger than ½ and hard axis anisotropy this leads to non-trivial magnetization curves, see Figure 2 for an example.

The above discussed setting corresponds to a single-impurity single-channel problem. It is solved by mapping the physical Hamiltonian onto a semi-infinite chain, see Figure 3. The Hamiltonian of this so-called Wilson chain is then iteratively diagonalized in the respective Hilbert (Fock) spaces. We used openMP directives to adapt the program to the multicore environment of the nodes.

On-going Research / Outlook

Recent efforts aim at the deposition, investigation and manipulation of magnetic molecules with more than one magnetic center (spin). Figure 4 shows a successful realization of deposited short spin chains on a lead substrate [4].

In the experimental investigation the total magnetic moment of the effective spin chain of a single stack was measured for various stack sizes, i.e. numbers of spins in the chain. In our recent investigations [5] we aim at a theoretical determination of the magnetization as a function of field for small stack sizes. One interesting question concerns the critical value of the interaction between lowest spin and substrate, above which the lowest spin is fully screened and the chain is effectively reduced by one site.

Our continuing theoretical research is going to complement the huge efforts on the experimental side to design such small deposited quantum systems, e.g. as deposited magnetic molecules or as quantum dots.

References and Links

[1] http://obelix.physik.uni-bielefeld.de/~schnack
[2] K. G. Wilson. 1975. Rev. Mod. Phys. 47, 773
[3] M. Höck, J. Schnack. 2013. Phys. Rev. B 87, 184408
[4] Xi Chen, Ying-Shuang Fu, Shuai-Hua Ji, Tong Zhang, Peng Cheng, Xu-Cun Ma, Xiao-Long Zou, Wen-Hui Duan, Jin-Feng Jia, and Qi-Kun Xue. 2008. Phys. Rev. Lett. 101, 197208.
[5] H.-T. Langwald, F. Kaiser, J. Schnack. In preparation.

Advanced Finite-Temperature Lanczos Method for Magnetic Molecules

Research Institution
Universität Bielefeld, Fakultät für Physik

Principal Investigator
Jürgen Schnack

Researchers
Oliver Hanebaum, Christian Heesing, Jürgen Schnack

Project Partners
–

SuperMUC Project ID: pr63fa

Introduction

Quantum spin systems constitute an important platform to realize next-generation ultra-small magnetic data storage devices as well as devices for quantum computation and quantum simulation. The fundamental problem consists in the massive Hilbert space dimension of this quantum many-body problem. To provide an impression, a quantum spin system of *N* identical spins with spin quantum number s amounts to a Hilbert space dimension of $(2s+1)^N$. From the numerical point of view this leads to matrices that are quadratic in this dimension. These matrices need to be diagonalized in order to determine the quantum mechanical energy eigenvalues and thus thermodynamic properties. But nowadays computers, even supercomputers, allow only to diagonalize matrices of order 100,000x100,000 (or maybe somewhat bigger). This corresponds to only about 17 spins s=1/2. A simulation of the magnetic properties of a bigger molecule or an extended solid is thus impossible via direct matrix diagonalization.

Therefore, our group [1] employs powerful Krylov space methods in order to approximate thermodynamic observables such as magnetization or specific heat. Our experience is that our method of choice – the Finite-Temperature Lanczos Method (FTLM) [2] – delivers quasi exact quantum results which enable us to better understand magneto-chemical correlations in nanoscopic quantum magnets [3].

Results and Methods

Besides modeling a large amount of magnetic molecules that could potentially be used for magnetocaloric purposes, i.e. for designing new coolants for sub-Kelvin cooling, we addressed the problem of calculating the magnetic properties of a very large anisotropic molecule, Mn_{12}-acetate [4]. This very important (and famous) molecule, see Figure 1 for a schematic view, had never before been modeled in a realistic spin model that contained all 12 spins. A side result of our investigation was that we could for the first time quantify the accuracy of parameterizations that are derived from Density Functional Theory (DFT).

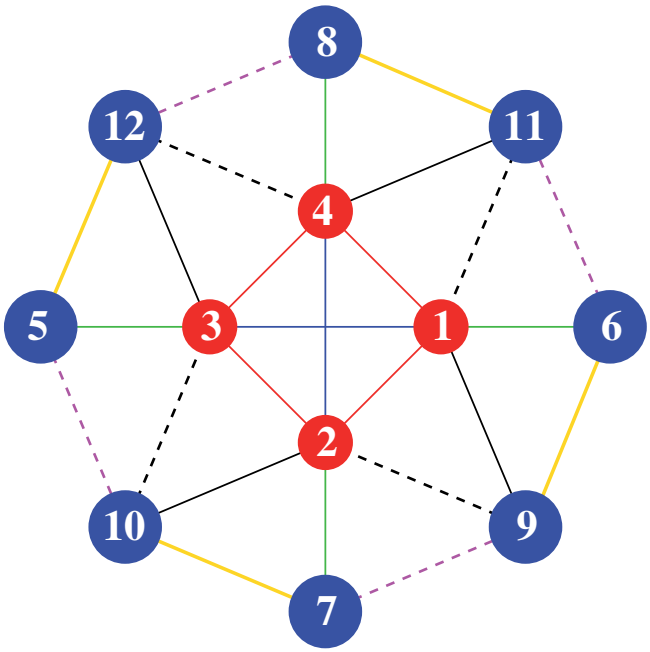

Figure 1: Schematic structure of Mn_{12}-acetate; $Mn^{(IV)}$ ions (1-4) are shown as red circles, $Mn^{(III)}$ ions (5-12) as blue ones. An S_4 symmetry of the molecule is assumed.

From the physics point of view we achieved two results: 35 years after its synthesis and 22 years after the first measurements of Mn12-acetate the Finite-Temperature Lanczos Method puts us in a position to evaluate thermodynamic functions of really large magnetic molecules. It thus complements DFT and other calculations of spin-Hamiltonian parameters for such big systems in so far that one can now use the DFT results in order to understand thermodynamic observables. DFT alone cannot deliver thermodynamic quantities.

In addition, one can now assess the quality of parameterizations for bigger spin systems. In the present case it turns out that although the range of addressed terms of

the spin-Hamiltonian is really impressive, the agreement with magnetic observables is not yet optimal, see Figure 2. It turns out that DFT underestimates the strength of some antiferromagnetic interactions as well as the strength of the anisotropy tensors of the $Mn^{(III)}$ ions. Our hope is that such investigations help to systematically improve DFT schemes for correlated electron systems with anisotropy.

For the evaluation of the magnetic observables we had to work in a Hilbert space of dimension 100,000,000. The FTLM as many Lanczos procedures works with three vectors of this length. It was modified to use only two vectors. In addition, the Hamiltonian which is a sparse matrix of dimension 100,000,000x100,000,000 was never stored. Its non-zero matrix elements were evaluated on the fly whenever needed. The huge effort arises from the fact that the procedure has to be executed for each magnetic field value (magnitude and direction) and a nearby field for calculating numerical derivatives.

We used a combination of MPI and openMP, where openMP was employed on each node to run the necessary matrix vector multiplications of the method and MPI was used to parallelize for the various magnetic fields.

In order to cover all parameter sets as well as the relevant temperature and field ranges, we had to use several hundred thousands of CPU hours. The numerical problem is so big, that a fit procedure, i.e. deriving the parameters of the spin Hamiltonian from experimental data is possible, but so far too expensive.

On-going Research / Outlook

The investigation of quantum many-body systems would be impossible without HPC support. This statement holds in general, and thus also in particular for correlated electron systems. The reason lies in the very nature of the problem, which is an eigenvalue problem (Schrödinger equation) in a large and potentially infinitely dimensional Hilbert space with non-trivial eigenfunctions.

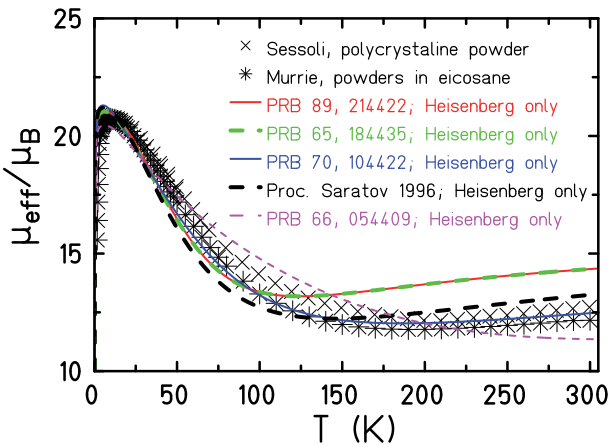

Figure 2: Effective magnetic moment of Mn12-acetate at B=0.1 T. Data of Sessoli and Murrie are given by symbols. Observables employing the Heisenberg part of parameterizations are displayed by curves. See [4] for details.

Our progress rests on two developments. On one hand we further advanced the Finite-Temperature Lanczos Method for anisotropic magnetic systems by sampling over pairs of time reversed random vectors instead of just random vectors. This improved the convergence for small sets of random vectors dramatically – and thus saves computing time. On the other hand we further improved our numerical schemes. We reduced the number of Lanczos vectors from three to two and restructured the program for a combined use of MPI and openMP.

We are convinced that this direction of research will be followed in the foreseeable future since quantum many-body systems are at the heart of many new quantum devices such as quantum computers and quantum simulators. In addition to static properties discussed in this paper we expect that time-dependent problems such as transport in low-dimensional systems will become progressively interesting in the near future. This could for instance include fundamental questions of thermal equilibration as well as spin transport phenomena currently investigated for spintronics applications. We definitely plan to continue our investigations and to start new projects along the outlined ideas.

References and Links

[1] http://obelix.physik.uni-bielefeld.de/~schnack
[2] J. Jaklič and P. Prelovsěk. 1994. Phys. Rev. B. 49, 5065–5068.
[3] (a) J. Schnack and O. Wendland. 2010. Eur. Phys. J. B 78, 535; (b) J. Schnack and C. Heesing. 2013. Eur. Phys. J. B 86, 46; (c) O. Hanebaum and J. Schnack. 2014. Eur. Phys. J. B 87, 194.
[4] O. Hanebaum and J. Schnack. 2015. Phys. Rev. B 92, 064424.

Study of complex microstructure evolution in ternary eutectic alloys with massive parallel large-scale phase-field simulations

Research Institution
Karlsruhe Institute of Technology, Karlsruhe University of Applied Sciences and Friedrich-Alexander University Erlangen-Nürnberg

Principal Investigator
Britta Nestler

Researchers
Johannes Hötzer, Marcus Jainta, Martin Bauer, Philipp Steinmetz, Michael Kellner, Harald Köstler, Ulrich Rüde, Britta Nestler

Project Partners
–

SuperMUC Project ID: pr84qi (Gauss Large Scale project)

Introduction

The development of optimized and new components for special applications requires high performance materials with defined properties. In ternary eutectic alloys, a wide range of different microstructures forms and allows to study the broad variety of the underlying physical processes. During the solidification of these alloys, a melt transforms into three solid phases at a defined temperature and concentration. The three solid phases arrange in various patterns, which strongly influence the mechanical properties. Due to the high technical importance, directional solidification experiments of ternary eutectics are in focus of current research. Especially the complex interplay of the process parameters such as temperature gradient and pulling speed is still not completely understood. Directional solidification experiments of the ternary eutectic system Al-Ag-Cu are even conducted at the international space station (ISS) as part of the SETA program to give insight into the pattern formation without the influence of convection induced by gravity [2]. Already today, high performance computing allows to treat such experiments in sufficiently large volume elements. The phase-field method has been established in the last two decades to simulate these kind of processes. By incorporating thermodynamic databases into the phase-field approach, different alloy systems with specific physical parameters and process conditions can be considered.

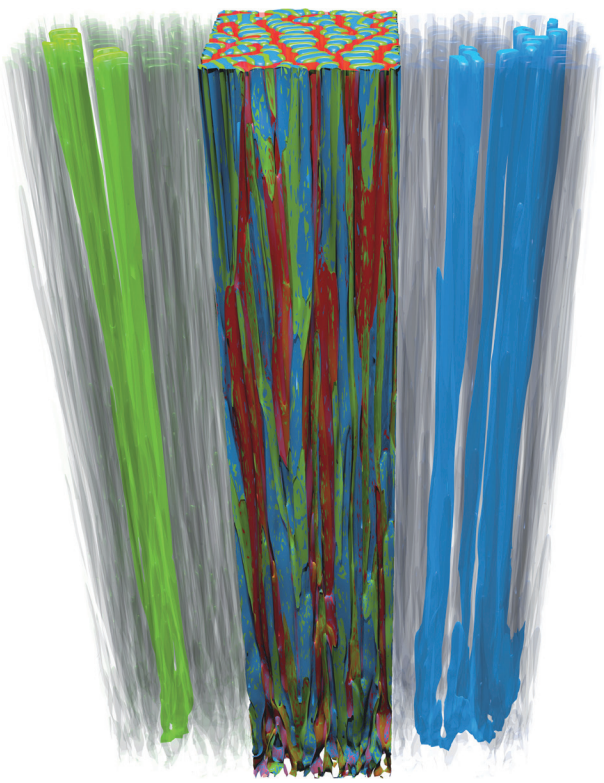

Figure 1: A 800x800x~4000 cell simulation of the directional solidified ternary eutectic system Al-Ag-Cu. In green and blue the evolving rods are exempted.

Results and Methods

For studying directional solidification of different ternary eutectics, a thermodynamic consistent phase-field model based on the grand potential approach is used [3]. This model is implemented in the highly parallel waLBerla framework [4]. The solver was optimized on various levels, like parameters, model, numerics, implementation, parallelization and explicit vectorization of the kernels. The systematic node level performance engineering results in 25% peak performance on a single SuperMUC node. The solver has proven to scale excellent on all three German Tier0/1 supercomputers: SuperMUC, JUQUEEN and Hazel Hen [5].

As initial setting, a randomized Voronoi tesselation for the nuclei in the melt is set at the bottom of the domain. To model an infinite domain, periodic boundary conditions are applied at the four side surfaces of the domain and Neumann boundary conditions are used at the top [6]. From the nuclei, the three distinct phases evolve dur-

ing the directional solidification in different patterns parallel to the growth front.

Ruggiero and Rutter predicted five different idealized patterns from experiments and geometrical considerations, in sections parallel to the solidification front [7]. Four of these five patterns are found in the high performance simulations [8].

By systematically varying the domain size in the plane perpendicular to the solidification direction, the influence on the pattern formation was investigated [9]. The variation induces a different lamellar spacing and hence leads to different undercoolings at the solidification front. A Jackson-Hunt analysis for a three dimensional hexagonal structure was conducted, showing qualitative accordance with the analytical predictions. An increase of the domain size reveals different aligned hexagonal structures which can not be observed in smaller domains. The results intend the necessity of large-scale simulations to capture pattern formation not effected by the domain boundaries. The diversity of microstructural ordering of the three phases of the ternary eutectic system Al-Ag-Cu was investigated in [6], using large scale simulations with a domain size of 800x800x~4000 cells on 13600 cores for ~16h. A good visual accordance with experimental micrographs for different process conditions was achieved. In Figure 1, a simulated 3D structure of the directional solidified ternary eutectic Al-Ag-Cu system is depicted showing a typical brick-like pattern as reported in experimental images.

A good quantitative comparison between the simulations and experimental micrographs of Al-Ag-Cu with principle component analysis based on two-point correlations is demonstrated in [10]. In this work, the necessity of large-scale simulations to establish large enough volume elements for the derivation of statistically meaningful conclusions is shown. Depending on the used parameter set and the arising patterns, a domain size parallel to the solidification front of at least 800x800 cells is required.

In an experimental cross-section of Al-Ag-Cu parallel to the growth direction, evidence of spiral growth was reported in [11]. Using large-scale simulations based on optimized parameters, this kind of spiral growth has been detected in computations for the first time [12]. Two spiraling rods are highlighted within the three dimensional structure in Figure 2. It is shown, that tilted growth in 2D as well as spiraling growth in 3D are driven by a defined spatial arrangement of the neighboring phases and the same physical parameters.

The introduced framework was further used to investigate the solidification behavior and microstructure formation in the high-temperature alloy Ni-AL-Cr.

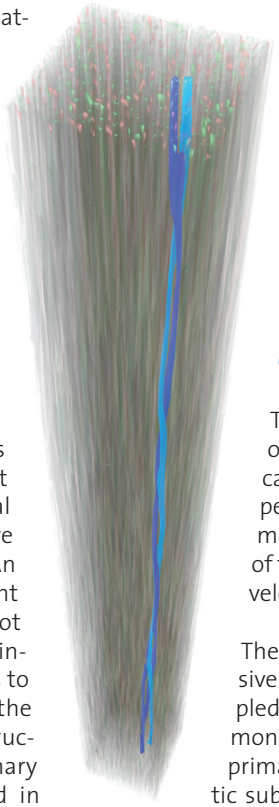

Figure 2: Spiraling growth within a simulation of a ternary eutectic system.

Highly optimized code and the computational power of current supercomputers like the SuperMUC, enable to investigate the microstructure evolution for various multiphysical process parameters and complex multicomponent material systems, analogue to experiments on the ISS.

On-going Research / Outlook

The ongoing research focuses on the influence of further process parameters, like the solidification velocity, the slope of the imprinted temperature gradient and the concentration of the melt for different systems. An implementation of the solver to exploit accelerator cards is in development.

The aim of forthcoming work is to perform massive multiscale simulations to explore the coupled dendritic and eutectic growth as it is commonly observed in experiments. The evolution of primary dendrites with a fine interdendritic eutectic substructure is an inherently multiscale microstructure and directly builds on the presented achievements of the pure eutectic growth.

References and Links

[1] http://www.gauss-centre.eu/gauss-centre/EN/Projects/MaterialsScienceChemistry/2015/hoetzer_TEDS.html

[2] http://www.dlr.de/dlr/desktopdefault.aspx/tabid-10337/1346_read-10047/#/gallery/14806

[3] Choudhury, A. & Nestler, B. (2012). Grand-potential formulation for multicomponent phase transformations combined with thin-interface asymptotics of the double-obstacle potential. Physical Review E, 85(2), 021602

[4] Feichtinger, C. et al. WaLBerla: HPC software design for computational engineering simulations. Journal of Computational Science, 2(2), 105-112. (2011).

[5] Bauer, M. et al. Massively parallel phase-field simulations for ternary eutectic directional solidification. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (p. 8). ACM. 2015

[6] Hötzer, J. et al. Large scale phase-field simulations of directional ternary eutectic solidification. Acta Materialia, 93, 194-204. (2015)

[7] Ruggiero, M. A. & Rutter, J. W. Origin of microstructure in the 332 K eutectic of the Bi-In-Sn system. Materials science and technology, 13(1), 5-11.(1997).

[8] Hötzer J. et al. Die Vielfalt der Musterbildung in Metallen. Horizonte, 45 (2015)

[9] Steinmetz et al. Large-scale phase-field simulations of ternary eutectic microstructure evolution. Submitted to Computational Materials Science

[10] Steinmetz, P. et al. Analytics for microstructure datasets produced by phase-field simulations. Acta Materialia, 103, 192-203. (2016).

[11] Genau, A. & Ratke, L. Morphological characterization of the Al–Ag–Cu ternary eutectic. International Journal of Materials Research, 103(4), 469-475. (2012).

[12] Hötzer J. et al. Phase-field simulations of spiral growth during directional ternary eutectic solidification. Acta Materialia, Volume 106, (2016)

The inverse problem of electronic structure

Research Institution
Martin-Luther University of Halle-Wittenberg

Principal Investigator
Miguel Marques

Researchers
Tiago Cerqueira

Project Partners
–

SuperMUC Project ID: pr84ra

Introduction

Due to continued advances in theory, scientific software, and high-performance computers, the calculation of the electronic band structure of materials is, by now, a routine problem. Starting from a given crystal structure we can easily obtain, often to a very good approximation, its band structure and all the properties derived from it. Unfortunately, the inverse problem is much more complicated, and led to one of the most exciting developments in condensed matter physics over the past years: materials design.

This new discipline aims at solving this inverse problem: given a certain desired property (or properties), discover (design) the material that possesses this property under a given set of constrains. These constrains can be related to the mechanical or chemical stability of the compounds, their price, their availability, etc. Several groups are actively working on this topic, with many projects related to energy materials including lithium batteries, photovoltaics, etc.

In this context we developed recently a technique based on genetic algorithms (GA), that can perform the inverse problem in an efficient way. This is done without resorting to databases or to any experimental input besides the periodic table of the elements and the laws of quantum mechanics. In a way, our approach fulfills the dream of any theoretician to be able to provide predictions with as little experimental information as possible (ideally none).

The main objective of this project was to prove the feasibility of this approach, by trying to invert the densities of states (a key property related to the electronic structure) of a few emblematic materials.

Results and Methods

Our approach is composed of three parts (see Fig. 1). We start with a list where each entry represents the chemical composition of a unit cell. For example, one of the entries could be "HHCaSbTe", which means that a unit cell contains two atoms of hydrogen and one of Ca, Sb, and Te. This initial list is completely random, and we only have to choose which chemical elements may enter our runs, and the maximum number of atoms allowed for each unit cell. Then, for each of the entries we perform a global structural prediction calculation [1] to obtain its ground-state crystal structure and several low-lying meta-stable states. We are currently using the minima hopping method [2] (MHM) for this task. These structures are then characterized, and its (electronic, mechanic, thermodynamic, etc.) properties are evaluated with density-functional theory [3, 4] (DFT) or related approaches. Finally, these values are passed to the (multi-objective) genetic algorithm (GA) that will determine the better-fit parents that will generate the offspring for the following generation (list of compositions). The cycle is then iterated until some convergence criterion is satisfied.

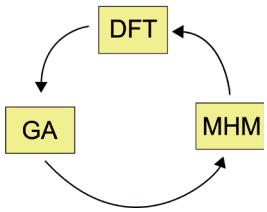

Fig. 1: Scheme of our unbiased optimization algorithm.

The key ingredient that allows us to perform the inverse problem on-the-fly, without the use of any pre-computed database, are state-of-the-art global structural prediction methods [1]. The mathematical problem of structural prediction is very easy to explain: given a certain composition of the system one wants to find the lowest energy (or enthalpy, or free-energy) crystal structure. There are several techniques used within this field, such as genetic algorithms, particle-swarm methods, random search, minima hopping, etc., that use different strategies to explore the minima of the potential energy surface. Beneath this machinery usually lies an ab initio DFT code that provides a reliable, efficient, and unbiased evaluation of energy and forces. It is true that structural prediction is a very computational intensive task, but with the recent advances in the algorithms, codes, and computers we can now explore the phase diagram as a function of pressure, temperature and composition for compounds containing up to 20-30 atoms in the unit cell.

Our approach is easily paralellizable, and scales extremely well with the number of computer cores. This is, in fact, one of the main advantages of the GA: the fitness function can be calculated independently for all individuals (material compositions) of a certain generation. On the other hand, the evaluation of the fitness function for a certain material is by far the limiting step of the approach, and it is limited to 16-24 cores due to scalability issues of the underlying DFT code. For around 100 individuals per generation we can therefore typically use up to 1600-2400 computer cores simultaneously.

Ideally, we would like to be able to give a target band-structure, and to find the material that possesses that band-structure. However, a band-structure is a complicated quantity that is defined inside the three-dimensional Brillouin zone. As the Brillouin zone depends on the unit-cell of the material and on its symmetry, it is not immediate to compare the whole band-structures of two different systems. Moreover, the band-structure usually has much more information than the one necessary to our problem.

We therefore decided to turn to a simpler quantity, namely the density of electronic states. This quantity includes much of the information of the band structure in an integrated form. Furthermore, as it is one-dimensional it is trivial to define a distance between two densities of state (such as the L^2 norm, or any similar metric), the basic ingredient to calculate the objective function for our genetic algorithms. Finally, it is trivial to generalize this distance to use the partial density of states.

Results for the evolution of the error when one tries to invert the density of states of MgB_2 are shown in Fig. 2. A typical run, such as the one depicted in the picture, requires about 2 million core-hours in SuperMUC. Storage requirements are relatively mild (less then 1 Tb), but a large number of files (tens of thousands) are generated in each run.

MgB_2 is the superconducting material with the highest transition temperature among conventional, electron-phonon driven superconductors at ambient pressure. Every column in the plot is a generation composed

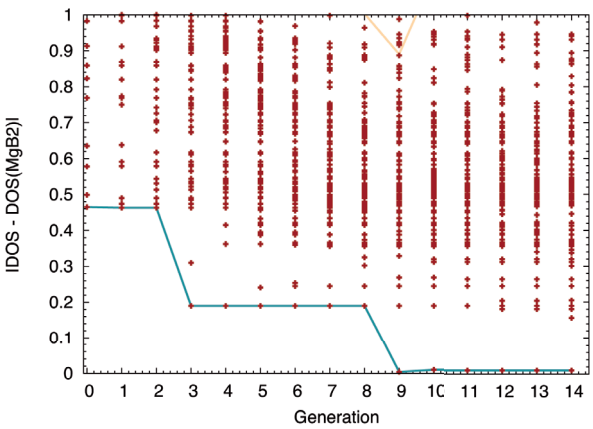

Fig. 2: Evolution of the error in the density of states for MgB_2.

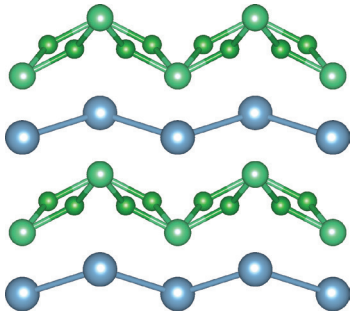

Fig. 3: Crystal structure of BeB_2Al.

of 100 different materials. Generation 0 was created with 100 random materials. As we can see, the evolutionary scheme we used quickly optimized the DOS. The target material MgB2 was found in generation 9. We emphasize that this result is highly non trivial: this is probably the first time that a general, practical solution for the inverse problem of electronic structure was proven.

Besides this result, we also obtained a series of materials with a DOS similar to MgB2, the most notable example being BeB_2Al (see Fig. 3). The similarities between the (partial) DOS of these two materials are indeed striking. Moreover, it seems that BeB_2Al is, like MgB_2, also a superconductor with a remarkably high transition temperature. This can be understood from the similarity between the Fermi surfaces of these two materials, that are mirrored in the respective DOS.

On-going Research / Outlook

SuperMUC was pivotal for our project, due to the large numerical cost involved that can only be provided by a large HPC facility. From the technical point of view, the largest problem we had to resolve was related to our unusual workflow, that involves several steps using multiple codes that had to be adapted to the SuperMUC configuration. Another issue was related to the very large number of small files generated (while the SuperMUC filesystem is optimized to few, large files) and that required the consolidation of the output.

After having proved that our method is capable of inverting the density of states of an existing material, the next step is to engineer a material with an artificial, human created, density of electronic states. This is a challenge that could open new ways for the optimization and the discovery of new materials with improved technological properties.

References and Links

[1] A. R. Oganov (ed.). 2010. Modern Methods of Crystal Structure Prediction. Wiley-VCH, Berlin.
[2] M. Amsler and S. Goedecker. 2010. Crystal structure prediction using the minima hopping method. J. Chem. Phys., 133, 224104. DOI: 10.1063/1.3512900.
[3] C. Fiolhais, F. Nogueira, and M. Marques (eds.). 2003. A Primer in density functional theory. Springer, Berlin.
[4] G. Kresse and J. Furthmüller. 1996. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. Phys. Rev. B, 54, 11169. DOI: 10.1103/PhysRevB.54.11169

The chemistry of porphyrins adsorbed on metallic surfaces

Research Institution
Department of Molecular Nanoscience and Chemical Physics of Interfaces

Principal Investigator
Wilhelm Auwärter

Researchers
Marie-Laure Bocquet, Tangui Le Bahers, Marie Lattelais, Torsten Houwaart

Project Partners
Ecole Normale Supérieure de Lyon and Massachusetts Institute of Technology.

SuperMUC Project ID: pr85la

Introduction

The mechanism for binding oxygen to metalloporphyrins is a vital process for oxygen breathing organisms. Understanding how small gas molecules are chemically bound to the metal complex is also important in catalysis or the implementation of chemical sensors. For investigating these binding mechanisms, porphyrin rings with a central cobalt or iron atom are supported on copper or silver surfaces, leading to non-planar deformation – a saddle geometry. In 2011, the collaborative team in ENS Lyon and TUM have deciphered, how the poisonous carbon monoxide gas attaches to one type of supported metal-porphyrin, by means of detailed Scanning Tunneling Microscopy experiments complemented by density functional theory calculations [1]. Indeed, two CO molecules dock between the central metallic atom and the two opposite nitrogen atoms of the macrocycle (see figure 1) and this dual ligation had never been observed elsewhere.

The initial idea of the simulation project was to tentatively generalize at the single molecule level an anomalous binding mode for CO molecules – called the *rider mode* – onto cobalt porphyrins adsorbed on well-defined surfaces. The straightforward generalization was to test the binding of other gaseous molecules like NO and O_2 or solid particles like Fe adatoms, to try other metal porphyrins like Fe porphyrin and to model other "weaker" but planar substrates like epitaxial graphene on Cu and graphene on Ir.

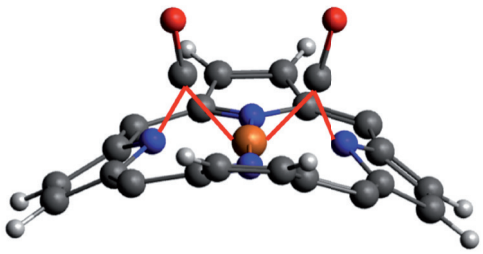

Figure 1: DFT model of dicarbonyl bonding at saddle-shaped cobalt tetraphenyl-porphyrins anchored on well-defined coinage metal substrates like Cu and Ag. The Ag(111) surface considered here and the phenyl substituents have been removed for clarity.

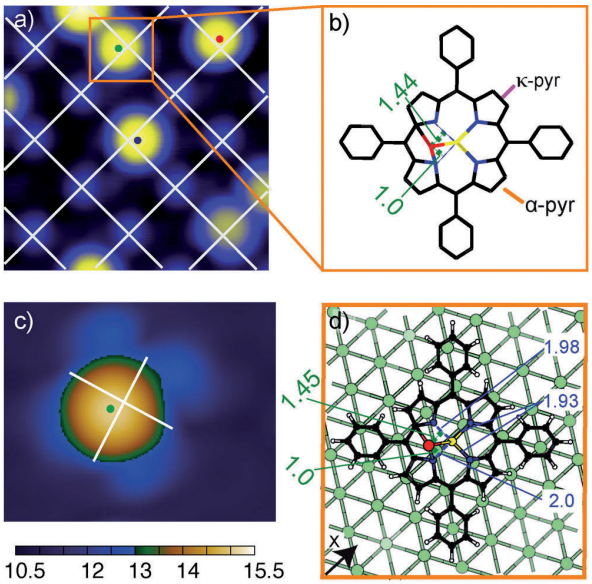

Figure 2: a) High-resolution STM image revealing a minute off-center adsorption of the Fe atoms on the porphyrin lattice. (I = 0.1 nA, V_b = -0.1 V, 55.4 Å x 55.4 Å). b) Schematic model extracted from the STM data. c) Simulated constant-current STM image based on the structure shown in d) (Vb = -0.2 V, 32 Å x 25 Å). The given distances are relative to the bare Ag substrate and the overall corrugation amounts to 5 Å. d) DFT optimized Fe/Co-TPP complex on Ag(111) in a 11x5sqrt(3) cell. The four Co-N distances are quoted. Green, black, blue, white, yellow and red balls depict respectively Ag, C, N, H, Co and Fe atoms.

Multiple tests have been simulated within the computer grant and a few of them have failed like the binding of O2 and CO molecules on FeTPP and the study of FeP on graphene on Cu(111). But some adjustments in the project have resulted in important findings and we will describe them in the following.

Results

Fe adatom – CoTPP (see Ref [2]):
In conjunction with high-resolution STM images acquired after Fe deposition (TUM in Auwärter's & Barth's groups), an Fe adatom is shown to bind in a bisector position from two Co-N directions (see figure 2). 4 positions are equivalent in the deformed porphyrin macrocycle and Fe is switching easily between these four equivalent sites in the macrocycle (fluxional ligand even at low temperature). We have estimated the activation energy barrier for the two diffusion paths between equivalent adsorption sites through Nudged Elastic Band (NEB) calculations. In addition DFT shows that the Fe adatom keeps part of its magnetism and even induces some magnetism in the Co center. This finding fits nicely to XMCD measurements (collaboration with MPI, Stuttgart) that evidence a ferromagnetic coupling between the Co central metal and the Fe ligand atoms. [2]

From FeP on graphene/Cu(111) to cycloadducts of FeP on graphene/Ir(111)

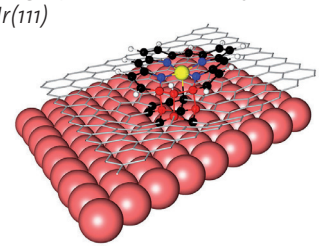

Figure 3: an exothermic cycloadduct from FeP on graphene/Ir as predicted by DFT. Upon cycloaddition, the graphene area is bending down to the Ir surface making strong covalent bonds.

The theoretical study of the physisorption of FeP on graphene and copper supported graphene (lattice-matched, atop-fcc) showed that for both surfaces the molecule-substrate interaction is weak. The adsorption energy of FeP is about 0.8-0.9 eV on freestanding graphene, where as it is 0.5 eV on copper supported graphene, when dispersion is accounted for the top copper layer. Similarly the adsorption distance is respectively 3.3-3.4 Å on pristine graphene and 3.1-3.3 Å on graphene-covered metal.

Because of the weakness of the interaction on copper supported graphene the porphine molecule will be too mobile to be properly investigated by low temperature STM. This is one obstacle to propose this system for an STM investigation. Hence we propose to covalently anchor Fe porphine to epitaxial graphene via concerted reversible cycloadditions. A single cycloadduct was found on graphene/Cu but was located at +2 eV above the initial porphine reactant. We have then searched for a more active metallic substrate and we found that both Re and Ir substrate under graphene do activate the cycloaddition process. And on graphene on Ir an exothermic cycloadduct has been even identified (see the structure on Figure 3). These theoretical results published in J. Phys. Chem C. (Ref. [3]) serve now as a basis for future STM investigations.

Side project with the STM group at TUM (Klappenberger&Barth): a new kind of organo complex on Ag(111)

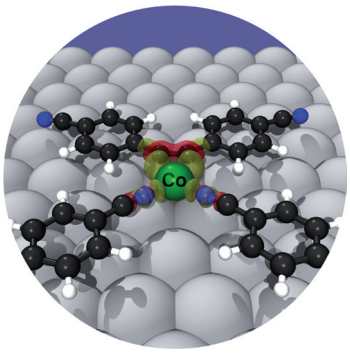

Figure 4: the organo Co complex formed on Ag after annealing as predicted by DFT. Shaded areas mark the electron transfer between the ligands and the Co adatom.

Our DFT calculations have helped to decipher the formation of a new complex on a Ag(111) surface resolved by high resolution STM and XPS measurements. By sequentially depositing, first a ethyne molecule substituted in each end by one benzonitrile group, second Co adatoms and third by annealing the surface, a new phase forms. We have demonstrated that this new complex results from the association of three organic molecules around the Co atom, two molecules remaining intact and bonded via their ending N atoms and one molecule being bent and bonded via its central alkyne unit. This work was recently published Angewandte Chemie (Ref. [4]).

Conclusion

In this project we have dealt with very complex molecules (metal porphyrins and organo complex) on either metallic surfaces or graphene-covered metallic surfaces. In addition to the adsorption of these molecules we have studied their reactivity (ligation of Fe adatoms or cycloaddition to the substrate). The project was very challenging knowing the fact that these entities carry a molecular spin. The grant has permitted to envisage numerous ways for the reactivity that could be also probed by STM investigation. We are convinced that the cycloaddition results, yet only DFT findings, will receive a large echo in the graphene community aiming at the chemical functionalization of graphene.

References

[1] "Cis-dicarbonyl binding at cobalt and iron porphyrins with saddle-shape conformation", K. Seufert, M.L. Bocquet, W. Auwärter, A. Weber-Bargioni, et al. Nature Chemistry, 3 (2011) 114-119.

[2] "Restoring the Co Magnetic Moments at Interfacial Co-Porphyrin Arrays by Site-Selective Uptake of Iron", S. Vijayaraghavan, W. Auwärter, D. Ecija, K. Seufert, S. Rusponi, T. Houwaart, P. Sautet, M.-L. Bocquet, P. Thakur, S. Stepanow, U. Schlickum, M. Etzkorn, H. Brune, J.V. Barth, ACS Nano, 9 (2015) 3605-3616.

[3] "Cycloaddition of Metal Porphines on Metal-Supported Graphene: A Computational Study", M. Lattelais, M.L. Bocquet, , Journal of Physical Chemistry C, 119 (2015) 9234-9241.

[4] "Surface-guided formation of an organocobalt complex », Peter B. Weber, Raphael Hellwig, Tobias Paintner, Marie Lattelais, Mateusz Paszkiewicz, Pablo Casado Aguilar, Peter S. Deimel, Yuanyuan Guo, Yi-Qi- Zhang, Francesco Allegretti, Anthoula C. Papageorgiou, Joachim Reichert, Svetlana Klyatskaya, Mario Ruben, Johannes V. Barth, Marie-Laure Bocquet*, and Florian Klappenberger*, Angewandte Chemie, 128 (2016) 5848-5853.

Adsorption of molecules in *giant* Metal-Organic Frameworks: An *ab initio* study

Research Institution
Department of Chemistry, University of Torino, Italy

Principal Investigator
Bartolomeo Civalleri

Researchers
Maddalena D'Amore, Roberto Orlando, Silvia Casassa, Piero Ugliengo, Guillaume Maurin

Project Partners
Institut Charles Gerhardt, Université de Montpellier 2, France

SuperMUC Project ID: pr85qu (PRACE project)

Introduction

Metal-organic frameworks (MOFs) have recently attracted a lot of attention for their high porosity due to the crystalline structure which shows different topologies (cages, channels, ...) and chemical versatility that allows one to easily modify their surface properties [1]. Their adsorption capacity can then be tailored to enhance their performance in the capture, separation and storage of small molecules, like CO_2 [2] as well as for advanced applications such as drug delivery [3].

Although, ab initio modeling offers a powerful tool to shed some light on the interaction between molecules and MOFs, at present, the target has been limited to small-to-medium size frameworks. In this project, we aimed at making a step forward by investigating the adsorptive capacity of the so-called giant MOFs, also known as mesoporous MOFs. The most representative one is probably MIL-100 [4] (see Figure 1) which is comprised of trimeric units of a trivalent metal octahedrally connected with 1,3,5-BTC (benzene-1,3,5-tricarboxylic acid). MIL-100 has cages of different dimensions from micro- (6.5 Å) to mesopores (25-30 Å) and 2788 atoms in the primitive cell. It has a surface area of 3340 m2/g, which is 3 times larger than the values measured for the MCM-41 inorganic mesoporous materials. MIL-100 is characterized by the presence of a large number (i.e. more than 200) of coordinatively unsaturated metal atoms that are exposed at the inner surface of the pores. Because of the huge size of the unit cell (74 Å), it then represents a tremendous challenge for current ab initio calculations.

In the project, the adsorption of CO_2 in MIL-100 with different metals (i.e. Al, Sc, Cr, and Fe) has been studied by considering the metal atom as the primary adsorption site. Nowadays, the removal of carbon dioxide from environment is very important to reduce its concentration in the atmosphere. In this respect, porous materials like MOFs have been identified as effective adsorbent to capture CO_2. In particular, MIL-100(Cr) have been shown to have one of the highest CO_2 adsorption enthalpy.

Results and Methods

The examined system is very large and very challenging from the computational point of view. Calculations were carried out within a periodic ab initio approach by using atom-centered Gaussian-type basis functions and hybrid HF/DFT theoretical methods. In particular, the B3LYP functional augmented with an empirical correction to include dispersion interactions (i.e. +D*) in combination triple-zeta quality basis sets (up to 50000 basis functions) were used. All calculations were performed with the massive parallel version of the CRYSTAL code [5,6]. On average, calculations were run on 1000-2000 cores.

Main results of the project include: (i) the prediction of the basic properties (structure, electrostatic potential, spin density) of the examined system with four different metals, namely: Al, Sc, Cr and Fe; (ii) the study of the interaction with CO_2 on different adsorption sites and coverages.

First, we have investigated structural features of the isoreticular MIL-100(M) family (M=Al, Sc, Cr and Fe). Although the ideal crystal structure of MIL-100 is cubic, a lowering

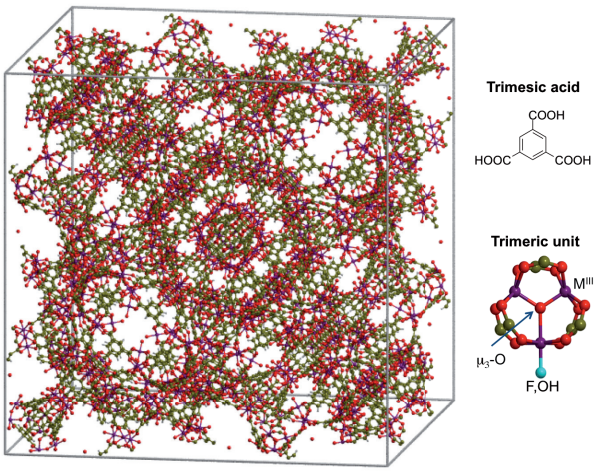

Figure 1: Unit cell of MIL-100 (on the left) along with the two secondary building units (i.e. organic linker and inorganic cluster): the trimesic acid and trimeric unit (on the right).

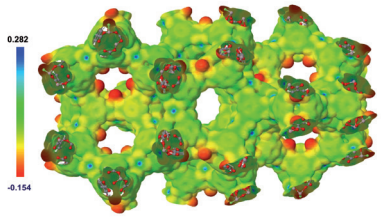

Figure 2: Electrostatic potential of MIL-100(Al) mapped on top of a charge density isosurface (0.003 e). Open metal sites are clearly visible as blue (positively charged) spots.

of the symmetry from cubic to tetragonal had to be taken into account since a counterion must be included to neutralize the positively charged trimeric unit (see Figure 1). Calculations were still affordable, but more costly than expected at the proposal submission stage. According to experiment, either fluorine atoms or hydroxyl groups can be used as counterions. For simplicity, we decided to model MIL-100(M) with fluorine. We have then optimized the structure of MIL-100 with M=Al, Sc, Cr and Fe. For all systems, the deviation from a cubic lattice is very small with the c/a ratio ranging from 0.9995 for Al to 1.0090 for Cr. Computed lattice parameters nicely agree with experimental data, even though for MIL-100(Al) and MIL-100(Sc) they refer to MIL-100(M) with OH groups as counterions. For MIL-100(Al), the computed average lattice parameter (<*a*>) is 71.52 Å to be compared with 71.69 Å, while, for MIL-100(Sc), <*a*>=74.73 Å vs 74.35 Å.

In MIL-100(Cr) and MIL-100(Fe), the two transition metals present unpaired electrons. Different spin state configurations have been investigated for the inorganic secondary building unit. For MIL-100(Fe), available experimental data show that iron is in a high spin state and there is an antiferromagnetic coupling among the atoms. Results for two spin state configurations show that, indeed, a ferromagnetic interaction that tends to align the spin vectors in a parallel fashion is disfavored with respect to a ferrimagnetic solution in which an antiferromagnetic interaction is present. For MIL-100(Cr) a similar ferrimagnetic solution has been found as the most stable. Predicted structures for MIL-100(Cr) and MIL-100(Fe) are also in very good agreement with experimental data.

Secondly, on the basis of the optimized structures of MIL-100(M=Al,Sc,Cr,Fe), the study of the adsorption of CO_2 has been carried out. By symmetry, nine non-equivalent open metal sites are present in the framework (i.e. 136 sites in the primitive unit cell). A close inspection of the structure shows that they are not exposed at the inner surface of giant cages, rather they are located in the channels (Figure 2). Two different adsorption sites have been investigated that correspond to two different coverages, namely: a first site that shows the smallest loading (i.e. 6%) of adsorbed molecules in the unit cell (i.e. 8 molecules) and a second site that permits to adsorb 16 molecules in the unit cell (loading: 12%). Results show that adsorption sites are rather homogeneous while experimental data suggest some heterogeneity. Computed interaction energies are underestimated with respect to experimental values but they confirm that the interaction of CO_2 decreases along the series: Cr > Sc > Fe > Al (Figure 3). Underestimation of the adsorption energies could be explained by the presence of defects in the real structure of MIL-100.

On-going Research / Outlook

Several open issues remain to be addressed as the use of OH groups as counterions instead of fluorine or improving the quality of computed adsorption energies through a multilevel approach that combines periodic and cluster calculations with different methods of increasing accuracy. Even more interesting, it would be to understand the role of defects, as lower coordinated metal sites, on the adsorption capacity of MIL-100.

Beyond the fundamental interest to investigate this giant MOF in interaction with small molecules as CO_2, there is a considerable interest in the adsorption of drugs for application in drug delivery [3]. This is an even more challenging task. By starting with the optimized structure of MIL-100(Fe), we tried to model the adsorption of Busulfan, an antitumoral. Different arrangements of the molecule in interaction with the metal sites were explored to dock it in the pores. Unfortunately, we could not obtain conclusive results. It would then be interesting to complete such a study in next projects. In addition, the water-drug competition would deserve to be explored, because the drug must remove water molecules before interacting with the metal sites. In this case, tier-0 supercomputing resources, as the SuperMUC Next Generation, would be the only way to tackle such calculations at an ab initio level.

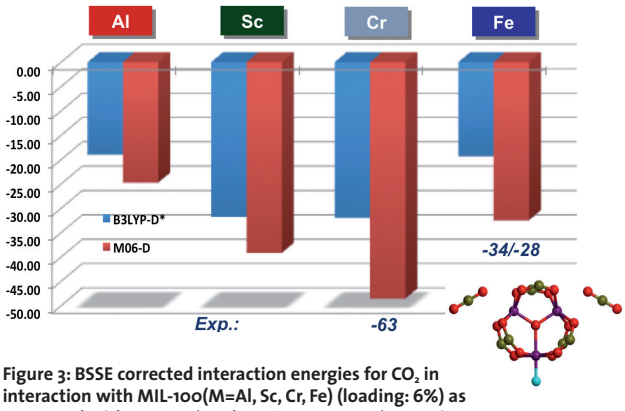

Figure 3: BSSE corrected interaction energies for CO_2 in interaction with MIL-100(M=Al, Sc, Cr, Fe) (loading: 6%) as computed with B3LYP-D* and M06-D//B3LYP-D*. Experimental data are reported for comparison.

References and Links

[1] Special issue on "Metal-Organic Frameworks". 2012. Guest Editors: H.-C. Zhou, J.R. Long and O.M. Yaghi. Chem. Rev. 112, 673–1268.
[2] K. Sumida, et al. 2012. Carbon Dioxide Capture in Metal-Organic Frameworks. Chem. Rev. 112, 724.
[3] P. Horcajada, et al. 2010. Porous metal-organic framework nanoscale carriers as a potential platform for drug delivery and imaging, Nature Mater. 9, 172.
[4] G. Férey, et al. 2004. A Hybrid Solid with Giant Pores Prepared by a Combination of Targeted Chemistry, Simulation, and Powder Diffraction, Angew. Chem., Int. Ed. 43, 6296.
[5] R. Dovesi, R. Orlando, A. Erba, C.M. Zicovich-Wilson, B. Civalleri, et al. 2014. CRYSTAL14: A Program for the Ab Initio Investigation of Crystalline Solids. Int. J. Quantum Chem. 114, 1287-1317.
[6] http://www.crystal.unito.it

"Hot" adatoms hopping: Equilibration dynamics from first-principles

Research Institution
Lehrstuhl für Theoretische Chemie, TU München

Principal Investigator
Karsten Reuter

Researchers
Vanessa Jane Bukas, Jörg Meyer, Christoph Scheurer

Project Partners
–

SuperMUC Project ID: pr85wa (Gauss Large Scale project)

Introduction

Due to its central role in technologically relevant processes such as heterogeneous catalysis, sensing, corrosion or epitaxy, the adsorption and dissociation of molecules on solid surfaces has been a most active field of research for many decades. When breaking things down to elementary processes on the atomic scale, however, our understanding of such phenomena remains rather limited. At present this receives even further stimulus from the basic energy science perspective which adds fundamental questions like the conversion of energy forms at interfaces to the agenda.

A particular example is the conversion of chemical energy to heat which arises as a consequence of a molecule's exothermic interaction with the surface and may easily amount to several electron volts of energy. Though sizable in view of potential microscopic dissipation channels, the prevalent picture in models of chemical kinetics is that this energy is quasi-instanteneously thermalized, ultimately into substrate lattice vibrations, i.e. phonons. Quite on the contrary, for a prototypical model reaction like the dissociative adsorption of oxygen molecules scanning-tunneling microscopy has suggested the formation of so-called "hot" adatoms on several metal surfaces. The latter refers to dissociation products characterized by a non-thermal transient mobility and thus strongly alludes to delayed heat dissipation of the initially released chemical energy.

As the experimental quest to generate molecular movies of such reactions is still ongoing, theory has been challenged to elucidate a full picture of the equilibration dynamics through predictive-quality simulations based on computationally demanding quantum-mechanical (QM) treatments: Here, the breaking and making of chemical bonds dictates high-level theoretical approaches such as Density Functional Theory (DFT) to accurately describe the adsorbate-surface chemical interaction. This already poses a considerable computational challenge for evaluating kinetic quantities from molecular dynamics (MD) simulations that require sufficient statistical averaging. Nevertheless, continuous increases in computer power and the development of efficient algorithms for electronic structure calculations are nowadays rendering such treatments increasingly accessible for periodic simulation cells containing already several tens of metal (Me) atoms. Complimenting this with a quantitative first-principles account of substrate phonons, however, remains an elusive hallmark and has been cited as one of the major conceptual problems in contemporary gas-surface dynamical modeling.

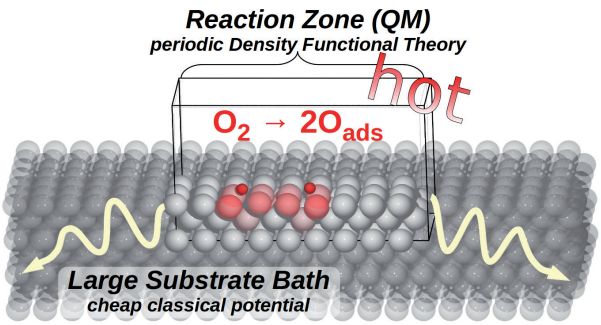

Figure 1: Schematic illustration of the QM/Me embedding scheme: Energy is dissipated out of a quantum-mechanically (QM) described "hot" reaction zone and into the extended heat bath of the metal (Me) substrate which is efficiently modeled by a classical interatomic potential.

Given this situation, in the present project [1] we employ a novel embedding scheme for metallic substrates (QM/Me [2]) that allows for energy to be dissipated out of a DFT-described reaction zone and into a computationally undemanding extended bath. The latter accounts for long-range phonon propagation which is reliably provided by a classical interatomic potential (CIP), thus coupling the localized QM-described chemical interaction with heat dissipation at the larger-scale. We apply QM/Me through large-scale MD simulations to investigate the equilibration of prominent catalytic processes, comparing to experimental studies and providing trend understanding into the influence of surface symmetry. Additionally, we advance QM/Me to include the concurrent treatment of multiple reaction zones centered around individual adatoms and dynamically following their motion. Standing

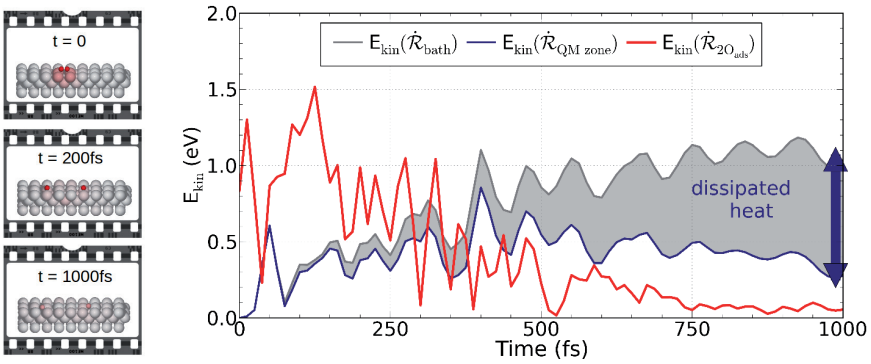

Figure 2: Left: Snapshots along a dynamical trajectory depicting the motion of "hot" oxygen adatoms on a Pd(100) surface (color-coded to kinetic energy). Right: Profiles of kinetic energy as a function of time. Note that already within 1 picosecond after the initial O_2 bond dissociation, phonons have propagated the larger fraction of the initially released chemical energy outside the QM-described reaction zone and into the extended substrate bath.

at the forefront of multi-scale materials' modeling, we thereby provide a new first-principles perspective into the intriguing propositions brought forth by such "hot chemistry", ultimately extending into fundamentals of heat management and sustainability at the larger scale.

Results and Methods

Classical MD simulations were performed according to the employed embedding scheme by interfacing calculations performed independently on the QM and classical levels of theory. Complicating the problem further, the novel QM/Me approach separates the chemical and elastic contributions in the QM interaction potential and thus dictates the computation of forces to be performed twice for each time-step (~1fs) within the considered embedded region(s): once with and once without the presence of the adsorbate. Relying on a generalized gradient approximation (GGA) of DFT, these calculations were performed within an all-electron basis set and dominated in terms of computational demand: Typical jobs of the "general" class were parallelized over 1,536 cores to propagate the system over ca. 50 fs. Overall, long trajectories of several ps were evaluated for different sets of initial conditions and for each of the investigated systems, thus consuming the larger fraction of the total ca. 20 mio CPUhs used (incl. extensive testing of the implementation). A many-body CIP was employed to (serially) evaluate an extended heat bath of 125,000 Me atoms which led to the generation of large output files (of ca. 100MB per 50fs) to contain all necessary phase-space trajectory information. The employed QM (Fortran) and CIP (C++) codes are listed in Ref. [3] and were dynamically coupled through an efficient in-house Python interface.

In the application to O_2 dissociation over Pd(100) we first establish the relevance of phononic dissipation through dynamical simulations that rely on QM forces (only) within a rigid slab model and augmented with various effective accounts of surface mobility [4]. Focusing in a second step on the microscopic details of equilibration, QM/Me predicts "hot" adatoms traveling ballistically over several lattice constants as a consequence of non-immediate energy transfer to the underlying substrate. The interfacial conversion of energy and its dissipation to the bulk occurs within a few picoseconds after the initial O_2 bond dissociation, thus indicating that thermalization is not instantaneous on the time-scale of the elementary process itself (cf. Fig.2) and clearly influences the actual adsorbate dynamics. The ensuing transient mobility thus intricately couples the elementary reaction steps of dissociation and diffusion; a notion hitherto not considered in prevalent kinetic models in catalysis.

Despite being similarly exothermic, a much shorter transient mobility has been measured experimentally for the O_2/Pd(111) reaction. We explain the small net displacement here by revealing randomized trajectories of hyperthermal hops between neighboring adsorption sites due to scattering from the strongly corrugated QM potential, thereby invalidating the simplistic physical picture implying a minimal transient mobility for small adatom separations. In both the Pd(100) and Pd(111) systems, however, a quantitative analysis of the underlying phonon excitations identifies certain groups of localized surface modes as the dominant dissipation channels which questions prevalent assumptions about energy sinks made in commonly used model Hamiltonians.

Outlook

High-end HPC resources enabled the formidable task of obtaining the required QM forces at each time-step of our MD simulations. Relying on the efficient parallel performance of the ELPA eigenvalue solver and SuperMUC's limited communication overhead (optimized in Phase 2), we fully benefited from scaling speed-up through massive parallelization over several hundreds of CPUs. Overall, the newly established methodological tools now finally set the foreground for a detailed atomistic understanding of the influence of surface temperature on the underlying phonon dynamics through extensive statistical sampling; an, overall, ideal follow-up project for "SuperMUC Next Generation".

References and Links

[1] https://www.elitenetzwerk.bayern.de/elitenetzwerk-home/aktuelles/meldungen/2015/enb-juli-2015/lindau-vortrag-bukas/

[2] Jörg Meyer and Karsten Reuter. 2014. Modeling Heat Dissipation at the Nanoscale: An Embedding Approach for Chemical Reaction Dynamics on Metal Surfaces. Angew. Chem. Int. Edit. 53, 18 (May 2014), 4721-4724. DOI:http://dx.doi.org/10.1002/anie.201400066

[3] Volker Blum et al. 2009. Ab initio molecular simulations with numeric atom-centered orbitals. Comp. Phys. Comm. 180 (June 2009), 2175 -2196. DOI=http://dx.doi.org/10.1016/j.cpc.2009.06.022; Steve Plimpton. 1995. Fast Parallel Algorithms for Short-Range Molecular Dynamics, J. Chem. Phys. 117, 1 (March 1995), 1-19. DOI:http://dx.doi.org/10.1006/jcph.1995.1039

[4] Vanessa J. Bukas, Shubhrajyoti Mitra, Jörg Meyer and Karsten Reuter. 2015. Fingerprints of energy dissipation for exothermic surface chemical reactions: O_2 on Pd(100). J. Chem. Phys. 143, 034705 (July 2015) 8 pages. DOI=http://dx.doi.org/10.1063/1.4926989

Shape and Catalytic Mechanism of RuO_2 Particles at CO Oxidation Reaction Conditions: First-Principles Based Multi-Scale Modeling

Research Institution
Lehrstuhl für Theoretische Chemie, TU München

Principal Investigator
Karsten Reuter, Lehrstuhl für Theoretische Chemie, TU München

Researchers
Christoph Scheurer, Lehrstuhl für Theoretische Chemie, TU München,
Tongyu Wang, Lehrstuhl für Theoretische Chemie, TU München

Project Partners
–

SuperMUC Project ID: pr86de

Introduction

For model catalyst studies on low-index single-crystal surfaces close agreement between detailed measurements and quantitative microkinetic modeling can increasingly be achieved. However, for 'real' catalyst particles, such structure–morphology–activity relationships are only scarcely established. This is prototypically reflected by the situation for RuO_2, as a most active catalyst for CO oxidation. Here, existing first-principles kinetic modeling is restricted to just one facet, namely the RuO_2(110) surface, which is not able to fully account for activity data obtained from polycrystalline RuO_2 powder catalysts. The overarching objective of this project was correspondingly to close this gap and demonstrate that similarly close agreement as for individual single-crystal model catalysts can also be achieved for catalyst particles. Specifically, we addressed experiments where an intact RuO_2 bulk structure is conserved, and establish the atomic-scale structure and reactivity of other RuO_2 low-index facets under the gas-phase conditions characteristic for catalytic CO oxidation.

Results and Methods

As a starting point we employed ab initio Wulff construction to obtain a first thermodynamic estimate of the crystal habit at different feed conditions. Overall, the simulation results fit the existing experimental data quite well, whereby remaining deviations can be ascribed to complex surface reconstructions not considered in this first approach. We therefore extended the project specifically to obtain a detailed structural model of the c(2x2)-RuO_2(100) reconstruction using a basin hopping approach. Simultaneously, we established first-principles kinetic Monte Carlo (1p-kMC) models to address kinetic effects on those facets identified by our Wulff construction study as dominating the habit at oxidizing conditions, i.e., RuO_2(111) facets. Here we provide further details on the three strands of activity pursued in the project:

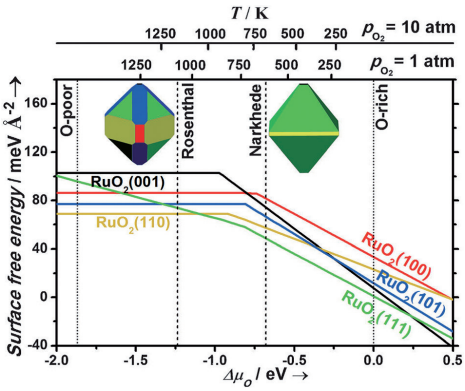

Figure 1. Minimized surface free energies for the five low-index RuO_2 surfaces and the corresponding particle shapes obtained from the constrained Wulff construction, in which the colours of the different facets match those of the corresponding surface free energy lines.

Ab Initio Wulff Construction for RuO_2 Nanoparticles under Oxidising Conditions [1]:
We pursued a DFT-based ab initio thermodynamic approach to obtain the surface structure and composition of single-crystal low-index RuO_2 surfaces in contact with O_2. Combining the corresponding surface free energies within a Wulff construction then yielded particle shapes with explicit dependence on temperature and reactant partial pressures (shown in Figure 1).

The obtained particle shapes are consistent with the two powder catalyst studies for the shape of the cross sections. However, the experimental shapes are much more columnar. For the Rosenthal crystals, we attribute this difference to the previously reported restructuring of (110) facets into c(2x2)-(100) microfacets, which is not included in our model. In the Narkhede experiments, the more likely explanation is that the experimental crystals are dominated by kinetic limitations of the growth process.

This part of the project was published in ChemCatChem [2]. All DFT calculations in this work were performed with the CASTEP package, which is fully parallelized based on the MPI standard. The CPU-cost for a single-point DFT

calculation at a production-level 9-layer (1x1) RuO_2(111) slab (250 electrons) in optimum SuperMUC operation mode amounted to ~21 CPUh. Including preliminary tests on the number of layers, cut-off energy, k-points, etc, we used about 12% of our available CPU-time (~0.2M CPU hours) for this part of the project.

Surface Structure Determination of c(2x2)-RuO_2(100) through First-principles Global Geometry Optimization:
We performed DFT-based global geometry optimization using a basin hopping approach to address the detailed structural model of c(2x2)-RuO_2(100). So far, the ongoing global geometry optimization has already found three structures, which exhibit a lower surface free energy than unreconstructed RuO_2(100) in some range of O chemical potential (see Figure 2). These candidates are all oxygen-rich terminations, but with different Ru and O ratios. Common to these three candidates is a trend to form RuO4 on top of all the surfaces, which is consistent with the well-known tendency of RuO_2 to form volatile RuO4 upon oxidation. However, the surface free energies of the three candidates found are still not low enough to induce microfacetting on RuO_2(110) according to the facet formation conditions indicated by Jacob et al.

For the search we have performed so far, we used up to 38% of our total CPU time (~0.8M CPU hours). Considering the huge consumption of computational resources, rather than continuing this project using the current global geometry optimization method, we decided to improve this methodology by applying more chemically intuitive displacements so as to improve the efficiency.

First-Principles Kinetic Monte Carlo Simulations of CO Oxidation on RuO_2(111) [2]:
In this part of the work we focus on RuO_2(111) facets and compare extensively to the established 1p-kMC model for CO oxidation at the hitherto primarily investigated RuO_2(110) facet. Figure 3 displays the 1p-kMC calculated steady-state average surface coverages and CO oxidation TOFs over a range of reaction partial pressures around ambient conditions. At the lowest pCO shown (~10^{-5}) the catalytic activity starts to die out. At increasing CO pressures catalytic activity sets in. The higher CO impingement increases the probability for ER reactions and enables LH reactions due to the increased stabilization of CO at the surface.

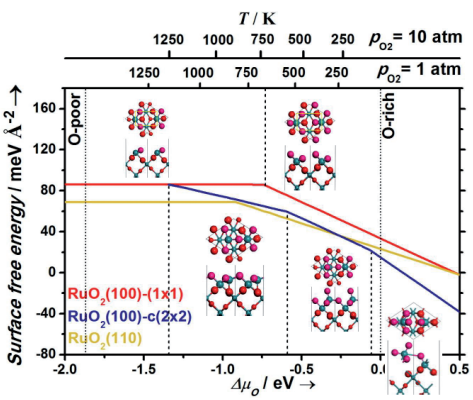

Figure 2. Same as Figure 1 with three c(2x2)-RuO_2(100) candidates identified in the global geometry optimization added with their respective surface free energies in blue lines.

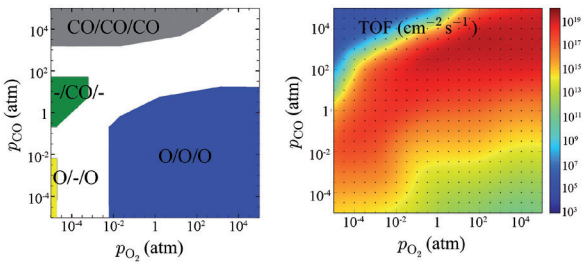

Figure 3. 1p-kMC computed steady-state O and CO surface coverage (left panel) and CO turnover frequency, TOF (right panel). Shown are data for a range of CO and oxygen partial pressures at 600 K.

We identified qualitative differences in the site arrangement, interplay of elementary processes and even the underlying individual elementary processes, demonstrating that CO oxidation at RuO_2(111) and RuO_2(110) facets is a structure sensitive reaction par excellence. Since nevertheless the calculated TOFs of both facets are almost identical, our study showed that the traditional classification of a reaction as structure insensitive based merely on the macroscopically observed catalytic function (typically explored only over a small set of feed conditions) can be a dangerous concept that does not adequately capture the underlying micro- to mesoscopic complexities.

This part of the work has been published in the Journal of Chemical Physics [3]. The main inputs are transition rates, which depend on energy barriers between states. A typical CI-NEB run to determine one transition state and therewith kinetic barrier required ~21 kCPUh, i.e. around 1000 single-point calculations to obtain a converged minimum energy path (forces < 0.05 eV/Å). For larger (1x2) surface unit-cells to extract lateral interactions (environment-dependent barriers) this cost increased to ~50 kCPUh. Combined with all the geometry optimization, we used about 50% percent of our CPU time (1M CPU hours) for this project.

Outlook

Under the SuperMUC grant our project has already established important insight into the CO oxidation at RuO_2 particles, in particular into the crystal habit during the induction period, the catalytic mechanism of CO oxidation on the apical facet (111), and possible avenues to the deactivation mechanism. The above success motivates further studies on the explicit structure of the RuO_2(100)-c(2x2) reconstruction and the activity of RuO_2(111) at lower temperatures, for which an improved estimate for the prefactor for Eley-Rideal type of reactions is required. As to the long-term steady-state activity, the obtained different composition and activity patterns of the RuO_2(110) and RuO_2(111) facets hint at interesting mass transport effects over the facet edges of RuO_2 nanoparticles. Such studies would then establish a comprehensive approach to CO oxidation at RuO_2 nanoparticles, significantly advancing the state-of-the-art in the field.

References and Links

[1] T. Wang, J. Jelic, D. Rosenthal, K. Reuter, Chem. Cat. Chem. 2013, 5, 3398.
[2] T. Wang, K. Reuter, J. Chem. Phys. 2015, 143, 204702.

New environmentally friendly thermoelectric materials

Research Institution
Max Planck Institute for Polymer Research

Principal Investigator
Davide Donadio

Researchers
Daniele Selli, Claudia Mangold, Joerg Behler, Pascal Pochet

Project Partners
Ruhr University, CEA Grenoble

SuperMUC Project ID: pr87bi (Gauss Large Scale project)

Introduction

The thermoelectric (TE) effect, discovered by Seebeck in the early 1800s, enables the conversion of thermal energy into electric power. In principle it would be a valuable renewable energy source, which would allow one to recover waste heat, e.g. from engines or chemical processes, or to use a broader range of the solar spectrum, thus complementing photovoltaics. An example of a typical TE device is shown in Figure 1.

While the TE effect occurs in a broad range of materials, the efficiency of TE conversion is usually low, with efficiency of few percent, thus relegating thermoelectric devices to niche applications. Furthermore, even if relatively low conversion efficiency may be tolerable, the best performing TE materials below 450 K known so far are semiconducting nano-structured transition metal chalcogenide alloys, such as lead and bismuth telluride, which contain rare and toxic elements.

To extend the use of thermoelectric technology beyond niche applications, it is therefore necessary to develop new Earth abundant and environmentally friendly thermoelectric materials. In this project we investigate the optimal structure of *ultra-thin silicon membranes and nano-structured manganese-doped germanium as viable TE materials.*

The efficiency of TE materials is determined by their dimensionless figure of merit $ZT=\varsigma S^2 T/\varkappa$, where T is the operating temperature, ς is the electric conductivity, S the Seebeck coefficient and $\varkappa$ the thermal conductivity. Current bismuth telluride TE devices operating at room temperature with $ZT<1$. Optimal ZT is obtained for doped semiconductors with high electrical conductivity – as in crystals - and low thermal conductivity – as in glasses, hence the paradigm of electron crystal phonon glass[2]. Since bulk crystalline silicon and germanium have large thermal conductivity, a strategy for TE optimization is to engineer them in such a way to block thermal transport without hindering electron conduction. The heat carriers in semiconductors are mostly phonons, which are quantized vibrations in periodic systems. Since the wavelength of heat-carrying phonons is typically of the order of tens of nanometers, TE optimization proceeds through nanoscale phonon engineering.

Results and Methods

Atomistic modeling is a powerful tool for nanoscale materials design. In this project we have employed a variety of complementary numerical methods to predict and optimize electron and phonon transport in ultrathin silicon membranes and nanostructured Mn-doped germanium. To compute heat transport we have employed molecular dynamics and lattice dynamics, whereas electronic properties were computed either at the level of density functional theory or by accurate semi-empirical methods.

Former experiments showed that the thermal conductivity of silicon membranes is largely reduced, compared to bulk, especially when the thickness of the membranes is of the order of few tens of nanometers or less. Molecular dynamics simulation of thermal transport in silicon membranes showed that such significant reduction of thermal conductivity is due to the presence of a native oxide layer that grows spontaneously at the surfaces when they are exposed to air. The reduction of thermal conductivity is proportional to the surface to volume ratio: the thermal conductivity of a 7 nm thick membrane can be 40 times lower than that of bulk silicon [4].

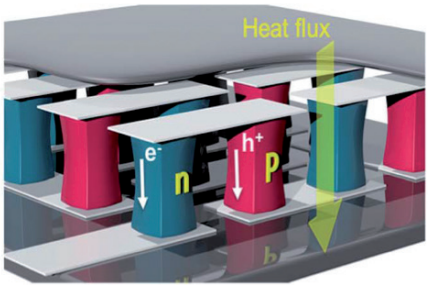

Figure 1: Scheme of a thermoelectric device in which p- and n-doped legs are connected electrically in series and thermally in parallel from [3].

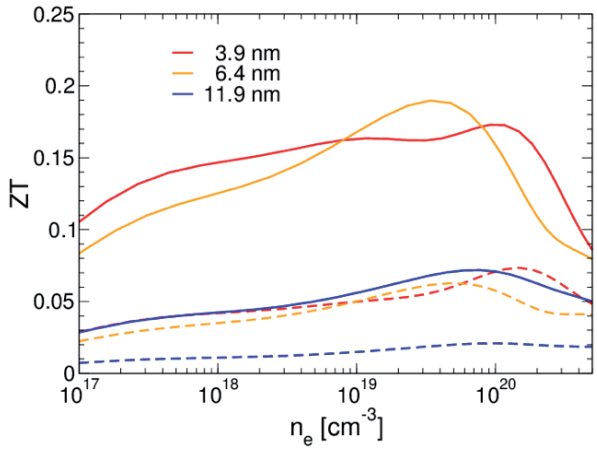

Figure 2: Thermoelectric figure of merit of ultra-thin silicon membranes of different thickness (3.9, 6.4 and 11.9 nm) as a function of the concentration of carriers (electrons) considering enhanced mobilities due to confinement (solid line) and bulk mobilities (dashed lines).

Electronic transport calculations, performed using density functional theory and using the results as input of the Boltzmann transport equation, confirm that doped ultra-thin silicon membranes, with thickness in the range between 4 and 12 nm, retain a good electrical conductivity and may even exhibit a slightly improved Seebeck coefficient, due to electron confinement effects.

The resulting thermoelectric figure of merit (Figure 2) peaks at about 0.2, which is a factor 40 higher than that of bulk silicon, for membranes 6.4 nm thick.

Electron transport calculations of membranes in open-system device-like configuration, performed using the non-equilibrium Green's functions (NEGF) approach, confirm similar values of ZT.

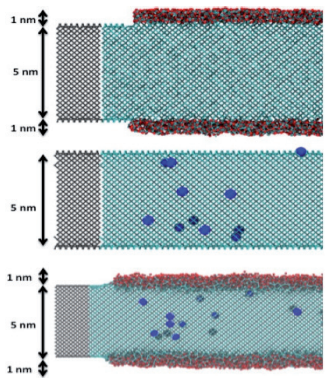

Figure 3: Models of membranes used for open-system device calculations: a 5 nm thick silicon membrane with amorphous native oxide at the surface (top), and a fully crystalline silicon membrane with 3% of substitutional defects (bottom panel). Only part of the systems is shown, next to the left lead (grey).

Lattice dynamics calculations of silicon membranes in device-like configurations (Figure 3) suggest that phonon scattering due to mass disorder in the membranes further reduces thermal transport thus improving the TE figure of merit. Mass disorder may be induced by either p or n doping or alloying, e.g. with germanium. A detailed analysis of the phonon transmission function, made possible by using the elastic scattering kernel method (ESKM) [5], shows that low frequency vibrations are mostly affected by the presence of surface features, such as native oxide and roughness, whereas mid-range frequency phonons are efficiently scattered by substitutional impurities with different mass. A further reduction of thermal conductivity, up to a factor 4 with respect to the previously reported data, can be achieved by introducing a 5% of germanium or arsenic, with the possibility of reaching ZT~1.

Applying the above mentioned atomistic techniques to realistic materials means that one needs to simulate system containing from hundreds to hundreds of thousands atoms, at the limit of the possibility of supercomputers. To obtain the results briefly reported above we exploited the efficient parallel architecture of SuperMUC. Large scale DFT calculations typically required the use of 512 cores per run. Pools of 1024 CPUs were used both for NEGF and ESKM calculations. In total these calculations took about 8 million CPU-hours.

On-going Research / Outlook

Given the chemical complexity of Mn-Ge systems, which exhibit a large variety of phases and stoichiometry, the study of nanostructured Mn-doped Ge for thermoelectric applications, required the fitting of a neural-network potential on density functional theory data. A neural network (NN) is a learning algorithm that establishes a mathematical relation between an input data set and an output data set. These two data sets are the configurations of compounds (input) – in our case the binary alloys of Mn-Ge plus bulk Ge – and their corresponding energies and forces (output) respectively. The more extensive the data base, the more transferable the resulting NN potential. About 11,000 structures (see Table 1 for details) were calculated to fit an accurate and transferable NN potential.

Table 1: Construction of the training set for the neural network potential

Structure	Type	# of entries
Ge	Bulk crystalline	2522
MnGe	Bulk crystalline	1937
Mn_5Ge_3	Bulk crystalline	2396
Ge/Mn_5Ge_3	Superlattice	4794

Accuracy and transferability of the NN potential were verified on the structural and vibrational properties of Ge, crystalline Mn_xGe_y and Ge/Mn_5Ge_3 heterostructures. Future work will consist of a full vibrational and thermal characterization of crystalline Ge containing Mn_5Ge_3 quantum dots.

References and Links

[1] www.merging.eu
[2] G. J. Snyder and E.S. Toberer, Nature Mater 7, 105 (2008).
[3] O. Bubnova and X. Crispin, Energy Environ. Sci. 5, 9345 (2012).
[4] S. Neogi et al. ACS Nano 9, 3820 (2015).
[5] I. Duchemin and D. Donadio, Phys Rev B 84, 115423 (2011).

Optical Simulation of Innovative Thin Film Solar Cells

Research Institution
Chair for Systemsimulation – FAU Erlangen-Nürnberg

Principal Investigator
C. Pflaum

Researchers
Birhanu T. Abebe, Julian Hornich

Project Partners
Energy Campus Nürnberg (EnCN), Fraunhofer COMEDD, ALANOD GmbH & Co.KG, WE-Forschungszentrum für Energietechnologie e. V. Oldenburg

SuperMUC Project ID: pr87fe

Introduction

Within this project the goal is to study and develop novel approaches to boost the perfomance of thin film solar cells. For this, 3D optical simulation of the photovoltaic devices is performed by discretizing Maxwell's equations. A sophisticated light management is important to construct thin-film solar cells with optimal efficiency. The light management is based on suitable nano structures of the different layers and materials with optimized optical properties. The design, development and test of new solar cell prototypes with respect to an optimal light management are a time consuming processes. For this reason, suitable models and simulation techniques are required for the analysis of optical properties within thin-film solar cells.

Method

A rigorous analysis by a discretization method for Maxwell's equations is needed, in order to predict optical properties of thin film solar cells. Rigorous EMF simulation methods are more accurate because they include physical effects like wave interference, reflection, scattering as well as plasmonic effects. Suitable discretization methods are the finite integration technique (FIT) and the finite difference time domain method (FDTD). These simulation techniques are computationally intensive because the random textures at the interface of composite layers are difficult to simulate.

We have developed a simulation tool, based on finite integration technique FIT, for calculating external quantum efficiency and short circuit current density of thin film solar cells. We use FIT, because it can accurately model curvilinear interfaces and is less computationally intensive compared to FEM. The program is parallelized using MPI and OpenMP and can be used on supercomputers with several thousand processors.

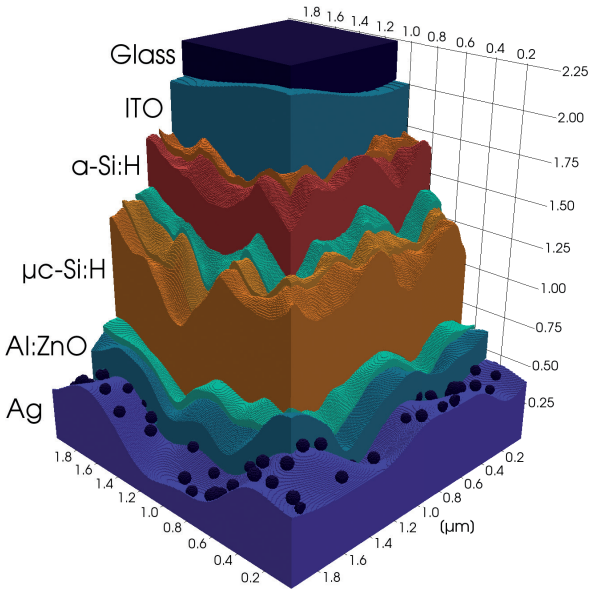

Figure 1: Cross-section of an exemplary simulation setup of a tandem thin-film solar cell. The amorphous (a-Si:H) and mi-crocrystalline silicon (µc-Si:H) layers have textured surfaces to increase the light trapping ability of the cell. SiO 2 nanoparticles are incorporated to further increase light scattering at the bottom electrode (Ag).

Simulation Results

Flexible Silicon Tandem Solar Cells
The Silicon based Thin Film Flexible Solar cells project (SiSOFlex) was envisioned to use the advantages of thin film solar cells and work on different weak aspects of the solar cell to achieve a stable and working solar cell with >11% stable efficiency. To achieve this goal simulation, modeling and optimization of different ideas and concepts plays an integrated role along with experiments before going to production. Figure [1] shows the crossection of such a solar cell. The roughness of the surfaces and the integration of scattering particles results in an optimized harvesting of the incident light. Optimization of those structures requires many simulations but allows for an accurate prediction of the performance of the resulting solar cells. More detailed results can be found in Ref. [2].

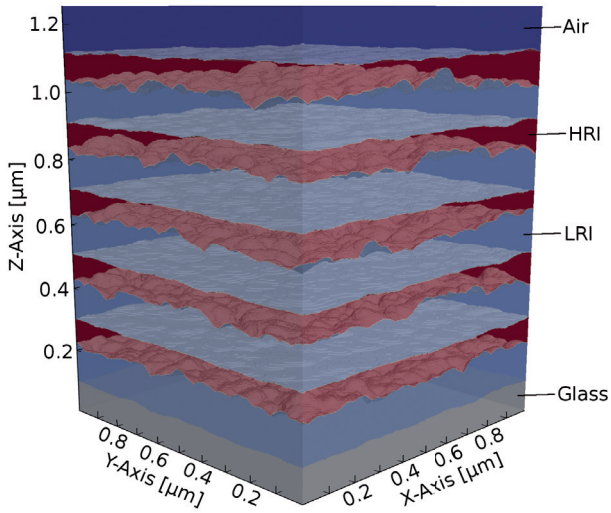

Figure 2: Setup of a dielectric (Bragg) mirror with five alternating high/ low refractive index stacks. See Ref. [3]

Wavelength Selective Dielectric Mirrors

Organic photovoltaics in combination with dielectric mirrors (DM) are a potential candidate for building integrated solar cells as they promise high efficiencies in parallel to the possibility to adjust the color and the transparency of the whole device. Wavelength selective filters which are also known as 1D photonic crystals, Bragg mirrors or DM are based on constructive or destructive interference in thin layers. For this purpose, a high refractive index (HRI) and a low refractive index (LRI) material have to be arranged alternatingly. Comparison to experimental measurements have shown that a good agreement can be reached with our simulation code. Figure [2] shows a simulation setup of a DM with five alternating HRI/LRI layers. The characteristic transmission spectrum of a DM depends not only on the used materials but also on the surface roughness between the HRI and LRI layers. This makes 3D optical simulations necessary.

Plasmonic Absorption Enhancement

The excitation of localized surface plasmons on metallic nano particles results in an enhancement of the electromagnetic field in the direct vicinity of the particles. The origin of this effect is the oscillation of the electron cloud in the metallic particles caused by the electric field of the incident light. The resulting field enhancement is perpendicular to the polarization of the incident electric field, see Figure [3]. Experimentally, it is extremely difficult to measure the near field absorption enhancement of such nano particles. Hence we are carrying out simulations to study the impact of several particle types and arangements incorporated into thin film organic photovoltaic devices.

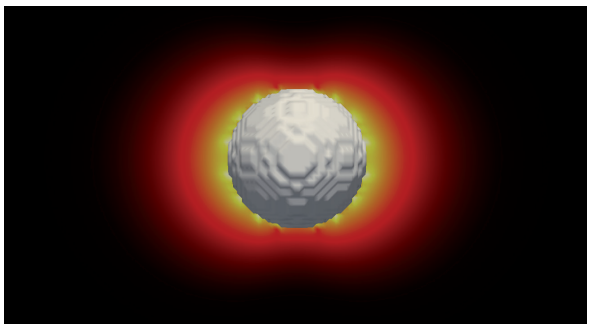

Figure 3: Visualization of the enhancedelectric field around a metallic nanoparticle dueto the excitation of localized surface plasmons.

Performance Optimization

In order to optimize the parallel efficiency of the used simulation tool we recently carried out a KONWIHR project at the RRZE, FAU Erlangen-Nürnberg. Within this project we were able to reduce the time spent on data exchange between processes by an improved mapping of the processes to the simulation grid and by more sophisticated overlapping/asynchronous data exchange algorithms. The benchmarks, see Figure [4] and [5], show that we were able to improve the parallel efficiency of the simulation code from ~70% up to >90% for more than 100 compute nodes.

On-going Research / Outlook

The projects presented above are ongoing and more detailed simulations will be performed. Since this will require more computational resources we are planning a further code optimization project with KONWIHR to enhance the single core performance of our code.

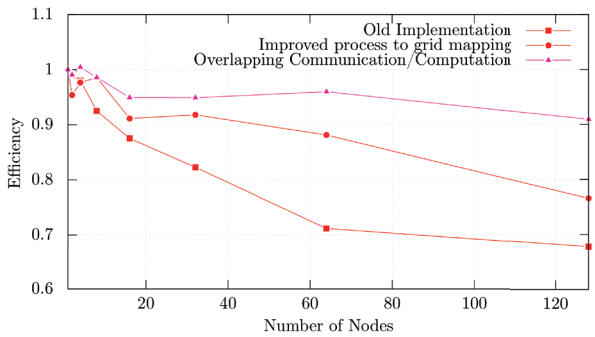

Figure 4: Comparison of the parallelefficiency with the methods implementedduring the KONWIHR project.

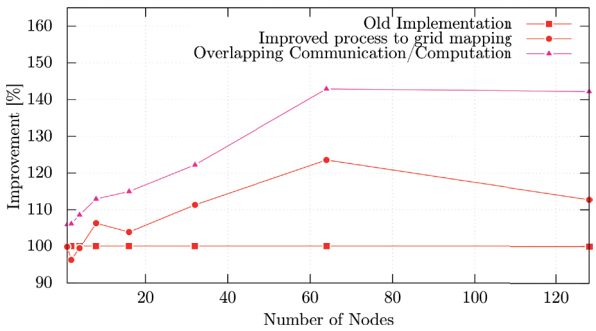

Figure 5: Comparison of the (time to solution) improvement with the methods implemented during the KONWIHR project.

References and Links

[1] www10.cs.fau.de/wissenschaftliches-rechnen
[2] www.eupvsec-proceedings.com/proceedings?fulltext=abebe
[3] Bronnbauer, C., Hornich, J., Gasparini, N., Guo, F., Hartmeier, B., Luechinger, N. A., Pflaum, C., Brabec, C. J. and Forberich, K. (2015), Printable Dielectric Mirrors with Easily Adjustable and Well-Defined Reflection Maxima for Semitransparent Organic Solar Cells. Advanced Optical Materials, 3: 1424–1430. doi: 10.1002/adom.201500216

Modeling, simulation and optimization of thin film silicon solar cells on flexible Aluminum substrate

Research Institution
University of Erlangen, Computer Science Department, Chair of System Simulation

Principal Investigator
Christoph Pflaum

Researchers
Birhanu Tamene Abebe, M.Sc

Project Partners
NEXT ENERGY EWE-Forschungszentrum, ALANOD® GmbH & Co. KG, Fraunhofer COMEDD, Dresden, Germany

SuperMUC Project ID: pr87fe, h0672

Introduction

Silicon based Thin Film Flexible Solar cells project (Si-SOFlex) was envisioned to use the advantages of thin film solar cells and work on different weak aspects of the solar cell to achieve a stable and working solar cell with > 11% stable efficiency. To achieve this goal simulation, modeling and optimization of different ideas and concepts plays an integrated role along with experiment before going to production. The process of combining experiment with simulation first saves time and resource and second provides information about the behavior of the solar cell that are difficult to measure experimentally. Overall, factors that affect the performance of thin film solar cells including but not limited to the materials used at each layer, processes used for the manufacturing, sequence of the layers, deposition order of the layers in the solar cell are investigated in detail.

Results and Methods

To enhance the efficiency of solar cells, the light coming in has to be trapped inside the solar cell. As a light trapping mechanism, the most efficient technique for thin film solar cells is textured interface between the layers. The solver takes surface data like the one shown Figure 2.1a, as an input to define the interface roughness between layers as shown on Figure 1.1b. The surface data can be artificially generated or experimentally measured

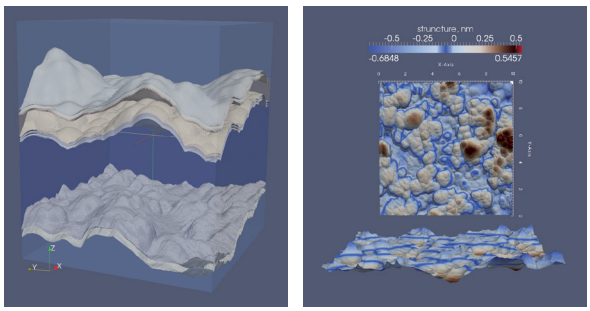

(a) (b)

Figure 1.1: (a) AFM scan integration (b) measured interface data using AFM scan method

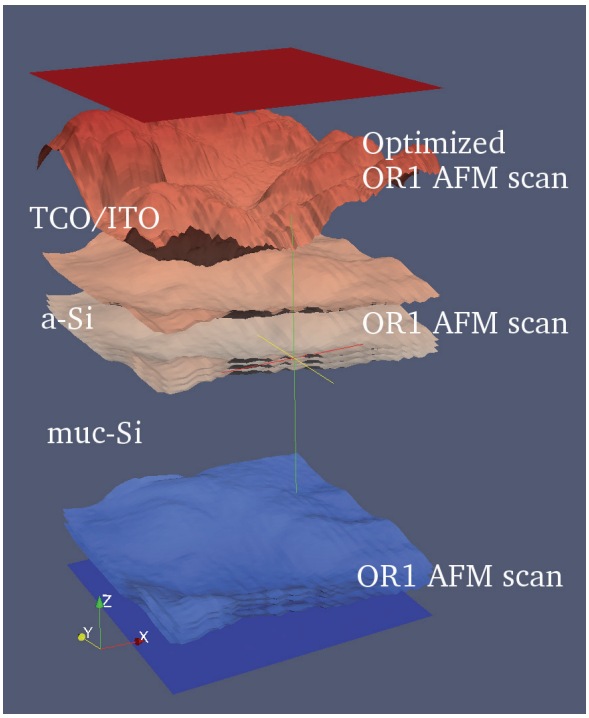

Figure 1: Thick Front Contact

using Atomic Force Microscopy (AFM). Simulation of AFM scan requires a relatively large simulation domain to accurately simulate the effect of the statistical data on the AFM scan. The other way is to simulate a number of smaller simulations that covers the AFM scan data adequately. Both ways need substantial computational resource in terms of memory and time. To satisfy the memory and time demand, the simulations are performed on HPC at LRZ. Some of the simulations performed are discussed below.

Thick front TCO
The idea behind thick front Transparent Conductive Oxide (TCO) is to make the front thickness thick enough to introduce roughness in addition to the roughness generated as a result of the back contact roughness. To this

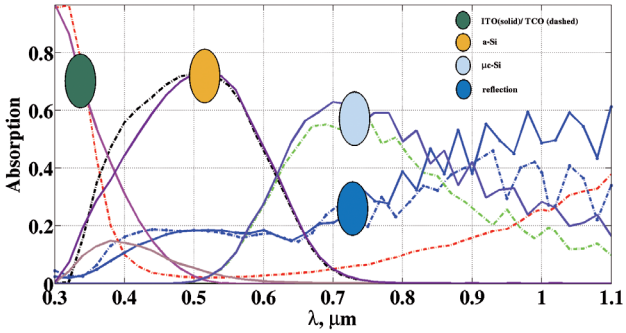

Figure 2: Thick front contact (TCO) result

end, the thickness of the front TCO layer is simulated with thickness in the range of 400 - 800nm. This allows more rough interface at the front contact.

Materials like Indium Tin Oxide (ITO) and Aluminum doped Zinc Oxide (ZnO:Al) were used as a front contact as shown on Figure 2.1. The result shown on Figure 1.2 and Table 3.1 indicates that the improvement resulted using this technique is not significantly better as a result of the high absorption in the thick TCO layer.

Intermediate Reflector

A closer look at simulation results in the previous sections shows that the top aSi sub-cell limits the performance of the solar cell. To improve the absorption at the top aSi layer, intermediate reflector placed between the top aSi sub-cell and the bottom mcSi sub-cell is simulated. The idea is discussed in [1] although in the first case the solar cell is superstrate configuration. The intermediate reflector will reflect the light back to the aSi sub-cell thereby increasing the probability the light will be absorbed in

Table 1. Intermediate Reflector setup

Glue	Infinity
ITO	80nm
p-aSi	10nm
i-aSi	300nm
$n-\mu$cSi	40nm
TCO	100nm
$p-\mu$cSi	40nm
$i-\mu$cSi	1600nm
$n-$aSi	40nm
TCO	80nm
Silver	300nm

The result for given on Table 3.4 and Figure 3.4 shows that the intermediate reflector results the desired outcome by increasing the absorption of the top aSi layer. There is also a slight decline on the performance of the mcSi sub cell. In general with intermediate reflector, it is shown that it is possible to achieve a solar cell with efficiency of 11%. Overall the use of intermediate reflector shows a promising result but Figure 3 shows that reflection from the solar cell increases with introduction of the intermediate reflector in the longer wavelength region. This is the reason behind the decline in the absorption of the bottom mcSi sub cell.

As a result further optimization is needed to find out better material that acts as a reflector for the top aSi sub-cell while at the same time it is transparent for the bottom mcSi sub cell.

Optimized Solar cell

After the optimizations discussed in the previous sections are done, complete solar cells where simulated with optimized parameters. The comparison is given on Table 2.2. The result shows that ITO as a front contact

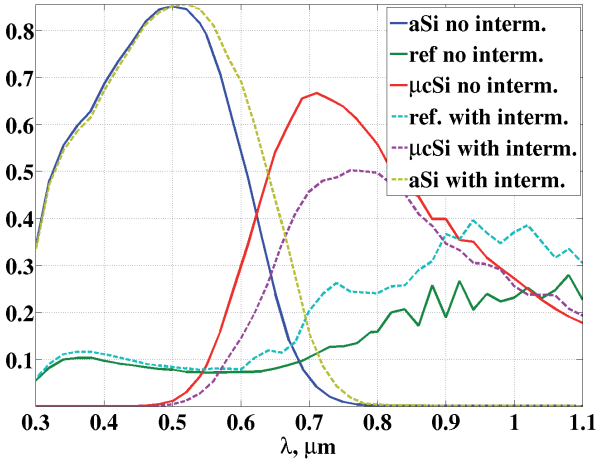

Figure 3: Intermediate reflector with ZnO:Al material

gives better result. Optimizations of the interface roughness + active layer thickness generally improves the solar cell but the more rougher the back contact roughness, shunting problem occurs. In addition, the roughness at the back contact greatly affects the quality of the layer being deposited [1]. As a result, the most viable solution to increase the performance of the solar cell is to use ITO as a front contact.

Conclusion

All the milestones set at the beginning of the project were successfully met. Various ideas originated from different partners of the project were simulated, multiple experiments were verified by simulation, the simulations provide insight on things that are difficult to see by experiment, case in point - parasitic absorptions in different layers. The optical simulation of University of Erlangen was successfully coupled with the electrical simulation from Next Energy.

References

[1] B. Abebe and C. Pflaum, "Simulation and optimization of flexible thin film silicon solar cells," in WAFER-BASED SILICON SOLAR CELLS AND MATERIALS TECHNOLOGY, EU PVSEC submitted, 2015.

[2] S. Geissendorfer, M. Theuring, T. Titz, S. Mogck, C. Pflaum, B. Abebe, F. Schutze, D. Wynands, U. Kirstein, A. Schweitzer, V. Steenhoff, A. Neumuller, K. Borzutzki, R.-E. Nowak, A. Philipp, P. Klement, O. Sergeev, M. Vehse, and K. von Maydell, "The sisoflex project: Silicon based thin-film solar cells on flexible aluminium ubstrate," EU PVSEC Conferences 29, pp. 1667–1670, 2014.

[3] B. T. Abebe, K. Hertel, and C. Pflaum, Modelling and comparison of light trapping caused by textured interfaces and nanoparticles in thin film solar cells. Proc. SPIE 8818, Nanostructured Thin Films VI, 2013. 8818: p. 88180Q-88180Q-11.

Tuning the electronic reconstruction at oxide surfaces and interfaces

Research Institution
Fakultät für Physik, Theoretische Physik und Center of Nanointegration (CENIDE), Universität Duisburg-Essen

Principal Investigator
Rossitza Pentcheva

Researchers
David Doennig, Markus Müller, Hamidreza Hajiyani

Project Partners
–

SuperMUC Project ID: pr87ro

Introduction

Transition metal oxide (TMO) interfaces host a remarkably rich electronic behavior distinct from the one of the bulk compounds, which opens perspectives for new electronics and spintronics devices as well as for energy conversion applications. In pr87ro large-scale material-specific density functional theory (DFT) calculations (Wien2k [1] and VASP code [2]) are carried out to achieve a fundamental understanding of the phenomena emerging at complex oxide surfaces and interfaces.

Results and Methods

The interface between the band insulators $LaAlO_3$ (LAO) and $SrTiO_3$ (STO) comprises one of the most intensively studied systems, due to the variety of functional properties that arise such as a two-dimensional conductivity, superconductivity and signatures of magnetism. One intriguing feature is the transition between insulating and conducting behavior in thin LAO films on STO(001) at a critical thickness of 4 ML (monolayers) LAO [3]. DFT calculations show that the thickness-dependent insulator to metal transition is a result of the interplay of an internal electric field that develops within the polar LAO film

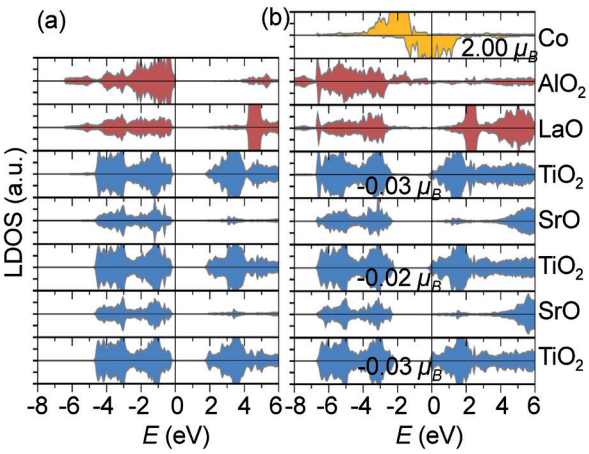

Figure 1: Layered density of states of 1 ML LaAlO3/SrTiO3 (001) (a) without and (b) with a Co layer on top [5].

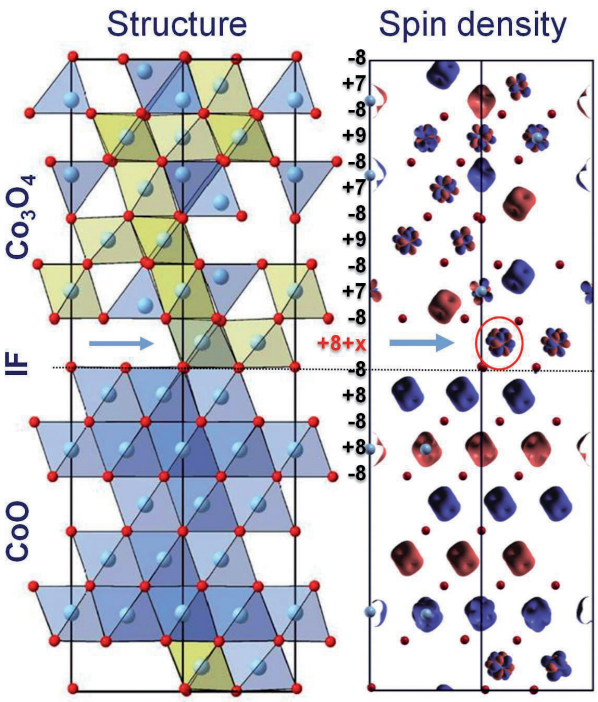

Figure 2: Electronic reconstruction at a polar CoO/Co_3O_4 interface. While octahedral Co^{3+} in Co_3O_4 bulk is in low spin (LS) state, the spin density from the DFT+*U* calculation indicates a change in the oxidation state of octahedral Co at the interface to Co^{2+} in LS state [6].

(see layer-resolved DOS of 1LAO/STO(001) in Fig. 1a) and a counteracting lattice polarization [4].

A recent combined experimental and theoretical study demonstrated that the critical thickness can be suppressed by a Co overlayer. The DFT calculations indicate that the metallic Co overlayer modifies the electrostatic boundary conditions and induces a charge transfer towards the Ti 3d bands. (Fig. 1) [5].

The electronic reconstruction at oxide interfaces is not restricted to the much studied perovskite structure. Recently, we have explored the origin of magnetism at the interface of two antiferromagnets CoO (rocksalt struc-

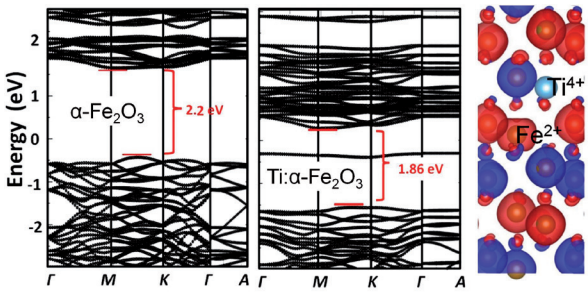

Figure 3: DFT+U band structure of α-Fe_2O_3 (left), Ti- doped α-Fe_2O_3 (middle) and the corresponding spin density of the latter (right), indicating that the substitution by Ti^{4+} is compensated by the formation of Fe^{2+}.

ture) and Co_3O_4 (spinel). The DFT+*U* calculations show that to compensate the polar discontinuity at the CoO/Co_3O_4(111) interface, octahedral Co^{3+} (d^6, $0_{\mu B}$) within the interfacial Co_3O_4 layer is partially reduced to Co2+. Moreover, the latter is in the LS state (0.4-0.8$_{\mu B}$) [6], in contrast to the high spin Co^{2+} (2.6 $_{\mu B}$) in CoO. The DFT+*U* calculations indicate that these defect moments are likely the origin of a stable above-room temperature ferrimagnetism at this interface [6].

On-going Research / Outlook

Due to their high thermal and chemical stability, non-toxicity as well as broad tunability of properties, TMO are gaining importance also in the field of energy conversion, e.g. in the light-assisted splitting of water into hydrogen and oxygen. Thereby, the oxygen evolution reaction (OER) at the anode comprises the transfer of four holes to obtain O2 and poses high requirements on the anode material including good visible light adsorption, conduction and valence band edges straddling the reduction and oxidation potentials of water, rapid charge transfer and suppression of carrier recombination. Hardly any current material fulfills all requirements. Within pr87ro we are currently investigating the optimization of anode materials such as hematite by doping as a promising route. Fig. 3 shows the influence of doping on the bulk band structure of α-Fe_2O_3: the band gap of the pure material of 2.2 eV is reduced to 1.86 eV by Ti substitution. Additionally, a localized level is formed in the band gap, due to the formation of Fe^{2+} to compensate for the Ti^{4+} incorporation.

In the framework developed by Nørskov, Rossmeissl and coworkers [e.g. [7] and Refs. therein] water oxidation is subdivided into four reaction steps comprising four reaction intermediates *OH_2, *OH, *O and *OOH, where asterisk denotes a species adsorbed on the surface. Fig. 4 shows the energetics of these intermediates at the α-Fe_2O_3(0001) surface (see also [8]). Currently we are exploring the influence of the surface termination and dopands to reduce the overpotential for OER at this surface. Funding by the DFG within SFB/TR80 and SPP1613 is gratefully acknowledged.

References and Links

[1] P. Blaha, K. Schwarz, G. Madsen, D. Kvasnicka, and J. Luitz, (Karlheinz Schwarz, TU Wien, Austria, 2001).
[2] G. Kresse and J. Hafner, Phys. Rev. B 47, 558 (1993); G. Kresse, D. Joubert, Phys. Rev. B 59, 1758 (1999).
[3] S. Thiel, G. Hammerl, A. Schmehl, J. Mannhart, Science 313, 1942 (2006)
[4] R. Pentcheva and W.E.Pickett, Phys. Rev. Lett. 102, 107602 (2009).
[5] E. Lesné, N. Reyren, D. Doennig, R. Mattana, F. Choueikani, H. Jaffrès, V. Cros, F. Petroff, P. Ohresser, R. Pentcheva, A. Barthélémy, and M. Bibes, Nature Commun. 5, 4291 (2014).
[6] Zi-An Li, N. Fontaíña-Troitiño, A. Kovács, S. Liébana-Viñas, M. Spasova, R. E. Dunin-Borkowski, M. Müller, D. Doennig, R. Pentcheva, M. Farle and V. Salgueiriño, Sci. Rep. 5, 7997 (2015).
[7] I. C. Man et al., Chem. Cat. Chem. 3, 1159 (2011).
[8] P. Liao, J. A. Keith and E. A. Carter, J. Am. Chem.Soc.134, 32 (2012).

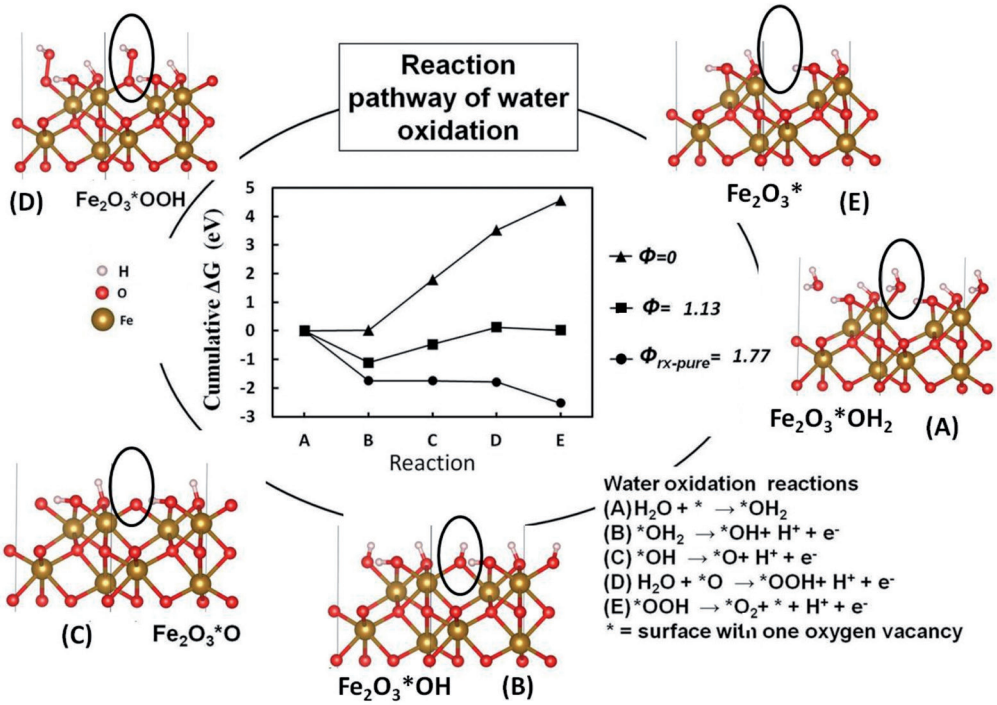

Figure 4: Cumulative free energies of the intermediates for different values of the reaction potential during OER at the α-Fe_2O_3(0001) surface obtained within DFT+*U*.

Hydrogen-bond symmetrization and the spin transition in ε-FeOOH

Research Institution
Fakultät für Physik, Theoretische Physik und Center of Nanointegration (CENIDE), Universität Duisburg-Essen

Principal Investigator
Rossitza Pencheva

Researchers
Carmen E. Quiroga

Project Partners
–

SuperMUC Project ID: pr87ro

Introduction

The behavior of water-bearing compounds at high pressures and temperatures is naturally related to the state of water in the deep Earth. Iron oxyhydroxides are abundant on the surface of the Earth, where they occur naturally in soils, aquifers and sediments in the crystal forms of goethite (α-FeOOH), akaganeite (β-FeOOH) and lepidocrocite (γ-FeOOH). A fourth polymorph, ε-FeOOH, can be synthesized as a high-pressure phase. Understanding the stability and properties of these oxyhydroxides at high pressures can reveal essential information about the influence of water on the mantle redox state and represents an important step towards quantifying more complex hydrogen-bearing compounds relevant to the Earth's interior.

Results and Methods

We have performed density functional theory (DFT) calculations including a Hubbard *U* term to study the structural and electronic properties of iron oxyhydroxides at high pressures, with particular focus on the spin state of iron as pressure is increased.

Analysis of the structural degrees of freedom as a function of pressure in ε-FeOOH reveals a complex interplay between the symmetry of the hydrogen bonds and the spin transition of Fe^{3+} (see Fig.1) [1].

As shown in Fig. 2, a second order structural phase transition takes place at ~ 40 GPa in high-spin ε-FeOOH, which is clearly identified by the freezening of the *P*nnm Γ4-displacive mode. Further evidence of the crystallographic phase transition is seen in the symmetrization of the hydrogen bonds as H reaches an equidistant position along the O...O line and the off-centering of the Fe cation with respect to the center of mass of the iron-oxy-hydroxyl octahedra vanishes. Above the phase transition, the O-H distances become nearly insensitive to pressure, which manifests in the lattice compression behavior as the stiffening of the lattice directions containing the hydrogen bonds (see Fig. 3). As a result, high-spin ε-FeOOH acquires a significantly larger bulk modulus above ~ 40

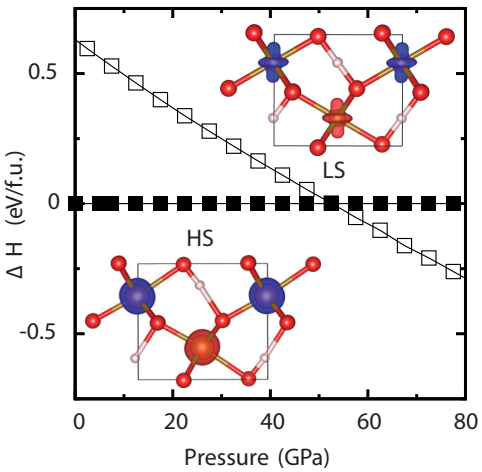

Figure 1: The spin transition in ε-FeOOH, showing the enthalpy difference between the high-spin (HS) and low-spin (LS) states. The spin density distributions serve to identify the spin state of Fe^{3+} in both configurations. Note a nearly spherical spin-density indicative of the HS state vs. the anisotropic spin density of the LS state.

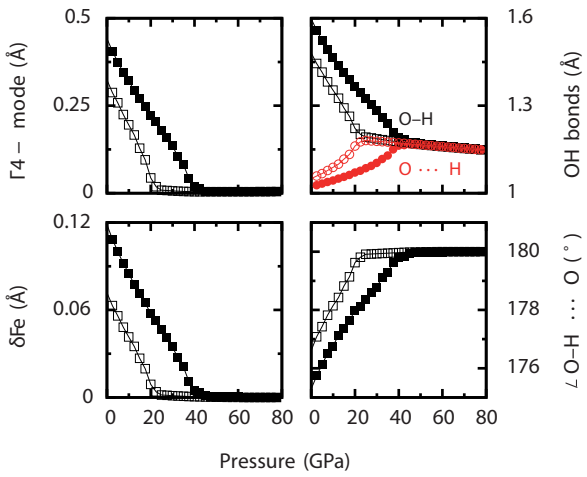

Figure 2: Pressure- induced symmetrization of the hydrogen bonds in ε-FeOOH and related quantities, showing the amplitude of the Γ4-mode identifying the structural phase transition, the pressure behavior of the O-H, O...H bonds, the O-H...O bond angle and the displacement of the Fe cation with respect to the center of mass of the iron-oxy-hydroxyl octahedra in the high-spin (solid squares/circles) and low-spin (open squares/circles) configurations.

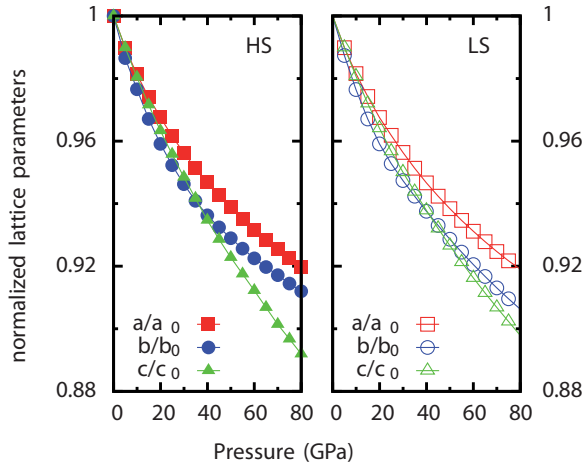

Figure 3: Normalized lattice parameters in the HS (right) and LS (left) spin configurations showing the stiffening of the ab plane in both the HS and LS configurations above the P21nm to *P*nnm phase transition shown in Fig. 2.

GPa, while low-spin ε-FeOOH allows greater compressibility due to its iron-oxy-hydroxyl octahedra being ~ 6% smaller. That, together with the indication from Gleason et al. XES data that the spin transition starts at ~ 40 GPa and completes after the volume reduction, suggests that the spin transition is assisted by the symmetrization of the hydrogen bonds [1].

Similar stiffening behavior has been recently observed in other distorted rutile-type oxy-hydroxides, including δ-AlOOH(D), β-CrOOH(D), β-GaOOH and InOOH, but among those only CrOOH contains a transition metal ion which is however not susceptible to a spin transition. Thus, our study establishes for the first time a relation between the effect of hydrogen bond symmetrization and spin crossover at high-pressure.

The theoretical description of transition metal oxides represents a significant challenge for ab-initio methods based on DFT. The commonly used approximations for the exchange and correlation energy, rely on some parametrization of the homogeneous electron gas, fail dramatically to describe the properties of strongly correlated materials, where valence electrons are localized in atomic-like orbitals and many-body terms of the electronic interactions acquire relevant significance. Several schemes have been introduced throughout the past decades to overcome this difficulty. Among them, the DFT+*U* [2] is one of the simplest, yet effective, methods and remains the most reasonable choice for systems of realistic complexity. Borrowed from the Hubbard model, the DFT+*U* functional incorporates a correction that selectively operates on the localized states to disfavor fractional orbital occupations. The strength of the "+*U*" correction is controlled by the effective Coulomb (*U*) and exchange (J) interactions, whose values are not known a priori. However, and as shown in Fig. 4 for ε-FeOOH, the spin transition pressure depends strongly on the chosen values of *U* and J entering the DFT+*U* functional. Fig. 4 also depicts that without the "+*U*" correction the spin transition pressure is significantly underestimated, consistent with the common trend of standard DFT to overstabilize the low-spin state configuration. It is clear then that in order to allow a meaningful prediction with DFT+*U*, the study of the spin transition requires a careful evaluation of both *U* and J.

To accomplish this, we have implemented in VASP [3,4] the linear response approach proposed by Cococcioni and Gironcoli (2005) [5] and tested it against the implementation in the Quantum Espresso package. Overall consistent results are obtained between the two codes and among the four studied polymorphs. With the calculated values (U = 5.98 eV and J = 1.36 eV) the spin transition is predicted to take place at 53 GPa in ε-FeOOH (see Fig. 1), in remarkable agreement with the experimental observation [1]. Using the same *U* and *J* values we are able to predict the spin transition pressure for α-, β- and γ-FeOOH at 49, 63 and 40 GPa, respectively [6].

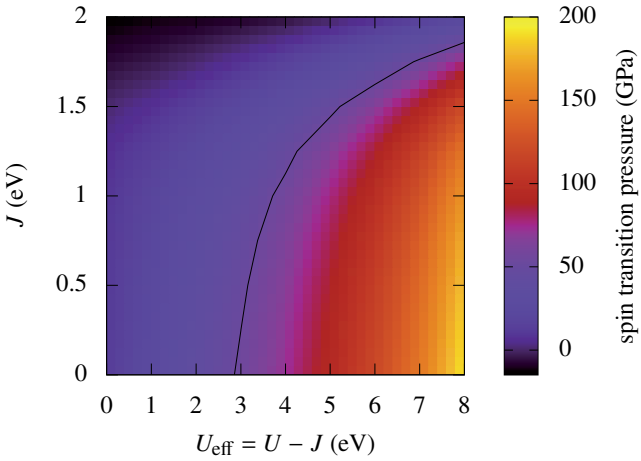

Figure 4: Dependence of the spin transition pressure in ε-FeOOH on the chosen values of U_{eff}=*U-J* and *J*. The contour line in solid black at 54 GPa marks the measured spin-transition pressure [1].

On-going Research / Outlook

As the presence of water, even in small quantities, can affect phase relations, melting temperature, rheology, and other key properties of the Earth's mantle, our study suggest a possible connection between water (hydroxyl) content and the spin-transition pressure of Fe^{3+} in the Earth's interior. This study is currently being extended to generalize the trends to all FeOOH polymorphs [6].
Funding by the DFG, SPP1236 PE883/8-1 is gratefully acknowledged.

References and Links

[1] A.E. Gleason, C.E. Quiroga, A. Suzuki, R. Pentcheva & W.L. Mao. Earth. Planet. Sci. Lett. 379, 49 (2013)
[2] A. I. Liechtenstein, V.I. Anisimov, J. Zaanen, Phys. Rev. B 52, R5467 (1995).
[3] G. Kresse and J. Hafner, Phys. Rev. B 47, 558 (1993).
[4] G. Kresse, D. Joubert, Phys. Rev. B 59, 1758 (1999).
[5] M. Cococcioni & S. de Gironcoli. Phys. Rev. B 71, 035105 (2005).
[6] C.E. Quiroga & R. Pentcheva, in preparation.

Phase transition based control of friction at the nanoscale

Research Institution
Empa Swiss Federal Laboratories for Material Science and Technology
Principal Investigator
Carlo Pignedoli
Researchers
Andrea Benassi
Project Partners
Institute for Practical Computation, LMU Munich

SuperMUC Project ID: pr89mi (PRACE project)

Introduction

The ability to control and manipulate frictional forces at the nanoscale is extremely important for technology, being closely tied to progress in transportation, manufacturing, energy conversion, and lubricant consumption, impacting on innumerable aspects of our health and environment. In recent years a lot of effort has been devoted to gain control of friction at both the macroscopic and microscopic scale. However, most of the employed techniques cannot be straightforwardly extended to the nanoscale, where a flexible and almost cost-free way to dynamically tune friction forces is still lacking.

The flexibility of selected physical properties of the sliding bodies, necessary to actuate a dynamical control of friction, might be provided by the occurrence of a phase transition. A few evidences of this possibility can already be found in literature, and are related to the development of atomic force microscopy (AFM), a quickly improving experimental tool that allows to study, using nanosized oscillating tips, surface properties by sampling contact and non-contact forces down to the atomic scale.

Recently Benassi and coworkers demonstrated the possibility to control nanofriction by switching the order parameter of a structural phase transition [1]. Friction force microscopy (FFM) experiments on a model ferro-distortive substrate have been simulated, showing a non-monotonic behavior of friction as a function of the substrate temperature, broadly peaking at T_c. Besides this unusual feature (stick-slip friction of a single contact on ordinary substrates is known to decrease monotonically unless multiple slips occur), below T_c, the frictional response is found to depend strongly on the substrate distortive order parameter: different values of the substrate distortion can give rise to a very different friction force. Acting now with an external stress field the distortive order parameter of the substrate can be changed reversibly and dynamically, increasing or decreasing the frictional properties of the substrate.

The goal of our PRACE project was to identify a real material hosting a structural phase transition in a range of parameters accessible to the standard experimental techniques. To this aim, a practical structural phase transition to look at is the rotational melting occurring near 260 K in pristine C_{60} molecular crystals (aka Fullerite, Figure 1). Such first order phase transition leads from an ordered structure with locked fullerene at low temperatures, to an angular disordered phase with almost freely rotating fullerenes at high temperatures. Recent measurements have indicated a sharp drop of tip adhesion and of AFM sliding friction in connection with this bulk transition [2], see figure 1 (a). However the explanation given for this strong frictional drop is in open contradiction with thermodynamics, as it is immediately visible from a comparison with the cohesive energy of the Fullerite crystal available in the literature. Despite the great importance that the measured frictional drop can have in the control of friction, no theoretical explanation or simulation data are available for the tribological properties of this system.

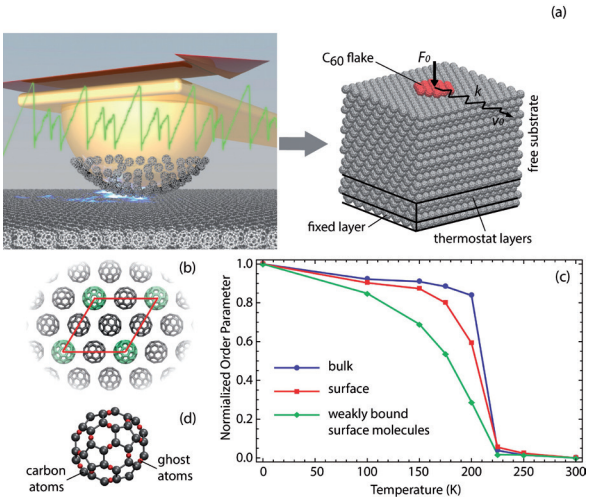

Figure 1: (a) Sketch of the C_{60} sliding physical system (left) and of its implementation in our MD simulations (right). (b) Surface cell for the low temperature phase, the weakly bound surface molecules are highlighted in green. (c) Variation of the rotational order parameter with temperature (see [3]) for bulk molecules (blue), all surface molecules (red) and weakly bound surface molecules (green) of panel (b). (d) Interaction centers for the intermolecular potential of ref. [4]. Reproduced from ref. [3] with permission from the Royal Society of Chemistry.

Results and Methods

Using classical molecular dynamics we simulated the pull-off and sliding friction experiments in presence of the rotational melting phase transition [3].

We started reproducing the bulk and surface rotational melting transition, to this aim the potential by Sprik et al. [4] proved to be best solution. This potential is constituted by a Lennard-Jones short range term plus a long range coulombic interaction.

Every C_{60} molecule is treated as a rigid body with 90 interaction centers placed on every carbon atom and on every double bond, to mimic the charge delocalization, see figure 1 (d).

Finite size effects are known to inhibit the occurrence of the phase transitions thus, to have a satisfactory description of rotational melting, we need to work with large systems, our typical simulation box contains 16x16x17 fullerenes, i.e. ~400000 interaction centers. As we simulate pull-off and sliding friction experiments we act with an external driving force on the systems. The work done by this force increases the internal energy and, in order to reach a steady state and to prevent the system from blowing up, the excess energy must be disposed off with a thermostat.

As demonstrated in reference [5], a good thermostatation procedure, which does not interfere with the simulated non-equilibrium phenomena, requires the thermostat to be applied far away from the sliding region (fullerite surface) and to have a large portion of solid between them in order to favor the thermalization of the excess energy before reaching the bottom of the simulation box. The simulation box details are given in figure 1(a-b). After the characterization of the phase transition we simulated pull-off and sliding friction experiments. In our reference experiment the tip was coated by a monolayer of C_{60} molecules thus, instead of simulating the full silicon nitrate tip, we simply slide or pull the red C_{60} flake of figure 1 (a) which is in direct contact with the Fullerite surface. While the order of magnitude of pull-off and friction forces is the same as in the experiments, the drop at the critical point is much smaller than the measured one, namely only 20 % instead of a factor two. After investigations along different lines we concluded that the only way to produce such a huge drop in both the pull-off and friction forces is through a rotation of the C_{60} flake modifying its degree of commensurability with respect to the surface. When the tip approaches the surface at $T > T_c$ the C_{60} flake attached to the tip will more likely contact the surface in an incommensurate configuration. This results in a small friction and pull-off force. Conversely, when the tip approaches the surface at $T < T_c$ the molecules are not rotating and the surface is ordered, this allows the C_{60} flake to rotate and get in registry with the substrate. In this configuration pull-off and friction forces can increase by an order of magnitude. In conclusion, the phase transition, i.e. the rotation of the C_{60} molecules, is not directly responsible for the drastic change in the frictional properties of the Fullerite surface, it rather triggers a change in the contact geometry between the tip and the surface. The possibility to control the contact geometry, the commensurability and thus the frictional properties of a surface through the occurrence of phase transitions deserves further investigation being of great potential impact in the field of nanomanipulation.

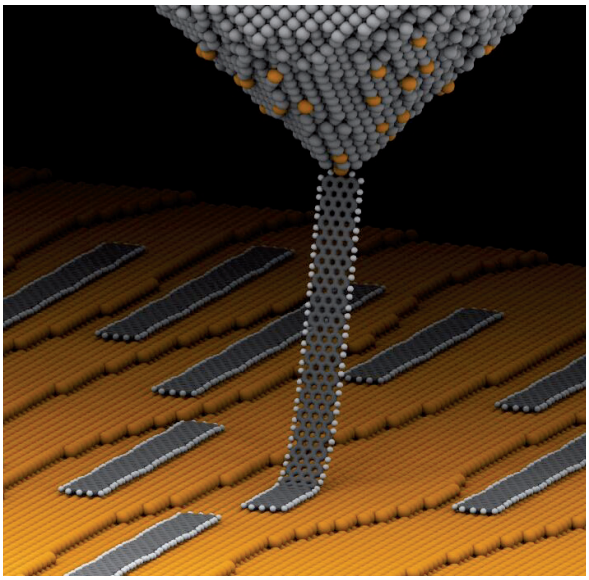

Figure 2: Pictorial representation of the manipulation procedure adopted in our combined experimental theoretical approach to investigate the tribological properties of graphene nanoribbons sliding on gold [7]

The molecular dynamics simulations have been performed with the LAMMPS code [6] an open source state of the art code specifically designed for HPC purposes. The typical production run required 4096 cores for 24 hours.

On-going Research / Outlook

New experiments to probe the dissipation of Fullerite across its rotational melting transition with non-contact AFM are ongoing at Empa and from the theoretical side new simulations will be carried out to understand how to promote or inhibit the occurrence of the phase transition thus gaining a dynamical control of the friction coefficient, this could be in principle achieved applying an external pressure. More recently we demonstrated graphene nanoribbons superlubricity when sliding on gold (Figure 2 and ref. [7]).

References and Links

[1] A. Benassi, A. Vanossi, G.E. Santoro, E. Tosatti Phys Rev Lett. 106, 256102 (2011)
[2] Q. Liang, H. Li, Y. Xu, X. Xiao J. Phys. Chem. B 110, 403 (2006)
[3] A. Benassi, A. Vanossi, C.A. Pignedoli, D. Passerone, E. Tosatti Nanoscale 6, 13163 (2016)
[4] M. Sprik, A. Cheng, M.L. Klein J. Phys. Chem. 96, 2027 (1992)
[5] A. Benassi, A. Vanossi, G.E. Santoro, E. Tosatti Phys. Rev. B 82, 81401 (2010)
[6] http://lammps.sandia.gov/
[7] S. Kawai, A.Benassi, E. Gnecco, H. Söde, R. Pawlak, X. Feng, K. Müllen, D. Passerone, C. A. Pignedoli, P, Ruffieux, R. Fasel, E. Meyer, Science 351, 957 (2016)

Stabilization of ferroelectric properties in Hafnia and Zirconia

Research Institution
Modeling and Simulation Lab, Department of Applied Sciences and Mechatronics, University of Applied Sciences Munich

Principal Investigator
Alfred Kersch

Researchers
Christopher Künneth, Robin Materlik

Project Partners
NamLab gGmbH Dresden, Electronic Materials Research Lab RWTH Aachen

SuperMUC Project ID: pr89pe

Introduction

Ferroelectricity is a property of a few crystalline insulators with partial ionic bonds. In such a crystal the positively charged metal ions and negatively charged oxygen ions can be arranged in structures of different symmetry the possible crystal phases. In a crystal the symmetry of atom arrangement and symmetry of charge arrangement has to be distinguished. In most cases the existing crystals, which are the energetically most favorable structures have a reduced atomic symmetry but the charge distribution is still symmetric. In a ferroelectric crystal a less symmetric phase with asymmetric charge distribution is energetically most favorable. Figure 1 shows the unit cell of the ferroelectric phase of HfO_2 with positively charged Hafnium and negatively charged Oxygen with an asymmetry in the Oxygen subsystem. As a result the crystal has an electric polarization without the application of an external electrical field.

A permanent crystal polarization implies piezoelectric and pyroelectric material properties which are the change of polarization – measurable as electric current or potential – with a relative shift of the positive and negative charges caused by mechanical stress or thermal expansion. Ferroelectrics have a wide range of applications in modern technology starting with the use of permanent polarization in FeRAM cells for data storage applications. As strong piezoelectrics they enable electro-mechanical transducers, as strong pyroelectrics they enable thermo-mechanical transducers including sensors for heat or motion detection and generators for electric energy from waste heat.

This Project is part of a research collaboration of the NamLab gGmbH in Dresden, the Electronic Materials Research Lab at RWTH Aachen, and the University of Applied Sciences Munich [1] with the goal to investigate Hafnia (HfO_2) and Zirconia (ZrO_2) ferroelectrics and their possible applications. This recently discovered [2] new ferroelectric material class is especially attractive for many applications because of their compatibility to silicon microelectronics, their biocompatibility and especially their strong pyroelectric properties. It some applications, it may substitute lead containing compounds.

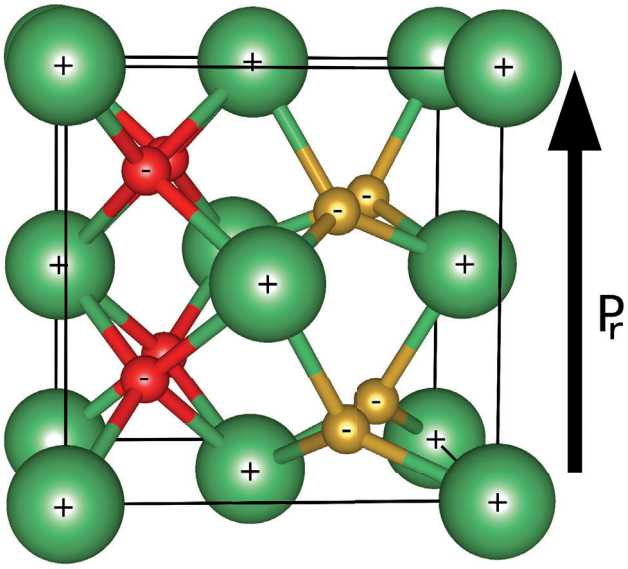

Figure 1: The ferroelectric Pca21 crystal structure of HfO_2. This structure is only one of several competing crystal phases. Green: metal ions, Red: oxygen ions, Gold: stable asymmetric positioned oxygen ions responsible for ferroelectricity. Arrow marks the polarization vector Pr.

Both Hafnia and Zirconia show no ferroelectric behavior in nature and only become ferroelectric under certain experimental conditions. This simulation project aims to understand what those conditions are, why they promote ferroelectric behavior in the mentioned materials, and how they can be improved.

Results and Methods

The crystal structures are simulated on the atomic scale, on the level of quantum mechanics with the open source simulation tool Abinit [3]. The electron distribution in the crystal is calculated with density functional theory, then the forces between the atoms are derived and a new po-

sition of the atoms closer to the equilibrium is searched, iteratively until a convergence is reached. The results include the position of the atoms, the crystal structure, and the energy of the crystal. Different initial positions lead to different crystallographic phases. The energies are finally compared to find the most stable phase depending on the applied conditions. These conditions include pressure, temperature, strain, doping, and defects. For doping and defects on the percent level, structures up to 100 atoms have to be investigated.

Each calculation is scaled on up to 1000 cores and runs for 12 h – 48 h. Due to the large amount of competing crystallographic structures and several relevant conditions that may contribute to the stabilization of the phases,six millions of CPU-hours have been granted.

The first major result is the identification of surface energy as a driving force for ferroelectric phase stabilization and the development of a crystal / surface energy model for the prediction of stable phases of undoped materials [4].

In thin film applications Hafnia and Zirconia are often manufactured with a nano crystalline morphology with grain sizes of approximately 10 nm radius. In such nano grains surface stress and surface energy play a major role for phase stability. Figure 2 illustrates the mechanism: in bulk or in large grains the crystal phase with the lowest energy has typically a reduced symmetry. Such low symmetry phases have disordered surfaces with high surface energy. If the grain size is reduced, the surface to volume ratio increases. Then, a core with a higher symmetric phase and higher crystal energy together with a less disordered surface with lower surface energy may prevail, when the sum of both energies becomes comparatively smaller. In HZO a mixture of HfO_2 and ZrO_2 a phase transition from the monoclinic phase to the ferroelectric phase happens below 15nm. Further reduction of the grain size leads to a further phase transition from the ferroelectric phase to the higher symmetric tetragonal phase below 8nm.

Our model contains the crystal energies under various conditions of temperature and stress/strain, obtained by ab initio calculations at SuperMUC. The influence of temperature and stress/ strain is calculated using approximately one million CPU-hours. This was done for pure Hafnia and Zirconia as well as various mixtures of both.

So far, the values of the surface energies are derived from experimental results, until they are calculated in a later work package. With the model it was possible to predict a new ferroelectric window which was subsequently found in experiment [5]. Further results concern oxygen defects and dopants and how they can promote sweet spots of ferroelectricity in Hafnia. Oxygen vacancies have a stabilizing effect on the ferroelectric tetragonal and cubic phase. However, to stabilize any of those phases with defects alone requires too large concentration. Nonetheless, vacancies widen the window of stability, which is opened by surface energy. Furthermore, vacancies explain very likely ferroelectric phase degradation in applications, because oxygen defects are mobile under an applied electric field.

Substitutional doping of Hf or Zr with dopants like Si, Y, Sr or Gd on the percent level provides annother knob to create windows for the ferroelectric phase. Indeed, the first ferroelectric thin film in Hafnia was 3 % Si: HfO_2. The stabilizing effect of dopants on the ferroelectric phase is very much dependent on concentration and valence. We find, that earth alkaline group elements very well. Together with the surface energy they open the ferroelectric window in a wide range of concentration and grain size. But due to their 2-valent nature they strongly interact with oxygen vacancies. This requires the calculation of a large number of structures to find realistic energies. For the realistic dopant concentrations, large structures with 96 atoms are necessary.

On-going Research / Outlook

Experiments by our collaborators show that not only earth alkaline metals but also lanthanides and boron group metals are capable of inducing ferroelectric behavior in Hafnia. These dopants are known to form more complex defect structures due to their 3-valent nature. This requires further simulation. In some cases computationally more expensive density functionals will be required.

References and Links

[1] http://www.fb06.fh-muenchen.de/fb/index.php/de/ labhome.html?labid=63
[2] Böscke et al., Appl. Phys. Lett. 99, 102903 (2011)
[3] http://www.abinit.org/
[4] R. Materlik et al., J. Appl. Phys. 117, 134109 (2015)
[5] P. Polakowski and J. Müller, Appl. Phys. Lett. 106, 232905 (2015)

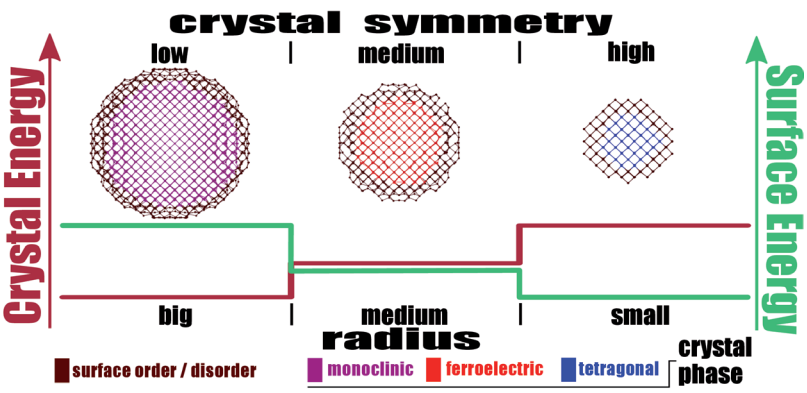

Figure 2: The diagram shows the crystal energy (red line) and the surface energy (green line) of three grain sizes. The surface energy of less symmetric phases is higher because of disorder. The core of the grains is either in the monoclinic, ferroelectric or tetragonal phase. The stability of the phase is determined by the sum of the energies.

Atomistic Insights into Novel Materials and Old Catalytic Problems

Research Institution
Chair for Theoretical Chemistry, TU Munich
Principal Investigator
Karsten Reuter
Researchers
Chiara Panosetti, Mie Andersen
Project Partners
–

SuperMUC Project ID: pr94sa

1. Subproject: Novel Nanoform of Silicon [1]

Endohedrally doped silicon clusters have attracted much interest over the recent years due to their cage-like structures in which the geometry is stabilized either by strong dopant-cage interaction or by external hydrogen passivation of the cage [2]. These structures represent *an intriguing novel nanoform of silicon with great promise for magneto-electronic applications*, such as data storage or optoelectronic devices. The present goal of cluster studies is to move from the gas-phase description to the investigation of their interaction with non-trivial environments such as extended surfaces. Experimentally, the fabrication of chemisorbed silicon structures at surfaces can be realized either through the controlled deposition of pre-formed cages ("soft landing") or via direct silicide formation. Part of this project is carried out in collaboration with Prof. Dähne's group at TU Berlin, investigating the synthesis of rare-earth Si_xM_y clusters on reconstructed Si(111)–7x7 [3]. Among others, a Si_3Tb_3 cluster has turned out particularly intriguing due to the existence of two stable equivalent conformations, rotated by π/3 with respect to each other. Flipping between the two can be induced by applying small voltage biases with the tip of a Scanning Tunneling Microscope (STM), while controlled voltage biases beyond a threshold of ±2 V appear to fix the cluster in either conformation.

Results and Methods
Metal-silicon clusters are stabilized through complex electronic mechanisms combining charge transfer and hybridization. The theoretical investigation of their interaction with an environment requires approaches able to account for such chemical subtleties. It is generally accepted that the standard semi-local formulation for exchange-correlation functionals in Density Functional Theory (PBE-DFT) possesses the desired features and provides an excellent description of systems of this size. We employ the FHI-aims package [4], which uses a localized numeric-atomic orbital (NAO) basis set, which allows for highly efficient calculations of both finite and periodic systems within the same numerical framework, with accuracy up to the order of meV. Despite the efficiency of the electronic structure method, systems of this size can at present only be simulated through highly demanding computing resources such as those available at SuperMUC.

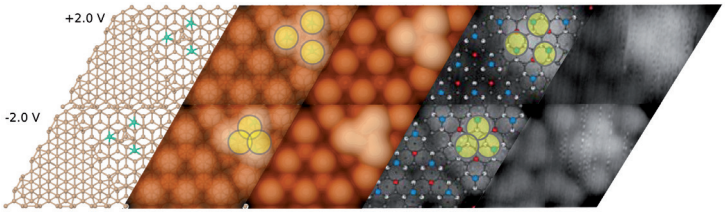

Fig. 1: Computed (colour) and experimental (black and white) STM images of a Si3Tb3 cluster on Si(111)–7x7.

We carried out local structural optimization at the PBE-DFT level, based on the structure assumed to best match the experimental STM images (c.f. Fig. 1). The resulting optimized structure differs from the initial guess in that the Si_3 core is found to be protruding over the Tb3 decoration. This feature was not directly visible experimentally and its determination could only be possible through theoretical calculations. To verify the structure, we simulated the corresponding STM image at both negative (filled-state, sensitive to Si) and positive (empty-state, sensitive to the metal) bias voltages.

We further carried out transition state search to compute the activation barrier for the flipping process. The computed barrier of 1.291 eV is consistent with the observed fix-and-flip behaviour. The simulation also reveals the existence of a symmetric intermediate, again a subtlety which could only be uncovered by means of simulations.

Outlook
Our ultimate aim is to theoretically predict how, and to what extent, the stoichiometry of rare-earth silicides at the Si(111)-7x7 surface can be controlled by finely tuning the environmental conditions during the synthesis. This synthesis can be modelled theoretically by providing a reservoir of gas-phase silicon and metal atoms, which is consistent with the experimental synthesis technique whereby the atomic components are made available by

laser vaporization from the bulk. By varying the availability of the atomic components in the gas-phase (i.e., their chemical potentials), we can determine which conditions lead to the formation of endohedral cages at the surface. This kind of simulation can be addressed by global optimization using the free energy of formation (as a function of the chemical potentials of gas-phase silicon and metal atoms) as the target quantity to minimize. Our hypothesis is that there exists some appropriate tuning of these two chemical potentials during the growth process such that the cages are naturally stabilized at the surface.

2. Subproject: kMC modeling of methane and higher alcohol synthesis [5]

Scaling-relation based kinetic Monte Carlo modeling of methane and higher alcohol synthesis microkinetic modeling has always been a cornerstone of modern heterogeneous catalysis research. A microkinetic model draws on kinetic information of individual elementary processes such as adsorption, desorption and reaction of chemicals at the catalyst surface and analyzes the interplay of these processes within the catalytic cycle to yield the intrinsic catalytic activity or selectivities. Most general microkinetic models thereby aim to explicitly resolve a maximum number of elementary processes, of active site types and even the geometric arrangement and local coverage of chemicals around these sites at the catalyst surface. This requires a very extensive database of energetic input, which is often obtained using density functional theory (DFT) [4]. The enormous computational costs involved in obtaining all required input data from DFT severely limits the system level that can be addressed. A way to overcome this problem is the identification of scaling relations that relate the adsorption strengths of reaction intermediates to those of their constituting base elements or the activation barriers to (thermochemical and computationally much less intense) reaction energies. This allows to express all relevant energetics in terms of typically only one to two key adsorption energies, the so-called descriptors.

Results and Methods

In the present work we focus on methane and higher alcohol synthesis from carbon monoxide and hydrogen gas (synthesis gas). *The conversion of synthesis gas to higher alcohols, remains challenging, as there are no catalysts that have a high enough selectivity and activity to make this process commercially viable.* Initially, we focused on the methane synthesis subnetwork consisting of 10 reaction steps, namely CO and H_2 adsorption, H-assisted CO splitting, C and CH_x hydrogenation steps up to CH_4 (followed by immediate desorption), as well as hydroxyl association and O and OH hydrogenation. A main result of the present project is that the activation energies of these reaction steps have been identified to follow the same scaling relations at different active sites and at different coverages. This shows that the effect of local coverage and of the different sites on the reaction barriers can be predicted from the effect on the adsorption energies alone. This reflects exactly the enormous

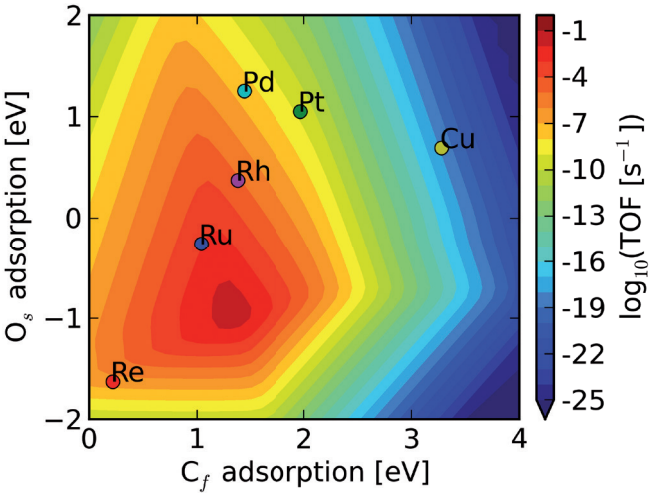

Fig. 2: Calculated catalytic activity measured as turnover frequencies (TOF) for methane production as a function of the C binding on the f step site and the O binding on the s step site as obtained with the mean-field model without adsorbate-adsorbate interactions. Reaction conditions are 523 K and 1 bar with a gas composition of 1 % CO, 97 % H_2, 1 % CH_4, and 1 % H_2O.

reduction in computational cost achieved through the exploitation of such scaling relations.

On the level of the microkinetic model, a great simplification can be obtained by making use of the mean-field approximation (MFA), where the spatial distributions of the chemicals at the surface is averaged to account only for the mean coverage. The presently available scaling-relation based DFT database has been used to compute the theoretical catalytic activity for methane production in the MFA as a function of the carbon and oxygen binding energies at stepped metal surfaces, c.f. Fig. 2. In good agreement with experiments, Ru is predicted as the most active metal.

Outlook

While our model has already yielded good agreement with experiments, an explicit account of the correlations, fluctuations, and detailed spatial distributions of the chemicals at the catalyst surface requires more extensive kinetic Monte Carlo (kMC) simulations. In future work the theoretical catalytic activity with and without effects of local coverage will be compared in the MFA and from kMC simulations. The insight obtained for the methane synthesis sub-network will provide important guidance when extending to the full reaction network up to higher alcohol formation.

References and Links

[1] http://www.th4.ch.tum.de/index.php?id=852
[2] Dennis Palagin and Karsten Reuter. $MSi_{16}H_{16}$ aggregates: from simple building blocks to highly magnetic functionalized materials. ACS Nano, 7:1763–1768, 2013.
[3] M. Franz, S. Appelfeller, M. Rychetsky, and M. Dähne. Formation and atomic structure of self-assembled Dy silicide clusters on the Si(111)–7x7 surface. Surf. Sci., 609:215–220, 2013.
[4] Volker Blum, Ralf Gehrke, Felix Hanke, Paula Havu, Ville Havu, Xinguo Ren, Karsten Reuter, and Matthias Scheffler. Ab initio molecular simulations with numeric atom-centered orbitals. Comp. Phys. Commun., 180:2175–2196, 2009.
[5] http://www.th4.ch.tum.de/index.php?id=853

Numerical simulations of topological and correlated quantum matter

Research Institution
Institut für theoretische Physik und Astrophysik, Universität Würzburg

Principal Investigator
Fakher F. Assaad

Researchers
Florian Goth, Giorgio Sangiovanni, Ewelina Hankiewicz, Michael Karolak, Gang Li, Martin Edelmann, Andreas Hausoel, Thomas Lang, Zi Yang Meng, Domenico Di Sante, Yasir Iqbal

Project Partners
–

SuperMUC Project ID: pr94vu

Introduction

The complexity of the solid state does not allow us to carry out simulations of correlated materials without adopting approximation schemes. In this project we are tackling this daunting task with complementary techniques. On one hand one can start with density functional theory in the local density approximation and then add dynamical local interactions using the so called dynamical mean-field approximation. This approach has the merit of being material dependent in the sense that it is possible to include the specific chemical constituents of the material under investigation. Progress in this domain will be described below. Another venue is to concentrate on phenomena occurring in a class of materials. Here, the strategy is to define models which one can simulate in polynomial time on supercomputing architectures, and which reproduce the phenomena under investigation. This route has been remarkably successful, and we are now in a position to provide controlled model calculations which can cope with antiferromagnetic fluctuations in metals, or nematic instabilities of fermi liquids. Both phenomena are crucial for our understanding of high temperature superconductivity in the cuprates and the pnictides. Access to the LRZ supercomputing center was imperative during the current grant period to do the relevant simulations on a wide range of topics on correlated electrons. In all cases access to supercomputing facilities allows to carry out simulations on larger and larger system sizes so as to be able to extrapolate to the thermodynamic limit relevant for the understanding of experiments and collective phenomena.

Nematic instabilities in Dirac metals

Very recently, there has been considerable progress in quantum Monte Carlo simulations of fermion systems addressing the question of electronic nematic phase transitions. For a square lattice with four fold rotation and translation symmetries, the nematic transition reduces the four fold rotational symmetry to a two fold one but preserves translation invariance. Simple models of SU(2) symmetric fermions coupled to a transverse Ising field model (TIFM) are amenable to sign-free quantum Monte Carlo simulations [1] and allow to study this transition numerically. In the absence of the so-called negative sign problem, the computational effort of quantum Monte Carlo simulations scales as the the cubed power of the number of lattice sites, N, and is proportional to the inverse temperature, beta.Hence, given access to supercomputing resources, detailed information on the nature of the quantum criticality can be achieved. The understanding of this quantum phase transition may shed light on our understanding of underdoped cuprates. In this project, our aim is to understand the nematic phase transition in Dirac metals using the approach described in [1]. In our model calculation, Dirac fermions are captured by a so call π-flux tight binding model, where each plaquette of the square lattice is threaded by half a flux quantum. The coupling between the fermions and Ising variables of the TIFM is such that in the ordered ferromagnetic (FM) phase the fourfold rotational symmetry is broken down to a two fold one, and that in the disordered paramagnetic (PM) phase it is restored. The nature of this transition has been discussed analytically, but with contradicting results. On one hand, a so called ε-expansion around three dimensions was carried out [2] and the results (runaway flow) were interpreted in terms of a first oder phase transition. On the other hand, a so-called large-N approximation carried out in Ref. [3] suggest a

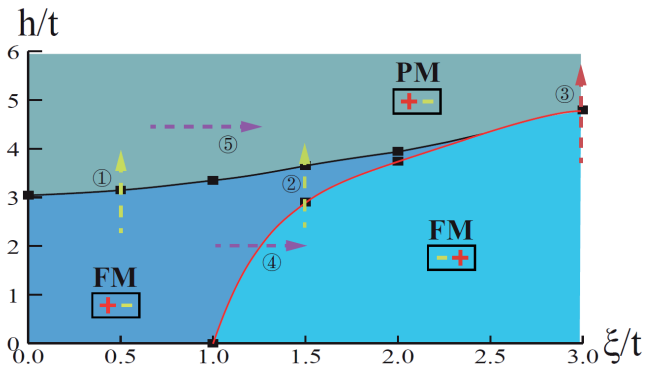

Figure 1: Zero temperature phase diagram in the transverse field versus coupling strength plane.

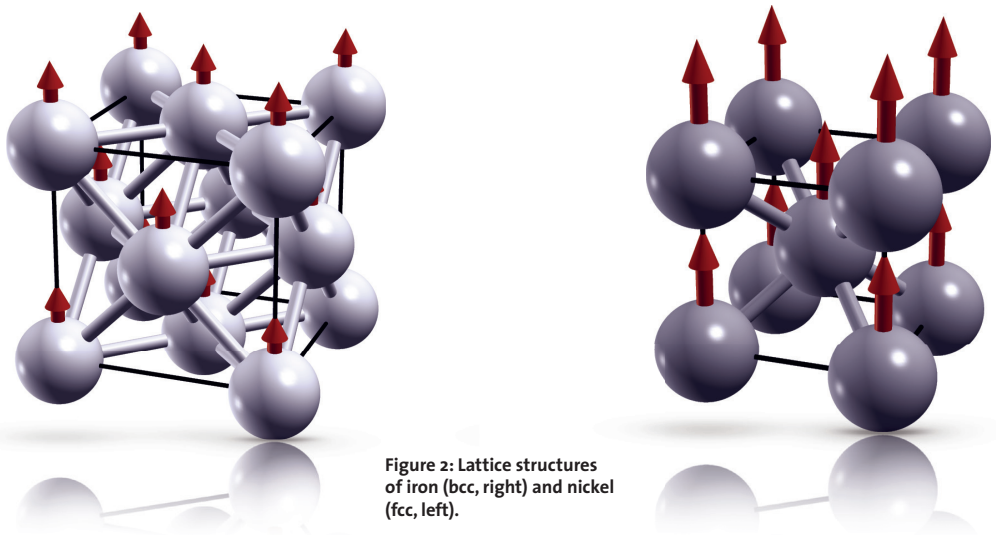

Figure 2: Lattice structures of iron (bcc, right) and nickel (fcc, left).

continuous transition above a critical coupling, where the anisotropy of the nodal quasiparticle dispersion strongly influences the interplay between the critical fluctuations and the quasiparticles. To shed light onto these contradicting analytical results we are presently carrying out large scale numerical simulations. Our results to date are shown in the phase diagram of Fig. 1 in the transverse field, h, versus coupling strength, ξ, plane.

For small values of the coupling strength, it turns out that the anisotropy in the velocity is irrelevant, and essentially renormalizes to zero (Path 1 in the figure). On the other hand, at strong coupling, we observe a first order transition between states where the Berry phase of the Dirac cones switches (Path 3). At h=0 and in the ferromagnetic phase, the switch of the Berry phases between the two cones takes place when one of the velocities vanishes. This yields a Fermi-liquid phase which is expected to be unstable due to the underlying particle-hole symmetry. Further simulations are required prior to publication so as to pin down the details of the phase diagram.

Ab-initio study of the finite temperature magnetism in iron and nickel

Recent developments in electronic-structure numerical algorithms allow to calculate very accurately the band structure of solids as well as the effect of strong interactions among the conduction electrons. The consequences of the Coulomb repulsion can be so dramatic that materials with partially filled d- and f-shells turn insulating even without the "help" of magnetic long-range order. These are called Mott insulators and several realistic studies have helped so far to reveal their intriguing properties. A less explored direction is to study the transition from Mott insulators to magnetically ordered ones.

The interest is to see to what extent our classical picture of ferromagnets, i.e. the Curie-Weiss theory or Pauli theory, can still be applied to strongly-correlated quantum magnets. Curie-Weiss theory describes localized electrons as they are present in Mott-insulators, whereas in the Pauli theory the electrons are delocalized i.e. they retain their Bloch nature.

Hovever real solids are in neither of the two limits, and up to now there exists no way to describe these two aspects in one analytic theory.

By combining the state-of-the-art band structure methods (density functional theory, implemented in the VASP package) and many body methods such as Dynamical Mean Field Theory, we attack this problem in the exemplary cases of bcc-Iron and fcc-Nickel with heavy numerics. Despite the fact that the lattice structures of iron and nickel are very simple (see figure 2), these materials have partially filled valence d-shells that hybridize with the valence p- and s-shells. Thus they are strongly correlated and have a very large low-energy Hilbert space.

Our code package *W2DYNAMICS* (Würzburg/Wien strong coupling solver), a Continuous Time Quantum Monte Carlo impurity solver developed by us, can tackle this NP-hard problem [4,5]. We highly optimized it over several years to treat systems that have a size at the border of feasibility. The massively parallel structure of SUPERMUC allows us to explore with unprecedented accuracy -- and for parameter regions never calculated before -- the electronic and magnetic properties of iron and nickel emerging from electronic correlations. The magnetism is of particular interest, because it is still to some extent not completely understood. This is quite remarkable as these ``old" magnets are very well characterized experimentally and have a very simple lattice structure (figure 2).

References and Links

[1] Y. Schattner, S. Lederer, S. A. Kivelson, and E. Berg, "Ising nematic quantum critical point in a metal: a Monte Carlo study", ArXiv e-prints (2015).

[2] M. Vojta, Y. Zhang, and S. Sachdev,"Quantum Phase Transitions in d-Wave Superconductors," Phys. Rev. Lett. 85, 4940 (2000).

[3] E.-A. Kim, M. J. Lawler, P. Oreto, S. Sachdev, E. Fradkin, and S. A. Kivelson, "Theory of the nodal nematic quantum phase transition in superconductors," Phys. Rev. B 77, 184514 (2008).

[4] Nicolaus Parragh, Alessandro Toschi, Karsten Held, and Giorgio Sangiovanni, "Conserved quantities of SU(2)-invariant interactions for correlated fermions and the advantages for quantum Monte Carlo simulations," Phys. Rev. B 86, 155158 (2015)

[5] P. Gunacker, M. Wallerberger, E. Gull, A. Hausoel, G. Sangiovanni, and K. Held, "Continuous-time quantum Monte Carlo using worm sampling," Phys. Rev. B 92, 155102 (2012)

Earth and Environmental Sciences

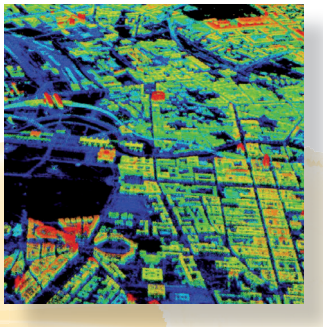

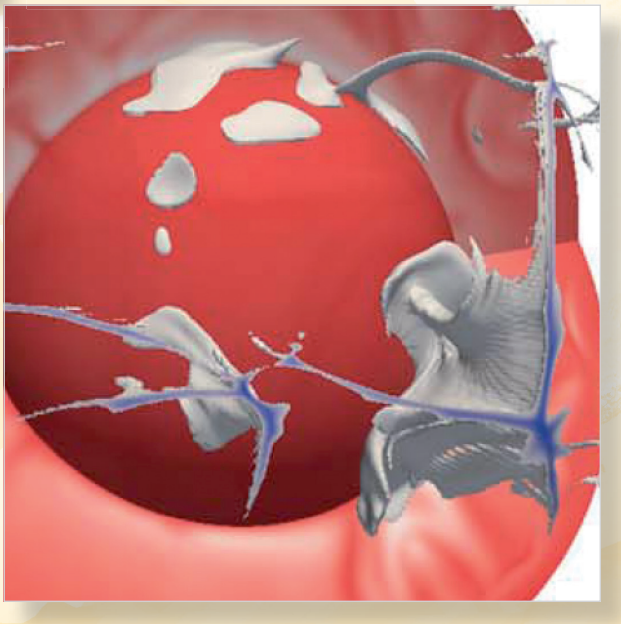

Synthetic Earth Models – Taming Chaotic Mantle Flow and Exploiting Seismic Wavefield Effects to Constrain Buoyancy

Research Institution
Geophyics Section, Department for Earth and Environmental Sciences, Ludwig-Maximilians-Universität München

Principal Investigator
Bernhard Schuberth

Researchers
Lorenzo Colli, Hans-Peter Bunge, Christophe Zaroli and Guust Nolet

Project Partners
Géoazur (Observatoire de la Côte d'Azur, Université de Nice - Sophia Antipolis, France),
Institut de Physique du Globe Strasbourg (Université de Strasbourg/EOST, France)

SuperMUC Project ID: pr32mu

Introduction

Much of the geological activity of the Earth arises from convective processes within the mantle that transport heat from the deep interior of our planet to the surface. Accurate knowledge of the associated forces that drive plate tectonics is of fundamental importance for studying natural hazards such as earthquakes. However, we still lack a quantitative understanding of the buoyancy distribution in Earth's mantle that ultimately drives all surface motions. In this multi-disciplinary project we combine simulations and post-processing tools from geodynamics, mineral physics and seismology, in order to improve conceptual models of Earth's deep interior. The general work-flow involves the generation of new mantle convection models with either updated sets of input parameters or increased resolution and improved numerical representation. The crucial step is then to comprehensively assess the quality of these models against a range of observations of the Earth system (including geologic, seismological and geodetic information). This involves two additional sets of simulations: First, the inverse modelling of mantle convection that allows one to track mantle motion back into the past. This way, one can test unknown parameters of the geodynamic models explicitly against the geologic record. Second, the simulation of 3-D global seismic wave propagation through the geodynamic models, which enables us to test the models directly against the huge amount of seismic data available nowadays.

Results and Methods

Our numerical simulations of mantle flow and seismic wave propagation place formidable demands on parallel computing resources, as they require extremely high resolutions in space and time in order to allow for models with earth-like physical parameters. Forward problems require the solution of the governing equations on computational grids with billions of degrees of freedom and thousands of time steps. Production runs on SuperMUC were performed using between 512 and ~4000 cores and required around 50,000 CPU-h each. The inverse problem of mantle convection is much more CPU-time intense as it requires a number of iterations on the forward-inverse cycle. In addition, very large parallel storage and I/O capacity is needed, as one must preserve all degrees of freedom per time step over the entire simulation period (~10 TB temporary disk space per simulation).

Taming the Chaotic Nature of Mantle Flow
Although stronger than steel and capable of transmitting seismic shear waves, the mantle can be treated as a fluid on geologic time scales (i.e., millions of years). Thus, mantle convection is governed by the hydrodynamic field equations (mass, momentum and energy conservation). Our simulations in 3-D spherical shell geometry are performed using the MPI parallelised finite-element code *TERRA*. The code avoids indirect memory addressing and allows for efficient data-structures leading to excellent

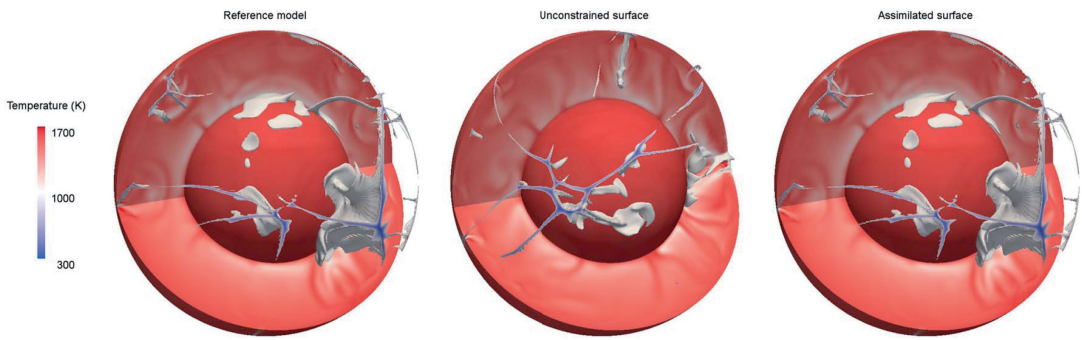

Figure 1: 3-D view of the temperature field of a MCM with isosurfaces at 1000° C. Left: reference run; middle; unconstrained twin experiment; right: twin experiment evolved with assimilation of surface velocities from the reference run. The thermal fields of reference run and twin with assimilated surface velocities are nearly indistinguishable. From [2].

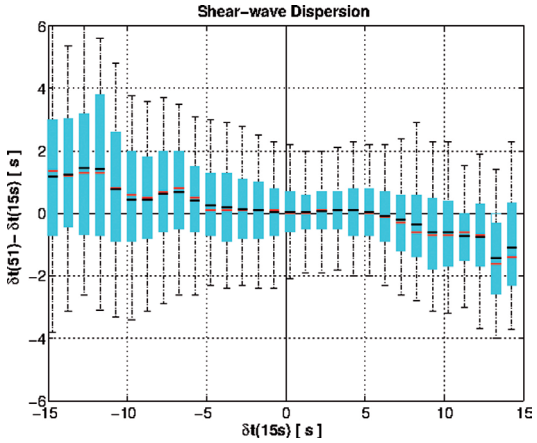

Figure 2: Diffraction-induced traveltime dispersion in our MCM (i.e., distributions of the long-period–short-period differences as a function of the short-period residual; Box-plots show (red) median, (black) mean and 25–75 and 5/95 percentiles). Differences between long-period and short-period residuals tend to increase with increasing short-period signal. From [3].

performance. Currently, we model mantle flow with 80-800 million finite elements and construct socalled mantle circulation models (MCMs), akin to the circulation models used in ocean and climate simulations. MCMs integrate mantle flow forward in time from an assumed initial state some time in the past using reconstructions of past plate motions as surface boundary condition for velocity. The problem is that initial condition and model parameters are not well known. Different assumptions and approximations lead to vastly different predictions. Alternatively, an initial condition for the present day can be estimated from seismic tomography of Earth's mantle. It is thus possible to evolve it into the future in accord with the governing equations. Unfortunately, as mantle convection develops over timescales of millions of years, these predictions and the underlying assumptions cannot be tested. Modeling past states of Earth's mantle, on the other hand, and relating them to geologic observations such as continental-scale uplift and subsidence, is an effective method for testing MCMs explicitly in time. However, mantle convection is chaotic and two identical models initialized with slightly different temperature fields diverge exponentially in time until they become uncorrelated, thus limiting retrodictions (i.e., reconstructions of past mantle states using present information) to the recent past. Given the uncertainty with which we know the present-day thermodynamic state of Earth's mantle, the limit of predictability for mantle convection models was estimated to be around 100 Ma at best. Using 3-D spherical mantle convection models, we recently showed that retrodictions can be extended significantly if knowledge of the surface velocity field is available [2]. Assimilating surface velocities produces in some cases negative Lyapunov times (i.e., e-folding times of the discrepancy between predicted and true state), implying that even a severely perturbed initial condition may evolve toward the reference state (Fig. 1). A history of the surface velocity field for Earth can be obtained from plate motion reconstructions for time periods of a mantle overturn (~200 Ma), suggesting that mantle flow can be reconstructed over comparable times.

Exploiting Seismic Wavefield Effects to Constrain Mantle Buoyancy

As an alternative to inverse flow simulations, we had demonstrated the great potential of computing 3-D wave fields in MCMs for testing the underlying assumptions. We use the spectral-element code SPECFEM3D_GLOBE to solve the seismic wave equation, and only as of late, HPC resources allow for simulations at relevant seismic frequencies. In a recent study, we have investigated the strength of diffraction-induced dispersion of traveltimes of seismic waves in MCMs [3]. This type of observation is expected to help constrain the length-scales and magnitude of buoyancy forces in the mantle. The synthetic multi-frequency measurements obtained for our MCM allowed us – for the first time – to characterize the wavefield effects to be expected in a mantle with earth-like structural length scales and realistic magnitudes of seismic heterogeneity. We found that our synthetic data do indeed show significant diffractioninduced dispersion between the longest and shortest period considered (see Fig. 2). The dispersion, defined here as the difference between traveltime residuals measured in different frequency bands, can also be regarded as a new seismic datum: a differentialfrequency time residual. In our analysis we concluded that this datum is to some extent insensitive to shallow structure and that it may potentially prove useful to improve tomographic models of the lower mantle. Overall, the dispersion of traveltime residuals introduced by virtue of diffraction is quite substantial and significant not only with respect to the magnitude of the residuals themselves, but also with respect to the dispersion observed in real data. If the information contained in the observed dispersion is fully exploited through multifrequency inversions, Earth's deep interior can likely be imaged at unprecedented resolution and accuracy.

Outlook

Our results on retrodiction of mantle flow showed that the assimilation of surface velocities prevents the chaotic long-term drift of unconstrained convection. Past states of mantle flow can thus be reconstructed for periods comparable to the available time span of past plate motion models. We next plan to systematically scan the parameter space of the inverse geodynamic simulations. Our aim is to push the limits of resolution and to solve problems with 10^{12} unknowns. Current computing resources still prevent us from running models with a number of relevant strength profiles. With SuperMUC "Next Generation", it will be possible to overcome this limitation and to perform a comprehensive systematic parameter space search. Furthermore, we will finally be able to fully resolve the extremely high Rayleigh numbers relevant for Earth's mantle.

References and Links

[1] www.geophysik.uni-muenchen.de/research/geodynamics
[2] Colli, L., H.-P. Bunge, and B.S.A. Schuberth, (2015), Geophys. Res. Lett., 42, 8341–8348.
[3] Schuberth, B.S.A., C. Zaroli and G. Nolet, (2015), Geophys. J. Int. 203, 2099–2118.

High resolution gravity field modeling

Research Institution
Institute of Astronomical and Physical Geodesy

Principal Investigator
Thomas Gruber

Researchers
Thomas Fecher

Project Partners
–

SuperMUC Project ID: pr89qu, pr32qu

Introduction

The static gravity field of the Earth is one of the key parameters for the observation and measurement of a number of processes and flows in the dynamic system of the living planet Earth. Its knowledge is of importance for various scientific disciplines, such as geodesy, geophysics and oceanography. For geophysics the gravity field gives insight into the Earth's interior, while by defining the physical shape of the Earth it provides an important reference surface for oceanographic applications, such as the determination of sea level rise or modelling of oceanographic currents. Moreover this reference surface is a key parameter on the way to a globally unified height system.

The scientific goal is to estimate the static gravity field as precise and detailed as possible. As the gravity field in general is represented by a spherical harmonic series, the parameters to be estimated in gravity field modelling are spherical harmonic coefficients. There exist various techniques to observe the gravity field, which have different advantages and complement each other. The observation of the Earth gravity field from dedicated satellite missions delivers high accurate and globally homogenous gravity field information for the long to medium wavelengths of the spherical harmonic spectrum (corresponding to spatial resolutions down to roughly 100 km). However, due to the large distance between the satellite and the Earth's surface, the gravity field signal is damped in satellite height. Therefore short wavelengths of the spherical harmonic spectrum (smaller than 100 km) cannot be observed from space. To complement the satellite information, terrestrial gravity field measurements over land and satellite altimeter observations over the oceans, which need to be converted to gravity field quantities, are used as additional data. As these observations are taken at the Earth surface (land and ocean) they contain the full undamped signal. The scientific challenge is to combine the different types of gravity field observations in the way that all data types keep their specific strengths and are not degraded by the combination with other information in specific spherical harmonic wavelength regimes. As mentioned,

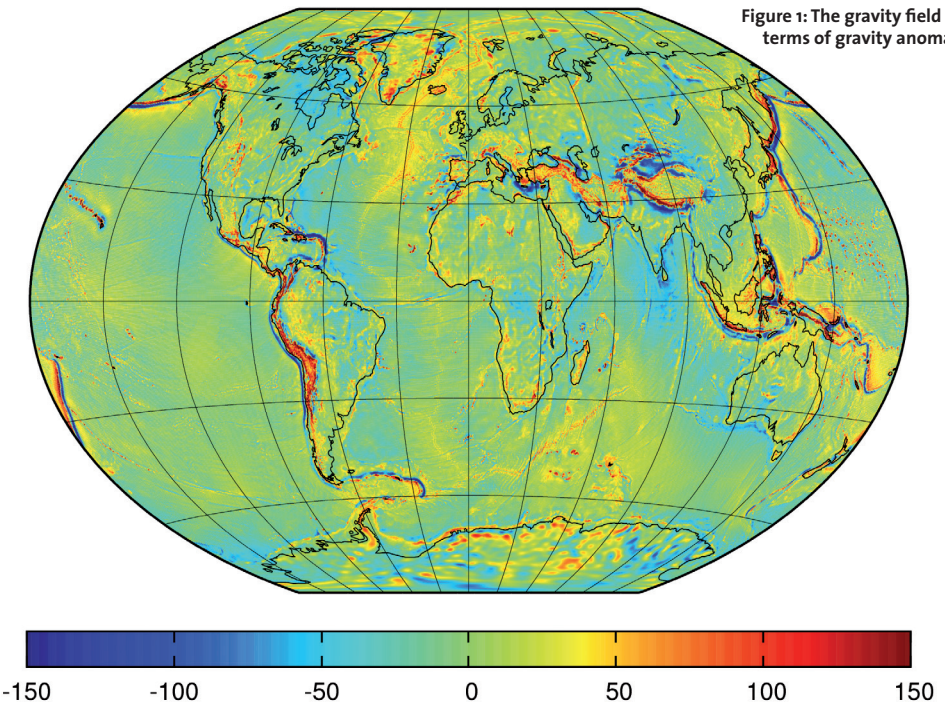

Figure 1: The gravity field of the Earth in terms of gravity anomalies [mgal].

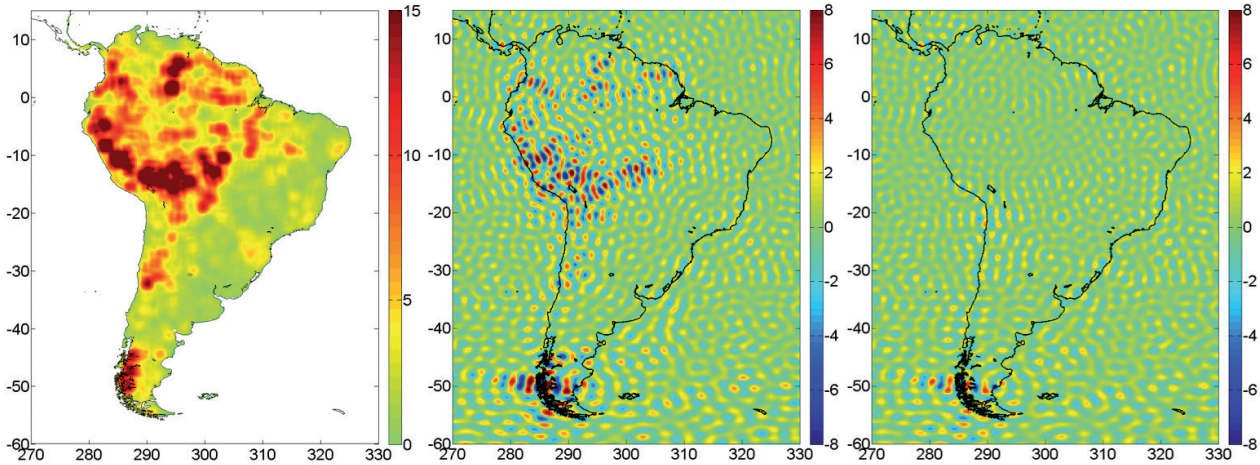

Figure 2: Accuracy map of terrestrial observations in South America [mgal] (left); difference between the block-diagonal (middle) and full normal (right) equation solution to the satellite information [mgal].

this procedure shall result in a set of spherical harmonic coefficients representing the global Earth gravity field up to highest possible resolution.

Strict full normal equation approach and the need for supercomputing

Our approach to determine the spherical harmonic coefficients is based on a strict least squares adjustment with a Gauß-Markov model. By that, the different data types can be combined optimally on normal equation base. The approach enables the ideal relative weighting between different data sets as well as the individual weighting of every single observation, what is important as especially the quality of terrestrial measurement data differs significantly (observations are in general quite accurate in regions such as Europe or North America, whereas they are worse in Africa or South America). Due to the high correlation of the unknown spherical harmonic coefficients, the corresponding normal equation system is a dense matrix. As the number of unknowns increases quadratic with the spherical harmonic degree, and again the number of elements of the full normal equation matrix is quadratic to the number of unknowns, full normal equation systems become quite large. This is a computational challenge, which can only be solved with supercomputers such as SuperMUC. With the available data we currently estimate the coefficients of spherical harmonic expansion up to degree and order 720, which corresponds to more than 500 000 unknowns and a normal equation system of 2 TByte. With availability of denser ground data coverage further extensions to even higher resolutions are planned for the near future.

Results

Below the benefit of the full normal equation approach is demonstrated exemplary for a gravity field solution in the area of South America. As observations data sets from the US-German satellite mission GRACE (Gravity and Climate Recovery Experiment), the satellite mission GOCE (Gravity Field and Steady State Ocean Circulation Explorer) of the European Space Agency (ESA), and terrestrial observations provided by the National Geospatial Intelligence Agency (NGA) are available. The quality of the terrestrial data differs significantly, e.g. it is quite high for example for coastal areas in Brazil and low for areas in the Amazon area. The major challenge is to find the optimal combination of the different data sets.

For that purpose, we derive first an accuracy map of the terrestrial observations by comparison with the satellite information in the low to medium wavelengths (Fig 2 left). In the following, we calculate two different gravity field solutions. One is based on a reduced block-diagonal normal equation approach, which requires strong data constraints like equal data weights, but can be executed on a single node computer system. The second approach is our full normal equation approach applying individual weights according to a predetermined accuracy map and using supercomputing facilities. One can investigate the difference between the two results and the original satellite information in the low to medium wavelengths, which delivers a measure of how good the estimation procedure is able to recover the field in this frequency range. The parameter estimation process can be regarded as good, when the differences are small, as that proofs, that the good performance of the satellite information was not degraded by the combination with terrestrial data. It is clearly visible, that the differences are quite high for the reduced block diagonal approach (Fig 2 middle). Bad data in the Amazon area and in the Andes affected the combination solution. With full normal equations and supercomputing approach, nearly no differences are visible (Fig 2 right). The combination is ideal for that case, what demonstrates the power of using full normal equation systems and supercomputing faciltiies for this purpose.

Details about the procedure and results can be found in [1].

References and Links

[1] Fecher, T.: Globale kombinierte Schwerefeld-modellierung auf Basis voller Normalgleichungssysteme; Dissertation, Ingenieurfakultät Bau Geo Umwelt (BGU), TU München, 2015.

EXtreme PREcipitation and Hydrological climate Scenario Simulations (EXPRESS-Hydro)

Research Institution
Leibniz-Rechenzentrum, Garching bei München, Germany

Principal Investigator
Dieter Kranzlmüller

Researchers
Jost von Hardenberg[1], Antonio Parodi[2], Alexandre Pieri[1], Antonello Provenzale[1]

Project Partners
[1]Institute of Atmospheric Sciences and Climate - National Research Council (ISAC-CNR) and
[2]CIMA Research Foundation, Italy

SuperMUC Project ID: pr45de (Gauss Large Scale project)

Introduction

Predicting weather and climate and its impacts on the environment, including hazards such as floods, droughts and landslides, continues to be one of the main challenges of the 21st century with significant societal and economic implications. At the heart of this challenge, lies the ability to have easy access to hydrometeorological data, to share predictive models, and to facilitate the access to High Performance Computing facilities supporting leading edge hydro-meteorological simulations. Advances in the science and observation of climate change are providing a clearer understanding of the inherent variability of Earth's climate system and its likely response to human and natural influences. Numerical simulations at the global scale of future climate in different emission scenarios, produced in the framework of the fifth climate model intercomparison project (Coupled Model Intercomparison Project Phase 5, CMIP5), have been typically obtained using hydrostatic global climate models (GCM) and are available only at quite coarse spatial resolutions which do not allow an accurate representation of intense precipitation events over complex topography areas such as Europe. An improved generation of regional climate models (RCM) is being developed and applied in the framework of the CORDEX (COordinated Regional climate Downscaling Experiment) initiative, with the aim to produce regional climate change projections worldwide for input into impact and adaptation studies, considering multiple forcing GCMs from the CMIP5 archive. Recent results for the European branch of the CORDEX initiative (Kotlarski et al. 2014) confirm, with simulations on grid-resolutions up to about 12 km (0.11°), the ability of RCMs to capture the basic features of the European climate for the period 1989-2008, but also show non-negligible deficiencies of the simulations concerning selected metrics, certain regions and seasons: for example seasonally and regionally averaged temperature biases are mostly smaller than 1.5 °C, while precipitation biases are in the ±40% range.

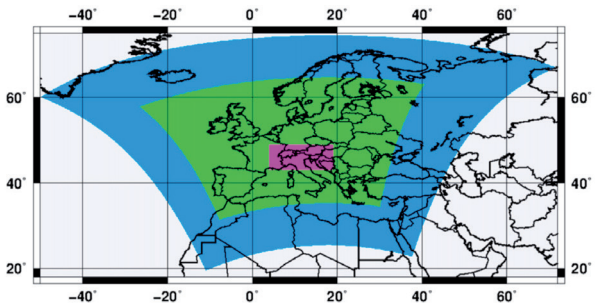

Figure 1: European domain defined in CORDEX (0.118; blue) and the IER (0.0378; green) used for the high-resolution integration. The Great Alpine Region (GAR) used for some of the diagnostics is displayed in purple (courtesy of Pieri et al. 2015).

Results and Methods

This reality, together with the fact that the European region is exposed to intense Atlantic synoptic perturbations and potentially their north shift due to climate change, make it critical to carefully understand all model sensitivities before drawing conclusions and assessing climate change impacts from the numerical simulations outputs. Starting from these ideas, the EXtreme PREcipitation and Hydrological climate Scenario Simulations (EXPRESS-Hydro) project moves one step further, performing very high-resolution regional dynamical downscaling of historical climate scenarios produced by the ERA-Interim reanalysis using the state-of-the-art non-hydrostatic Weather Research and Forecasting (WRF) regional climate model.

To the best knowledge of the current literature, EXPRESS-Hydro performed for the first time long climate simulations (1979-2008, Pieri et al. 2015) over the European domain (Inner European Region, IER, Figure 1.) at a very fine cloud-permitting resolution of about 4 km (0.037°) with explicitly resolved convection and a sharp representation of orography, thanks to the possibility of running very computationally and data storage demanding simulations (about 700 Tb of outputs produced) on the state-of-the-art SuperMUC Petascale System, one of the fastest supercomputers in the world.

EXPRESS-Hydro explored the WRF regional climate model capability in reproducing observed precipitation extremes, over Europe and in detail over the Greater Alpine Region, comparing simulations results with available high-resolution gridded observational data sets. Overall our results indicate that increased resolution with explicitly resolved convection helps to obtain a closer representation of the precipitation field, reducing the overestimation of precipitation (about 25% on average over the European domain) and allowing to better reproduce the distribution and the statistics of the rainfall rate, particularly over the Alps (Figure 2).

EXPRESS-Hydro complemented the present climate with a future scenario projection at 0.11°, in which the WRF model is forced with CMIP5 output from the EC-Earth global climate model, with RCP4.5 greenhouse gas and aerosol concentrations. Furthermore part of the EXPRESS-Hydro team undertook also high-resolution simulations, using again SuperMUC Petascale System, down to the cloud-resolving range (200 m), for high impact weather events (HIWE) occurring in the Mediterranean region, such as the recent Genoa 2011 and Genoa 2014 flash-flood events (Fiori et al. 2014, Fiori et al. 2015, Hally et al. 2015). Fiori et al. (2015) Taking advantage of the availability of both observational data and modelling results (WRF-ARW runs) at the micro-α meteorological scale (2 km – 0.2 km and 1 hour or less, Orlanski, 1975), have provided new insights about the triggering mechanism and the subsequent spatio-temporal evolution of Genoa 2014 HIWE.

Outlook

The post-processed datasets produced in this project will be used as inputs for regional impact studies, in particular forcing hydrological models and ecosystem dynamics models over selected study areas. The data will be distributed in the framework of major international projects such as DRIHM, and the Italian Project of Interest NextData.

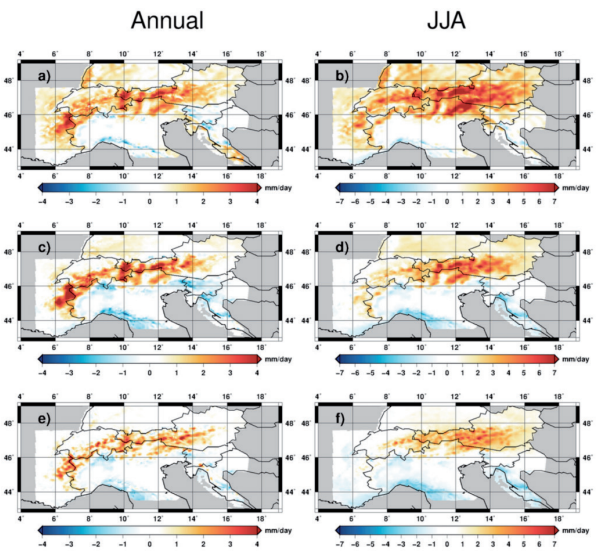

Figure 2: Differences in the (left) annual and (right) june-july-august (JJA) precipitation climatology over the GAR between WRF and EURO4M-APGD for (a),(b) 0.11° run using Kain-Fritsch convection parameterization; (c),(d) 0.11° run using Betts-Miller-Janjic convection parameterization; and (e),(f) 0.037° run using explicit convection (courtesy of Pieri et al. 2015).

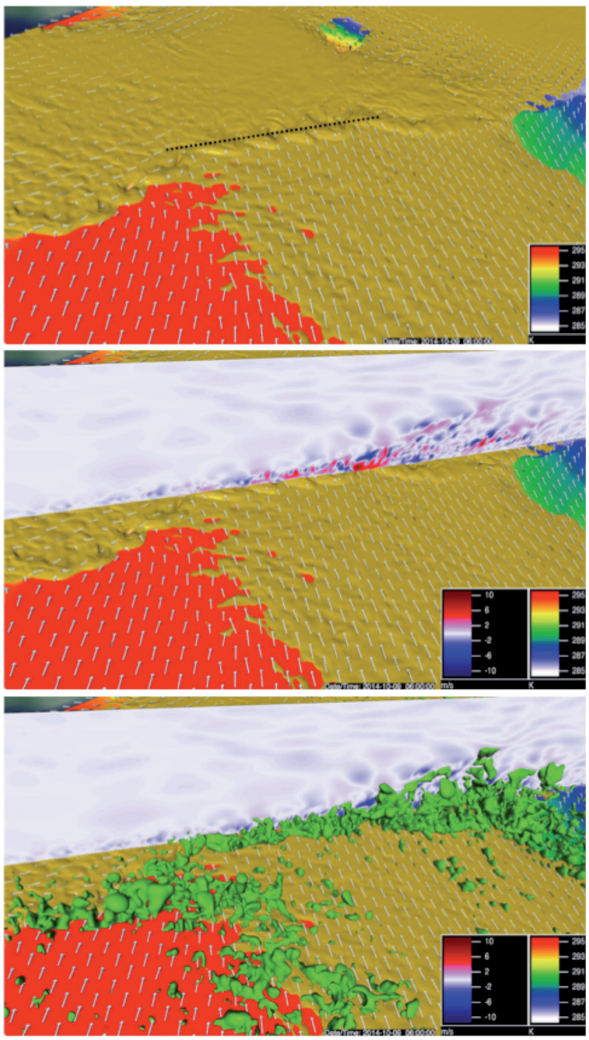

Figure 3: Genoa 2014 WRF-ARW 200 m run. Upper panel, T 2m temperature field and potential temperature isosurface (θ=293.5 K) at 6:00UTC; middle panel, T 2m temperature field, potential temperature isosurface (θ=293.5 K), and vertical velocity field cross section, at 6:00UTC; lower panel, T 2m temperature field, potential temperature isosurface (θ=293.5 K), vertical velocity field cross section, and vertical velocity isosurface (W=1 m s-1) at 6:00UTC.

References

[1] Fiori, E., Comellas, A., Molini, L., Rebora, N., Siccardi, F., Gochis, D. J., ... & Parodi, A. (2014). Analysis and hindcast simulations of an extreme rainfall event in the Mediterranean area: The Genoa 2011 case. Atmospheric Research, 138, 13-29.

[2] Fiori, E., Ferraris, L., Molini, L., Siccardi, F., Kranzlmueller, D, Parodi, A. (2015). Morphology of the triggering and evolution of a micro-α convective system in the Mediterranean Sea. Quarterly Journal of the Royal Meteorological Society, submitted.

[3] Hally, A., Caumont, O., Garrote, L., Richard, E., Weerts, A., Delogu, F., ... & Clematis, A. (2015). Hydrometeorological multi-model ensemble simulations of the 4 November 2011 flash flood event in Genoa, Italy, in the framework of the DRIHM project. Natural Hazards and Earth System Science, 15(3), 537-555.

[4] Pieri, A. B., Hardenberg, J. V., Parodi, A., & Provenzale, A. (2015). Sensitivity of precipitation statistics to resolution, microphysics and convective parameterization: a case study with the high-resolution WRF climate model over Europe. Journal of Hydrometeorology, (2015).

Dynamic Rupture Simulations at Petascale

Research Institution
Technische Universität München, Ludwig-Maximilians-Universität München

Principal Investigator
Michael Bader, Alice Gabriel

Researchers
Alexander Breuer, Alexander Heinecke, Sebastian Rettenberger

Project Partners
Intel Parallel Computing Lab

SuperMUC Project ID: pr45fi, pr83no

Introduction

The simulation of earthquakes and of propagating seismic waves can target quite a variety of questions, such as the estimation of earthquake hazards for particular regions or to study how seismic waves propagate through the Earth's crust. In *dynamic rupture simulations*, the focus is on simulating the key process that initiates strong ground motions: assuming a known fault system and certain initial stresses, we simulate how the rupture process propagates along the fault, how seismic waves evolve from this rupture process and even how seismic waves interact with the rupture process. The dynamic rupture code SeisSol simulates such events in a coupled way and together with the resulting wave propagation, which allows for physics-driven simulation of strong ground motions during earthquakes.

As a physical process, dynamic rupture covers various length scales. While on the fault plane the dynamic rupture process happens on a meter-scale, fault systems reach a couple of kilometers deep into the crust, and wave propagation and destructive ground shaking affects regions that can extend over hundreds of kilometers. To capture these processes as accurately as possible, simulations with high resolution are necessary. Our largest simulation so far used a mesh with close to 200 million grid cells, approximately 100 billion degrees of freedom and more than 200,000 time steps. Simulations of this size ask for petascale performance. As productive simulations need to study a wide range of parameters and thus have to be repeated again and again, optimizing the software for best-possible time-to-solution and achieving performance as close as possible to the machine peak is imperative. In our project, we optimized SeisSol for the SuperMUC phase 1 and 2 platforms, such that it achieved a performance of nearly 1.5 PFlop/s in production runs.

Methods and Results

Turning SeisSol into a petascale-ready code required improvements and optimizations along the entire simulation pipeline. To perform large simulations on the full SuperMUC machine required a redesign of the input/output routines. Substantial improvements in node-level computational performance were achieved by optimization of all computational kernels based on a code-generation approach for the element-local matrix operations. Finally, parallel scalability greatly profited from moving to a hybrid MPI+OpenMP parallelization.

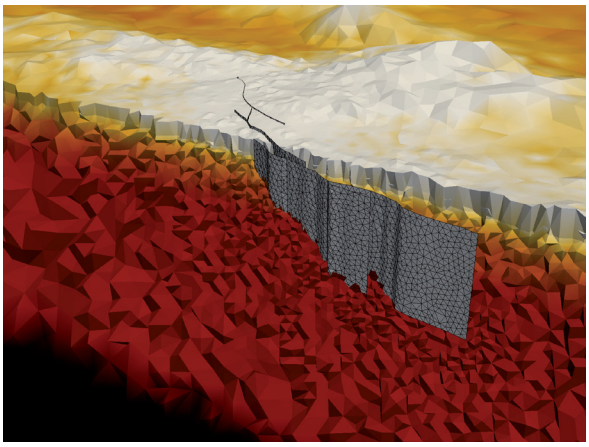

Figure 1: Illustration of the setup for the simulation of the Landers 1992 Earthquake. The branched fault system (in grey) is discretized with a resolution of approx. 200m. The coloring (yellow to red) indicates the increasing velocity of seismic waves in the crust.

Scalable I/O for Petascale Simulations

SeisSol works on unstructured tetrahedral meshes to account for complicated geometries of topography and complex branched fault systems. The traditional workflow – generate such meshes from CAD models, use standard graph partitioning packages to compute balanced partitions and read mesh file and partitioning file

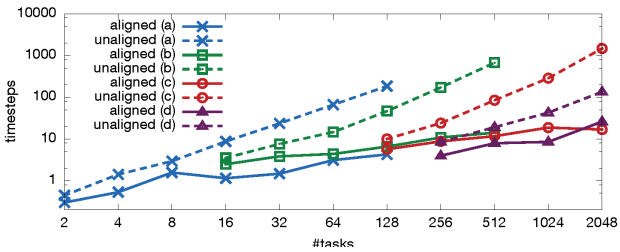

Figure 2: Overhead caused for writing the wave field output in SeisSol (relative to the cost of a single timestep). Notice the substantial reduction in writing times when aligning output sizes to the block size of the file system (by combining output from several ranks). [4]

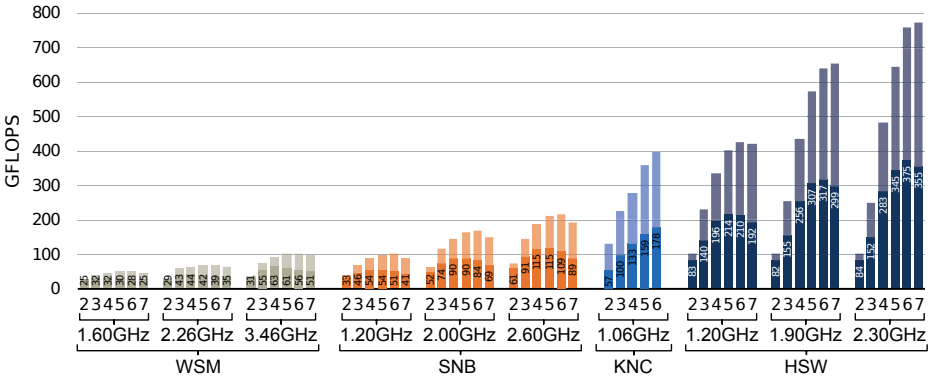

Figure 3: Performance (in GFlop/s) achieved by SeisSol for the LOH.1 benchmark – given are results for several Intel architectures (Westmere, Sandy Bridge, Knights Corner, Haswell), disretization orders 2–7, and varying clock frequencies. The dark-colored part reflects the amount of operations without padding effects. [3]

into SeisSol – did not scale beyond a few thousand MPI processes and a few million elements. We therefore designed a novel two-step workflow that combines mesh generation and partitioning into a preprocessing tool, PUMgen. SeisSol can read PUMgen files using fast parallel I/O, such that meshes with hundreds of millions of grid cells are now read on the full SuperMUC machine in less than a minute. Similarly, SeisSol's output routines were optimized by automatically reducing the number of writing nodes to speed up data transfer – see Figure 2.

Optimization of Node-Level Performance

SeisSol uses the Arbitrary-high-order-DERivative Discontinuous Galerkin (ADER-DG) method for discretization and time stepping. Its implementation can be expressed via element-local multiplication of small matrices. While discretization matrices are sparse with fixed sparsity pattern, the matrix of quantities is dense, but with a small fixed number of columns (corresponding to the nine quantities of the underlying PDE system). SeisSol relies on code generation for these sparse-dense matrix operations to obtain highest-possible performance. For each sparse-dense operation and each occurring sparsity pattern, either a sparse implementation (using unrolling of operations) or a tailored dense matrix kernel is selected via an autotuning approach. Figure 3 illustrates the achieved performance for the LOH.1 benchmark on a range of Intel platforms.

Scalability on SuperMUC

SeisSol relies on a hybrid MPI+OpenMP parallelization. Especially on full-machine runs, the substantial reduction of MPI ranks resulting from OpenMP parallelization reduced the scalability demands on the MPI system and led to strongly improved performance. Recently, overlapping of computation and communication was implemented. As a result, a recent simulation of the Landers 1992 scenario (with 191 million grid cells and 96 billion unknowns) was executed [5] with a performance of 1.4 PFlop/s on SuperMUC phase 2. The strong scaling from 128 to 3,072 nodes (with a parallel efficiency of 94%) is shown in Fig. 4.

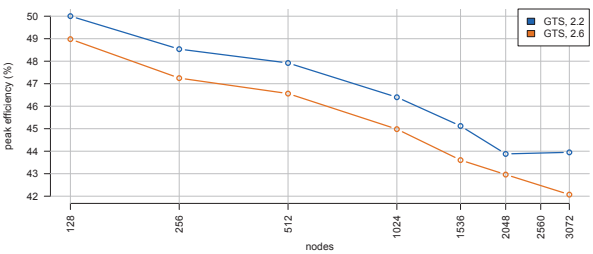

Figure 4: Percentage of peak performance obtained in a strong scaling experiment of the Landers scenario on 128 to all 3072 nodes of SuperMUC phase 2. 44% and 42% of peak performance where reached when setting the clock frequency to 2.2 GHz or 2.6 GHz, respectively. [5]

On-going Research / Outlook

We are currently working on establishing large-scale simulation of tsunamigenic earthquakes, for example the 2004 Sumatra-Andaman earthquake, with the goal to obtain physically consistent, time-dependent displacements of the ocean floor during the earthquake. We plan to pipe this displacement data into highly resolved tsunami simulations.

On the modelling side, SeisSol is being extended to include effects of plasticity and viscoelastic attenuation during wave propagation. In addition, a scalable local-time-stepping algorithm has been implemented to mitigate the effects of small time steps forced by very fine or even degenerate grid cells [5]. To combine local time-stepping for wave propagation with dynamic rupture will be our next step.

References and Links

[1] www.seissol.org

[2] A. Heinecke, A. Breuer, S. Rettenberger, M. Bader, A.-A. Gabriel, C. Pelties, A. Bode, W. Barth, X.-K. Liao, K. Vaidyanathan, M. Smelyanskiy, P. Dubey: Petascale High Order Dynamic Rupture Earthquake Simulations on Heterogeneous Supercomputers. Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis SC14, p. 3–14. IEEE, 2014.

[3] A. Breuer, A. Heinecke, L. Rannabauer, M. Bader: High-Order ADER-DG Minimizes Energy- and Time-to-Solution of SeisSol. High Performance Computing, 30th International Conference, ISC High Performance 2015, Lecture Notes in Computer Science 913, p. 340–357. Springer, 2015.

[4] S. Rettenberger, M. Bader: Optimizing Large Scale I/O for Petascale Seismic Simulations on Unstructured Meshes. 2015 IEEE International Conference on Cluster Computing (CLUSTER), p. 314–317. IEEE Xplore, 2015.

[5] A. Breuer: High Performance Earthquake Simulations. PhD Thesis, Department of Informatics, Technische Universität München, 2015.

4D City – Space-time Urban Infrastructure Mapping by Multi-sensor Fusion and Visualization

Research Institution
Signal Processing in Earth Observation, TU München
Principal Investigator
Xiaoxiang Zhu
Researchers
Yuanyuan Wang, Gerald Baier, Yilei Shi
Project Partners
German Aerospace Center

SuperMUC Project ID: pr45ne

Introduction

Static 3D city models are well established for many applications such as architecture, urban planning, navigation, tourism, and disaster management. However, they do not represent the dynamic behavior of the buildings and other infrastructure (e.g. dams, bridges, railway lines). Very high resolution spaceborne Synthetic Aperture Radar (SAR) Earth observation satellites, like the German TerraSAR-X, provide for the first time the possibility to derive both shape and deformation parameters of urban infrastructure on a continuous basis.

Therefore, this project is aimed to generate 4D (space-time) city models and their user specific visualizations to reveal not only the 3D shape of urban infrastructures but also their deformation patterns and motion.

The research envisioned in this project will lead to a new kind of city models for monitoring and visualization of the dynamics of urban infrastructure in a very high level of detail. The deformation of different parts of individual buildings will be accessible for different users (geologists, civil engineers, decision makers, etc.) to support city monitoring and management and risk assessment.

Results and Methods

The main technique employed in the project is the so-called differential SAR tomography (TomoSAR). For retrieving the 3D position and the deformation parameters of one pixel, we solve an inversion problem with typical dimension of 100x1,000,000 (the forward model matrix). Considering the large size of spaceborne SAR images, the computation over a large area is very time consuming.

With the resource of LRZ, we are so far the only team in the world that can produce the 3D reconstruction and deformation in city-scale using TomoSAR. In total, this project consumed 6 mio core-hours and 10 TB of storage. For each processing, over 500 cores are required. So far, we have processed the following datasets:

Table 1. Processed datasets and size

Dataset	# of images	Image size	resolution
Las Vegas	180	11kx6k	0.6x1.1m2
Berlin	550	11kx6k	0.6x1.1m2
Shanghai	29	25kx55k	1.2x3.3m2

Among them, the Berlin dataset was reprocessed multiple times with different parameters setting. In the following content, some representative results are shown.

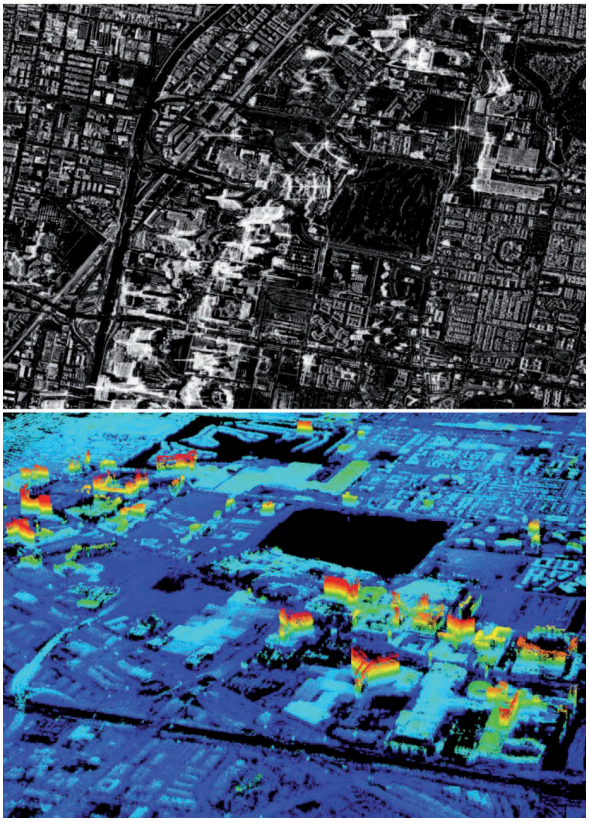

Figure 1. Upper: TerraSAR-X high resolution spot light image of Las Vegas, and lower: 3D point cloud of Las Vegas reconstructed using our algorithm. The color represents the high [2]-[4].

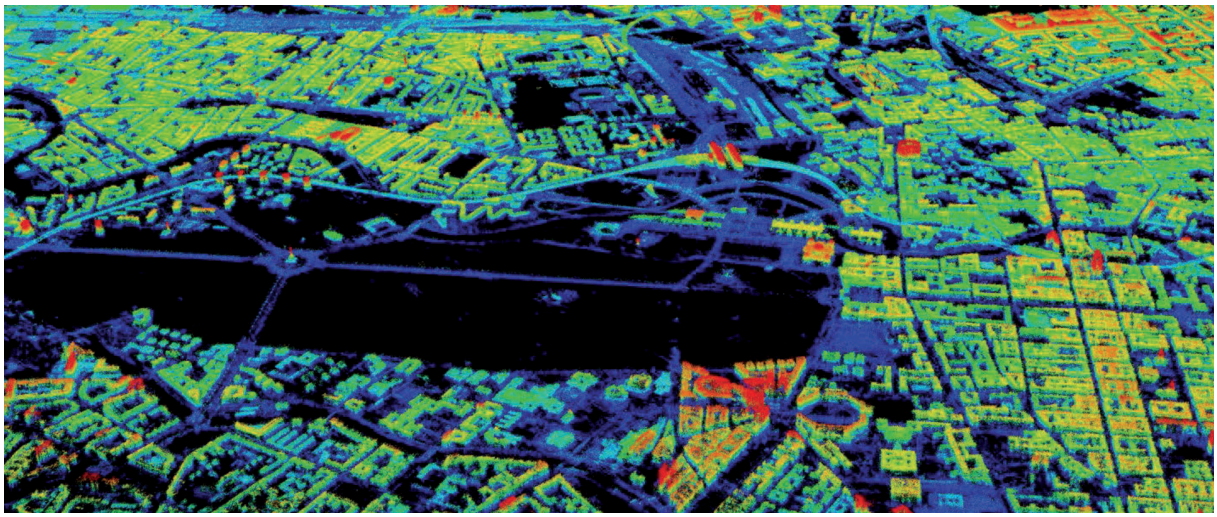

Figure 2. Fusion of two reconstructed 3D point clouds. The combined point cloud contains over 40 million points.

Las Vegas
The following upper subfigure is one of the input TerraSAR-X images of Las Vegas. By applying the TomoSAR algorithm on tens of such images, a 3D point cloud was reconstructed (lower subfigure). This point cloud contains around 10 million points. Most importantly, each point contains not only the 3D position information, but also its deformation information (so-called 4D, but not shown here).

Berlin
By fusing two point clouds from different viewing angles, we obtain a complete coverage over an entire city. Figure 2 is the example of Berlin. As always, each point is associated with its movement information. The combined point cloud contains about 40 million points. The number of points exceeds 100 million, if we combine all six reconstructed point clouds of Berlin.

4D Building Model
Figure 3 illustrates the potential of these point clouds. It depicts the facade model of the Bellagio hotel in Las Vegas reconstructed from the TomoSAR point cloud [5]. The points are overlaid onto the model with the color indicating the estimated motion parameter (here is the amplitude of seasonal motion caused by thermal dilation). This information can be used for developing dynamic building models from spaceborne SAR data that can help to monitor individual buildings and even the whole city.

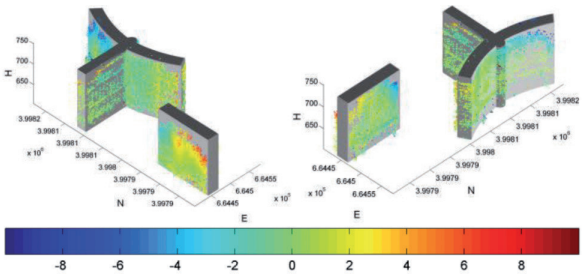

Figure 3 Reconstructed 4D building façade model, amplitude of seasonal motion is color-coded. [unit: mm]

On-going Research / Outlook

Based on the achieved results from previous years, we would like to conduct further research on large area object reconstruction from the resulting TomoSAR point clouds exemplified in Figure 1 and Figure 2, as well as SAR image filtering for improving the resolution of the current 3D digital elevation model (DEM). This would require computational resource for:

- Object reconstruction algorithm developments and tests based on the TomoSAR point clouds;
- Scientific visualization of the reconstructed dynamic city models.
- Non-local filtering of SAR images

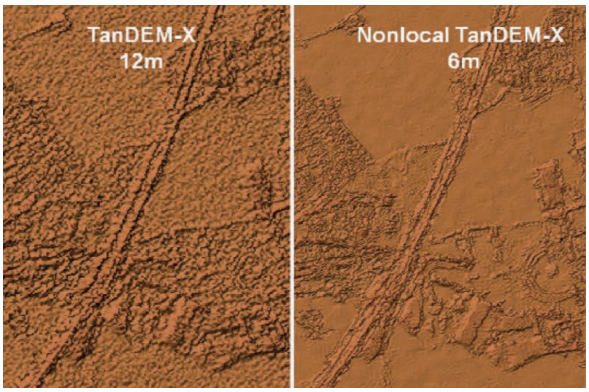

Figure 4. The standard 12m TanDEM-X DEM (left), and the non-local DEM with 6m resolution (right).

Figure 4 is a comparison of the standard 12m TanDEM-X DEM and the nonlocal filtered TanDEM-X DEM. Figure 4. The standard 12m TanDEM-X DEM (left), and the non-local DEM with 6m resolution (right).

References and Links

[1] http://www.sipeo.bgu.tum.de/
[2] Xiaoxiang Zhu. 2011. Very High Resolution Tomographic SAR Inversion for Urban Infrastructure Monitoring: A Sparse and Nonlinear Tour, Deutsche Geodätische Kommission.
[3] Xiaoxiang Zhu, Yuanyuan Wang, S. Gernhardt, and R. Bamler. 2013. Tomo-GENESIS: DLR's Tomographic SAR Processing System. In Urban Remote Sensing Event, 2013 Joint. 159–162.
[4] Yuanyuan Wang, Xiaoxiang Zhu, and Richard Bamler. 2014. An Efficient Tomographic Inversion Approach for Urban Mapping Using Meter Resolution SAR Image Stacks. IEEE Geosci. Remote Sens. Lett. 11, 7, 1250–1254.
[5] Xiaoxiang Zhu and Muhammad Shahzad. 2014. Façade Reconstruction Using Multiview Spaceborne TomoSAR Point Clouds. IEEE Trans. Geosci. Remote Sens. 52, 6, 3541–3552.

Computational Wave Propagation

Research Institution
Department für Geo- und Umweltwissenschaften, LMU Munich

Principal Investigator
Heiner Igel

Researchers
Moritz Bernauer, Stefanie Donner, Michael Dumbser, Alice-Agnes Gabriel, Celine Hadziioannou, Kasra Hosseini, Lion Krischer, Elizabeth Madden, Anne Obermann, Simon Stähler, Thomas Ulrich, Stephanie Wollherr

Project Partners
Department of Earth Sciences, ETH Zurich, Switzerland; Department of Civil, Environmental and Mechanical Engineering, University of Trento, Italy; University of Oxford, UK; King Abdullah University of Science and Technology, Saudi Arabia; Princeton University, USA; University ofToronto, Canada; University of Nice – Geoazur, France; Oak Ridge National Laboratory (ORNL), USA; Leibniz-Institute for Baltic Sea Research, Warnemünde; Chair of Scientific Computing, Technical University Munich; Department of Statistics, LMU Munich; Intel Corporation, Santa Clara, USA

SuperMUC Project ID: h019z, pr63qo

Introduction

In 2015, the Computational Seismology group of LMU Munich intensively exploited SuperMUC HPC infrastructure in a variety of international research projects covering computational wave propagation across spatial and temporal scales. A high number of small- to midscale runs as well as certain large-scale simulations were performed in order to support on-going research projects of the several subgroups. State-of-the-art modeling software based on high-order accurate SEM and ADER-DG methods were used to gain insight for example in industry related projects, as e.g. full-waveform inversion experiments, for physically based seismic hazard assessment in high-resolution earthquake scenarios as well as in fundamental geophysics, as e.g. in understanding earthquake source physics and geophysical signals in noise recordings of seismometers. Given the variety of methods and topics covered in the project we summarize in the following results, impact and outlook of the following subprojects in detail:

- **Large-scale Full Waveform Inversion**
- **3-D probabilistic sensitivity kernels to image medium changes in the Earth crust**
- **New Computational Tools for Seismology**
- **Forward simulations: Seismic wave propagation and earthquake dynamics**
- **A-posteriori sub-cell limiting of Discontinuous Galerkin methods**
- **Inversion for seismic moment tensors combining translational and rotational ground motion data**

1. Large-scale Full Waveform Inversion

Researchers
Lion Krischer, Heiner Igel (both LMU)

Collaborators
Andreas Fichtner, Christian Böhm, Saule Simute, Michael Afanasiev (all ETH Zürich)

Approx. Resources
1 million CPU hours, 50 million files on SCRATCH and WORK , 10 TB temporary storage, typically 3000 cores per job

This on-going project strives to image the crustal and upper mantle subsurface structure beneath the North American continent and the Northern Atlantic. We are employing physically meaningful anelastic waveform simulations through three-dimensional complex

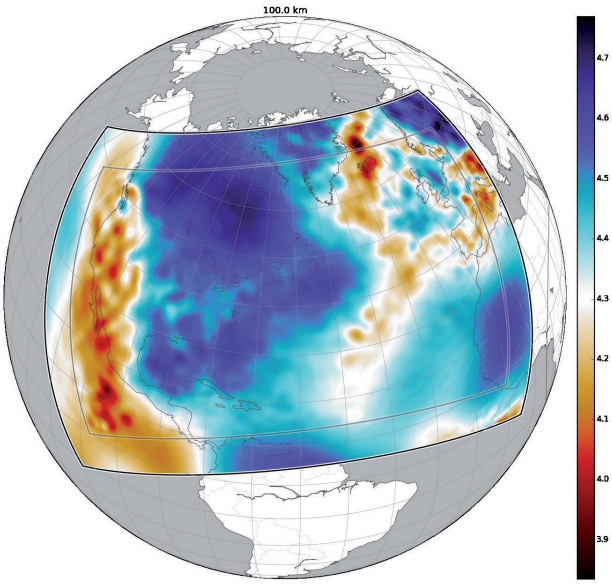

Figure 1: Preliminary vsh velocity model at 100 km depth.

media for the forward as well as the adjoint problem. The non-linear iterative L-BFGS inversion schema incorporates large amounts of data giving rise to interesting challenges in I/O and workflow optimization. We developed a framework to carry out these inversions in an efficient, reproducible, and less-error prone fashion [1]. A preliminary result is shown in Fig. 1. In an effort to build upon existing work we integrated existing models part of the Collaborate Seismological Earth Model [2]. After completion the model derived in this study will flow back into this collaborative model benefiting future studies. All waveform simulations and parts of the optimization procedure were carried out on SuperMUC.

2. 3-D probabilistic sensitivity kernels to image medium changes in the Earth crust

Researchers
A. Obermann (ETHZ), C. Hadziioannou (LMU)

Approx. Resources
400,000 CPU hours on typically 20 cores/run, 250,000 files generated, 30 TB temporary on SCRATCH and 200 GB long-term storage requirement

Temporal changes in the later, scattered part of the seismic signal, referred to as the "coda", can be used to monitor perturbations of mechanical (velocity, pressure, etc) or structural (change of scatterer position, for instance due to fracturing) properties of the crust and of engineering structures. Besides the detection of the temporal changes, an important aspect is the localization of the changes. Several studies to date have shown that spatial localization of medium changes on the horizontal plane (2-D). This method has been successfully applied to forecast the location of upcoming volcanic eruptions at Piton de la Fournaise volcano, La Réunion Island.

In order to extend the imaging method to 3-D, we have run several 3-D simulations in complex media to determine the properties of the mixture of surface and body waves at different times in the coda, for different scattering media. We demonstrated that we can use the lapse-time dependent behavior of the velocity changes to discriminate a change that occurs at the surface from a change at depth. This work is presented in a paper which has been recently submitted [3].

The results of our simulations also form a necessary first step to construct 3-D sensitivity kernels for the diffuse coda waves. More simulations are needed to fully characterize wave propagation in scattering media, which will be pursued in future research. We will simulate wave propagation in complex media for different scattering properties, and for several realizations of each medium. The resulting information on wave type partition in the late part of our seismograms will be applied to construct sensivity kernels for the seismic coda, which in turn can be used to locate subsurface changes at depth, which opens new opportunities in seismic imaging and monitoring applications.

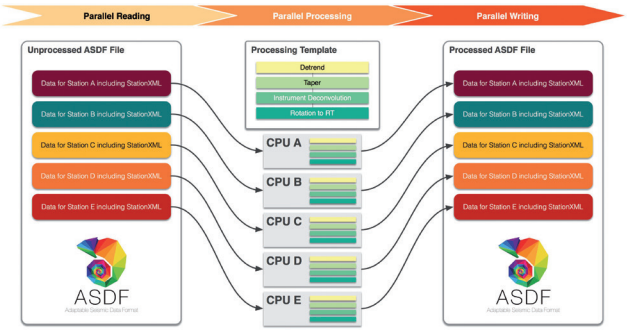

Figure 2: Parallel processing of seismological data by defining a processing template in the form of a script that is automatically applied in parallel to all data contained in a single file.

3. New Computational Tools for Seismology

3.1 A New Efficient and Parallel Data Format for Seismology

Researchers
Lion Krischer (LMU)

Collaborators
James Smith[1], Matthieu Lefebvre[1], Wenjie Lei[1], Youyi Ruan[1], Elliott Sales de Andrade[2], Norbert Podhorszki[3], Ebru Bozdag[4], and Jeroen Tromp1

[1]Princeton, [2]University of Toronto, [3]Oak Ridge National Laboratory (ORNL), [4]University of Nice – Geoazur

Approx. Resources
200 000 CPU, 10 TB storage requirements on WORK

The full waveform inversion study of section 1 and requirements of various persons and groups in seismology made it abundantly clear that our community requires a new and exchangeable data format and associated tools. Some problems require data from a very large number of stations. Using existing data formats potentially results in the creations of millions of files which are cumbersome, error-prone, as well as slow and inefficient to work with. The developed format also solves other problems like data provenance and easier exchange of information. The ASDF format (http://seismic-data.org, [4]) has been designed with the goal to enable fast and parallel processing. We developed an adapter for the widely used waveform propagation code SPECFEM 3-D GLOBE to write and read the ASDF format. Additionally, we developed a Python API granting parallel data processing abilities to users without needing to know anything about parallel programming or MPI, see Fig. 2. These programs

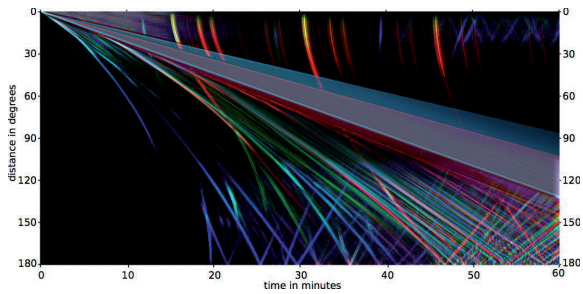

Figure 3: A stack of global seismograms accurate to 2 seconds computed with Instaseis. Color encodes the data component and automatic gain control was employed to balance amplitude variations.

and certain designs decisions of the format have been influenced by tests run on SuperMUC and we have the hope that they benefit the whole seismological community for times to come.

3.2 Instaseis

Researchers
Lion Krischer (LMU), Simon Stähler (LMU), Kasra Hosseini (LMU)

Collaborators
Martin van Driel (ETH Zurich), Tarje Nissen-Meyer (Oxford)

Instaseis (http://instaseis.net, [5]) enables the rapid, almost instantaneous, calculation of fully three-dimensional seismograms for arbitrary source-receiver geometries on the whole globe. The trade-off is that the structural model has to be spherically symmetric, a property which we explicitly use. Special wavefields are precomputed by using the axisymmetric spectral element code AxiSEM [6]. These can reach sizes of more than 10 TB requiring advanced I/O strategies and are stored in NetCDF files. SuperMUC has been used to calculate a range of these databases which in turn already enabled a couple of other studies [7], [8].

4. Forward simulations: Seismic wave propagation and earthquake dynamics

Approx. Resources
5 million CPUh, on average 500 cores/run, 1 million files, 20 TB storage requirements on SCRATCH and WORK

4.1 Statistical analysis of ground motion and source dynamics of a 3-D dynamic rupture scenario for the Northridge 1994 earthquake

Researcher
A.-A. Gabriel (LMU)

Collaborator
T. Kühn, A. Stöcker, F. Scheipl (Statistics, LMU)

200 high-resolution simulations of dynamic earthquake scenarios for the 1994 Northridge event based on [16] were performed with the ADER-DG based software package SeisSol (www.seissol.org) on SuperMuc and subsequently statistically analyzed to find computationally efficient approximations of the non-linear relations between earthquake source processes and ground shaking as well as of macroscopic earthquake source parameters with each other (see Fig. 4). The scenarios varied in initial conditions in terms of fault friction and strength as well as tectonic background stresses. The statistical analysis required automatized memory-intensive post-processing of large data amounts and were performed on SuperMuc as well. The employed statistical analysis included gradient boosting for parameter selection and yielded in a profound analysis of flexible statistical (GAM) models compared to simple 2-D analytical models previously developed.

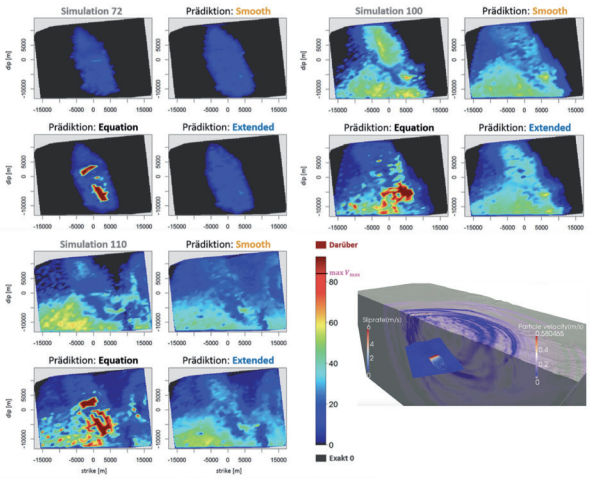

Figure 4: Particle velocity across the curved earthquake source of 3 typical realizations of a Northridge event scenario fitted by 3 statistical models of increasing complexity.

4.2 Development of a high-frequency 3-D dynamic earthquake scenario for petascale simulations on heterogeneous supercomputers

Researchers
C. Pelties, A.-A. Gabriel (LMU)

Collaborators
A. Breuer, S. Rettenberger, M. Bader (TUM), A. Heinecke (Intel)

Modeling dynamic earthquake rupture (on meter scale) coupled to seismic wave propagation (on scales of hundreds of kilometers) requires a multiscale and non-linear multi- physics approach. Realistic model setups should acknowledge topography, 3-D geological structures, rheology, and fault geometries with appropriate stress and frictional parameters (Fig. 5), all of which contribute to complex ground motion patterns. A fundamental challenge hereby is to model the high frequency content of the three-dimensional wave field. Given the lack of in-situ observations from the seismogenic zone, tens of lower resolution test scenarios (up to 1 Hz) were performed with SeisSol in the scope of this project varying fault strength and tectonic initial conditions. The model outputs were

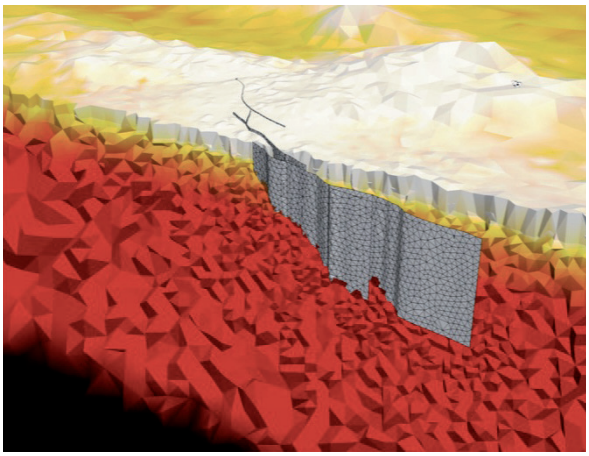

Figure 5: Tetrahedral discretization of the Landers fault system including all ruptured segments embedded in a realistic geological structure including topography.

compared to indirect earthquake observations yielding in the preferred model which was ran as a multi-petaflop simulation enabling the observation of high-detail rupture evolution and ground motion frequencies up to 10 Hz relevant for civil engineering purposes [17].

4.3 Off-fault plasticity with SeisSol

Researchers
S. Wollherr, A.-A. Gabriel (all LMU)

Collaborator
P.M. Mai (KAUST)

Rock and other solids undergo permanent deformation when stresses reach some critical value. High stress concentration during dynamic rupture simulations are typically observed around the fault and may be in reality large enough to cause failure of off-fault materials. In order to create realistic earthquake scenarios, we included off-fault plastic yielding into SeisSol. We successfully implemented the Drucker-Prager yield criterion with viscous regularization. The effectiveness of plastic yielding is very sensitive to the viscoplastic relaxation coefficient which is itself dependent on mesh discretization and S-wave speed. In order to investigate its effect with respect to the mesh discretization and order of accuracy, we ran several simulations for parameter studies.

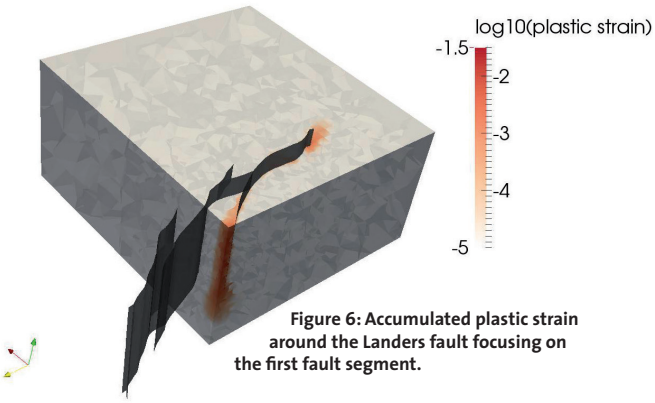

Figure 6: Accumulated plastic strain around the Landers fault focusing on the first fault segment.

The ADER-DG method is especially suited for complex geometries by enabling unstructured tetrahedral meshes. In this context, we investigate in a large-scale earthquake scenario the effects of off-fault-plasticity on the branched fault zone structure of the 1992 Landers earthquake. Compared to a purely elastic setup, plastic rock deformation does not only influence the step-over times but also the location of the rupture jumps in depth (Fig. 6). Previous studies already investigated the influence of off-fault yielding on the activation of branched faults in two-dimensions; however we find important 3-D effects of plasticity on the rupture behavior in depth. Therefore we plan to study the activation of branches and the probability of rupture to jump to adjacent fault segments especially in three dimensions. These simulations will be computationally expensive, but will help to constrain the conditions under which rupture is able to jump or to propagate into branches and as a consequence to constrain the dimension of a potential earthquake on a complex fault system.

4.4 Earthquake ruptures on self-similar fractal faults

Researchers
T. Ulrich, A.-A. Gabriel (both LMU)

The waviness of faults' surfaces has been observed over large ranges of length-scales, from the microscopic to regional scales. We would like to better understand the effects of this geometric complexity on earthquake rupture processes by simulating earthquakes on complex faults using SeisSol. A snapshot of a rupture propagating on such a fault is presented in Fig.7. Rough fault simulations are computationally intensive. In fact, the smallest features of the fault geometry have to be captured by the mesh in order to efficiently assess the effect of them on the rupture. Typically, a 10 second simulation of an earthquake propagating across a 40 km long fault, characterized by a 500 meter roughness scale requires a 20 million element mesh (50m mesh discretization on the fault) and runs in approximately 3 hours on 256 nodes on SuperMuc 2 (order 4 in space and time) [18]. We plan a dozen of simulations of this kind, for getting answers to the following fundamental questions:

- Is the rupture process mostly influenced by a specific band of roughness scales?
- To which accuracy do we have to know the complex geometry to accurately simulate the rupture process?
- Would rupture scenarios presented in previous publications be profoundly affected if some small-scale random roughness would be added to the geometry?

The installation of SuperMuc 2 enabled larger hybridization of SeisSol and thus higher efficiency one large meshes. Furthermore, the output files are less divided and easier to postprocess, because we can use a smaller mesh partitions per core. A SuperMuc machine "next generation" will allow modeling small scales features in large scale scenarios, for example caused by wave scattering in the domain which will allow more realistic ground motion simulations with high complexity. Biggest limitations in this project is the visualization and post-processing of simulation results on highly partitioned unstructured meshes.

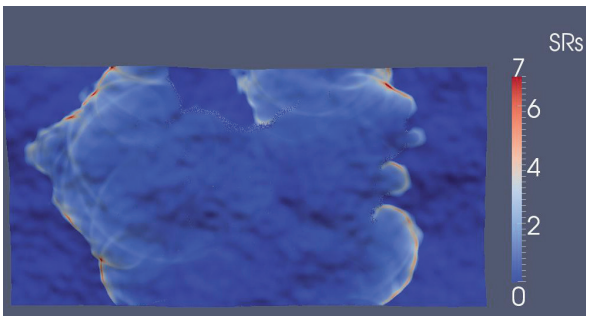

Figure 7: Slip-rate snapshot in the direction of slip of a typical rough fault. The initially circular rupture front has been split in several pulses by the geometrical barriers encountered along its way, leading to a complex but realistic rupture process.

4.5 Simulation of a 2004 Sumatra earthquake scenario

Researchers

T. Ulrich, B. Madden, S. Wollherr, A.-A. Gabriel (all LMU)

In the framework of the ASCETE project (www.ascete.de), we are developing physically realistic dynamic rupture scenarios of the 2004 Sumatra earthquake [19]. The ground displacements generated with SeisSol will be used by our ASCETE partners, to conduct large scale tsunami simulations. Dynamic simulation of this 1300 to 1500 km rupture is a computational and geophysical challenge. In addition to capturing the large-scale rupture, the simulation must resolve the process zone at the rupture tip, whose characteristic length is comparable to smaller earthquakes and which shrinks with propagation distance. Thus, the fault must be finely discretized. Moreover, previously published inversions agree on a rupture duration of ~8 to 10 minutes, suggesting an overall slow rupture speed. Hence, both long temporal scales and large spatial dimensions must be captured. As a consequence, we expect even the first test simulations, in which we will try different stress conditions, fault geometries and friction properties, to be already computationally intensive (probably a few hours on 128 nodes each, on phase 2). The most realistic scenario will be simulated using a larger mesh and higher space-time order in a petascale simulation. The setup combines the detailed geometry of the subducting fault proposed by Slab1.0 to the south and aftershock locations to the north, with high-resolution topography and bathymetry data. Due to the lack of near-field observations, these simulations will be constrained by the overall characteristics of the rupture, including the magnitude, propagation speed, and extent along strike. The possibility of inhomogeneous background stress, resulting from the curved shape of the slab along strike and the large fault dimensions, will be investigated, as will the possible activation of thrust faults splaying off the megathrust in the vicinity of the hypocenter (Fig. 8).

4.6. Simplified subduction zone scenarios for coupled dynamic earthquake rupture and tsunami simulations

We are currently validating this new coupled framework between a dynamic earthquake within Earth's crust and the resulting tsunami wave within the ocean using a simplified model setup. The SeisSol earthquake simulation will be coupled to an adaptive mesh discretizing the shallow water equations with a Runge-Kutta discontinuous Galerkin (RKDG) scheme subsequently allowing for an accurate and efficient representation of tsunami evolution and inundation at the coast. The simplified earthquake model of Fig. 9a incorporates the most fundamental features of a subduction zone earthquake, but also is as simple as possible, to facilitate setting-up the workflow between the codes, and to understand the main features of the chain of events. To allow faster runs, a small fault model is preferred. Nevertheless, the respect of some fundamental features of the subduction earthquakes requires that the fault is not too small. In fact, subduction zone faults are characterized by low dipping angles and large depth spans: their average dip angle is about 20 degrees, and their average seismic depth range spans from 10 to 45 km. In addition, slip can occur in the shallowest 10 km of the fault, where the material rheology theoretically do not favor rupture propagation (e.g. the 2011 Tohoku earthquake in Japan). As a consequence, we use an along-dip dimension of 130 km that intersects with the sea floor in order to capture these fundamental megathrust characteristics in preparation for more complex models.The resulting seafloor displacements along an initially planar (and later realistic) bathymetry profile will be transferred to the tsunami setup with an initially simple coastal run-up profile. Once validated in this simplified setup, we will constrain the earthquake initial stress and strength conditions from realistic and physically consistent seismo-thermo-mechanical modeling on long timescales. Fig. 9B compares the simulated surface displacements profiles, considering elastic and plastic constitutive laws. The vertical surface displacements are found to be significantly influenced by off-fault plasticity newly implemented into SeisSol, which may have important implications for tsunami generation and propagation.

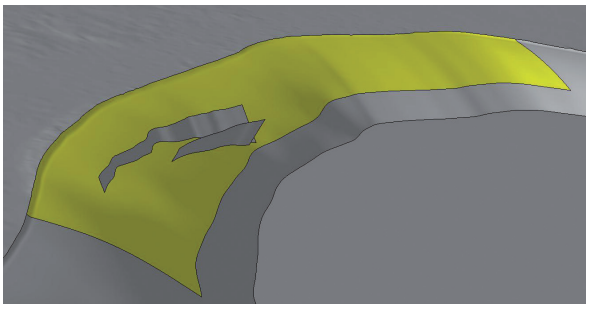

Figure 8: Snapshot of the current state of the Sumatra model, where is detailed the main fault (yellow), part of the complex ground surface, and two splay fault that may or may not have slipped.

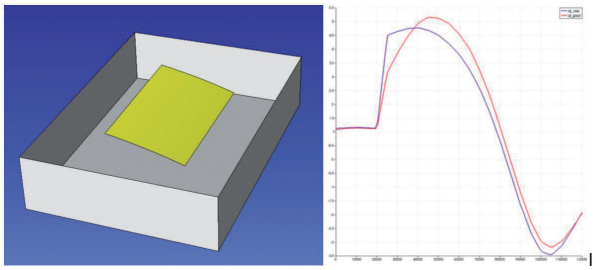

Figure 9: (a) Model set-up featuring a curved, shallowly dipping fault, in homogeneous elastic half space. (b) Vertical surface displacements profiles, transversely to the center of the fault, are presented for the elastic (blue) and plastic (case). In case of plasticity, the surface displacement is shifted laterally from the trench. The peak of vertical surface displacement is slightly higher in case of plasticity.

4.7 Benchmarking of the ADER-DG method for advanced dynamic rupture scenarios

Researchers

T. Ulrich, S. Wollherr, A.-A. Gabriel (all LMU)

The goal of simulating rupture on a fault and the resulting wave propagation in a medium is to get a better insight into the complex physical processes during an earthquake. Ideally, we could compare the results of our simulations with real earthquake data in order to verify

the numerical approach. But in most cases there is a lack of observation close to large earthquakes. In addition to that, there doesn't exist an analytical solution even for the simple most dynamic rupture problem.

Extending the recent verification of SeisSol [20] and in order to verify for example the implementation of new features in SeisSol such as off-fault plasticity or the use of rough fault geometries it is crucial to compare our results with the results of other dynamic rupture codes. We use therefore test problems of the Dynamic Earthquake Rupture Exercise (http://scecdata.usc.edu/cvws/) created by the South California Earthquake Center (SCEC) and compare amongst others the synthetic ground motions as well as the slip and stresses on the fault for numerically simulated earthquakes.

- *Benchmark tests using off-fault plasticity*

Off-fault plasticity was verified by two benchmark setups: We achieved very good accordance in comparison to other codes for a strike-slip scenario and a scenario with a 30° dipping fault. The later test case proves the ability of SeisSol to use the new implementation in subduction zone scenarios with dipping angles of even less than 30° where plasticity might influence the resulting floor-uplifting.

- *Benchmark tests using rough faults*

The ability of SeisSol to accurately simulate ruptures on complex fault geometry has been assessed. Some bias on the fault normal stress is observed in case of rough faults, particularly if the order of accuracy in space is high. We think it may be related to the low order of the mesh, which is composed of straight faces tetrahedron. Nevertheless, the ability of SeisSol to simulate such kind of rupture is not compromised, if some rules are respected. In particular, to increase accuracy, refining the mesh should be preferred, rather than using a higher space accuracy of the elements. In fact, a higher element order ensures a better efficiency of the code, and a better accuracy of the waves propagation, but could result in a lower accuracy of the rupture process in case of strong geometric complexity.

- *Benchmark tests using 1D velocity structure*

Two 1D layered benchmarks featuring a planar strike-slip fault was tested using SeisSol. In the first benchmark, the material properties were constant inside each layer. On the other hand, the second benchmark was a bit more demanding, as it featured linearly increasing properties inside some of the layers. The possibility of using varying properties inside an element is a feature that has been abandoned by SeisSol. As a matter of fact, the layers with linearly varying properties had to be meshed finely. The obtained results were in good agreements with the other codes.

- *Benchmark tests using a low-velocity zone around the fault*

The most recent benchmark features a relatively thin zone of low-velocity around a planar strike-slip fault. The low-velocity properties are accounting from the damage The low-velocity properties stemming from rock damage induced by previous earthquakes. The obtained results are in good agreements with the other codes.

4.8 Evaluating effects of detailed fault geometries on Landers earthquake dynamics

Researcher

B. Madden (LMU)

Understanding the influence of source structures on earthquake behavior is critical for earthquake forecasts and ground motion characterizations. Of primary importance is how segmented fault systems host earthquakes. After compiling slip date from 30+ earthquakes, [23] finds that 3-4 km is the upper limit of step distances across which an earthquake will continue to rupture. However, the scatter in the results suggest that step distance is not the only parameter affecting slip transfer. We consider the role of secondary faults in transferring slip across the step by modeling the M7.3 Landers earthquake, which occurred in the Mojave Desert of southern California, USA in 1992 (Fig. 10a). This earthquake ruptured segments of 5 different right-lateral, strike-slip faults, resulting in a complex surface slip and aftershocks (Figure 10b). One of these faults is the 4 km long Landers-Kickapoo Fault (LKF) within a 3 km releasing step between the Johnson Valley Fault to the southwest, where the earthquake nucleated, and the Homestead Valley Fault to the northeast. Both of these faults slipped only along a portion of their length. Quasistatic, boundary element method models incorporating 3-D, non-planar faults, suggest that slip along the LKF is required to transfer slip through the step [21]. These models incorporate a refined fault geometry along the southern part of the rupture that includes a dipping Johnson Valley Fault and a discontinuity between the LKF and the Homestead Valley Fault (Figure 11a). However, these models cannot evaluate the role of acceleration and triggering from seismic waves emitted from the rupture tip, which may allow the rupture to jump the step, with slip along the LKF following the initiation of slip on the far side of the step. Using SeisSol, we are evaluating under what conditions the rupture jumps the step and if and when the LKF is required to transfer slip. SeisSol accommodates the complex, 3-D fault shapes used in the quasistatic model, as well as a fully connected, vertical structural model

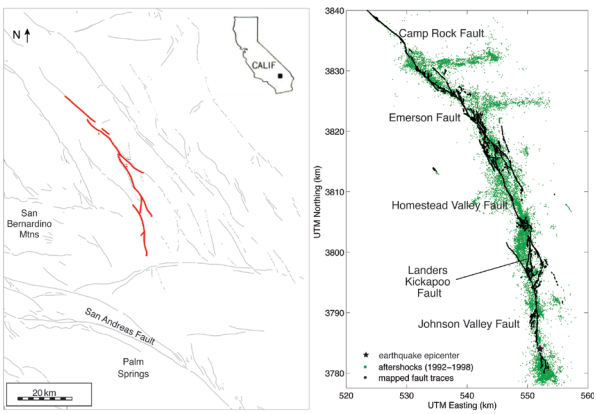

Figure 10: (a) Location and (b) surface rupture with aftershocks of the 1992 Landers earthquake.

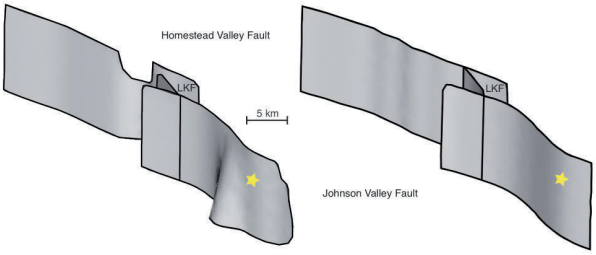

Figure 11:Structural models of the southern Landers rupture: (a) Dipping and disconnected, (b) fully connected and vertical. Only southern part of rupture where structures differ are shown. Star is earthquake hypocenter.

(Figure 11b, only southern part of rupture where structures differ are shown, the star is hypocenter). We evaluate how these different structural models affect rupture dynamics in order to address the role of the LKF in allowing this earthquake to move from the Johnson Valley Fault to the Homestead Valley Fault, and therefore become a much larger earthquake than was expected for the region based on individual fault lengths alone [22].

4.9 Oceanic Love wave excitation

Researchers

C. Hadziioannou, D. Ziane

Longuet-Higgins (1950) proposed that secondary microseismic noise can be attributed to oceanic disturbances by surface gravity wave interference causing non-linear, second-order pressure perturbations at the ocean bottom. In order to complete preliminary work on "Investigating the generation of Love waves in secondary microseisms using 3-D numerical simulations" (presented in [15]) numerous further simulations need to be performed to cover a meaningful parameter range. In particular, we will investigate and quantify the effects of the statistical properties magnitude of perturbation, correlation lengths and fractal dimensions for both, internal scattering and bathymetry variations.

5. A-posteriori sub-cell limiting of Discontinuous Galerkin methods

Researcher

M. Dumbser (Trento)

5.1 Introduction

The numerical solution of hyperbolic problems is required in several physical and technological applications involving fluid dynamics, multiphase flows, air flows around aircraft or cars, astrophysical flows, free surface flows, environmental and geophysical flows. Among numerical methods developed to solve hyperbolic problems, there are Discontinuous Galerkin (DG) methods. DG methods are very robust and they show high flexibility and adaptivity strategies in handling complex geometries. In addition, they also verify an entropy condition which confers them nonlinear L2 stability. Unfortunately, DG schemes suffer from a serious problem, which has negatively affected their popularity. Namely, they produce oscillations as soon as a discontinuity, such as a shock wave, appears in the solution (Gibb's phenomenon).

In this project, we have proposed and successfully applied a completely new strategy [9], which is based on the so-called MOOD paradigm [10]. In practice, we verify a-posteriori the validity of a discrete candidate solution against physical and numerical detection criteria after each time step and for those troubled cells that need limiting, our new limiter approach re-computes the discrete solution by scattering the DG polynomials at the previous time step onto a set of suitable finite volume

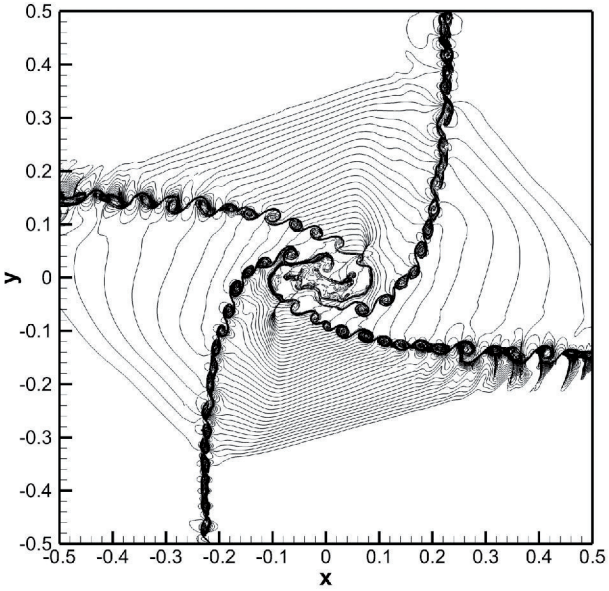

Figure 12: Two-dimensional Riemann problem solved with DG-P5

sub-cells per space dimension. In this way, the Gibbs phenomenon is essentially suppressed and the high accuracy of the solution is preserved.

When combined with Adaptive Mesh Refinement (AMR) [11], the new scheme allows for an unprecedented ability in resolving even the finest details in the dynamics of the fluid. The whole approach has been applied to

- The Euler equations of gas dynamics [9];
- The magnetohydrodynamics equations, both classical [11] and relativistic [12], with very interesting potential applications to astrophysical systems;
- The compressible Navier-Stokes equations [13], exploiting a novel hyperbolic formulation of continuum mechanics.

5.2 Results and Methods

The new a-posteriori DG limiter method presented by [9] and implemented in this project

1. Computes the solution by means of an unlimited ADER-DG scheme, which represents the discrete solution in terms of piecewise polynomials of degree N.
2. Detects a-posteriori the troubled cells by applying a simple discrete maximum principle and by checking the positivity of density and pressure,
3. Creates a local sub-grid within these troubled DG cells, each of which is composed by Ns=2N+1 sub-cells.
4. Re-computes the discrete solution at the sub-grid level via a more robust Total Variation Diminishing (TVD) or Weighted Essentially Non Oscillatory (WENO) finite

volume scheme. The final non-oscillatory DG solution on the main grid is then recovered from the sub-cell averages by means of a finite-volume reconstruction operator that acts on the cell averages of the sub-grid.

The choice of Ns=2N+1 sub-cells is specifically meant to allow a perfect matching of the maximum admissible time step of the finite volume scheme on the sub-grid with the maximum admissible time step of the DG scheme on the main grid, and to minimize the local truncation error.

The numerical scheme briefly described above has been implemented in Fortran 2003 and parallelized through standard MPI. The spectacular resolution properties of the new scheme have been shown through a wide number of test cases performed in two and in three space dimensions, for the Euler equations of compressible gas dynamics and for the magnetohydrodynamics (MHD) equations, the latter both in the Newtonian and in the special relativistic regime, as well as on a new unified model of continuum mechanics that is able to describe viscous fluids and elastic solids at the same time [13]. A limited sample of these results, fully available in references [9], [11], [12], [13] is reported in Fig. 12, Fig. 13 and Fig. 14. Collecting references [9], [11], [12] and [13], an approximate number of ~100 models have been evolved. Typically, excluding the simplest 1D models and the most expensive 3-D models, 1024 cores per job have been used, every model generating ~150 files on average.

At the heart of our DG scheme there is a local predictor step, which allows to perform the time update through a single time-step. This is particularly advantageous in terms of MPI communications, and it is also very rewarding when AMR is also activated. All these features have been fully exploited on SuperMuc. In three space dimen-

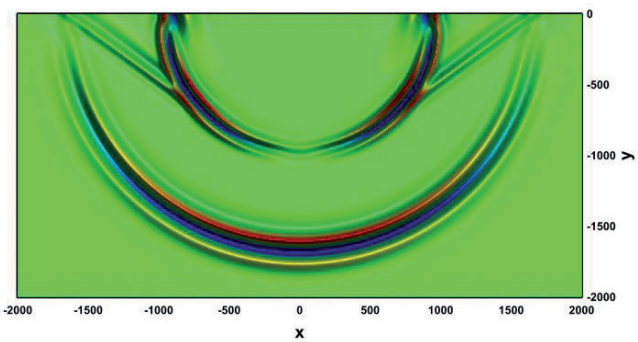

Figure 13: Lamb's problem, solved with ADER-DG P4 scheme using a new unified model of continuum mechanics that is able to deal with both, fluids and solids.

sions, and when the degree of the DG polynomial is very large (N > 6), the memory usage can become an issue, and we have found convenient to resort to the fat-node.

5.3 On-going Research / Outlook

SuperMUC computational resources have been fundamental for the development of this innovative project, providing an ideal infrastructure both in terms of computational efficiency and data storage.

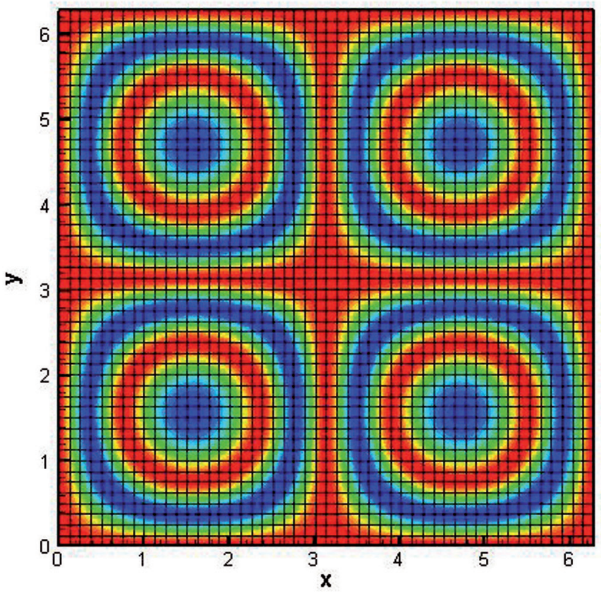

Figure 14: Distortion tensor in the Taylor-Green vortex problem, solved with ADER-DG-P3 using a new unified model of continuum mechanics that is able to deal with both, fluids and solids.

For the near future, we plan to further exploit our approach in a recently proposed numerical model of continuum mechanics [13], able to treat within a single framework the dynamics of inviscid and viscous compressible Newtonian and non-Newtonian fluids with heat conduction, as well as the behavior of elastic and visco-plastic solids.

6. Inversion for seismic moment tensors combining translational and rotational ground motion data

Researchers
S. Donner, M. Bernauer, H. Igel (all LMU)

Approx. Resources
1 TB storage requirement on WORK

Reliable earthquake source mechanisms are a fundamental key element in seismology for several applications. They do not only provide us with information on the processes within the source which are used to better understand how, where, and when earthquakes occur. They are also an important input for further analysis, such as seismic hazard studies or Earth tomography.

Especially on a regional scale the waveform inversion for seismic point-source moment tensors still faces several difficulties. These difficulties include e.g. a sparse and unfavorable data coverage, the necessity of a detailed structural model, and low signal-to-noise ratios, especially for smaller earthquakes (M<4.5). Thus, the resolvability of the moment tensor from inversion is reduced. Often, intense constrains are necessary to obtain a reliable solution and the full power of the moment tensor method is not exploited.

To increase the resolution of the moment tensor components and the centroid depth we included rotational ground motion data in addition to the conventional

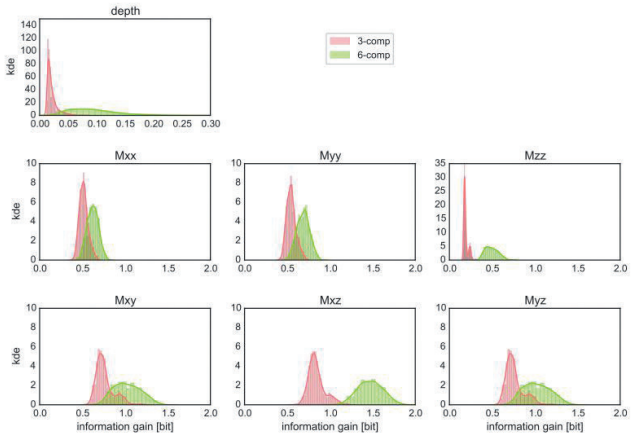

Figure 15: Gaussian kernel density estimations for the information gain corresponding to 1000 inversion runs. Coral and green distributions show the information gains from inverting 48 3-component and 24 6-component randomly chosen stations for the given parameters. The higher the information gain, the more the inversion benefits from using rotational motions. Same information gain means same inversion resolution using half the number of stations but the same amount of data.

translational ground motion data into waveform inversion for the moment tensor. In accordance with theoretical considerations we could show that the inversion of the combined data set provide substantially more information on the moment tensor. Especially, the resolution of parameters containing spatial derivatives with depths improved drastically. Moreover, our study demonstrated the possibilities of determining highly improved results with only half the number of stations, which is an important factor concerning e.g. maintenance (see Fig. 15). To face the difficulties of the inversion problem and obtain quantifications of uncertainties we applied a Bayesian, i.e. probabilistic, approach. The computational expenses of such an approach depend on the availability of modern high performance computing systems. The result will be published in [14].

References and Links

[1] Krischer, L., Fichtner, A., Zukauskaite, S., & Igel, H. (2015). Large-Scale Seismic Inversion Framework. Seismological Research Letters, 86(4), 1198–1207. doi:10.1785/0220140248

[2] Afanasiev, M., Peter, D., Sager, K., Simut, S., Ermert, L., Krischer, L., & Fichtner, A. (2015). Foundations for a multiscale collaborative Earth model. Geophysical Journal International, 204(1), 39–58. doi:10.1093/gji/ggv439

[3] A. Obermann, T. Planes, ` C. Hadziioannou, M. Campillo ; Lapse-time dependent coda wave depth sensitivity to local velocity perturbations in 3-D heterogeneous elastic media, submitted to Geophys. J. Int.

[4] Krischer, L., Smith, J., Lefebvre, M., Lei, W., Ruan, Y., Sales de Andrade, E., Podhorszki, N., Ebru Bozdag, E., and Tromp, J., An Adaptable Seismic Data Format, Journal of Geophysical Research, to be submitted, 2016.

[5] van Driel, M., Krischer, L., Stähler, S. C., Hosseini, K., & Nissen-Meyer, T. (2015). Instaseis: instant global seismograms based on a broadband waveform database. Solid Earth, 6(2), 701–717. doi:10.5194/se-6-701-2015

[6] Nissen-Meyer, T., van Driel, M., Stähler, S. C., Hosseini, K., Hempel, S., Auer, L., ... Fournier, A. (2014). AxiSEM: broadband 3-D seismic wavefields in axisymmetric media. Solid Earth, 5(1), 425–445. doi:10.5194/se-5-425-2014

[7] Sens-Schönfelder, Christoph, Roel Snieder, and Simon Stähler (2015), The lack of equipartitioning in global body wave coda, Geophysical Research Letters, 42, doi:10.1002/2015GL065108.

[8] Hosseini, K., & Sigloch, K. (2015). Multifrequency measurements of core-diffracted P waves (Pdiff) for global waveform tomography. Geophysical Journal International, 203(1), 506–521. doi:10.1093/gji/ggv298

[9] Michael Dumbser, Olindo Zanotti, Raphael Loubere, Steven Diot. A posteriori subcell limiting of the Discontinuous Galerkin finite element method for hyperbolic conservation laws, JCP 278, 47-75 (2015)

[10] Stephan Clain, Steven Diot, Raphael Loubere. A high-order finite volume method for systems of conservation laws Multi-dimensional Optimal Order Detection (MOOD), JCP 230, 4028-4050 (2011)

[11] Olindo Zanotti, Francesco Fambri, Michael Dumbser, Arturo Hidalgo. Space-time adaptive ADER discontinuous Galerkin finite element schemes with a posteriori sub-cell finite volume limiting, Computers & Fluids 118, 204-224 (2015)

[12] Olindo Zanotti, Francesco Fambri, Michael Dumbser. Solving the relativistic magnetohydrodynamics equations with ADER discontinuous Galerkin methods, a posteriori subcell limiting and adaptive mesh refinement, MNRAS 452, 3010-3029 (2015)

[13] Michael Dumbser, Ilya Peshkov, Evgeniy Romenski, Olindo Zanotti. High order ADER schemes for a unified first order hyperbolic formulation of continuum mechanics: viscous heat-conducting fluids and elastic solids, JCP (submitted, http://arxiv.org/abs/1511.08995)

[14] Donner, S., Bernauer, M., and Igel, H., Inversion for seismic moment tensors combining translational and rotational ground motions, submitted to Geophysical Journal International, 2015.

[15] Wenk, S., Hadziioannou, C., Pelties, C., and Igel, H., Investigating the generation of Love waves in secondary microseisms using 3-D numerical simulations, EGU 2014 General Assembly, Vienna, Austria, 2014.

[16] Gabriel, A.-A., Pelties, C., Atanasov, A., Sachdeva, V., Passone, L., Jordan, K. E., Ely, G., Mai, P. M., Large-Scale Earthquake Dynamic Rupture Simulations with the ADER-DG Method: Towards simulation based seismic hazard assessment, Oral presentation, 2013 SIAM Conference on Mathematical and Computational Issues in the Geosciences, Padova, Italy, 2013.

[17] A. Heinecke, A. Breuer, S. Rettenberger, M. Bader, A.-A. Gabriel, C. Pelties, A. Bode, W. Barth, X.-K. Liao, K. Vaidyanathan, M. Smelyanskiy and P. Dubey, "Petascale High Order Dynamic Rupture Earthquake Simulations on Heterogeneous Supercomputers", SC14 Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, 3–14, ACM Gordon Bell Prize Finalist.

[18] Thomas Ulrich, Alice-Agnes Gabriel, Influence of fault geometric heterogeneities on the dynamic rupture process, presented at NMEM (Numerical Modeling of Earthquake Motions, Slovakia, 07/2015)

[19] Thomas Ulrich, Elizabeth H. Madden, Stephanie Wollherr and Alice-Agnes Gabriel, Large scale dynamic rupture scenario of the 2004 Sumatra-Andaman megathrust earthquake,submitted to EGU General Assembly 2016

[20] C. Pelties, A.-A. Gabriel, and J.-P. Ampuero, "Verification of an ADER-DG method for complex dynamic rupture problems", Geosci. Model Dev., 7, 847– 866, doi:10.5194/gmd-7-847-2014; Geosci. Model Dev. Disc., 6(4):5981–6034, doi:10.5194/gmdd-6-5981-2013.

[21] Madden, E.H., F. Maerten and D.D. Pollard (2013), Mechanics of non-planar faults at extensional steps with application to the 1992 M 7.3 Landers, California, earthquake, Journal of Geophysical Research, v.118, 1-15, doi: 10.1002/jgrb.50237.

[22] Wesnousky, S. G. (1986), Earthquakes,Quaternary Faults, and Seismic Hazard in California. Journal of Geophysical Research, v.1, 12587-12631.

[23] Wesnousky, S. G. (2008). Displacement and geometrical characteristics of earthquake surface ruptures: Issues and implications for seismic-hazard analysis and the process of earthquake rupture, Bull. Seismo. Soc. Am. 98, no.4, 1609-1632.

Life Sciences and Biology

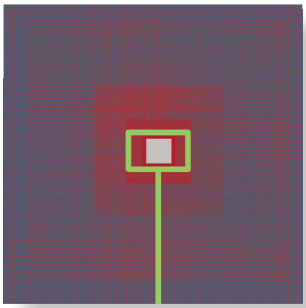

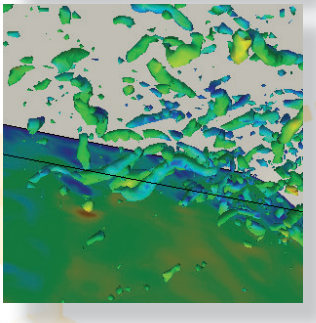

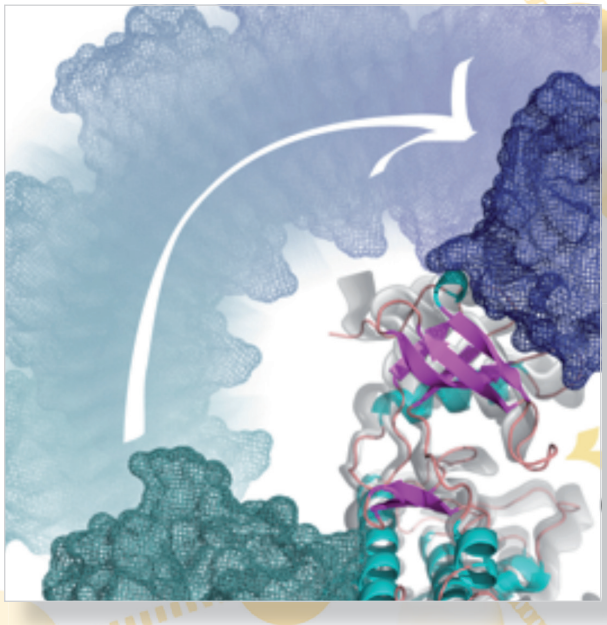

Dynamics of Transmembrane Domains: Impact on Complex Membrane Processes

Research Institution
Chemie der Biopolymere, WZW, TUM, and Physics Department, TUM

Principal Investigator
Dieter Langosch, Christina Scharnagl

Researchers
Norman Blümel, Alexander Götz, Markus Gütlich, Kai Libicher

Project Partners
–

SuperMUC Project ID: pr42ri

Introduction

The biological functions of all proteins are intrinsically linked to their structure, their interactions with other proteins, and their conformational dynamics. This is particularly true for integral membrane proteins that are involved in all kinds of communications between the inside of a cell and its exterior. Hence, a deep understanding of these molecular machines is of utmost importance. Unfortunately, membrane proteins are notoriously difficult to study by experimental methods. Computational approaches like molecular dynamics (MD) simulations provide those insights at a level of detail unreachable by experimental approaches. However, this level of detail is connected to a high demand on computational resources only offered by modern high performance (HPC) clusters.

Results and Methods

We applied MD simulations to study the membrane-embedded parts of certain membrane proteins, in order to understand the relationships between their structure, conformational dynamics, and functional properties [1]. The central questions and results are summarized in Fig.1.

In a first project, the conformational dynamics of the transmembrane helix of the amyloid precursor protein (APP) was investigated. APP is cleaved proteolytically by γ-secretase forming small toxic peptides that are thought to underlie the etiology of Alzheimer's disease. Understanding the molecular architecture of the APP transmembrane helix and its link to cleavage is therefore of utmost importance for understanding the first steps leading to this disease. To simulate the environment in the active site of the enzyme, the dynamics of the substrate helix was studied in a membrane mimicking solvent, containing a low percentage of water. Building on our previous work [2], we characterized bending

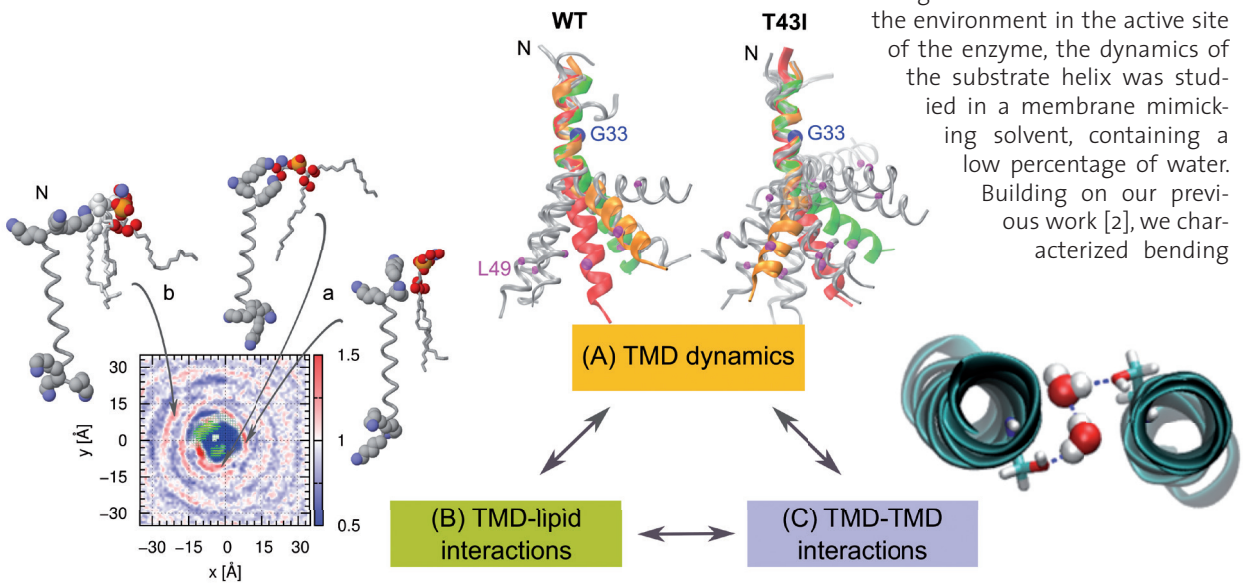

Fig. 1 Dynamics and interactions of single-span transmembrane domains (TMDs). (A) Large-amplitude backbone motions of the TMD of wild type (WT) APP and its T43I mutant, which is linked to heritable Alzheimer's disease. The removal of a stabilizing side-chain to main-chain hydrogen bond increases helix bending around a central hinge. (B) Bilayer restructuring around the charged N-terminus of a *de novo* designed highly flexible model TMD. Areas in red (blue) colour indicate increased (reduced) lipid density as compared to bulk. Shown are also snapshots of TMD-lipid arrangements from the first (a) and second shell (b). (C) Water (van der Waals spheres) linking the transmembrane helices of quiescin sulfhydryl oxidase II via hydrogen bonding to serine residues.

motions at a hinge within the transmembrane helix and their modification by primary structure changes, some of which are linked to heritable Alzheimer's disease (Fig. 1A). Our results [2-5] thus provide insights into the mechanism by which such mutations cause the disease. Moreover, our results suggest an entirely new model of intramembrane proteolysis were reaching a cleavage-competent state requires searching a complex energy landscape by the substrate/enzyme complex, rather than simply local unfolding of the substrate around its cleavable bond. This novel view prompted formulation of a number of computational and experimental approaches in a collaborative research program which is now funded by DFG (research unit 2290 "Understanding Intramembrane Proteolysis").

The computational challenge lies in the sufficient exploration of the protein's conformational space. We tackle this problem using multiple (> 150) independent copies of the simulation system, combined with Markov-State modeling. Up to now an aggregate time of 20 microseconds has been reached and will be further extended by adaptive sampling runs. To handle this large number of jobs in parallel we used a MPI based job farming approach. Therefore, we applied the Redisexec framework developed by the LRZ, which allows the implementation of complex dependency trees between simulations. We generally used 40k cores (peak 60k) without any major performance loss. Our work will be continued in project pr48ko which has been granted 35 mio core-h by the Gauss Collaboration in November 2015.

In a second project, we have been investigating a set of paradigmatic transmembrane helices de novo designed by mixing different types of amino acids in order to obtain different extents of conformational dynamics in the helices. While earlier simulations have revealed how primary structure dictates helix dynamics, we have now accumulated extensive simulations in membranes (aggregate time > 50 microseconds). These long simulation times are necessary to achieve sufficient convergence of the membrane environment. This project benefited from the enhanced performance of SuperMUC as compared to previous HPC systems. The production of >60 nanoseconds/day reduced the time necessary for such simulations from several months to only a few weeks. We used both single-component lipid membranes as well as complex lipid mixtures. While a detailed analysis is still ongoing, our preliminary results indicate that i) a membrane environment rigidifies a helix, ii) lipids preferentially interact with the peptides via electrostatic interactions at the membrane/lipid head group interface, and iii) peptide/lipid interaction is a two-way-process whereby the helices are influenced by lipids and vice versa (unpublished results) (Fig. 1B). We hope to clarify how helix dynamics and peptide/lipid interactions affect certain functional properties of the helices known from experiments, such as their ability to fuse lipid membranes and to flip individual lipid molecules across a bilayer. The results will shed light on both of these elementary, and currently ill-understood, processes in cell biology, the ability of certain proteins to fuse membranes in cellular secretion as well as in redistributing lipid molecules after their synthesis at intracellular membranes. We plan to extend this work in order to probe the relevance of our model studies for natural membrane proteins in physiological membrane environment (project pr92so).

In a third project we have investigated the self-interaction of the transmembrane helix of quiescin sulfhydryl oxidase II, an enzyme that is important in protein folding in different subcellular compartments. Originally inspired by bioinformatic analysis that identified a conserved face of its transmembrane helix, we have confirmed the relevance of the conserved residues for self-interaction using a bacterial reporter system. To better understand the molecular nature of the interface between the helices, we performed atomistic MD simulations in a bilayer system mimicking a cellular membrane. While the molecular model is fully consistent with the experimental data, it identified an interesting new feature. One, two, or three water molecules entered the helix-helix interface during the simulation via the apolar lipid membrane thus forming an interhelical hydrogen-bonding network that appears to stabilize this helix dimer (Fig. 1C). The entrance of water into the hydrophobic interior of the membrane is a rather slow process. In our simulations we could overcome the time barriers by running many simulations from different start conditions. In addition, the simulations revealed binding of certain lipid types to ionizable and polar residues of the transmembrane domain. It appears therefore as if the propensity of these helices to form a dimer might be modulated by interfacial water and/or by binding of lipid molecules to the helices (in print).

Conclusion

In conclusion, MD simulations provide an extremely valuable atomistic view onto the inner workings of molecular machines that support the life of cells in multiple ways. No other approach can provide the molecular detail attained by atomistic simulations. Clearly, however, the major current goal is to reach simulation times that come close to biologically relevant times.

References and Links

[1] http://cpb.wzw.tum.de/
http://users.physik.tu-muenchen.de/scharnagl/

[2] Pester, O., Barret, P., Hornburg, D., Hornburg, P., Pröbstle, R., Widmaier, S., Kutzner, C., Dürrbaum, M., Kapurniotu, A., Sanders, C. R., Scharnagl, C. & Langosch, D. (2013). The Backbone Dynamics of the Amyloid Precursor Protein Transmembrane Helix Provides a Rationale for the Sequential Cleavage Mechanism of γ-Secretase. J. Am. Chem. Soc. 135, 1317-1329.

[3] Pester, O., Götz, A., Multhaup, G., Scharnagl, C. & Langosch, D. (2013). The Cleavage Domain of the Amyloid Precursor Protein Transmembrane Helix does not Exhibit Above-Average Backbone Dynamics. ChemBioChem 14, 1943-1948.

[4] Scharnagl, C., Pester, O., Hornburg, P., Hornburg, D., Götz, A. & Langosch, D. (2014). Side-Chain to Main-Chain Hydrogen Bonding Controls the Intrinsic Backbone Dynamics of the Amyloid Precursor Protein Transmembrane Helix. Biophys. J. 106, 1318-1326.

[5] Langosch, D., Scharnagl, C., Steiner, H. & Lemberg, M. K. (2015). Understanding Intramembrane Proteolysis: from Protein Dynamics to Reaction Kinetics. Trends Biochem Sci 40, 318-327.

Gene mapping methods for complex diseases and integrative analysis of high dimensional 'omics data

Research Institution
Helmholtz Zentrum München, Institut für Genetische Epidemiologie und Abteilung für Molekulare Epidemiologie, Ludwig-Maximilians-Universität München, IBE – Lehrstuhl für Genetische Epidemiologie

Principal Investigator
Konstantin Strauch

Researchers
Brigitte Kühnel, Melanie Waldenberger, Christian Gieger

Project Partners
–

SuperMUC Project ID: pr45zi (Gauss Large Scale project)

Introduction

The identification of genetic factors that are responsible for complex diseases in humans is an important albeit challenging task. Finding the genes responsible for a certain disease opens the perspective to development of a causative therapy, for example by means of a new medication that acts on the corresponding drug target, or to enable approaches to individualized treatment or prevention. There are two major types of study designs used for the genetic mapping of diseases and related quantitative traits: Linkage analysis and association analysis. Linkage analysis is based on the co-segregation, i.e., joint inheritance, of a genetic marker locus and a trait locus (disease gene) through a family, also called pedigree. On the other hand, association analysis relies on assessing allele frequency differences between affected and unaffected individuals, or between individuals with different levels of a quantitative phenotype, either in unrelated individuals or nuclear families. Here, an allele refers to a particular genetic configuration at a certain position within the genome. Whereas association analysis is well suited to localize genetic variants that are common in the studied population (i.e., occur at a frequency of at least 5%), linkage analysis has good potential to map rare genetic variants, which can be enriched in the sample by using families instead of unrelated persons.

Along these lines, with our SuperMUC project we aimed at contributing to the dissection of complex human diseases by further development and application of gene-mapping analysis methods using disease phenotypes and large-scale 'omics data, including a substantial number of genetic as well as epigenetic markers. While an individual inherits the genetic information (coded within the deoxyribonucleic acid or DNA) from the parents, epigenetic marks can be either inherited or newly set during an individual's life. Here, with regard to epigenomics, we focus on methylation levels of DNA, which are not only important for gene regulation but also for the control of genomic imprinting [1]. This epigenetic effect leads to preferential expression of the copy of a gene that has been inherited either from the father or from the mother [2]. The first part of our SuperMUC project, which we describe in detail here, involved systematically evaluating the relation between genetic and epigenetic markers. In particular, a genome-wide association analysis was performed for methylation levels at many positions, so-called CpG sites, throughout the genome. This led us to gain important information regarding which genetic markers play a role in setting epigenetic patterns, and thereby ultimately help understand the mechanisms by which genetic markers contribute to determining gene expression and, if altered abnormally, how they contribute to the development of diseases. Our analyses are based on the population-based KORA study (Collaborative Research in the Region of Augsburg), for which disease phenotypes, life-style factors and large-scale 'omics data are available. Hence, the KORA study represents a unique epidemiological data resource that is well suited to address our research question.

Genotypes and DNA Methylation

There is considerable interest in the field of epigenetics, not only in developmental biology and cancer where it continues to uncover mechanistic insight and novel diagnostic and prognostic biomarkers, but more broadly in relation to many common complex diseases in humans. Epigenetics has gained prominence with respect to how the environment and lifestyle impact upon disease risk, in the field of developmental origins of adult diseases and in age-related functional decline. A major challenge facing the field of epigenetics is deciphering the contribution of genetic variation to epigenetic variation. Several examples of single-nucleotide polymorphisms (SNPs) associated with DNA methylation have already been reported suggesting that the inter-individual variability of DNA methylation is frequently shaped by genetic variation. A study on DNA methylation of cytosine-guanine dinucleotide (CpG) islands on chromosome 21 in leukocytes from healthy individuals observed a strong correlation between DNA methylation and genotype [3]. Along this line, a recent study in humans identified several cas-

es of allele-specific DNA methylation at nonimprinted gene loci, where the methylation status of each allele was likely controlled in cis by the local DNA sequence [4]. Therefore, allele-specific methylation represents an epigenetic pathway of how genetic polymorphisms may lead to phenotypic variability [3]. So far the relation between genotypes and DNA methylation has only been investigated for specific genetic regions.

We analyzed the influence of genetic variants on methylation intensity in cis and trans on a genome-wide scale. To this end, 1800 individuals of the KORA F4 study (Collaborative Research in the Region of Augsburg) with methylation data for 440,000 CpG sites as well as genetic SNP data were available.

Genetic SNP data for KORA F4 have been generated using the Affymetrix Axiom array and imputed using the 1000Genomes reference panel and were analyzed for genome-wide associations to methylation data obtained by the Illumina HumanMethylation 450K BeadChip. This array covers 96% of CpG islands (CGIs) with multiple sites in the annotated CGI, north or south shores (regions flanking island), and north or south shelves (regions flanking shores), 99% of RefSeq genes, with an average of 17 CpG sites per gene region distributed across the promoter, 5'UTR, first exon, gene body, and 3'UTR. Quality control and normalization of the 450k data has been performed as described previously for our association study between methylation levels and smoking [5].

The association analyses were based on current versions of 1000Genomes imputation data which yield 10 million SNPs throughout the genome with minor allele frequencies above 1%. For each of the 440,000 CpG sites we performed a genome-wide association analysis, i.e., we evaluated the effect of each of the 10 million SNPs on each of the 440,000 CpG sites. This computationally extensive analysis allows us to gain important information regarding which genetic markers play a role in setting epigenetic patterns, and thereby ultimately help understand the mechanisms by which genetic markers contribute to determining gene expression and, if altered abnormally, how they contribute to the development of diseases. With our analysis we found a high proportion of significant associations. These results lay the foundation for a deeper understanding of the mechanisms of gene expression and regulation. The Manhattan plot (Figure 1) shows the trans-chromosomal results. The cis results indicate an even higher amount of significant associations.

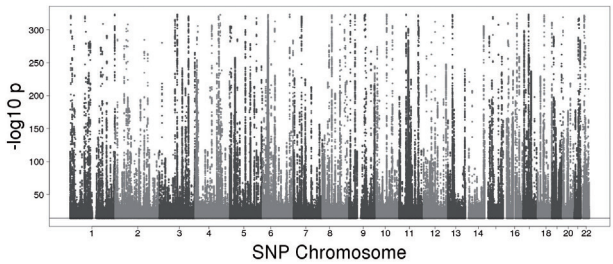

Figure 1: Manhattan plot for trans-chromosomal results

Computation

For the genome-wide association analysis of 440,000 CpG sites, we used the standard software packages OmicABEL and reshuffle, which are freely available and exactly suitable to this type of association analysis. The OmicABEL package for R allows rapid mixed-model based genome-wide association analysis; it efficiently handles large datasets, and both single trait and multiple trait ("omics") analysis.

The software is licensed under the GNU GPL v.3 license as part of the GenABEL project for statistical genomics at http://www.genabel.org/packages/OmicABEL.

The output produced by OmicABEL is kept in a compact binary format for performance reasons. The user can then use "reshuffle" to extract the desired information. Since the output may be very large, reshuffle allows the selection of specific ranges of SNPs and traits. We ran our OmicABEL computations on 40 cores, since it is limited to one node. Due to the SuperMUC requirement that every single job must not run longer than 48 hours we had to split the genome-wide analysis into 533 parts. Since OmicABEL produces very large output files we applied for more disk space. It was agreed that for a limited time frame of 6 weeks we had up to 263 TB of disk space at our disposal. In order to reduce the prohibitively large OmicABEL output files we used its accompanying utility program reshuffle to extract only a subset of significant pairs of SNPs and methylation sites. Our computations and subsequent data management used a total of 40,000 core hours.

References

[1] Sasaki, H., DNA methylation in epigenetics, development, and imprinting. Encyclopedia of Genetics, Genomics, Proteomics and Bioinformatics, 2005. 1: 1.3: 32.

[2] Strauch, K., Gene mapping, imprinting, and epigenetics. Encyclopedia of Genetics, Genomics, Proteomics and Bioinformatics, 2005. 1: 1.4: 49.

[3] Zhang, Y., et al., Non-imprinted allele-specific DNA methylation on human autosomes. Genome Biol, 2009. 10(12): R138.

[4] Kerkel, K., et al., Genomic surveys by methylation-sensitive SNP analysis identify sequence-dependent allele-specific DNA methylation. Nat Genet, 2008. 40(7): p. 904-8.

[5] Zeilinger, S., et al., Tobacco smoking leads to extensive genome-wide changes in DNA methylation. PLoS One, 2013. 8(5): e63812.

Binding Specificity and Allostery of Biomolecular Interactions

Research Institution
Center for Bioinformatics, Saarland University, Saarbrücken

Principal Investigator
Volkhard Helms

Researchers
Ozlem Ulucan

Project Partners
Matthias Engel

SuperMUC Project ID: pr58go

Introduction

Biological cells are filled to about 30% with proteins. Except for those proteins that are bound to some cellular structures, proteins diffuse freely in the cell according to the well-known Brownian motion. In the crowded cellular environment, this means that proteins constantly bump into other proteins.

In about some cases, such a collision actually leads to a biologically meaningful contact of the two proteins. Then, the binding partners may either remain bound or a chemical reaction may take place whereby, for example, an electron or a phosphate group is transferred from one protein to the other one before they unbind again. Such meaningful encounters are called "specific" interactions. In all other cases, the proteins only form short-lived "non-specific" contacts that are biologically not meaningful. Specific complexes often involve larger contact interfaces than non-specific complexes and are stabilized by favorable interactions such as hydrogen bonds or salt-bridges between the binding partners.

Our group has a particular interest in unravelling the biophysical principles underlying such biomolecular interactions [1]. In project pr58go that was also supported by DFG normal grant HE 3875/11-1, we therefore studied the dissociation of three different protein-protein pairs. The selected protein pairs all form specific complexes that are stabilized by hydrophilic binding interfaces. The three protein complexes are illustrated on the left side of Fig. 1.

Results and Methods

We employed atomistic molecular dynamics (MD) simulations that were performed with the GROMACS 4.5 software package [2] to characterize the free energy landscape governing the association and dissociation of the three protein-protein pairs. Precisely, we computed with the help of umbrella-potential restrained simulations the potential of mean force (PMF) between the two proteins along a one-dimensional reaction coordinate [3, 4]. Because we had to follow the unbinding pathway up to

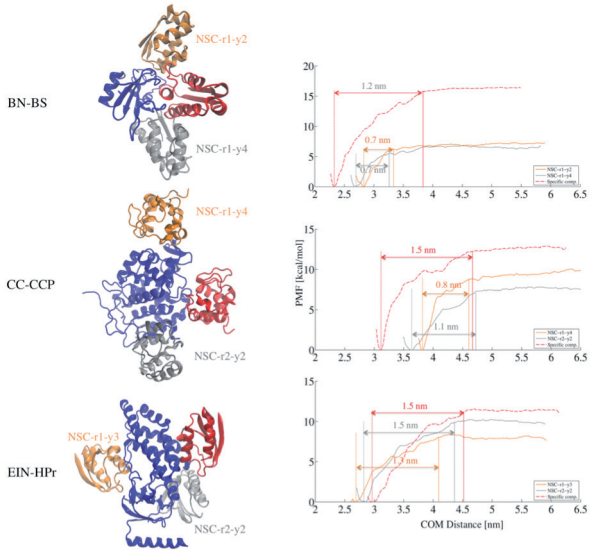

Figure 1: Pictures on the left show from top to bottom the protein-protein complexes barnase (blue) and barstar, cytochrome peroxidase c (blue) and cytochrome c, and the complex of the N-terminal domain of enzyme I (blue) with the histidine-containing phosphocarrier. The red colored orientations of barstar, cytochrome c, and histidine-containing phosphocarrier are those orientations observed in the specific complexes that were determined by X-ray crystallography. The silver and gold colored orientations are alternative binding orientations observed in MD simulations that led to transiently bound non-specific contacts. Pictures on the right illustrate the computed potentials of mean force (PMF on the y-axis) for the dissociation of the red, gold or silver orientations into the unbound state beyond separation distances beyond 4.0 - 4.5 nm between the centers of mass of the two proteins (x-axis).

distances of several nanometers - where the proteins do not attract eachother anymore – the simulation systems contained several hundreds of thousands of atoms. For each binding orientation, we computed 21 windows at different separation distances of 10 – 40 ns length each. In total, the two projects described here consumed 12 million core hours on SuperMUC.

The right side of Fig. 1 illustrates the computed free energy profiles. Interestingly, we found that all three protein complexes bind "downhill" into both the specific and non-specific complexes without passing over an activation energy barrier. This behavior is comparable to some fast-folding proteins that fold "downhill" in a steep folding funnel.

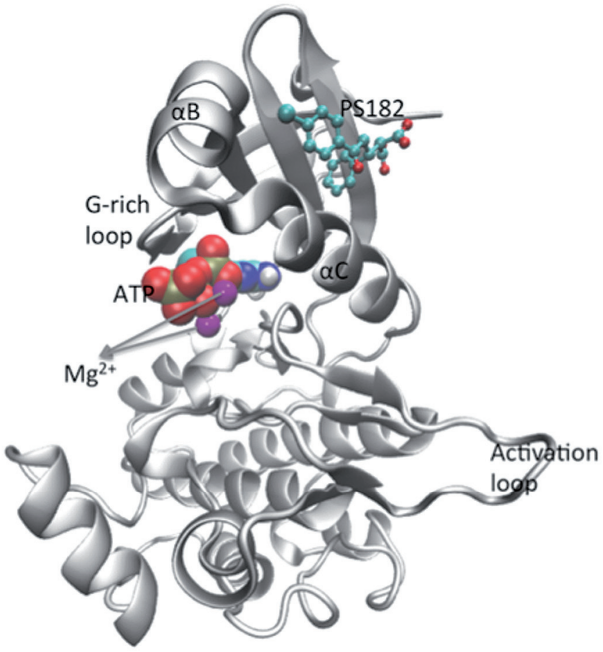

Figure 2: Crystal conformation of the small-molecule ligand PS182 and the co-factor ATP bound to the protein kinase PDK1.

The second interesting observation was that the attractive interactions between both proteins extend to 1.2 to 1.5 nm of separation and then become flat. This behavior was found for all three complexes, and may be a general behavior of specific hydrophilic protein complexes [3].

Studying such systems on a computer enables one to generate also starting conformations that cannot be studied so easily in experiment. To contrast the binding into specific complex orientations with the binding into non-specific contact orientations mentioned before, we performed unbiased MD simulations of the binding partners and monitored all contact orientations that differ from the specific complexes. Among all such contacts we selected the ones with largest binding interface and the longest-lived ones as starting positions for PMF calculations. In Fig. 1, these orientations are colored gold and silver for the three systems.

For these "non-specific" contacts, we then computed the same association free energy profiles [4] as for the specific orientations before. As expected, the difference in free energy between the bound state and the unbound state turned out to be much smaller than for the specific complexes. Interestingly, also the range of the attractive interaction is shorter (0.7 – 1.0 nm) - except for the EIN-HPr pair - and the interface formed between the proteins is significantly smaller than for the native specific complexes.

In a second project, we investigate in collaboration with Dr. Matthias Engel/Dept. of Pharmacy, Saarland University how the modulator PS182 alters the affinity of phosphoinositide kinase 1 (PDK1) towards the co-factor ATP:Mg2, see Fig. 2. This is an example of allosteric regulation of enzymatic activity. Allostery is found in practically all regimes of cellular processes, but its underlying principles remained largely unclear to date.

Using free energy perturbation calculations with the GROMACS package [2], we computed the binding free energy of ATP:Mg_2 in the presence and absence of PS182. Since PS182 induces conformational changes in PDK1 that favor a tighter packing of the protein around ATP, the difference between the two binding free energies can be attributed to as the allosteric contribution of PS182 to the binding free energy for ATP, see the thermodynamic cycle shown in Fig. 3.

This study is, to our knowledge, the first attempt to characterize such an allosteric contribution based on molecular dynamics simulations.

On-going Research / Outlook

Due to the lengthy and large-scale MD simulations, only access to SuperMUC made this project feasible at all. Fortunately, we did not face technical obstacles during the course of this project. Instead, both systems presented here required careful fine-tuning of the simulation protocols in terms of how to generate "good" starting conformations for the 21 windows (first project) and in terms of how two missing magnesium ions should be positioned in the active site of PDK1 where they helped to stabilize the geometry of its active site (second project). Our group will continue working towards a mechanistic understanding of biomolecular interactions. Dr. Mazen Ahmad at the nearby MPI of Informatics recently derived an intriguing theoretical model how conformational changes contribute to binding thermodynamics [5]. We plan to incorporate such developments in our further research in order to study interconnected intramolecular and intermolecular processes.

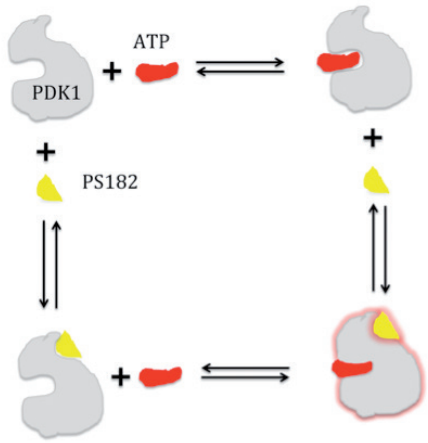

Figure 3: Thermodynamic cycle to characterize the contribution of the allosteric modulator PS182 (yellow) to the binding free energy of ATP (red) to the protein kinase PDK1. The allosteric contribution is defined as the difference of the two binding reactions shown at the top and bottom, respectively.

References and Links

[1] gepard.bioinformatik.uni-saarland.de
[2] http://www.gromacs.org
[3] Ulucan O, Helms V (2014) J Chem Theor Comp, 10., 3512-3524.
[4] Ulucan O, Helms V (2015) J Phys Chem B, 119, 10524-10530.
[5] Ahmad M, Helms V, Lengauer T, Kalinina O (2015) J Chem Theor Comp 11, 2945-2957

Disentangling Evolution on the SuperMUC

Research Institution
Heidelberg Institute for Theoretical Studies

Principal Investigator
Alexandros Stamatakis

Researchers
Alexey Kozlov, Lucas Czech, Diego Darriba, Tomas Flouri

Project Partners
1KITE initiative, Leibniz-Institut DSMZ, TU Kaiserslautern

SuperMUC Project ID: pr58te

Introduction

Phylogenetic tree reconstruction strives to infer the evolutionary relationships among a set of organisms (species, frequently also denoted as taxa) based on molecular sequence data. Recent advancements in sequencing technology, in particular the emergence of so-called next generation sequencers, have generated an avalanche of sequence data, that now makes it possible to use whole transcriptomes and even genomes of a large number of species for tree reconstruction.

Likelihood-based approaches (Maximum Likelihood and Bayesian Inference) represent an accurate and widely used, but at the same time also highly compute-intensive approach for reconstructing phylogenetic trees. Thus far, we were able to test and substantially improve the scalability and efficiency of our Bayesian and Maximum Likelihood based tree reconstruction tools on SuperMUC. We have also already published the first results of large-scale empirical data analyses.

Results and Methods

Toward the end of 2014 we published two papers in Science that both made it to the cover page. The papers presented phylogenetic trees for two of the largest datasets analyzed to date with likelihood-based methods. The bird tree of life contained data for about 50 bird genomes while the insect tree of life comprised roughly 150 insect transcriptomes. Perhaps more importantly, we freely made available as open source tools and also further developed and maintained our two tools for likelihood-based phylogenetic inference: ExaML (Exascale Maximum Likelihood) and ExaBayes (Exascale Bayesian inference). Interestingly, a recent critique of the bird phylogeny paper that was also published in Science actually used our ExaBayes code.

Since those two first empirical data analyses (birds & insects) had been completed in late 2014 and because we achieved substantial performance improvements for ExaML and ExaBayes because of the introduction of novel parallelization schemes, data distribution algorithms, etc., we have consumed less processor hours than initial planned.

This provided us the chance to (i) work on developing new computational methods and further improve our codes and (ii) conduct additional empirical data analyses on insects but also on two new projects for which the processor time budget would not have been sufficient, given the state of the software when we initiated this project.

In the following we quickly outline the software improvements as well as the new data analysis projects that were initiated in 2015.

With respect to software, one important contribution, not only regarding our own code ExaBayes, but also to Bayesian phylogenetic inference in general was the development of a novel branch length proposal mechanism that can substantially accelerate the convergence of Bayesian tree searches. The new method was recently published in *Systematic Biology* [4] (Impact factor: 14.3) and is fully implemented in ExaBayes.

Another important contribution was the release of ExaML version 3.0 [5] that features the novel data distribution algorithm, additional models of evolution, faster handling of I/O, and production-level support for the Intel Xeon PHI architecture. The Xeon PHI scalability experiments and adaptation were conducted on the respective LRZ system.

While the bird genome evolution project is completed, this is not the case for the 1000 insect transcriptome project (www.1kite.org). In 2015 we used SuperMUC to infer phylogenies for subsets/subtrees of the 150 insects analyzed for the 2014 Science paper, partially using different and simply more data. That is, after getting the big picture, we are now taking a closer look at individual parts of that tree. We expect the results of all these analyses to be published in 2016.

In addition to the above, we have also initiated two new collaborations. We started collaborating with the Leibniz-Institut DSMZ - Deutsche Sammlung von Mikroorganismen und Zellkulturen GmbH to infer the phylogeny of around 1000 full bacterial genomes using ExaML on the SuperMUC.

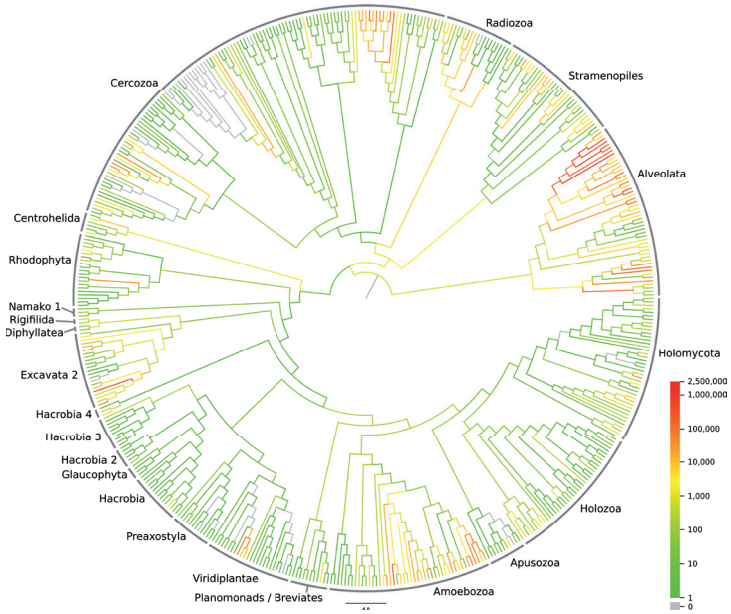

Figure 1. Protist diversity in new world tropical forest soils.

With collaborators at TU Kaiserslautern we are currently working on determining the protist diversity in new world tropical forest soils. Samples have been collected at 150 positions in three central American rain forests. Protists are Eukaryotes (i.e., contain a single or multiple cells with a nucleus). They form a diverse taxonomic group that can not be classified as animals, plants, or fungi. The protists include, for instance, amoeba, red algae, and dinoflagellates. This project does neither rely on ExaML nor on ExaBayes, but uses the so-called Evolutionary Placement Algorithm (EPA) we developed in our lab. We expect to publish the results of this project in early 2016. In Figure 1. we depict the protist diversity in the central American tropical forest soils by means of a reference phylogeny. Here, red branches indicate a high degree of diversity.

On-going Research / Outlook

As already mentioned, the 1KITE insect project is still on-going. In 2015 our collaborators sequenced approximately 1500 insect transcriptomes which we intend to analyze on SuperMUC in 2016 with ExaBayes and ExaML. This will yield a tree that is one order of magnitude larger with respect to the number of species than the phylogeny presented in late 2014.

While our software will be able to handle these analyses on SuperMUC, there is room left for making the codes run faster and hence in a moire energy-efficient way. In particular, there is a plethora of computational shortcuts available that can be used to further reduce memory footprints and calculation times of the so-called phylogenetic likelihood function that consumes more than 90% of the accumulated total run-time of ExaML and ExaBayes. However, not all tricks work efficiently on all types of datasets and for all types of computations. To this end, we will develop an auto-tuned likelihood kernel that will automatically determine the most optimal likelihood function implementation for the dataset at hand.

The implementation of these tricks will also require us to adapt and change the current data distribution algorithms of ExaML and ExaBayes, since some of the assumptions we have made thus far for optimally distributing data will not hold any more with the revised likelihood calculation approaches.

Another challenge will be to work on the scalability of the Evolutionary Placement Algorithm that is being used for analyzing Protist diversity in new-world tropical forest soils.

Initially, the code was not designed for deployment on supercomputers, since at the time of writing it, we underestimated the molecular data avalanche. While it can be deployed on SuperMUC as is, we are currently working on a complete re-design of the code, that will also make use of the aforementioned tricks for accelerating likelihood calculations. This re-designed code will be better suited for deployment on supercomputers.

Finally, we are also actively working on novel methods for the post-analysis of EPA results using various data mining and machine learning approaches. For instance, one of the key questions is how to compare samples. Assume that we take protist samples from two different tropical forests we want to be able to quantify and compare their degree of diversity. Note that, these techniques also have important medical applications, since the analysis of gut microbial diversity and associated pathogenicity boils down to exactly the same computational challenges.

References and Links

[1] lab web-site: www.exelixis-lab.org

[2] E.D. Jarvis, S. Mirarab, A.J. Aberer, B. Li, P. Houde, C. Li, S.Y.W. Ho, B.C. Faircloth, B. Nabholz, J.T. Howard, A. Suh, C.C. Weber, R.R. da Fonseca, J. Li, F. Zhang, H. Li, L. Zhou, N. Narula, L. Liu, G. Ganapathy, B. Boussau, Md.S. Bayzid, V. Zavidovych, S. Subramanian, T. Gabaldon, S. Capella-Gutiérrez, et al. Whole-genome analyses resolve early branches in the tree of life of modern birds". Science, 346(6215):1320-1331, 2014.

[3] B. Misof, S. Liu, K. Meusemann, R.S. Peters, A. Donath, C. Mayer, P.B. Frandsen, J. Ware, T. Flouri, R.G. Beutel, O. Niehuis, M. Petersen, F. Izquierdo-Carrasco, T. Wappler, J. Rust, A.J. Aberer, U. Aspöck, H. Aspöck, D. Bartel, A. Blanke, S. Berger, A. Böhm, T.R. Buckley, B. Calcott, J. Chen, F. Friedrich, M. Fukui, M. Fujita, C. Greve, P. Grobe, S. Gu, Y. Huang, L.S. Jermiin, A.Y. Kawahara, L. Krogmann, M. Kubiak, R. Lanfear, H. Letsch, Y. Li, Z. Li, J. Li, H. Lu, R. Machida, Y. Mashimo, P. Kapli, D.D. McKenna, G. Meng, Y. Nakagaki, J.L. Navarrete-Heredia, M. Ott, Y. Ou, G. Pass, L. Podsiadlowski, H. Pohl, B.M. von Reumont, K. Schütte, K. Sekiya, S. Shimizu, A. Slipinski, A. Stamatakis, W. Song, X. Su, N.U. Szucsich, M. Tan, X. Tan, M. Tang, J. Tang, G. Timelthaler, S. Tomizuka, M. Trautwein, X. Tong, T. Uchifune, M.G. Walzl, B.M. Wiegmann, J. Wilbrandt, B. Wipfler, T.K.F. Wong, Q. Wu, G. Wu, Y. Xie, S. Yang, Q. Yang, D.K. Yeates, K. Yoshizawa, Q. Zhang, R. Zhang, W. Zhang, Y. Zhang, J. Zhao, C. Zhou, L. Zhou, T. Ziesmann, S. Zou, Y. Li, X. Xu, Y. Zhang, H. Yang, J. Wang, J. Wang, K.M. Kjer, X. Zhou. Phylogenomics resolves the timing and pattern of insect evolution. Science, 346(6210): 763-767, 2014.

[4] A.J. Aberer, A. Stamatakis, F. Ronquist. An Efficient Independence Sampler for Updating Branches in Bayesian Markov chain Monte Carlo Sampling of Phylogenetic Trees. Systematic Biology, 65 (1): 161-176, 2016.

[5] A.M. Kozlov, A.J. Aberer, A. Stamatakis. ExaML Version 3: A Tool for Phylogenomic Analyses on Supercomputers. Bioinformatics, 31(15): 2577-2579, 2015.

High Performance Methods for Computational Fluid Dynamics including Multiphysics Scenarios

Research Institution
Institute for Computational Mechanics, Technische Universität München

Principal Investigator
Wolfang A. Wall, Martin Kronbichler

Researchers
Christian Roth, Benjamin Krank, Ursula Rasthofer, Francesc Verdugo

Project Partners
–

SuperMUC Project ID: pr83te

Introduction

Fluid flow is an essential aspect of many applications in multiphysics including the broad field of biomechanics, fluid-structure interaction, scalar transport as well as two-phase flows, which may occur both in laminar and fully turbulent state. We have therefore developed an efficient approach to computational fluid dynamics that is especially designed for application in such multiphysics scenarios. For example, our code allows us to make predictive simulations of the entire human lung using a fully coupled fluid-structure interaction approach [1]. One of the most challenging aspects of such multiphysics developments is the complexity of the software employed. We therefore develop our own multiphysics software platform BACI [2], which is based on continuous and discontinuous finite element methods. BACI is continuously advanced by our research group in an object-oriented C++ environment. Parallelization is based on domain decomposition methods using MPI. State-of-the-art solution techniques for nonlinear and linear systems of equations as well as for coupling of several physical fields are incorporated in BACI and are continuously developed in our group. For the implementation of efficient parallel sparse linear algebra operations, BACI makes use of the open-source software framework "Trilinos" (Sandia National Laboratories).

Simulations of the entire human lung

Mechanical ventilation is a life-saving treatment at the intensive care unit, but can also exacerbate acute and chronic lung diseases such as acute respiratory distress syndrome (ARDS) or severe asthma by locally overstraining healthy lung tissue. In recent years, advances have been made towards more protective ventilation strategies. However, there is still a lack of regional information on an individual patient's lung during ventilation which limits the effectiveness of personalized treatment.

Medical imaging techniques such as computed tomography (CT) can provide a patient's lung geometry at a specific state at relatively high resolution. However, regional information on airflow velocities, pressures and tissue strains are not visible in such images but of outmost importance in medical treatment. Advanced modeling and simulation is used to identify key parameters for an individual patient. We developed a volumetrically constrained nonlinear fluid-structure interaction algorithm that takes into account the full physical behavior of the human lung [3]. With this model it is possible to compute regional airflow velocities as well as local tissue strains and aeration at a spatial resolution of approximately 1.0mm and a temporal resolution of 10ms during several breathing cycles. Investigations of an exemplary patient show that distribution of air is relatively homogeneous during natural breathing (Figure 1). During mechanical ventilation, however, airflow in single smaller airways changes, leading towards an inhomogeneous shift of tissue ventilation compared to the healthy case (Figure 2).

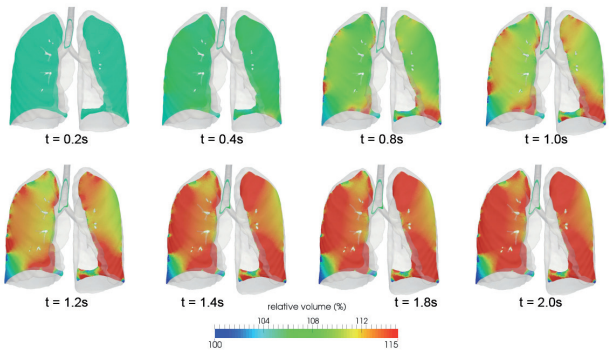

Figure 1: Regional volume increase of lung tissue during the inhalation phase of natural breathing.

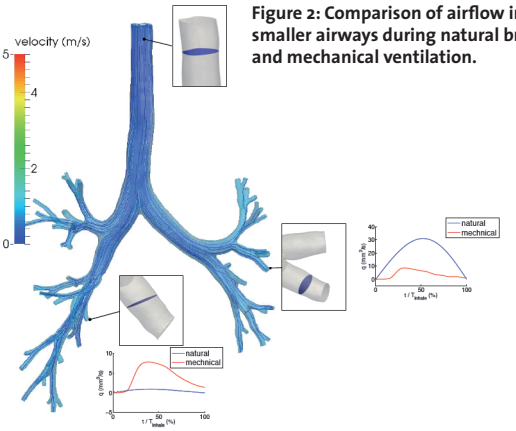

Figure 2: Comparison of airflow in single smaller airways during natural breathing and mechanical ventilation.

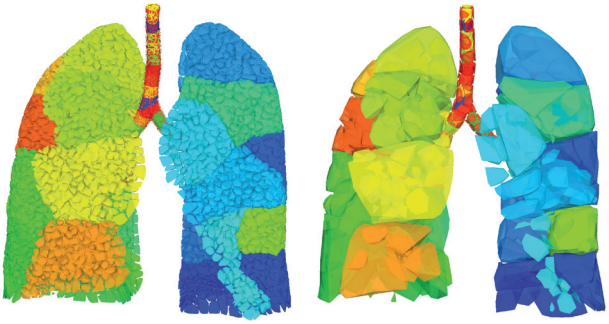

Figure 3: Algebraic multigrid aggregates for the fine level (left) and the coarse level (right) representation of the human lung model.

To conduct larger patient studies with several different ventilation strategies, large-scale numerical simulations have been performed in this project, using in total 3 million core hours (including uncertainty quantification). We designed an algebraic multigrid preconditioning technique for the highly-resolved human lung model [3] (Figure 3). An extension of a previous approach has been used to handle the volume constraint that links airflow and tissue deformation in the lung model. With these methods, an efficient parallel computation on more than one hundred processors is possible for the first time with significant speed-up of the highly-resolved human lung simulations (see [4]).

Large-scale turbulent single and two-phase flow simulations

Another focus of our work is the development of computational tools for large-scale simulations of turbulent flows, which forms the most important component in several multiphysics solvers, including the aforementioned lung simulations. As part of this work, we have elaborated an efficient turbulence modeling technique as well as a novel approach to wall modeling for large-eddy simulation (LES).

Many popular subgrid models for LES only model the net effect of the subgrid scales, i.e., the transfer of energy to consecutively smaller scales, by introducing a subgrid viscosity closure model. In contrast, our turbulence modeling approach, denoted by Algebraic Variational Multiscale-Multigrid-Multifractal Method (AVM4) [5], reconstructs the unresolved subscales directly and includes the actual non-linear interaction. For this reconstruction, a filtered velocity field is required, which is obtained from level-transfer operators of plain aggregation algebraic multigrid methods. This definition renders the method particularly suitable for large-scale computations. Our turbulence modeling approach has successfully been extended to other applications in multiphysics such as active and passive scalar mixing, low-Mach-number flow with variable density [6], and turbulent two-phase flow [7]. The large-scale simulations performed with the resources provided by this project enabled verification of the approach for difficult flow situations, using approximately 6 million core hours on a few hundreds to a few thousands of cores. The computational cost of LES of wall-bounded flows increases dramatically with the friction Reynolds number if near-wall turbulence is resolved, since the size of dominating eddies decreases towards the no-slip boundary. In order to circumvent this dependence, we have developed a new approach to wall modeling within the AVM4 [8]. The basic idea of our technique is to extend the finite element function space by a mean boundary layer profile that can be represented using very coarse meshes. The approach flexibly adapts to separated flows and settings with high adverse pressure gradients by automatically choosing the most appropriate solution among the standard polynomial part and the enrichment. We have found excellent results with this method for example for the flow over periodic hills visualized in Figure 4.

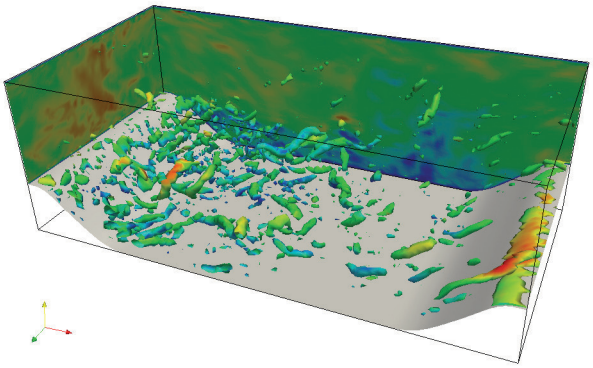

Figure 4: Wall-modeled large-eddy simulation over periodic constrictions.

Outlook

Ongoing work for the lung is focused on investigation of different protective ventilation strategies and their effects on a patient's individual pathophysiology, both in adults and in neonates. We hope that our simulations can help to predict an optimal patient-specific ventilation strategy to minimize lung damage and maximize oxygen uptake for these patients. Regarding turbulent flows, we will utilize our wall-modeled LES approach to investigate high-Reynolds-number fluid-structure interaction problems which have not been accessible with today's FSI solvers. The wall model will further be extended to a new framework including high-order discontinuous Galerkin methods.

References and Links

[1] Lung research group: http://www.lnm.mw.tum.de/ research/applications/biomedical-respiratory-system/

[2] W. A. Wall and M. W. Gee (2010). BACI: A parallel multi-physics simulation environment. Technical Report, Institute for Computational Mechanics. Technische Universität München.

[3] L. Yoshihara, C. J. Roth, and W. A. Wall (2015). Fluid-structure interaction including volumetric coupling with homogenized subdomains for modeling respiratory mechanics. International Journal for Numerical Methods in Biomedical Engineering, submitted.

[4] F. Verdugo, C. J. Roth, L. Yoshihara, and W. A. Wall (2016). Efficient solvers for coupled models in respiratory mechanics. International Journal for Numerical Methods in Biomedical Engineering, doi: 10.1002/cnm.2795.

[5] U. Rasthofer and V. Gravemeier (2013). Multifractal subgrid-scale modeling within a variational multiscale method for large-eddy simulation of turbulent flow. Journal of Computational Physics, 234, 79-107.

[6] U. Rasthofer, G. C. Burton, V. Gravemeier and W. A. Wall (2014). An algebraic variational multiscale-multigrid-multifractal method (AVM4) for large-eddy simulation of turbulent variable-density flow at low Mach number. International Journal for Numerical Methods in Fluids, 416-49.

[7] U. Rasthofer (2015). Computational multiscale methods for turbulent single and two-phase flows. Dissertation, Technische Universität München.

[8] B. Krank and W. A. Wall (2016). A new approach to wall modeling in LES of incompressible flow via function enrichment. Journal of Computational Physics, 316, 94-116.

Kinetics and thermodynamics of conformational changes upon protein association studied by molecular dynamics simulations

Research Institution
Lehrstuhl für Molekulardynamik, Physik-Department T38, TU München

Principal Investigator
Martin Zacharias

Researchers
Rainer Bomblies, Christina Frost, Florian Kandzia, Alexander Knips, Manuel Luitz, Giuseppe La Rosa, Maria Reif, Christina Schindler, Nadine Schwierz, Fabian Zeller

Project Partners
Aliaksei Krukau, Leibnitz Rechenzentrum, München

SuperMUC Project ID: pr84ko (Gauss Large Scale project)

Introduction

Proteins and peptides are essential components of basically all biological processes. During association protein molecules can undergo a variety of conformational changes essential for its function. Upon association such transitions can involve global domain motions for example of enzymes to switch between open accessible but inactive and closed active states. The enzyme adenylate kinase (ADK) consists of several semi-flexible segments or domains that rearrange from an open state to a closed enzymatically active state upon substrate binding (illustrated in Figure 1). The open conformation is easily accessible for the substrate and transitions to a closed state result in catalytic activation with the substrate largely buried in the protein molecule [1]. In order to elucidate how the lid-domain motion couples to binding of substrate and inhibitor molecules we employed Molecular Dynamics (MD) free energy simulations using the position of the two lid domains as independent reaction coordinates [1].

Besides of global motions association of proteins and peptides can also involve refolding of peptide segments. For example, certain peptides from stable aggregates in solution called also fibrils or amyloids [2]. Such amyloid fibrils play a major role in the several diseases of which the Alzheimer disease is probably the most well known example. Although the structure of several peptide amyloids has been elucidated by experimental methods the mechanism of formation is not very well understood. Using a similar methodology as for studying global motions in ADK we characterized the thermodynamics and kinetics of the propagation of Alzheimer $A\beta_{9\text{-}40}$ amyloid fibrils.

Results and Methods

Global domain motion in Adenylate Kinase
Besides of the two lid domains (termed ATP-lid and AMP-lid) the ADK enzyme contains a central core domain. In order to systematically investigate the lid domain mobility the distance of each lid to the central domain served as reaction coordinate. Using a two-dimensional (2D) umbrella sampling protocol coupled with replica exchanges it was possible to calculate a complete 2D-free energy landscape for the lid domain motions in response to different bound substrates and inhibitors. The coupling of umbrella sampling with replica exchange resulted in rapid convergence of the calculated free energy surface and showed very good scaling on a parallel super computer. The study is the first to investigate this free energy landscape for all possible substrate and inhibitor bound states of ADK and required >5 million SuperMUC cores hours. It demonstrated a dramatic dependence of the lid-motion on the binding state (see Figure 1). The simulation results helped to explain why for this type of enzyme two lid domains have evolved to efficiently bind and catalyze a reaction requiring two

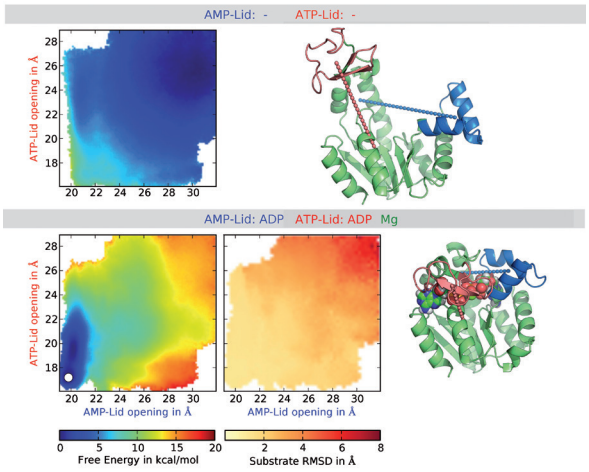

Figure 1: Upper panels indicate the calculated free energy landscape of the lid domain motion in adenylate kinase (ADK) versus the opening/closing of the two lid domains (left panel). The structure of the ADK enzyme in the apo form is illustrated as cartoon representation in the open state. The lid domains are colored red (ATP-lid) and blue (AMP-lid), respectively, and the two independent distance coordinates are shown as sticks. In the lower panels (from left to right) the free energy landscape for lid-domain motion in the presence of substrates (color-coded free energy), the fluctuation of the substrates (middle) and the bound ADK conformation (right panel) are indicated.

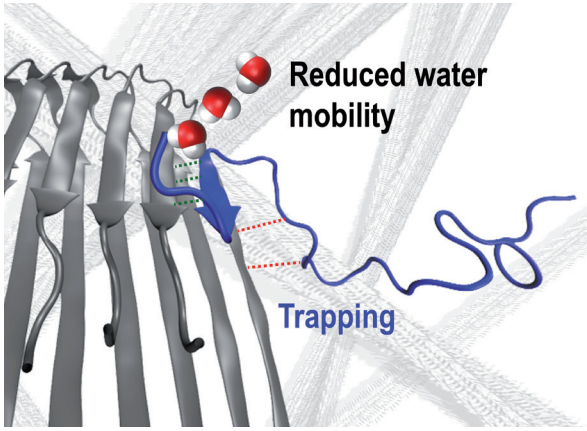

Figure 2: Illustration of an intermediate step during the association of an Alzheimer A$\beta_{9\text{-}40}$ amyloid (blue cartoon) to an already formed tip of an amyloid fibril structure. The state corresponds to a snapshot taken from extensive Molecular Dynamics simulations of the system. Transient non-native hydrogen bonds are shown in red and water molecules are included to indicate areas of reduced water mobility during the process.

substrate molecules and to elucidate the most likely order of binding events [1]. In collaboration with an experimental single-molecule group (Rief group, TUM Physics Department) the study also indicated the stabilization of a half open ADK state adopted in the presence of certain inhibitor molecules [3].

Propagation of Alzheimer A$\beta_{9\text{-}40}$ amyloid fibrils

In a second subproject we investigated the propagation step of the Alzheimer A$\beta_{9\text{-}40}$ amyloid fibril formation (illustrated in Figure 2). Using extensive Molecular Dynamics simulations combined with umbrella sampling we were able to study the thermodynamics and kinetics of fibril extension in good agreement with available experimental data [2]. The simulations identified the entropy change of water as a main driving force for association and interactions between the already formed fibril and a new monomer to overall disfavor association. The simulations also confirmed a dock/lock mechanism of association with a rapid docking phase characterized also by many non-native contacts followed by a slow locking step which finally leads to the native peptide association mode (Figure3). It was also possible to estimate the diffusion profile for the monomer approaching the fibril tip and to estimate the kinetics of fibril propagation. In order to achieve sufficient convergence of calculated free energies and other physical properties it was necessary to extend the simulations to several hundred nanoseconds per umbrella sampling window adding up to several microseconds total simulation time only feasible on the SuperMUC computer.

On-going Research / Outlook

The ADK enzyme is an excellent model system for studying the coupling between the enzymatic function of the protein and the global motion induced by substrate binding. Our coupled umbrella sampling and replica exchange methodology shows excellent scaling on parallel super computers. In future research we plan to study the influence of mutations on the energy landscape of domain motions in ADK. The simulation studies on amyloid formation and propagation will be extended to the study of inhibitors of amyloid propagation in collaboration with experimental groups at TUM. Both the systematic studies of global lid domain motions in the enzyme ADK but also the simulations of amyloid propagation were only possible by using the SuperMUC parallel computer facilities.

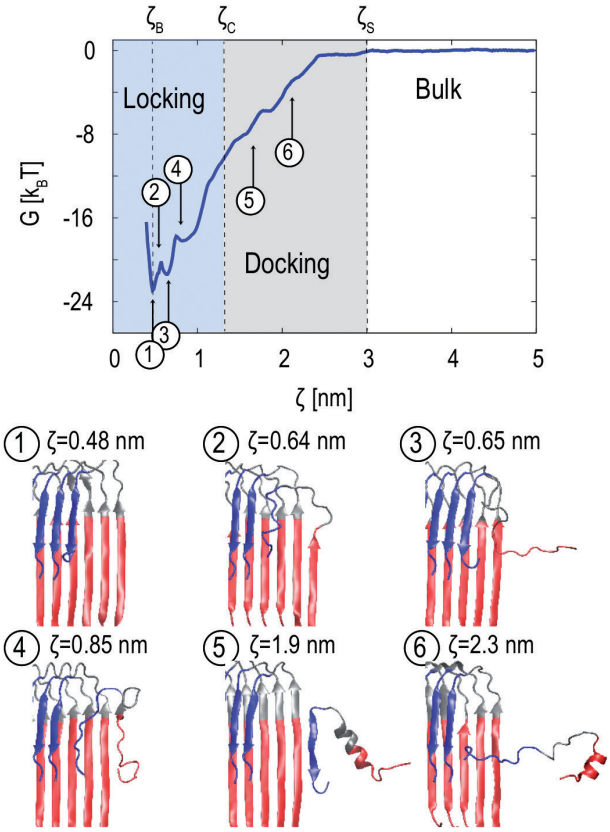

Figure 3: (Upper panel) Calculated free energy profile for the association/dissociation of a single A$\beta_{9\text{-}40}$ monomer to the end (tip) of an already formed amyloid segment. The reaction coordinate is the distance ζ between monomer and fibril fragment. The docking and locking phases are indicated and the snapshots shown in the lower panel are numbered according to the occurrence along the reaction coordinate. In the bulk regime the monomer adopts an unfolded conformation whereas in the locking regime it adopts a hairpin structure that associates with the tip of the amyloid fibril [2].

References and Links

[1] Zeller F, Zacharias M Substrate Binding Specifically Modulates Domain Arrangements in Adenylate Kinase. Biophys. J 109 (2015) 1978-85.

[2] Schwierz N, Frost CV, Geissler PL, Zacharias M. Dynamics of Seeded Aβ40-Fibril Growth from Atomistic Molecular Dynamics Simulations: Kinetic Trapping and Reduced Water Mobility in the Locking Step. J Am Chem Soc. 138 (2016) 527-39.

[3] Pelz B, Žoldák, Zeller F, Zacharias M, Rief M. Sub-nanometer enzyme mechanics probed by single-molecule force spectroscopy Nature Commun. 7 (2016) 10484.

Simulating transition to turbulence in hemodynamics of intracranial aneurysms at extreme scale

Research Institution
Simulation Techniques and Scientific Computing, University of Siegen, Germany

Principal Investigator
Sabine Roller

Researchers
Kartik Jain, Harald Klimach, Kent-Andre Mardal

Project Partners
Kent-Andre Mardal, Simula Research Laboratory and University of Oslo, Norway

SuperMUC Project ID: pr85mu

Introduction

An intracranial aneurysm (IA) is a pathological dilatation in a cerebral blood vessel characterized by a localized balloon-like enlargement of the blood vessel. Around 5% of the human population is reported to be suffering from un-ruptured IA. The rupture of an aneurysm, due to hemodynamic forces that act on its wall is an important cause of morbidity and mortality in the modern world. A major challenge that clinicians face today is the estimation of risk of rupture of a particular aneurysm, which consequently affects their decision for intervention.

Computational fluid dynamics (CFD) can, to a certain extent, provide patient specific information on the hemodynamic forces, namely the wall shear stresses, pressure and velocities that act on an aneurysm. Due to the low Reynolds number at which blood flows in a human artery, most CFD studies, however, have rested on the common premise of a laminar blood flow. With the help of computing resources on SuperMUC, we have been able to assess and quantify the occurrence of turbulence like phenomena in blood flow in various aneurysms. This phenomenon is exhibited by the presence of high-frequency fluctuations in velocity fields inside an aneurysm and resembles a flow regime that is not fully developed turbulence but rather transitional than laminar.

Results and Methods

To accurately assess the transitional hemodynamics in patient specific intracranial aneurysms, we employ the Lattice Boltzmann Method (LBM) for its computational efficiency, simple representation of complex geometries, and excellent scalability on modern supercomputers. Various LBM relaxation schemes are implemented in the *Musubi* [1] solver, which is part of the end-to-end parallel simulation framework, APES (adaptable poly-engineering simulator) [2].

MR images of patients are obtained from clinical collaborators, which are then segmented to obtain 3D models in the STL format. The models are discretized using an octree representation by our mesh generator *Seeder*. The meshes describe the geometrical layout for the flow computation in *Musubi*, and the results are then post-processed to obtain the relevant physical quantities. Spatial and temporal resolutions of 8µm and 1µs result in meshes that consist of up to ~ 1 billion cells and 1 million time steps per second. To gather sufficient statis-

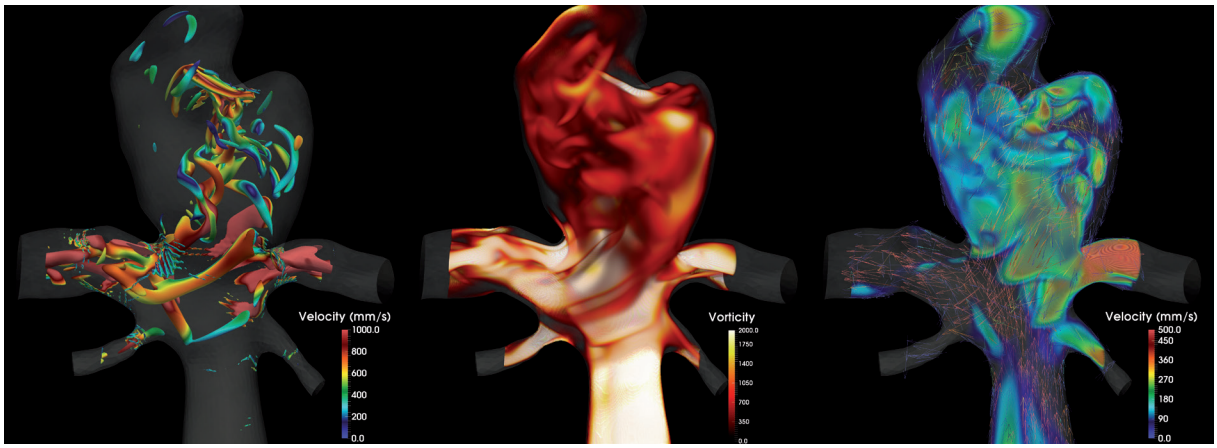

Figure 1: Turbulent characteristics in a patient specific intracranial aneurysm during peak systole. Velocity colored Q-isosurfaces (left), volume rendering of Vorticity magnitude (middle) and volume rendering of velocity magnitude overlaid by velocity vectors (right)

tics for the analysis of the transitional flow, we compute as many as 30 cardiac cycles. Simulations on SuperMUC are conducted using 32768 cores (4 islands) and computation of each cycle requires ~ 40 minutes. In a recent comprehensive study on space and time refinement [3], we have elucidated the role of high resolutions in a simulation of hemodynamics in intracranial aneurysms, and have explored the critical Reynolds number at which the flow in aneurysms leaves the laminar regime [4].

Figure 1 shows the velocity colored Q-isosurfaces (L), volume rendering of the vorticity magnitude (M) and velocity field overlapped with vectors (R) in a patient-specific aneurysm during peak systole. The morphology of the aneurysm at the bifurcation educes miniature vortices inside the aneurysm whereas the flow remains laminar in other parts of the vasculature. This shows that the manifestation of an aneurysm itself can act as an initiator of fluctuations in the flow. The velocity oscillates at around 10KHz in the center of the dome with a rapidly varying frequency during the cardiac cycle. During the cardiac cycle, the flow re-laminarizes during acceleration, fluctuations commence during peak systole and are enhanced by the large decelerative forces during deceleration [4].

Efforts for the validation of our results, particularly in the context of transitional physiological flows have included comprehensive comparisons against previously published data while extending those studies for quantification of Kolmogorov micro-scales and exploring the appropriate conditions when the flow in a complex anatomical geometry may or may not transition to turbulent like regime [5].

With the help of computing resources on SuperMUC, we were able to participate in the international aneurysm CFD-Challenge, an event in which a team of Neurosurgeons gave MR images of 5 aneurysms to the participants, and the task was to assess the rupture status of the aneurysms. No information about the aneurysms was provided, as the target was to explore what research teams would prescribe as boundary conditions. This event brought together results from 28 teams across the globe, and the results are in the process of analysis and publication by the organizers.

Our LBM framework *Musubi* is continuously improved and recent efforts include improvements for vector architectures and optimization of computing kernels. Figure 2 demonstrates the achieved performance of *Musubi* in TFLOP/s against the theoretical peak performance on four different computing architectures including the GCS systems SuperMUC, Hornet and Juqueen. The fourth (Kabuki) is a small NEC SX-ACE vector system installed at HLRS. A fixed problem size of 88 cells is chosen for this comparison of strong scalability, and a log-log scale is used to illuminate the wide scaling range. *Musubi* achieves nearly 20% of the peak performance on a single node of Kabuki while it achieves nearly 4% on other three systems. A super-linear speedup can be observed for the three scalar systems on large number of nodes as the problem per process fits into the cache of the processors.

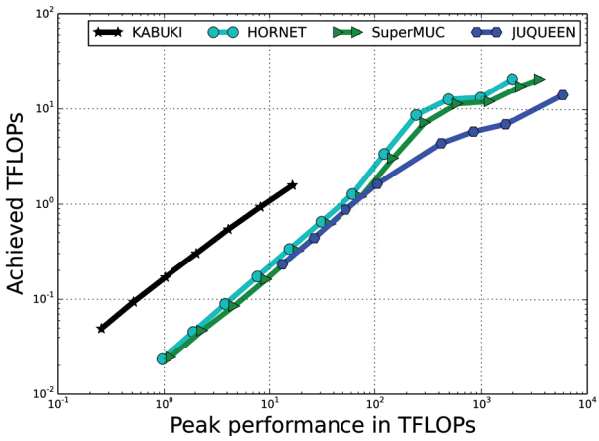

Figure 2: Performance comparison of *Musubi* on different computing architectures.

On-going Research / Outlook

Simulations of physiological flows conducted on SuperMUC have revealed the occurrence of turbulent like phenomena in aneurysms. Our on-going research efforts include investigations of the flow regime in aneurysms with different morphology, size, and location in the vasculature. The goal is to increase the tenacity of the presence of a transitional flow regime in aneurysms by exploring various factors. Moreover, we are investigating the impact of surrounding vasculature on the intra-aneurysmal flow field, which is usually truncated from the simulation domain. Such simulations are expected to increase the computational effort by two times, as the meshes will consist of up to 2 billion cells. Validation of physiological flow computations, in the absence of experimental data remains a challenge, and we are comparing the velocity fields obtained from simulations against that from MR-imaging to quantify the discrepancies from our results against the in vivo situation.

Musubi has been investigated on the phase 2 of SuperMUC, and our on-going work will include improvements in the sustained performance and investigations to improve the utilization of vector operations.

References and Links

[1] Manuel Hasert et al. (2014). Journal of Computational Science, 5(5), 784-794.
[2] Sabine Roller et al. "An adaptable simulation framework based on a linearized octree." High Performance Computing on Vector Systems 2011. Springer Berlin Heidelberg, 2012. 93-105.
[3] Kartik Jain, Sabine Roller, and Kent-Andre Mardal. Transitional flow in intracranial aneurysms – A space and time refinement study below the Kolmogorov scales using Lattice Boltzmann Method. Computers & Fluids, 127:36–46, 2016.
[4] Kartik Jain and Kent-Andre Mardal. Exploring the critical Reynolds number for transition in intracranial aneurysms - highly resolved simulations below Kolmogorov scales. 560 – 563. 2015 Computational and Mathematical Biomedical Engineering.
[5] Kartik Jain. (2016) Direct numerical simulation of transitional pulsatile stenotic flow using Lattice Boltzmann Method. PeerJ PrePrints 4:e1548v3

Modulation of voltage-gated potassium channel Kv1.2 by PIP_2 lipids

Research Institution
Universite de Lorraine, Boulevard des Aiguillettes, 54506 Vandoeuvre-les-Nancy, France

Principal Investigator
Mounir Tarek

Researchers
Marina A. Kasimova, Lucie Delemotte

Project Partners
–

SuperMUC Project ID: pr86ba (PRACE project)

Introduction

Voltage-gated potassium channel Kv1.2 is a transmembrane protein that enables the passive flow of potassium ions across a plasma membrane when the latter is depolarized. It consists of four peripheral voltage sensor domains, responding to the applied voltage, and a central pore domain that encompasses a hydrophilic path for passing ions [1]. Each voltage sensor domain is composed of four transmembrane helices, called from S1 to S4. When the electric field is applied, one of these helices, S4, moves across the membrane, promoting the transition between the most activated state of the channel (S4 is "UP") and the most resting state (S4 is "DOWN") through several intermediate states (see Figure 1A).

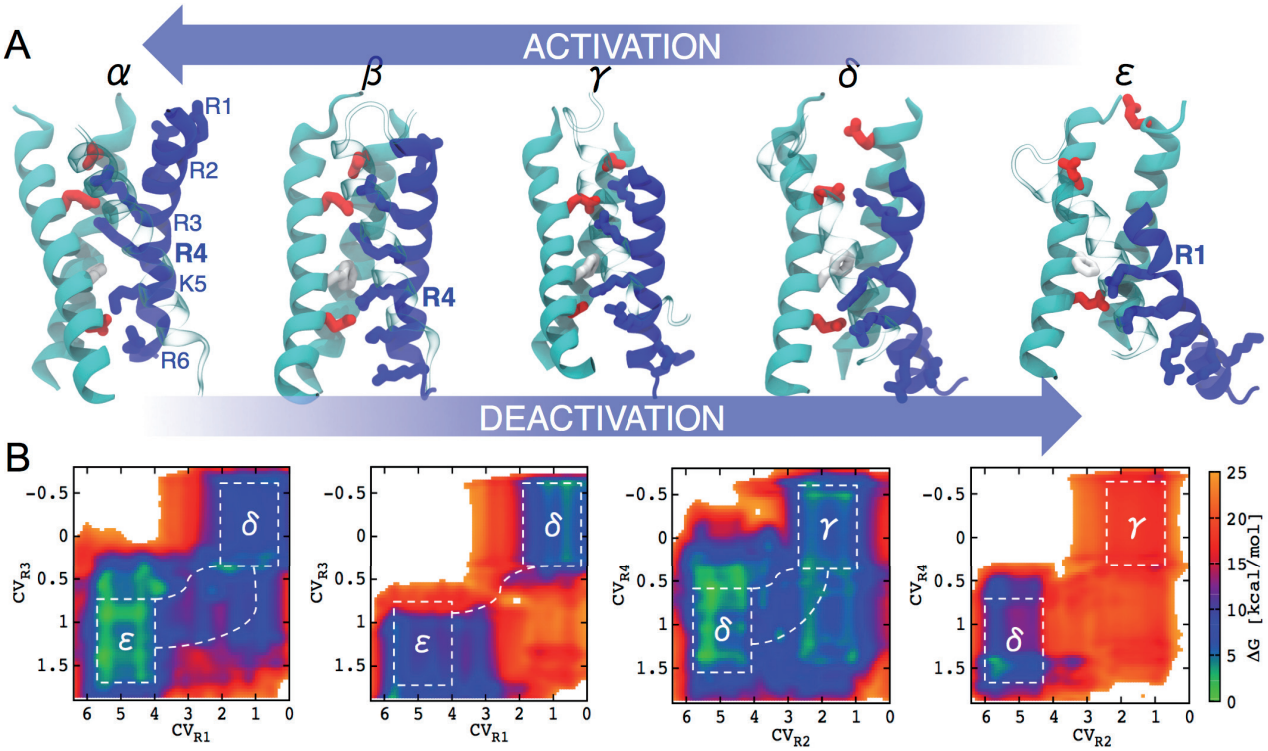

Figure 1: Conformational rearrangements of the Kv1.2 voltage sensor taking place during the channel activation. A. The activated to resting states transition of the Kv1.2 voltage sensor. The entire transition is split into 4 sub-transitions: ε-δ, δ-γ, γ-β and β-α. The most activated and most resting states are α and ε respectively; β, γ and δ are the intermediate states. The transition from the most activated to most resting states, from ε to α, is called activation; the opposite transition, from α to ε, is called deactivation. During the channel activation/deactivation, the S4 helix of the voltage sensor (shown in blue) moves across the membrane; in parallel, the positive residues of S4 (R1, R2, R3, R4, K5 and R6, shown as blue sticks) change their binding sites (negatively charged residues, shown as red sticks). In the four states of the voltage sensor, β, γ, δ and ε, these residues constitute a binding site for the lipids of the lower bilayer leaflet including PIP_2. B. Free energy surfaces underlying the two sub-transitions, ε-δ (first two plots) and δ-γ (last two plots), in the presence (right) and absence of PIP_2 (left). In the presence of PIP_2, the γ state is less stable compared to the δ state, while in the absence of this lipid, these two states have similar free energies. Also, PIP_2 affects the free energy barriers of the two sub-transitions: in the presence of this lipid, the free energy barriers are higher.

Phosphatidylinositol-4,5-bisphosphate (PIP_2) is a minor lipid of the inner plasma membrane leaflet. This highly negatively charged lipid modulates the Kv1.2 functioning: its application to the channel results in the delay of the motion of the voltage sensors [2,3]. So far, the mechanism underlying this effect remained unknown. In our study, we attempt to shed light on it using molecular dynamics simulations and metadynamics.

Results and Methods

Using unconstrained MD simulations, we have identified a potential PIP_2 binding site in Kv1.2. In this site, present only in the resting state of the channel, PIP_2 interacts with the lower positively charged residues of S4. Based on this finding, we hypothesized that PIP_2 stabilizes the resting state of the channel through the interactions with S4, which results in the effect of the delayed motion of the voltage sensors observed experimentally [2,3].

To verify our hypothesis we decided to estimate the free energy surface of the activated to resting states transition of the Kv1.2 voltage sensor in the presence and in the absence of PIP_2. To do so, we split the entire transition into 4 sub-transitions and explored the free energy surfaces of each sub-transition individually. To enhance sampling, we used metadynamics with "multiple walkers" scheme implemented in gromacs and plumed. Based on our benchmark estimations, we selected the number of walkers to be 42 and the number of processors involved by each walker to be 96. Using this setup, we were able to simulate ~ 0.8 us per day. We also chose high frequency for trajectory writing in order to further perform detailed post-analysis. Hence, the data generated during the project took ~ 15 TB of space. In total, 42.000.000 cpu-hours were used to generate ~ 350 us of molecular dynamics. To note, our system contained ~ 100.000 atoms and was 9x8x15 nm3 in its size.

Based on the obtained data, we calculated the free energy surfaces of all the sub-transitions along the activation path in the presence and absence of PIP_2 (see Figure 1B for the first two sub-transitions). We found that PIP_2 affects the relative stability of the voltage sensor states: in the presence of PIP_2, the so-called γ state is less stable compared to the more resting δ state, while in the absence of this lipid, the two states have similar free energies. We also found that PIP_2 affects the free energy barriers separating the states, indicating that it modulates the rate of the sub-transitions. The results of our work have been already published in the following manuscript [4] and the other manuscript is going to be submitted shortly.

On-going Research / Outlook

In the ongoing research, we will perform more detailed analysis of the generated trajectories in order to explore the role of each individual residue of the voltage sensor domain, first, in the activation process and, second, in the modulation by PIP_2.

References and Links

[1] Chen, X., Wang, Q., Ni, F., and Ma, J. Structure of the full-length Shaker potassium channel Kv1.2 by normal-mode-based X-ray crystallographic refinement. Proc Natl Acad Sci U S A, 107(25). 11352-11357.

[2] Abderemane-Ali, F., Es-Salah-Lamoureux, Z., Delemotte, L., Kasimova, M.A., Labro, A.J., Snyders, D.J., Fedida, D., Tarek, M., Baro, I., and Loussouarn, G. Dual effect phosphatidylinositol-(4,5)-bisphosphate PIP_2 on Shaker K+ channels. J Biol Chem, 287(43). 36158-36167.

[3] Rodriguez-Menchaca, A.A., Adney, S.K., Tang, Q.Y., Meng, X.Y., Rosenhouse-Dansker, A., Cui, M., Logothetis, D.E. PIP_2 controls voltage-sensor movement and pore opening of Kv channels through the S4-S5 linker. Proc Natl Acad Sci U S A, 109(36). E2399-2408.

[4] Delemotte, L., Kasimova, M.A., Klein, M.L., Tarek, M., and Carnevale, V. Free energy landscape of ion-channel voltage-sensor-domain activation. Proc Natl Acad U S A, 112(1). 124-129.

The key is in the movements: allosteric activation of an oncologically relevant kinase

Research Institution
Spanish National Cancer Research Center (CNIO), Madrid, Spain. University College London (UCL), UK.

Principal Investigator
Francesco L. Gervasio

Researchers
Nicole Dölker, Ludovico Sutto, Giorgio Saladino, Silvia Lovera

Project Partners
Research Center for Molecular Medicine of the Austrian Academy of Sciences, Ce-M-M

SuperMUC Project ID: pr86ga (PRACE project)

Introduction

Protein kinases are the key enzymes that control most cellular activities. A kinase that fails to work properly can cause severe damage to the organism, causing diverse diseases including cancer. It is therefore highly desirable to develop drugs that modulate the activity of specific protein kinases. However, this task is complicated by the fact that all protein kinases share a common architecture, so a drug tailored to one kinase is likely to affect others as well, causing serious side effects. The Abelson tyrosine kinase (Abl) is of special interest because of its importance as an anti-cancer drug target. A genetic defect affecting the gene that encodes Abl causes chronic myeloid leukemia (CML), and Abl inhibitors are the only known anti-CML drugs. The treatment is very effective, but a significant proportion of patients relapses due to drug resistance-causing mutations in Abl.

During activation, the catalytic domain of a kinase undergoes a number of distinct conformational changes, which are usually induced by phosphorylation. Furthermore, most kinases also require interactions with other domains or proteins for full activation. Activation of Abl involves a complete rearrangement of the domains of the protein, and thus the interface between the catalytic and the modulator (SH2) domains seems to be a promising target for new anti-CML drugs [2]. In our project, we set out to shed light on the effect of the SH2 domain on the intrinsic motions of the catalytic domain.

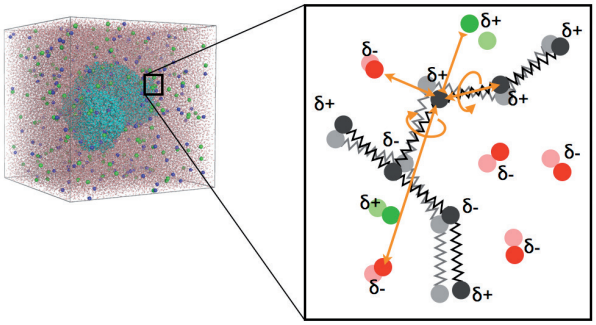

Figure 1: Schematic representation of MD simulations.

Methods

We used Molecular Dynamics (MD) simulations, in which all atoms of the system are represented by rigid spheres, and interactions between them are described by simplified models. Integration of Newton's equations of motion then yields a trajectory describing the movements of the system. The integration is carried out in tiny time steps in the range of femtoseconds, which have to be repeated millions of times. The simulation of meaningful protein motions therefore requires massive computer power.

Classical MD simulations reproduce protein dynamics in the nano to microsecond timescale, whereas the quantitative description of large-scale conformational changes requires the use of so-called "enhanced sampling methods".

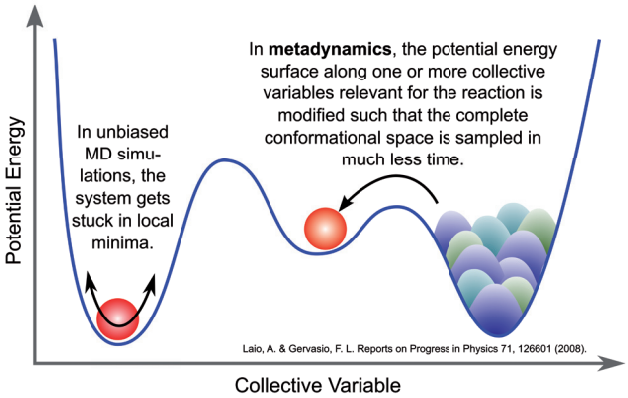

Figure 2: Enhanced sampling along a collective variable.

We used metadynamics with parallel tempering (PT-metaD) to reconstruct the free energy surface (FES) of the conformational change of a conserved Asp-Phe-Gly (DFG) motif at the active site of the catalytic domain. PT-metaD is an extremely efficient and intrinsically parallel method for the calculation of the free energy as a function of one or more variables. variables (CVs).

All calculations were carried out with the GROMACS4.5 molecular dynamics package and the PLUMED-plugin for free energy calculations. GROMACS is a standard

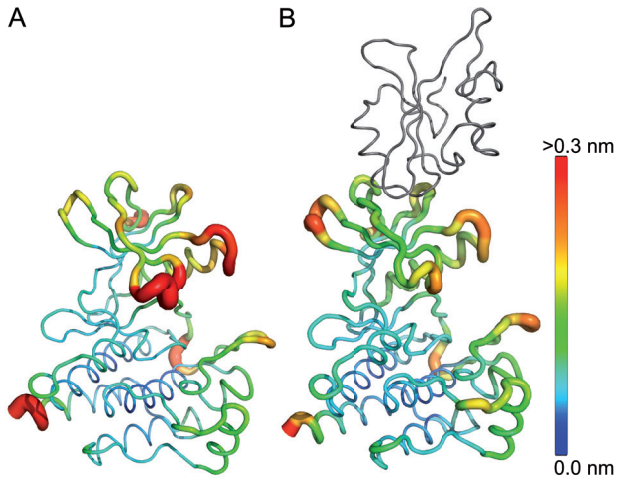

Figure 3: The regulatory SH2 domain (in grey) stabilizes catalytically important motifs of the kinase domain of Abl. The colouring and thickness of the tube-loke cartoon represent the differential flexibility of the residues of the kinase.

package for parallel molecular dynamics calculations of large biomolecular systems, which has been extensively tested and used on high performance computer systems, including a number of Tier-0 resources.

The project was awarded 25,000,000 core hours, which we used with an average of 4,000 CPU/job.

Results

Our simulations indicated that the SH2 domain modifies the motions of the catalytic domain of Abl in two ways. On the one hand, it stabilizes and positions some elements of the enzyme in an optimal arrangement, and on the other hand it channels thermal fluctuations into global movements involved in the catalytic process.

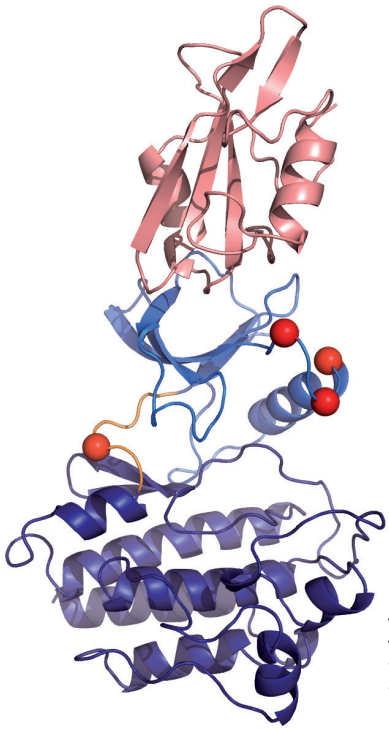

Figure 4: Residues crucial for the transmission of the activation from the SH2 domain to the active site (red balls).

In order to validate our model, in collaboration with the group of Prof. Giulio Superti-Furga at CeMM in Vienna, we designed a set of mutants we predicted to change the response of the Abl kinase to its regulatory SH2 domain and tested them experimentally for their catalytic activity. A number of the proposed mutants did in fact modulate the effect of the SH2 domain, thus supporting our computational results.

We expect that in the future our findings may open the door to the development of anti-CML drugs with a new mechanism of action and therefore effective for the treatment of patients who have become resistant to the current treatments.

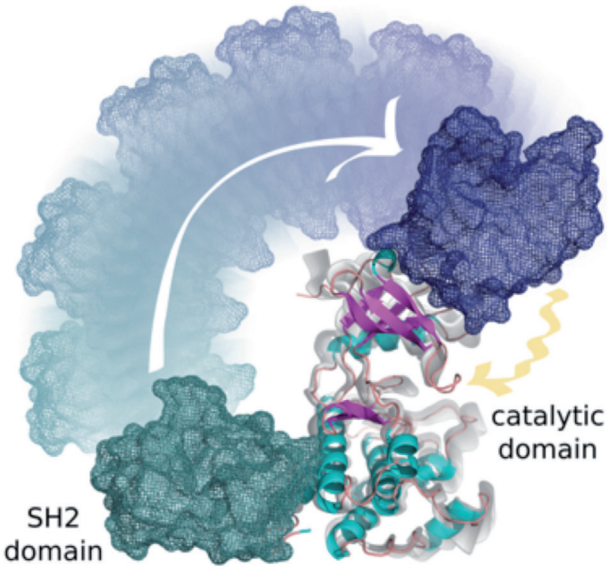

Figure 5: Abl kinase activation by a large-scale domain rearrangement, which leads to changes of the global alnd local dynamics of the catalytic domain.

On-going Research / Outlook

We were awarded computer time on MareNostrum at the Barcelona Supercomputing Center for a follow-up project, in which we carried out MD simulations of the mutants found to modulate the response to the SH2 domain during the first phase of the project. In this second phase, we could explain in detail how the mutants interrupted the interplay with the SH2 domain. The results of phases 1 and 2 of the simulations, together with the experimental data, gave rise to a publication [3]. Similar approaches are currently applied to other kinases with medical relevance, in order to elucidate their activation mechanisms and obtain data for drug design studies [4].

References and Links

[1] https://www.ucl.ac.uk/chemistry/research/ group_pages/prot_dynamics/research
[2] Filippakopoulos, P. et al. Structural coupling of SH2-kinase domains links Fes and Abl substrate recognition and kinase activation. Cell 134, 793–803 (2008).
[3] Lovera, S. et al. The different flexibility of c-Src and c-Abl kinases regulates the accessibility of a druggable inactive conformation. J Am Chem Soc 134, 2496–2499 (2012)
[4] Towards a Molecular Understanding of the Link between Imatinib Resistance and Kinase Conformational Dynamics. 11, e1004578 (2015).

From Biomolecular Structures to Thermodynamic Ensembles: Cellular Logistics Controlled by Disordered FG-Nucleoporins

Research Institution
Theoretical and Computational Biophysics, Max Planck Institute for Biophysical Chemistry

Principal Investigator
Helmut Grubmüller

Researchers
Sarah Rauscher, Carsten Kutzner

Project Partners
–

SuperMUC Project ID: pr86se, pr84ma (both Gauss Large Scale projects)

Introduction

The nuclear pore complex facilitates transport of macromolecules into and out of the nucleus
In eukaryotic cells, the nucleus is separated from the cytoplasm by the nuclear envelope, which is a double lipid bilayer. The maintenance of the integrity of this barrier is crucial to cellular viability, and therefore the passage of large macromolecules across the nuclear envelope is a tightly regulated process. The only passageway through which molecules larger than 40 kDa may enter and exit the nucleus is provided by the nuclear pore complex (NPC, Figure 1). Ions and small molecules pass through the NPC via passive diffusion. However, larger cargo (like protein and RNA) must be bound to nuclear transport receptors in order to pass through the NPC. The NPC spans the nuclear envelope and gates bi-directional selective transport. Because it must accommodate the passage of large molecules, such as the subunits of the ribosome, the NPC is an immense macromolecular machine. For example, the NPC of yeast contains 456 proteins and has a mass of approximately 50 MDa.

The constituent proteins of the NPC are called nucleoporins. They are classified into two types: structured and disordered. There is significantly more experimental data available for structured nucleoporins, including their locations within the NPC. The crucial selectivity, however, is conveyed by the disordered nucleoporins, for which, in contrast, only very little structural information is available. They are known to occupy the central region of the NPC, and are thought to form a dense, mesh-like structure, represented schematically in Figure 1. However, the precise arrangement and conformations of the nucleoporins in this region are not known, and is currently the subject of controversy. Importantly, the mechanism by which this "mesh" distinguishes between different cargo molecules is unclear. To obtain more detailed structural information on the disordered nucleoporins is, therefore, crucial.

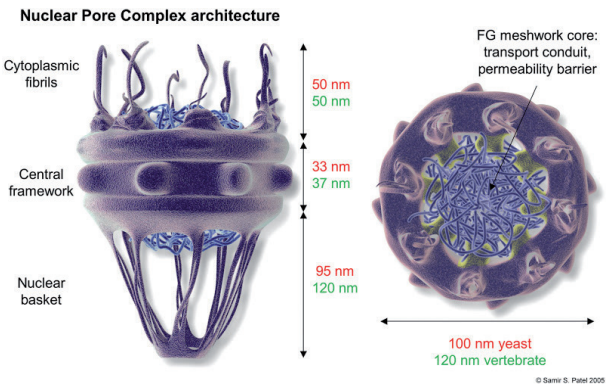

Figure 1. The Architecture of the Nuclear Pore Complex (image by Samir Patel, adapted from [1]). Left: side view, right: top view. The dimensions of both the yeast and vertebrate NPC are shown in red and green, respectively.

FG-nucleoporins are the gatekeepers responsible for selective transport through the nuclear pore complex
The disordered nucleoporins contain many phenylalanine-glycine (FG) repeat motifs, and are, therefore, called FG-nucleoporins (FG-nups). The FG-nups have been proposed to act as gatekeepers to the nucleus. It has been hypothesized that their thermal motion creates an entropic barrier to the passage of large molecules. The aggregation of FG-nups has been studied extensively *in vitro*. Polypeptide sequences derived from the FG-nucleoporin Nsp1p are able to self-assemble to form hydrogels. Remarkably, these hydrogels, which consist of only disordered nucleoporins, have been found to reproduce the selectivity of intact nuclear pore complexes. On the basis of these observations, Frey *et al* put forth a model for the selectivity region of the NPC which resembles a hydrogel (the "selective phase model").[2] In their model, the FG-nups form a mesh that is primarily held together by weak hydrophobic interactions to form a barrier to the diffusion of larger macromolecules. However, at present, the experimental data available for the structure of FG-nups cannot provide a high resolution description of their self-aggregation and structural properties. It is clear that we need to go beyond schematic illustrations of the interior of the NPC, if we are to understand how the NPC achieves such a high selectivity. The goal of our simulations is to obtain information at the atomistic level concerning the structure and aggregation properties of disordered FG-nups.

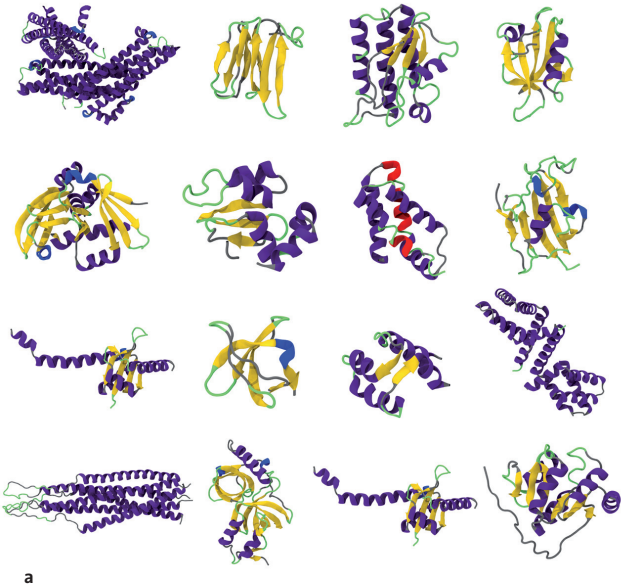

Figure 2a: 16 structures of globular proteins selected from the Protein Data Bank (PDB). The native state of these proteins is a folded, globular structure with well-defined secondary structure elements (purple: alpha-helix; yellow: beta-sheet; green: turn; grey: coil).

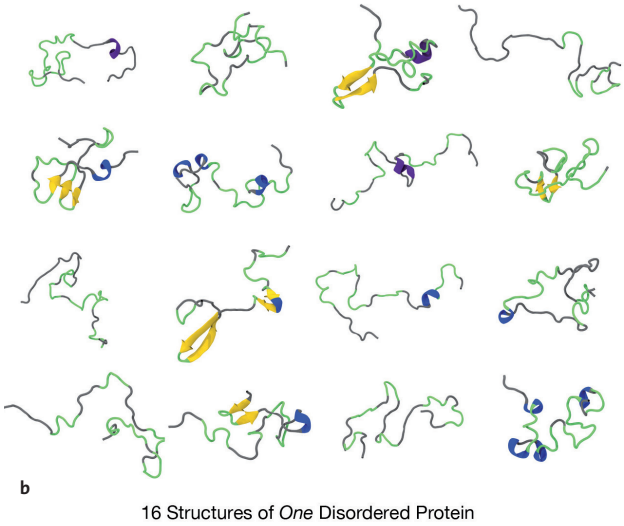

Figure 2b: In contrast to the native state of globular proteins (Figure 2a), the native state of a disordered protein consists of a collection of many possible structures (a "structural ensemble"). Here, we show a collection of 16 randomly chosen structures from the ensemble of a disordered FG-nucleoporin.

FG-nucleoporins are intrinsically disordered proteins
The FG-nups are disordered in solution, which means that they populate many conformational states at equilibrium. Thus, their structure cannot be described as a single, well-defined structure like globular, folded proteins (Figure 2a), but must instead be described as a structural *ensemble* (Figure 2b).

FG-nups are a prototypic example of the functional role of protein disorder in biological systems, and beyond their particular function are key model systems for disordered proteins. Intrinsically disordered proteins (IDPs) fulfill important biological roles including cell signaling and cell cycle regulation. Disordered proteins are highly abundant in all kingdoms of life: more than one-third of eukaryotic proteins and more than three-quarters of proteins linked to cancer are predicted to contain disordered regions. Despite their importance as potential drug targets, disordered proteins are poorly understood relative to the wealth of structural information available for folded proteins. Disordered proteins are notoriously difficult to study using experimental approaches due to their tendency to aggregate. Moreover, and on a very fundamental level, their description poses formidable challenges. Disordered proteins and disordered regions of proteins have many energetically-accessible conformational states. Thus, even with hundreds of experimental observables, it is not possible to obtain an unambiguously-determined ensemble of conformations. The structural characterization of disordered proteins is an inherently under-determined problem: a small number of restraints are insufficient to uniquely define the conformations of a system with thousands of degrees of freedom. Molecular simulations, with their empirical force fields, can offer the additional information required to obtain conformational ensembles for disordered states of proteins. However, these simulations must contend with a massive sampling problem, which we have successfully achieved using the SuperMUC.

Results and Methods

Numerical Methods and Algorithms
All simulations described here were performed using the GROMACS 4.6 [3] molecular dynamics package. GROMACS is one of the fastest molecular dynamics engines available today and is highly optimized for both single core performance and scalability. Recent benchmarks performed on the SuperMUC cluster show near perfect scaling for large (~12 million atoms) systems up to 4096 cores, and quite good scaling for smaller (~2 million atoms) systems comparable in size to the systems studied here up to 512 cores [4], Figure 3. Additionally, we have used generalized ensemble methods (the replica exchange algorithm, already implemented within GROMACS) in which the simulated system is allowed to perform a random walk in temperature space. These methods are designed to greatly enhance conformational sampling and to achieve near perfect scaling regardless of system size, as very minimal information (only temperature) must be communicated between the individual replicas.

An ensemble simulation on N nodes will then have high parallel efficiency as each of its R individual replicas running on N/R nodes. Analysis and processing of data generated by our simulations require large, long-term storage. Configuration snapshots are typically saved at 50 ps time intervals and simulation lengths generally span the microsecond to millisecond range. The simulations described here thus generate a significant volume of data consisting of the three-dimensional atomic positions of each saved system snapshot.

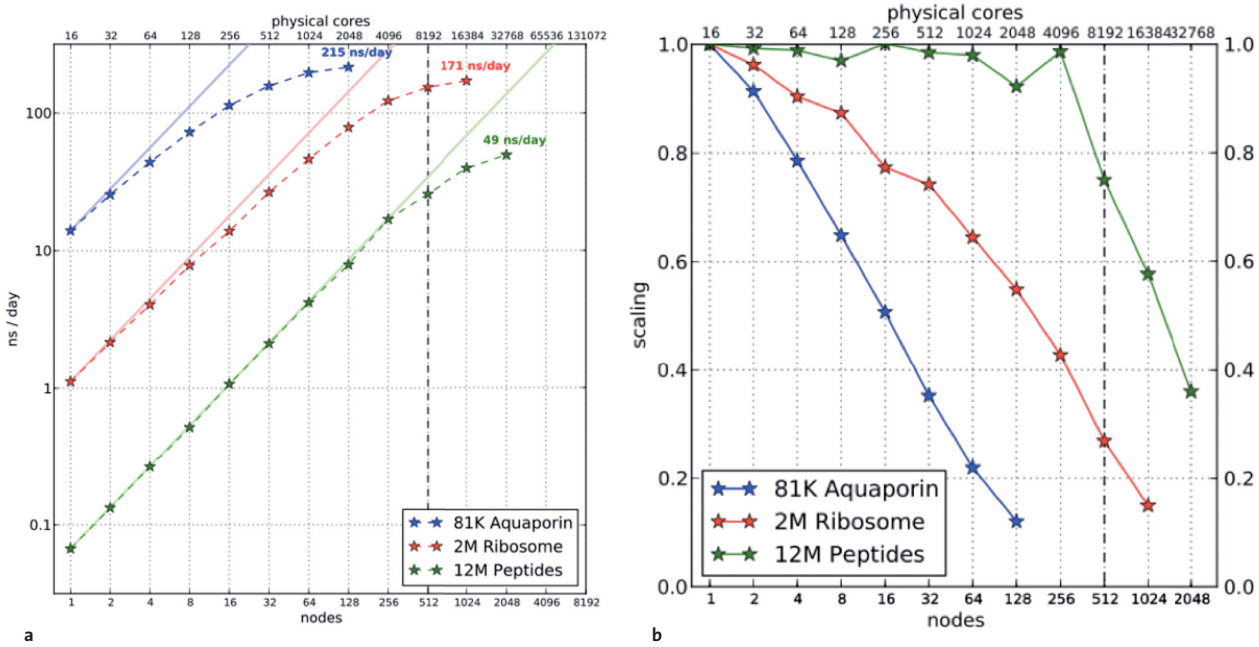

Figure 3. Scaling of GROMACS on SuperMUC for different system sizes. a, Performance of GROMACS 4.6 on the SuperMUC cluster, using typical benchmark systems of 81k atoms (blue, Aquaporin), 2 M atoms (red, the ribosome) and 12 M atoms (green, large aggregate of peptides) in size. Lines show perfect scaling. b, Scaling of the systems in a. (Figure taken from Kutzner et. al. 2014 [4])

Development and Benchmarking of the Simulation Approach

Molecular simulations are increasingly being used to obtain conformational ensembles of IDPs. However, there is currently no consensus on the accuracy of these ensembles, or the suitability of modern empirical force fields for this purpose. We have therefore carried out a rigorous evaluation of the accuracy of *de novo* intrinsically disordered protein (IDP) ensembles. [5] In our study, we assessed the accuracy of IDP ensembles obtained using state-of-the-art force fields by comparing to NMR and small angle x-ray scattering data. Carrying out such a comparative study presents a huge computational challenge, which was only possible with the large compute time allocation on the SUPERMUC; the accumulated simulation time for all simulations reported in this study is nearly one millisecond.

Our comparison of force fields also led to several unexpected results. First, the extent of the difference between ensembles is unexpectedly large, spanning the complete range from globule-like to highly expanded (Figure 4). The key finding of our joint experimental-computational study is that one single force field, CHARMM 22*, stands out in that it is consistent with small angle x-ray scattering and NMR data within experimental error. Thus, having obtained an accurate IDP ensemble, the potential long-term impact of this work extends far beyond an assessment of force field accuracy. This work has immediate application to the study of IDPs, which are currently of widespread interest in the fields of biochemistry, biophysics and computational chemistry.

Related Simulation Studies on Large Biomolecular Systems

We have recently carried out related work on another large biomolecular system, the ribosome. Our department has obtained atomic structures of ribosomes in 13 intermediate states of tRNA translocation by combining X-ray crystallography and cryo-EM data. [6] Translocation of tRNAs is coupled to large-scale rotation of the two ribosomal subunits. Extensive all-atom explicit-solvent MD simulations of the intermediate structures have shown that the transition rates for intersubunit rotation are surprisingly high, and that the movement of tRNAs is rate-limiting for the overall translocation process. A flexible element of the ribosome, the L1 stalk, was shown to pull on the tRNAs, thereby helping to overcome the barriers hindering tRNA motion. A recent follow-up study showed how the dynamic contact network between the two subunits adapts to keep the inter-subunit binding free energy at a constant level almost independent of the rotation angle. [7] This flatness of the free energy landscape rationalizes the rapid rotation on microsec-

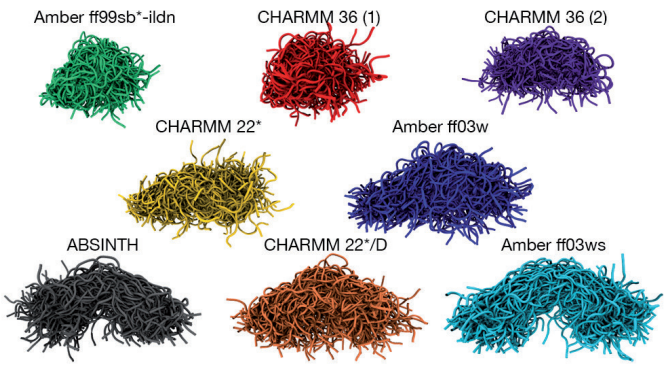

Figure 4. Ensembles of an intrinsically disordered protein (IDP) in eight different force fields (taken from Rauscher et al., 2015 [5]).

ond timescales. With state of the art microsecond-long simulations of the ribosome approaching experimental timescales, it was shown that the resolution of new emerging cryo-EM techniques is limited by the local atomic fluctuations. [8] With this knowledge we can now extract information about the dynamics of large molecular complexes from cryo-EM experiments. We have also carried out a joint experimental-computational study on the nuclear transport receptor that interacts with the FG-nucleoporins (importin β), which revealed a remarkable sensitivity of this protein to solution conditions. [9]

On-going Research

We are currently using our allocation on the SuperMUC to perform large-scale simulations of a variety of intrinsically disordered proteins that are implicated in human disease. We also are continuing to carry out comparisons between our simulated ensembles and primary experimental data to assess the accuracy of our approach, as well as to improve existing all-atom force fields.

References and Links

[1] S.S. Patel, B.J. Belmont, J.M. Sante, and M.F. Rexach. 2007. Natively Unfolded Nucleoporins Gate Protein Diffusion across the Nuclear Pore Complex. Cell. 129, 83–96.

[2] S. Frey, R.P. Richter, and D. Görlich. 2006. FG-Rich Repeats of Nuclear Pore Proteins Form a Three-Dimensional Meshwork with Hydrogel-Like Properties. Science. 314, 815–817.

[3] S. Pronk et al. 2013. GROMACS 4.5: a high-throughput and highly parallel open source molecular simulation toolkit. Bioinformatics 29, 845–854.

[4] C. Kutzner, R. Apostolov, B. Hess, and H. Grubmülller. 2014. Scaling of the GROMACS 4.6 molecular dynamics code on SuperMUC. In Parallel Computing: Accelerating Computational Science and Engineering (CSE); Bader, M. et al. ISO Press: Clifton, VA.; pp. 722– 730.

[5] S. Rauscher, V. Gapsys, M. J. Gajda, M. Zweckstetter, B. L. de Groot, and H. Grubmülller. 2015. Structural Ensembles of Intrinsically Disordered Proteins Depend Strongly on Force Field: A Comparison to Experiment. J. Chem. Theory Comput. 11, 5513–5524.

[6] L. V. Bock, C. Blau, G. F. Schröder, I. I. Davydov, N. Fischer, H. Stark, M. V. Rodnina, A.C. Vaiana, and H. Grubmülller 2013. Energy barriers and driving forces in tRNA translocation through the ribosome. Nat. Struct. Mol. Biol. 20, 1390-1396.

[7] L. V. Bock, C. Blau, A. C. Vaiana, and H. Grubmülller. 2015. Dynamic contact network between ribosomal subunits enables rapid large-scale rotation during spontaneous translocation. Nucleic Acids Res. 43, 6747–6760.

[8] N. Fischer, P. Neumann, A. L. Konevega, L. V. Bock, R. Ficner, M. V. Rodnina, H. Stark. 2015. Structure of the E. coli ribosome–EF-Tu complex at <3 Å resolution by Cs-corrected cryo-EM. Nature. 520, 567-570.

[9] K. Halder, N. Dölker, Q. Van, I. Gregor, A. Dickmanns, I. Baade, R. H. Kehlenbach, R. Ficner, J. Enderlein, H. Grubmülller and H. Neumann. 2015. MD Simulations and FRET Reveal an Environment-Sensitive Conformational Plasticity of Importin-β. Biophys. J. 109, 277–286.

Rapid and accurate calculation of ligand-protein binding free energies

Research Institution
Ludwig-Maximillians-Universität, Munich (LMU)

Principal Investigator
Dieter Kranzlmüller (LRZ/LMU, Munich)

Researchers
Peter Coveney, Shunzhou Wan, Agastya Bhati, Serge Jovanovic (UCL, London)

Project Partners
Shantenu Jha, Department of Computer Engineering, Rutgers University, USA

SuperMUC Project ID: pr87be

Introduction

Our present project [1], originally titled "computing clinically relevant binding free energies of Ableson kinase inhibitors", is to predict the strength of macromolecular binding free energies using computationally based molecular modelling. Rapid and accurate calculation of binding free energies is of major concern in drug discovery and personalized medicine. To perform modelling and calculation with optimal efficiency, we have developed the Binding Affinity Calculator (BAC), a highly automated molecular simulation based free energy calculation workflow tool.

Over the past few years, we have uncovered and developed two new ways of calculating the free energy of binding of ligands to proteins. One is ESMACS (enhanced sampling of molecular dynamics with approximation of continuum solvent); the other is TIES (thermodynamic integration with enhanced sampling). Emphasis has been placed on speed and reliability. These methods have considerable potential for uptake in the pharmaceutical industry; they are also likely to play a key role in the more forward-looking field of personalized medicine, underpinned by genomic analysis, for drug selection in a clinical context (Fig. 1).

The underlying computational method is based on classical molecular dynamics (MD). For purposes of reliability, ensembles of replica MD calculations are performed and we have found that ca 25 of these are required per MD simulation in order to guarantee reproducibility of predictions. This is due to the intrinsic sensitivity of MD to the initial conditions, since the dynamics is chaotic. On multicore machines such as SuperMUC, this plays into our hands because, in the time it takes to perform one such calculation, all of the members of an ensemble can be computed. The method is therefore fast, with free energies being determined within around 12 hours. Considerable automation is necessary to perform these calculations, which consist of a large number of steps, including model building, production MD and data analytics performed on the resulting trajectory files, all on SuperMUC. This workflow is a much evolved version of the original Binding Affinity Calculator (BAC) [2].

The first system of interest to us within the present project was the Abelson (ABL) kinase, implicated in breast cancer, to determine how effectively it binds to a number of drugs. In the initial stages of this work over the past year or so, we have been focused mainly on developing and testing the BAC workflow software, tools and services, alongside determining the performance of the system in production. In the past three months, we have converged on the optimum set of software and been able to move into full scale production. This includes a new FabSim package for supporting management of complex workflows and their execution on remote supercomputers [3], coupled to a user-friendly version of the Binding Affinity Calculator, uf-BAC [4].

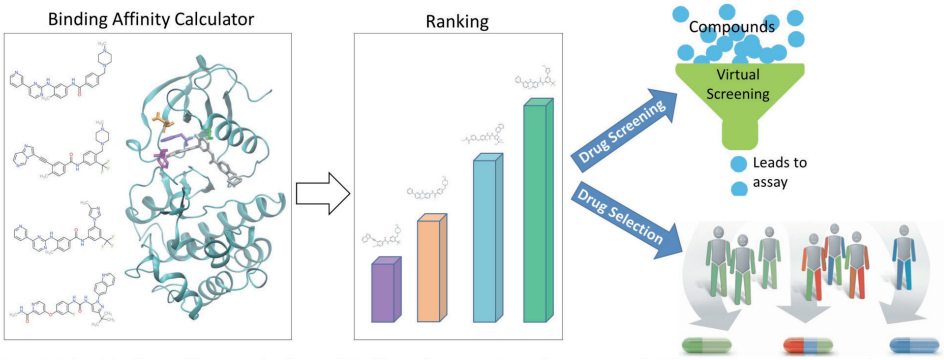

Fig. 1: The binding affinity calculator (BAC) ranks a series of compounds/drugs to a given protein target, and conversely that of a set of variant (mutant) proteins against a given ligand. BAC can be used as a virtual screening tool in pharmaceutical drug discovery for a large number of compounds, and as a key component in clinical decision support systems for personalized medicine.

Results and Methods

We have been able to make very important progress in our research based on the foregoing initially invested effort. At this time we could produce rapid, reliable, accurate and precise predictions of binding free energy ranking for the set of compounds originally proposed for the ABL kinase system using ESMACS. The same method has been used to perform studies on peptide-MHC (major histocompatibility complex), FGFR1 (fibroblast growth factor receptor 1) [5] and TrkA (tropomyosin receptor kinase A) target proteins. In the last two cases, we have also been able to apply TIES to compute relative free energies of binding within ±0.2 kcal/mol.

The speed, accuracy, precision and reproducibility of these results are unprecedented, in a field which has developed a poor reputation in respect of all these criteria. As a result, our work has attracted considerable interest from the pharmaceutical industry, for whom an automated, fast and reliable in silico method could dramatically transform their productivity.

On SuperMUC, we have been using NAMD and Amber packages for ensemble simulations of biological systems of interest. A typical calculation requires ca 9,800 cores using ESMACS and ca 12,740 cores using TIES. We can fill up a single phase of the machine (i.e. using ca 160,000 cores) and compute ca 16 binding affinities using ESMACS; or ca 12 relative free energies using TIES, within the space of 12 hours in each case; and double this turn around if we also make use of both phases of SuperMUC. A large number of simulations have been performed, involving studies of four drugs and five sequential variances of ABL kinase. A total of 1,480,315 CPU hours (96% of our allocation) has been consumed, and more than 250 GB data has been generated within the WORK directory, by this work in January 2016.

On-going Research / Outlook

Our predictions from ensemble simulations are in agreement with previously published experimental findings, and/or have been verified by subsequent experimental studies performed by colleagues in collaboration with us, including within leading pharmaceutical R&D laboratories worldwide [5]. Our findings have demonstrated that this approach is able to deliver an accurate ranking of ligand binding affinities quickly and reproducibly. Speed as well as reliability is essential in order to accelerate the drug discovery and development process, as well as to offer real-time clinical decision support in personalized drug selection. We are exploiting the multicore nature of SuperMUC to run many parallel jobs concurrently. It is possible for us to produce results within 12 hours. Then, with the very substantial amounts of data acquired (hundreds of gigabytes) residing on the same machine, we can perform the required data analysis, also in parallel, to produce the free energies and error analysis.

The purpose of the requested extension we are making here is to further explore, validate and apply the techniques we have developed to existing and emergent disease cases, and to better delineate the relative merits of our two approaches, ESMACS and TIES, for calculating binding affinities. Being able to make such predictions in a routine way using computation will also create many new opportunities for Europe's pharmaceutical sector, by reducing the time and cost of developing new medicines, as well as in healthcare delivery in the clinical context. This in turn will benefit wider society through improved treatments of an increasingly personalized form, through linking drug selection to a patient's genomic profile.

We plan to use the requested allocation of CPU time to investigate two groups of molecules. The first group focuses on drug selection in cancer treatment, which includes: estrogen receptor (ER), progesterone receptor (PR), aromatase, cyclin-dependent kinase (CDK) and human epidermal growth factor receptor 2 (HER2); these are important prognostic biomarkers and therapeutic targets in breast cancer; both ESMACS and TIES will be applied to these systems for selecting the best possible drug for a given patient. The second group focuses on drug development in the pharmaceutical sector, including: Jun N-terminal kinase 1 (JNK1), myeloid cell leukemia 1 (MCL1), p38 protein kinase, protein-tyrosine phosphatase 1B (PTP1B), thrombin, tyrosine kinase 2 (Tyk2); a large number of congeneric compounds are available to each of the above proteins, to which our TIES approach (and perhaps ESMACS too) will be applied for drug screening and compound optimization (Fig. 1). This makes up a total of ca 200 ligand-protein combinations we wish to calculate affinities for.

For the purpose of ranking a set of available compounds based on their efficiency of binding to a given protein, applicable both in the area of personalized medicine and lead optimization, we will need to run ensemble simulations for many compounds concurrently. As indicated above, a set of biomedical systems can be executed concurrently with a single submission script, ensuring optimal use of the future SuperMUC Next Generation (SuperMUC-NG).

Based on earlier estimates of number of core hours to compute a single ESMACS and a single TIES result, the studies we plan to perform require 30M core hours on SuperMUC for one year, preferably starting from April 2016. This proposal is intended to support work we are committed to perform within EU H2020 ComPat project (H2020-FETHPC-2014, 2015-18), UK MRC Medical Bioinformatics grant (MR/L016311/1, 2014-19) and special funding from the UCL Provost.

References and Links

[1] http://ccs.chem.ucl.ac.uk/supermuc

[2] S. Sadiq, D. Wright, S. Watson, S. Zasada, I. Stoica, P. Coveney. 2008. J. Chem. Inf. Model. 48, 1909–19.

[3] D. Groen, A. Bhati, J. Suter, J. Hetherington, S. Zasada, P. V. Coveney. Comput. Phys. Commun., in press (2016); also at arX-iv:1512.02194

[4] S. Wan, A. Bhati, S. Zasada, P. Coveney. "Rapid, accurate, precise and reproducible ligand-protein binding free energy prediction". Preprint (2016).

[5] D. Wright, B. Hall, O. Kenway, S. Jha, P. Coveney, 2014. J. Chem. Theory Comput. 10, 1228-41; S. Wan, B. Knapp, D. Wright, C. Deane, P. Coveney, 2015. J. Chem. Theory Comput., 11, 3346-56; T. Bunney, S. Wan, N. Thiyagarajan, L. Sutto, S. Williams, P. Ashford, H. Koss, M. Knowles, F. Gervasio, P. Coveney, M. Katan, 2015. EBioMedicine, 2, 194–204.

CAMEL Project Report

Research Institution
University of Bordeaux and the Medical University of Graz

Principal Investigator
Edward Vigmond / Gernot Plank

Researchers
Jason Bayer, Aurel Neic

Project Partners
–

SuperMUC Project ID: pr87na (PRACE project)

Introduction

This research deals with 2 projects concerning cardiac electrophysiology:

1) The most effective therapy for terminating ventricular fibrillation (VF) is a strong electric shock delivered by an implantable cardioverter-defibrillator. However, the high energy required for far-field defibrillation causes severe pain, tissue damage, and increased mortality. Since investigating new defibrillation technologies in vivo is extremely challenging, we utilized computational modeling to test wide-area low-energy tissue simulation for reducing defibrillation energy requirements.

2) The heart is a mechanical pump controlled by an electrical signal. Current ccmputer models of cardiac electromechanical function are limited in the sense that they cannot take into account the wealth of structural and functional data provided by state of the art experimental and clinical science. This is to be blamed to a large extent on the significant computational costs incurred in models that represent both cardiac anatomy and function with high geometric and biophysical fidelity. The most complex human heart simulations use meshes with millions of degrees of freedom, have behaviour at each node described by up to 100 differential equations, and must solve a linear and non-linear systems at each time step. Within the frame of this project we addressed several of these aspects to facilitate the simulation of a human heartbeat at an unprecedented level of detail.

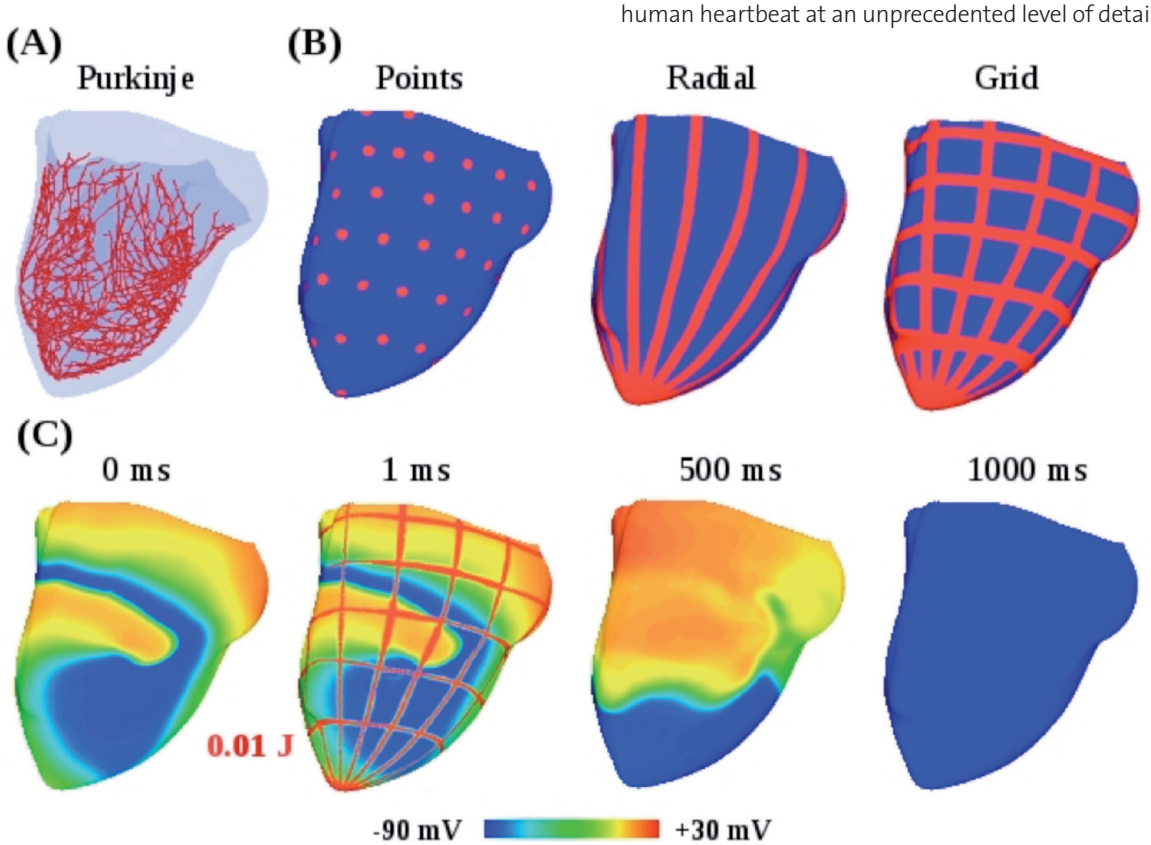

Figure 1: Wide area defibrillation. A) Cardiac conduction system. B) Defibrillation scenarios tested with electrode configurations shown in red C) Electrical activity shown with defibrillation shock applied at time 0. Colors represent transmembrane voltage.

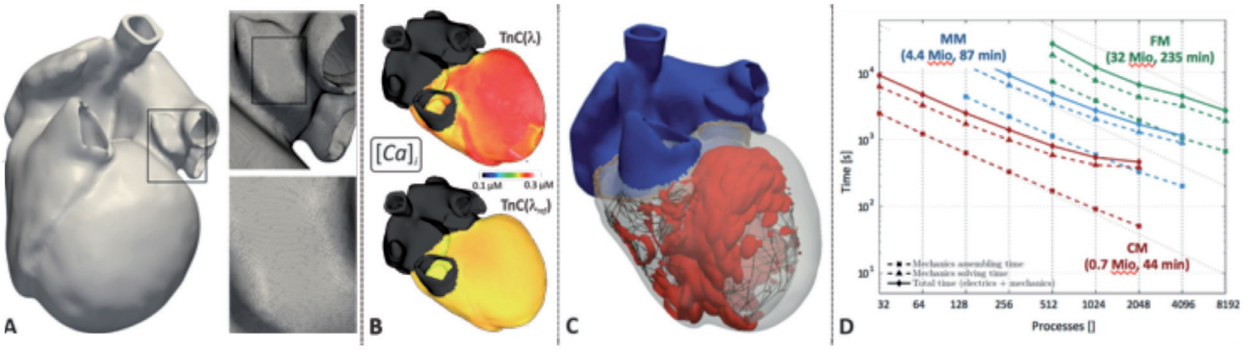

Figure 2: A) High resolution human four chamber heart geometry model. B) Impact of length dependent Calcium binding of TnC. Shown is Calcium with (top) and without (bottom) length dependence of binding affinity. C) Activation sequence (wavefront visualized as red isosurface) initiated by a topologically realistic model of the cardiac conduction system. D) Strong scaling experiments using three different spatial resolutions, a coarse mesh (CM), a medium mesh (MM) and a fine mesh. Degrees of freedom and minimum execution times are given in [1].

Methods and Results

Defibrillation
VF was initiated via rapid pacing in an image-based computer model of the failing human ventricles, complete with a Purkinje network (Fig 1A). Termination of VF with energy below the human pain threshold (<0.1J) was tested with the electrode configurations in Fig. 1B. Of the three configurations, only line electrodes on the ventricular surfaces terminated VF at energy levels <0.1J (Fig 1C).

Electromechanics
We developed the first strongly coupled electro-mechanical human four chamber heart model. All chambers are represented with full macroscopic scale geometric detail (Fig 2A) at high spatial resolutions in the range of 220 µm up to 880µm [1]. A biophysically highly detailed model of a human ventricular myocyte model, comprising state-of-the-art representations of both electrophysiology and mechanical force generation, was developed and used in organ scale simulations of a human heartbeat [2]. Strong coupling, achieved by accounting for length dependence of Calcium binding affinity of the Troponin C buffer, revealed that stretch heterogeneity throughout the heart induces a noticeable variability in cytosolic Calcium (Fig 1B).

The four chamber model was equipped with a topologically realistic model of the cardiac conduction system (Fig 2C). This allows the generation of physiologically realistic activation sequences, which serve as a trigger of the heartbeat.

Computational tractability was achieved by improving the scaling properties of all PDE solvers used (Fig 2D). In particular, a novel algebraic multigrid preconditoner for an iterative conjugate gradient solver was developed, optimized for deformation problems. Depending on the spatial resolution and the corresponding size of the problem, efficient strong scaling properties were achieved up to 2k to 8k cores.

On-going Research / Outlook

It was concluded that low-energy tissue stimulation with line electrodes effectively terminates VF, thus paving the way for a less damaging, pain-free alternative to traditional defibrillation therapies. Biological experiments are underway to confirm these results.

Overall, the developments of this project constitute a significant methodological break-through, which enables studying cardiac electro-mechanical function at an unprecedented level of detail. Specifically, we are going to study cardiac resynchronization therapy, wherein the goal is to improve pumping efficiency by electrically stimulating at several locations.

References and Links

[1] Crozier A, Augustin C M, et al. Image-Based Personalization of Cardiac Anatomy for Coupled Electromechanical Modeling. Ann Biomed Eng. 44(1):58-70, 2016.
[2] Augustin C M, Neic A, Liebmann M, et al. Anatomically accurate high resolution modeling of human whole heart electromechanics: A strongly scalable algebraic multigrid solver method for nonlinear deformation. J Comp Phys 305:622-646, 2016.
[3] J. D. Bayer, R. D. Walton, B. J. Boukens, I. R. Efimov, O. Bernus, E. J. Vigmond. Terminating lethal cardiac arrhythmia with low-energy line electrodes. Heart Rhythm, Vol. 12, No. 5, pp. S324-S325, 2015.
[4] J. D. Bayer, F. Vaddakkumpadan, E.J. Vigmond. Terminating ventricular fibrillation in the human ventricles with wide-area low-energy tissue stimulation. Heart Rhythm. Vol. 11, No. 5, pp. S5, 2014.

Iphigenie/CPMD: Accurate and Efficient QM/MM Molecular Dynamics

Research Institution
Ludwig-Maximilians-Universität München

Principal Investigator
Gerald Mathias

Researchers
Magnus Schwörer

Project Partners
KONWHIR

SuperMUC Project ID: pr89xe

Introduction

Chemical and physical properties of biomolecules are distinctly modified by the surrounding aqueous solvent. Thus, solute-solvent interactions are experimentally probed to elucidate biomolecular working principles. Accurate modeling of such solvation effects supports the interpretation of experiments down to atomic resolution. Here, a key application is the computation of vibrational spectra that have been experimentally probed during a biomolecular reaction.

Computational approach

Molecular properties of the solute such as vibrational spectra can only be accessed by high-level quantum-chemical methods, in our case the grid-based density functional theory (DFT) program CPMD [1], which is, however, computationally too demanding to employ it for large solute-solvent systems. Instead, we have developed a hybrid molecular dynamics (MD) approach, see Figure 1, in which we model the solvent on molecular mechanics (MM) force field level, whose accuracy is enhanced by additional polarizable terms (PMM) implemented the (P)MM-MD package IPHIGENIE [1]. The computation of

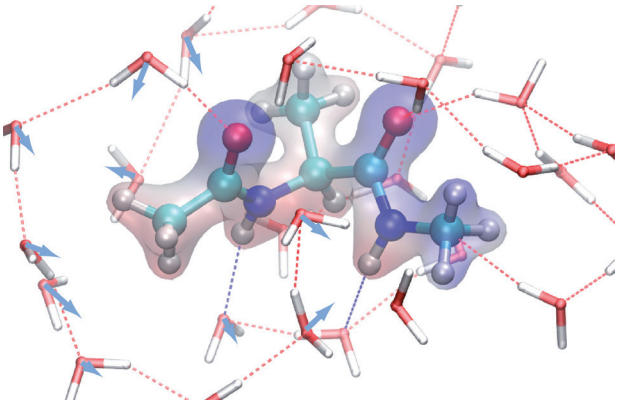

Figure 1: A alanine dipeptide (Ac-Ala-NHMe) treated by DFT immersed in PMM water. The electron density (shaded area) is colored according to the electrostatic potential generated by the solvent. PMM dipoles are indicated by blue arrows.

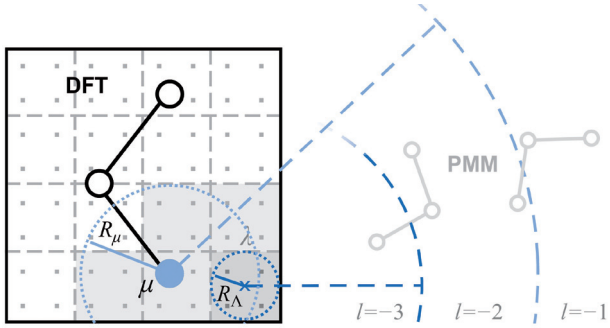

Figure 2: Sketch of the DFT/PMM FMM scheme. Close PMM atoms (within range *l*=-3) interact directly with the grid points. At larger distances (range *l*=-2) a PMM atom interacts collectively with a group of grid points (called voxels, dark gray square). A PMM atom within range *l*=-1 collectively interacts with all voxels associated with the DFT atom μ (light gray squares). For even larger distances both DFT and PMM atoms are grouped together (ranges *l*=0,1,2,..) in the FMM scheme of Iphigenie.

the interactions between the DFT and PMM fragments poses a considerable challenge because it involves on the order of 10^{12} pair interactions between the DFT charge density on ~10^7 grid points and ~10^5 PMM charges and dipoles of the solvent. Moreover, the DFT electron density and the PMM polarization terms have to be brought to mutual self-consistency, which requires multiple evaluations of DFT/PMM interactions during each MD integration step. This computational effort has become manageable by developing a nested hierarchical fast multipole expansion scheme (FMM), which computes DFT/PMM interactions in a linearly scaling fashion [2]. The approach is Hamiltonian and, thus, guarantees energy conservation and excludes artificial distortions of the electron density at the interface between the DFT and PMM fragments by using Gaussian distributed PMM charges and dipoles.

A second challenge when computing the properties of biomolecules in solution is the conformational sampling of the solute and of the solvent environment. Particularly sampling the conformational space of (bio-)molecules by plain molecular dynamics (MD) simulations is computationally inefficient. If one aims at an unbiased structural ensemble or has no prior knowledge of the free energy landscape, so-called generalized ensemble methods can alleviate the sampling problem. Here, we have combined

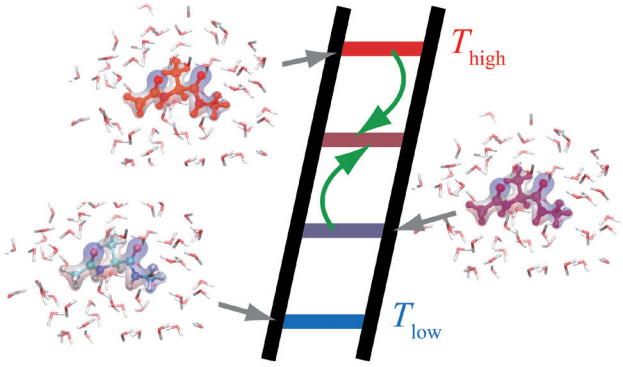

Figure 3: In SST simulations the system replicas frequently change the simulation temperature. At high temperatures enthalpic barriers can be crossed and newly visited molecular conformations can contribute to the ensemble of interest at low temperature after the replica has travelled through temperature space.

the DFT/PMM hybrid approach with the simulated solute tempering (SST) generalized ensemble method [4] sketched in Fig. 3. Because SST effectively heats up only the solute it requires only a few rungs on the SST temperature ladder and yields a particularly efficient sampling needed for the costly DFT/PMM method. The necessary SST parameters, so-called weights, can be obtained from inexpensive preparatory PMM simulations.

Results

The combination of the DFT/PMM and SST methods provides an efficient and highly accurate simulation setup that is scaling almost perfectly on SuperMUC Phase 2. Figure 4 shows the strong scaling of 128 replicas of the molecule alanine dipeptide (22 DFT atoms) shown in Fig. 1 solvated in PMM water (30000 PMM atoms). The inset shows the single replica scaling that demonstrates that the DFT/PMM coupling retains the excellent scaling behavior of CPMD for hybrid MPI/OpenMP parallelization. Particularly the optimal 7 OpenMP threads per MPI process reflect that CPMD fully profits from the Haswell microarchitecture of SuperMUC Phase 2. During the initial friendly user phase of Phase 2 we were able to scale the program to the full machine using up to 86k cores.

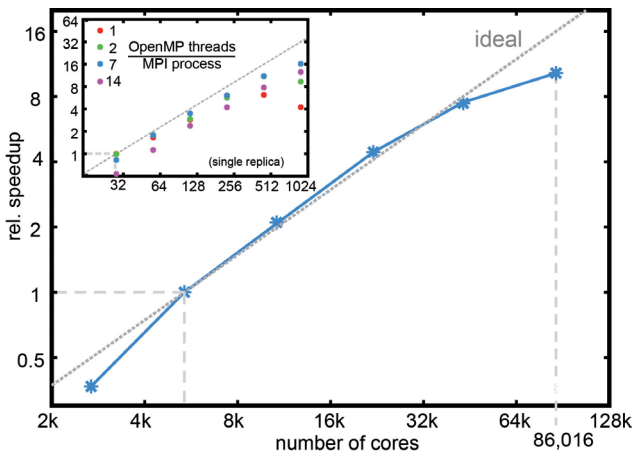

Figure 4: Strong scaling of a single replica (top inset) vs. the scaling of the SST generalized ensemble. The single replica setup is used to determine the optimal parallel setup (MPI+OpenMP).

As a first sample application we chose the free energy landscape of alanine dipeptide shown in Figure 5, which allows to identify typical molecular configurations. Here we find two prominent PP2 (ϕ, ψ) ≈ (-90°, 150°) and α-like (-90°, 0°) minima and the separating barriers.

On-going Research / Outlook

Currently we are analyzing further data obtained on the alanine dipeptide system. Particularly, the data refine the free energy landscape (Fig. 5) and compare free energies and barrier heights to other sampling methods. Furthermore, we have computed vibrational spectra based on

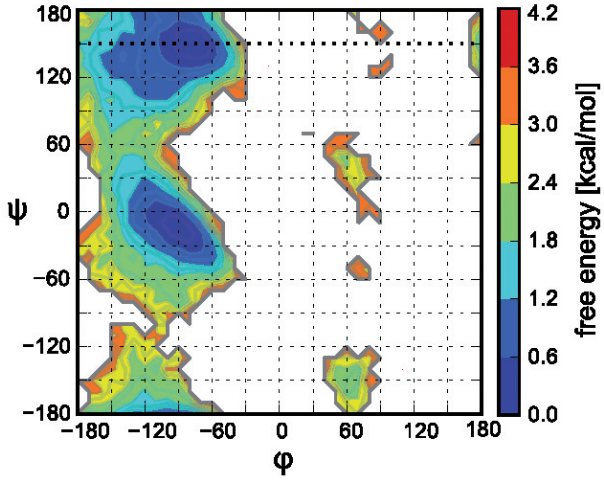

Figure 5: Free energy landscape of alanine dipeptide [4] obtained from DFT/PMM SST simulations.

the SST conformational sampling, such that we can map spectral features to the conformations of the molecule.

The CPMD - IPHIGENIE interface has been updated to the latest CPMD release 4.1 and has been verified for larger DFT subsystems.

References and Links

[1] https://sourceforge.net/projects/iphigenie/
http://www.cpmd.org
[2] M. Schwörer et al., J. Chem. Phys. 142, 104108 (2015)
[3] http://sc15.supercomputing.org/sites/all/themes/SC15images/tech_poster/tech_poster_pages/post177.html
[4] M. Schwörer et al., J. Chem. Theory Comput. 12, 992 (2016)

IPHIGENIE

Poster

The Molecular Mechanism of Cooperative Activation and Control of G-Protein Coupled Receptors

Research Institution
Computer-Chemie-Centrum and Interdisciplinary Center for Molecular Materials

Principal Investigator
Timothy Clark

Researchers
Noureldin Saleh

Project Partners
Francesco Gervasio (Chair of Biomolecular Simulations, University College London)

SuperMUC Project ID: pr94to

Introduction

The β2-adrenergic (ADRB2) and vasopressin receptors are members of the seven transmembrane receptor (7TMR) family, also known as G-protein coupled receptors (GPCRs), which represent the targets for more than 30% of marketed medications. The discovery of two different affinity-states of the ADRB2, controlled by the intrinsic activity and guanine nucleotide concentration, led to the proposal of a ternary complex model, in which both a ligand and an intracellular binding partner cooperate to achieve an active-state G-protein coupled receptor.[1] Cathecolamine binding to ADRB2 leads to heterotrimeric G-protein-dependent signaling, which results in the release of cyclic adenosine monophosphate (cAMP).[2] This eventually causes GPCR kinase- (GRK)-dependent phosphorylation followed by β-arrestin binding, a process known as desensitization.[3,4] It was subsequently discovered that some ADRBs can promote G-protein independent β-arrestin-signaling. The mechanism underlying the effect of an agonist ligand on the conformational space of the GPCR and the resulting conformational selectivity/cooperation toward coupling partners remains unresolved. In the present project, we have performed over 120 microseconds of atomistic simulations reveal the mechanism that underlies the cooperative effect between intracellular binding partners (G-protein, β-arrestin or protein nanobodies) and agonist ligands. These studies allowed the underlying mechanism to be extended to antagonists and inverse agonists and confirmed the cooperative arrestin-ligand activation of 7TMRs.

The project has resulted in an *in silico* protocol for predicting not only ligand-binding affinities, but also functional selectivity and intrinsic activity. This protocol can be extended to other GPCRs.

Results and Methods

We have developed a novel scheme based on combining atomistic molecular dynamics (MD) simulations, the well-tempered variant of metadynamics and the multiple-walker technique. The highly parallelized GROMACS suit with the PLUMED plug-in for the metadynamics simulations could be scaled up to one island of the Haswell nodes on SuperMUC. This allows extremely expensive metadynamics simulations of two-microseconds each to be converged in an overnight run.

Our simulations on a set of ligands with varying intrinsic activity and functional selectivity (toward both the G-protein and arrestin) showed that changes in the binding affinity caused by the intracellular binding partner (G-protein or arrestin) correlate with the ligands' intrinsic activities towards these coupling partners. The binding affinity is enhanced in the case of agonists, decreased for inverse agonists and remains unaffected for antagonists (see Figure 1). The changes found for the ligand-binding affinities were also found for the binding affinities of the corresponding G-protein and arrestin. These results, therefore, allow simulations to be used for the first time to predict and explain the intrinsic and functional selectivity of ligands on GPCRs.

Ongoing Research / Outlook

Molecular-dynamics simulations have so far not scaled well enough to allow sufficient sampling to derive thermodynamics parameters to predict experimental properties accurately. In this project, with the assistance and support of the LRZ, we were able to design a new sampling protocol that can derive these properties in simulations as fast as an overnight run on the 512 node of a phase 2 islands. Extending to more nodes or islands can potentially make these expensive calculations routine in the computer-aided drug-design process. We aim to be able in our next project to extend our calculations over several islands, which will allow us to address far more complex problems. This will require further method development to optimize the simulation protocol for this task. A new large-scale proposal has thus been submitted for the current Gauss-call.

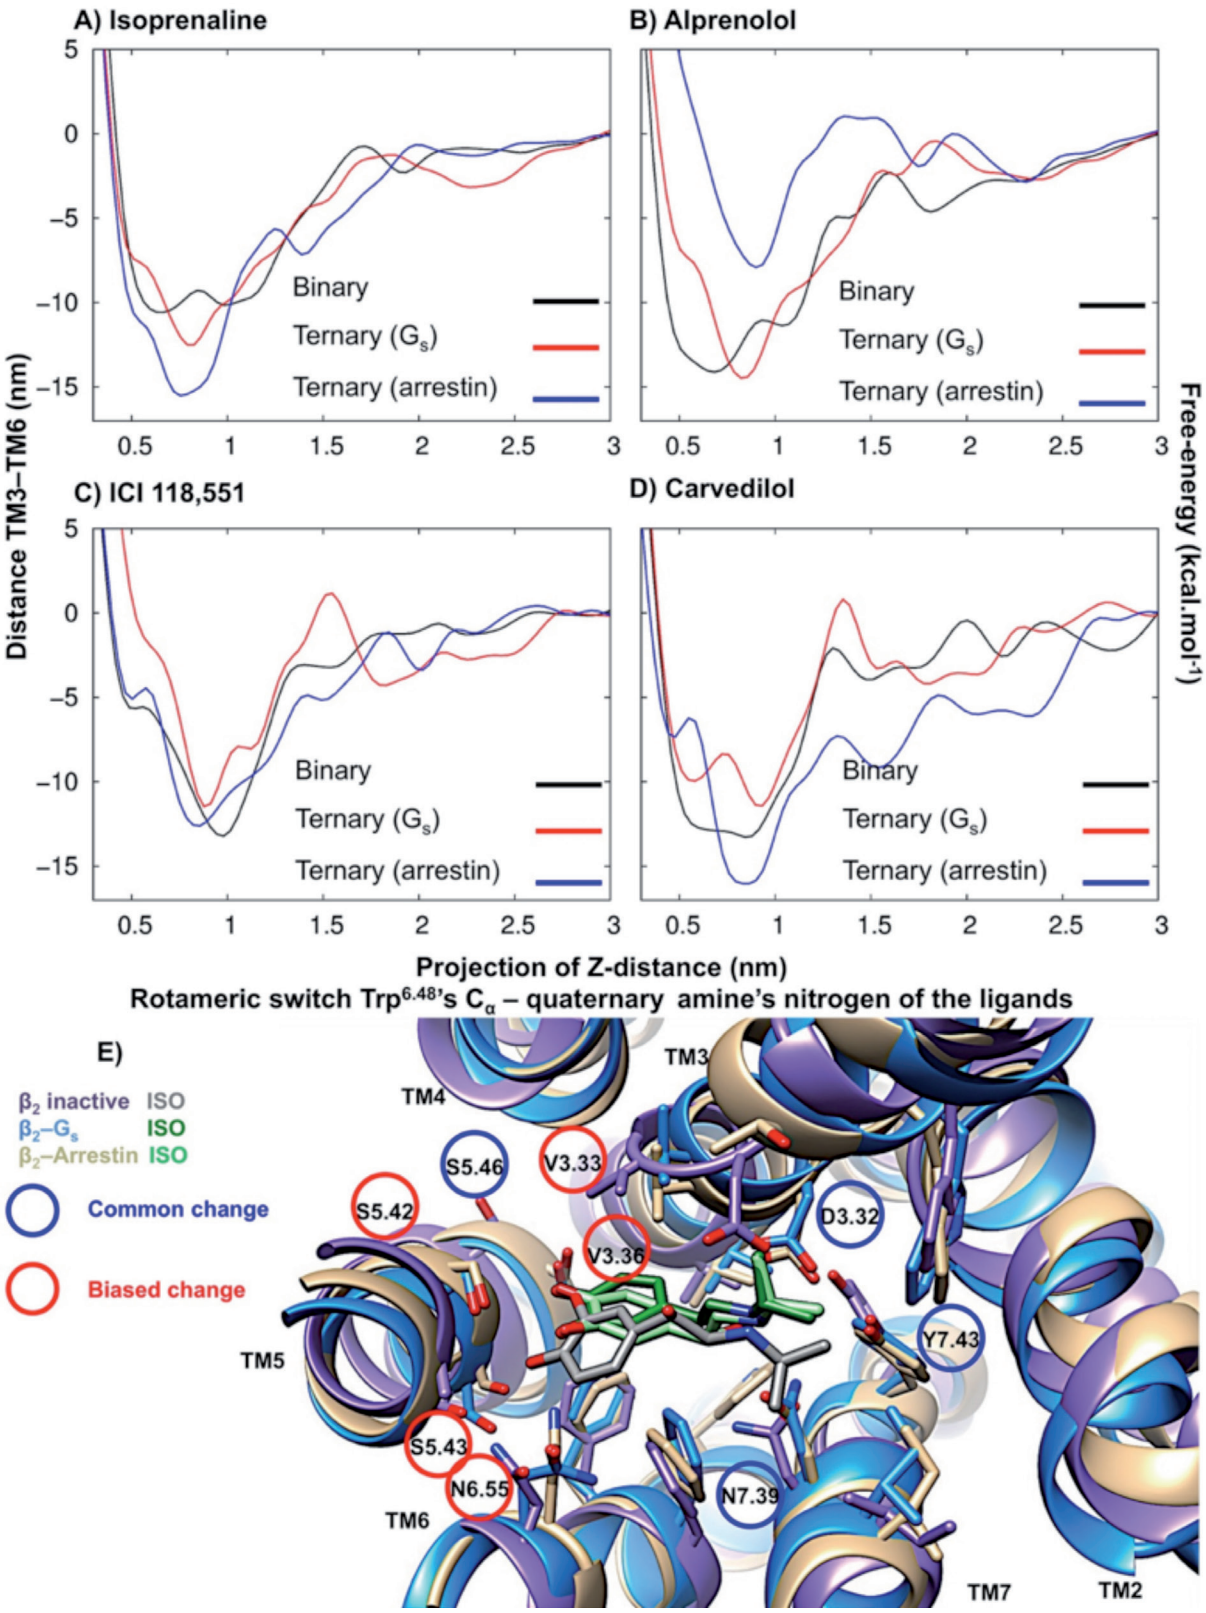

Figure 1: A-D: Free energy profiles for the ligand binding in the binary inactive-uncoupled ADRB2 and the ternary active-Gs coupled models. E) Comparison for the associated changes in the binding pocket for Isoprenaline.

References and Links

[1] De Lean, A., Stadel, J. M. and Lefkowitz, R. J. A ternary complex model explains the agonist-specific binding properties of the adenylate cyclase-coupled beta-adrenergic receptor. The Journal of biological chemistry, 255, 15 (Aug 10 1980), 7108-7117.

[2] Kenakin, T. Principles: receptor theory in pharmacology. Trends in pharmacological sciences, 25, 4 (Apr 2004), 186-192.

[3] DeWire, S. M., Ahn, S., Lefkowitz, R. J. and Shenoy, S. K. Beta-arrestins and cell signaling. Annual review of physiology, 692007), 483-510.

[4] Pitcher, J. A., Freedman, N. J. and Lefkowitz, R. J. G protein-coupled receptor kinases. Annual review of biochemistry, 671998), 653-692.

Targeting FtsZ assembly for the development of new antibiotics

Research Institution
Rocasolano Physical Chemistry Institute (IQFR-CSIC). Biological Physical Chemistry Department

Principal Investigator
Pablo Chacon

Researchers
Erney Ramírez-Aportela, José Manuel Andreu

Project Partners
Centro de Investigaciones Biológicas

SuperMUC Project ID: pr94ze (PRACE project)

Introduction

Cell division protein FtsZ is the organizer of the bacterial cytokinetic ring. FtsZ is a filament-forming GTPase that is thought to generate constriction force by a combination of filament bending, condensation and recycling. FtsZ and its eukaryote relative, tubulin, have molecular switches in their assembly-disassembly cycle, triggered by the presence or absence of the nucleotide gamma phosphate, which permit polymer regulation. Given its key role in division of the majority of bacteria, FtsZ has emerged as target for seeking new antibiotics to fight the widespread emergence of pathogens resistant to current antibiotics.

Only very recently the assembled sates of a straight FtsZ have been solved at atomic detail. Using this novel information, our goal is to perform all-atom simulations of long filaments bound to several biologically characterized modulators to better understand their inhibition mechanisms. In particular, we are interested in the relationship between the assembly molecular mechanism and the binding of modulators to help the rational design of new antibiotics.

We recently report a state of the art study of the FtsZ filament dynamics interpreted in the context of the assembly cycle of this essential cell division protein [2]. In contrast with all previous studies based on the inactive (not functional) closed-cleft FtsZ conformation studies, our large scale simulations studies disclose different filament curvatures supported by nucleotide-regulated interfacial dynamics. Moreover, we have monitored, for the first time, the relaxation from the active polymer conformation to the inactive closed-cleft conformation of FtsZ monomers. In agreement with experimental data, these groundbreaking results unravel the natural mechanism of the FtsZ assembly switch. Integrating this assembly switch and the nucleotide-dependent interfacial filament stability, our work offers a detailed molecular interpretation of the assembly-disassembly FtsZ cycle and its inhibition. Based on these results, we strongly think that the structure-based drug discovery efforts on this system can be only tackled by targeting the dynamics of the functional filament structure using atomistic simulations.

Results and Methods

Here we performed a molecular dynamics simulation study of FtsZ filament structures bound to several biologically characterized modulators to better understand their inhibition mechanisms. MD simulations were performed on FtsZ heptamers (~500,000 atoms, including water molecules). generated by crystallographic symmetry operations from the X-ray crystal structures of SaFtsZ bound to GDP (PDB ID 3V08) and to complex PC190723 (PDB ID 3V0B). The GDP and PC190723 crystallographic coordinates were replaced by optimized poses of the tested compounds in the corresponding binding site. For each filament-compound system, we carry out at least five MD simulations of 300-500 ns in length. Note that these FtsZ filaments are quite flexible and require long simulation times to achieve convergence. We typically employ 1024 cores to produce the long scale simulations that imply around 0.65 CPU hours/ns. Therefore the production time of one simulation system roughly takes several days to complete. The amount of sampling required to carry out all the simualtions planned would not be possible without the level of performance and scalability offered by SuperMUC.

There are a growing number of substances, some of them the product of large screens that have been reported to have some effect on FtsZ polymerization, FtsZ GTPases or bacterial cytokinesis [3]. There are two main binding sites 1) the nucleotide binding site located at the interface between dimers 2) interdomain cleft between C-terminal and the N-terminal domains. Therefore depending on the binding site we performed to type of studies:

1) *Targeting the polymerization interface between FtsZ monomers including the GTP binding site.* This includes two optimized analogs [4] of the polyhydroxy aromatic compounds, a Chrysophaentins fragment reported to be active [5], and two unpublished compounds obtained from an in house virtual screening study. All

these compounds that cover representative different scaffolds are GTP-replacing FtsZ inhibitors with antibacterial activity.

2) *Targeting the FtsZ activation switch.* First we revisited the filament stabilization mechanism of the effective antibacterial compound PC190723, which binds at a cleft between both FtsZ domains. Then we study more effective derivatives [6]. We also study a several fluorescent derivatives amenable for binding assays. Finally, we performed simulation on three novel compounds with antibacterial activities that also target the PC190723 binding site that we have identified from virtual screening approximation.

On-going Research / Outlook

Our results indicate the in the case of nucleotide site all compounds tested have a clear destabilization effect. By the contrary, in the case of PC190723 derivatives, the binding into the cleft between domains blocks the protein in the open filament-forming conformation. Conforming our previous simulations and in agreement with biochemical and structural studies, the maintenance of the cleft open avoids disassembly. In this context, we also explore some of the resistance mutations to these compounds (see Figure 1). We are currently analyzing the data in order to exploit these results for the structure-based drug discovery and desing. Molecular dynamics simulation was also used to test the stability of binding of three virtual screening compounds. Unfortunately, we encounter stabilization and convergence issues with such compounds depending on the initial configuration used in the simulation that we are still trying to overcome. However, we have no problems to simulate the fluorescent derivatives. In fact, thanks to the simulations performed on superMUC, we gathered enough information to guide the design of such compounds. Based on this information, and in collaboration with experimentalists, we have designed and synthesized several fluorescent probes to be used in ligand-binding assays. This probes are already been used in competitive binding assays and we are preparing a manuscript to be submitted within the next weeks.

In our ongoing research we currently analyze the effect of all these modulator compounds on the FtsZ assembly mechanism based on the simulations run on SuperMUC. In collaboration with experimentalist we are now confronting these results with experimental evidences. We hope that structural keys observed in the dynamics of the filaments will be very valuable to optimize and design new antibiotic compounds targeting Ftsz dynamics.

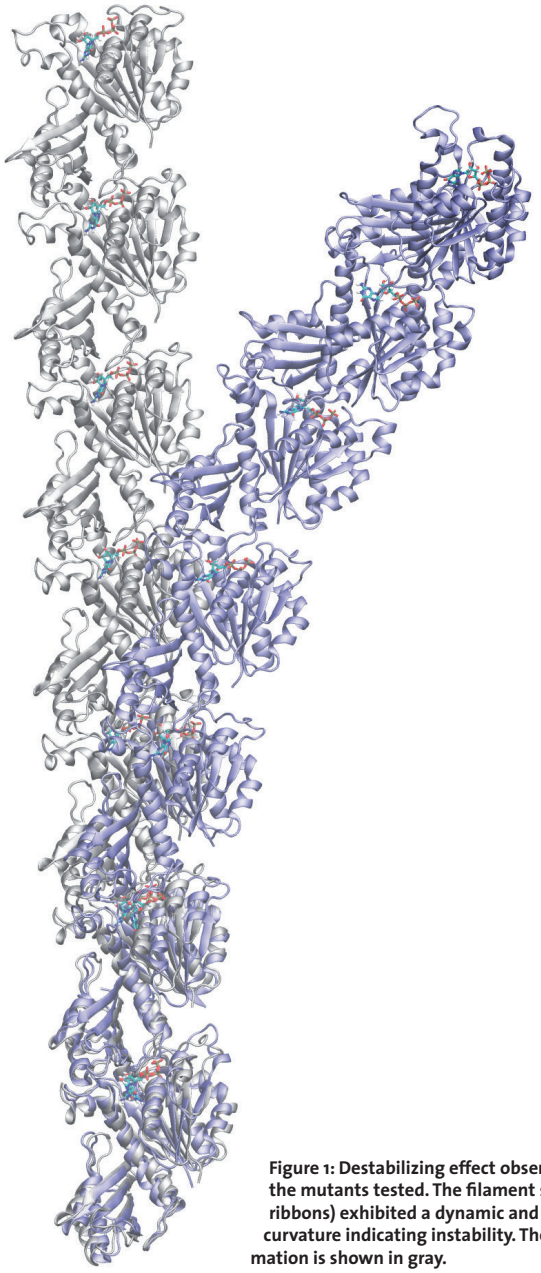

Figure 1: Destabilizing effect observed in one of the mutants tested. The filament structure (violet ribbons) exhibited a dynamic and heterogeneous curvature indicating instability. The initial conformation is shown in gray.

References and Links

[1] www.chaconlab.org

[2] Ramírez-Aportela E., López-Blanco J.R., Andreu J.M., and Chacón P. (2014). Understanding Nucleotide-Regulated FtsZ Filament Dynamics and the Monomer Assembly Switch with Large-Scale Atomistic Simulations. Biophys J. 107:2164–2176.

[3] Schaffner-Barbero C, Martín-Fontecha M, Chacón P, Andreu JM. (2011). Targeting the Assembly of Bacterial Cell Division Protein FtsZ with Small Molecules. ACS Chem Biol. 7:269-77

[4] M.Artola, L.B.Ruiz-Avila, A.Vergoñós, S.Huecas, L. Araujo-Bazán, M. Martín-Fontecha, H. Vázquez-Villa, C. Turrado, E. Ramírez-Aportela, A. Hoegl, M. Bruce Nodwell, I. Barasoain, P. Chacon, S. Axel Sieber, J.M. Andreu and M.L. López-Rodríguez. (2015) Effective GTP-Replacing FtsZ Inhibitors and Antibacterial Mechanism of Action. ACS Chem. Biol. 10:834–843.

[5] J. L. Keffer, S. Huecas, J.T. Hammill, P. Wipf,§ J.M. Andreu, and C.A. Bewley Chrysophaentins are competitive inhibitors of FtsZ and inhibit Z-ring formation in live bacteria Bioorg Med Chem. (2013). Bioorg Med 21(18): 5673–5678.

[6] Elsen N.L et al. (2012). J. Am Chem Soc 134, 12342-12345

Extreme Scaling on SuperMUC

Extreme Scale-out on SuperMUC Phase 2

In May 2015, LRZ started the reliability test of the new Peta-Scale System SuperMUC Phase2. In a 28-day block operation, the most demanding software applications were used to stress test the new parallel environment and MPI, the interconnect, and the I/O system, as well as the electrical and cooling infrastructure. Therefore, the LRZ invited 41 users to the "Extreme Scale-Out" block operation, to perform scaling tests on the new system using 14 different application codes. The codes as well as their measured power consumption are listed in Table 1.

Table 1: Applications from the Extreme Scale-out workshop on SuperMUC Phase 2. TDP stands for Thermal Design Point.

Software	Application Science	Watts/Node	Percent TDP	MPI tasks
SeisSol	Seismology	348	97,21	3072
Seven-League Hydro	Astropysics	254	70,95	86016
PSC	Plasmaphysics	249	69,55	86016
Musubi	Lattice Boltzmann	244	68,16	86016
BQCD	Quantum Chromodynamics	241	67,32	86016
CIAO	CFD, Combustion	237	66,20	86016
LS1-Mardyn	Material Science	234	65,36	172032
FLASH	Astro CFD	229	63,97	86016
waLBerla	Lattice Boltzmann	223	62,29	172032
Gadget	Cosmology	173	48,32	6144
ILBDC	Lattice Boltzmann	166	46,37	86016
Vertex3D	Stellar Astrophysics	159	44,41	6144
Iphigenie	Molecular Dynamics	108	30,17	12288
GPI-2 / GASPI	Global Adress Space Library	98	27,37	57344

The "Extreme Scale-Out" block operation was organized as a workshop. During a 1-day kick-off meeting, all teams presented their codes and could reserve SuperMUC Phase 2 for up to three full days. Each team was required to be present at LRZ during their reserved time slot. During regular office-hours, the team members would then get exclusive access to the LoadLeveler queue "special", where they could profile, debug and run their code up to the full system size of 86,016 cores. The team members would also have exclusive access to the new GPFS file system, where they could perform massively parallel I/O tests. In the special queue, jobs from 1 node up to the full system size of 3,072 nodes could be submitted and would start immediately since no users from other teams were allowed on the system.

This first part of the reliability test was designed to find bugs in the hardware, parallel environment and I/O system. Therefore, application experts and system administrators from IBM and LRZ, as well as the team members were present for the application runs, and only one application code was executed at a time. Indeed, several bugs and problems were found and could be fixed before the regular user operation of the new system started (see the end of this article for examples).

A second aspect of the reliability test was to simulate the regular operation of the new system. For this, a second LoadLeveler queue "general" was created that was only active at night. Users from all teams were allowed to continuously submit jobs to the general queue, from 1 node up to half of the system size, i.e. 1,536 nodes. The nightly runs using the general queue would therefore simulate the regular operation of the system.

In total, 63.4 million core-hours were available during the 28-day period for the participants of the workshop, of which 43.8 million core-hours were used by the applications, resulting in an overall utilization of 69%. The top 3 users were using 15.2, 6.4, and 4.7 million core-hours, respectively.

Some highlights from the workshop were:

- Largest cosmological simulation to date (10% of the visible universe simulated using the Gadget code)
- Largest pseudo-spectral simulation of interstellar turbulence ($10{,}000^3$ Cells using FLASH)
- Factor 100 better resolution for molecular spectra, and largest QM/MM replica exchange simulation (Iphigenie/CPMD)
- 2 Applications with sustained PFLOP/s Performance for more than 20 hours (SeisSol and LS1-Mardyn)
- Impressive strong scaling results for a seismic reconstruction problem using GPI-2 (computing time for one specific problem could be reduced from 16 hours to 55 seconds)

During the workshop, the performance and system monitoring tools registered a total of 6,751 jobs. There were 1,151 jobs that used at least 1,000 cores, 808 jobs that used more than 10,000 cores and 291 jobs using more than half the machine (43,008 cores). There were 9 full-machine jobs with a runtime of over 6 hours, and the

largest job consumed 1.2 million core-hours in a single run of 14.4 hours.

Each dual-socket Haswell EP node of SuperMUC Phase 2 contains 28 cores, resulting in a maximum of 56 tasks per node, if hyper threading is used. Only one software package utilized 56 tasks per node, most projects used 28 tasks per node. However, a significant amount of projects used a hybrid approach with varying numbers of OpenMP threads per node. In the hybrid OpenMP/MPI set-up, the applications used 1, 2, or 4 MPI tasks per node and 28, 14, or 7 OpenMP threads, respectively. For some codes, these factors of 7 instead of the more common powers of 2 for the thread-count were a challenge and required some modifications.

Figure 1 shows the electrical power in Watts consumed by SuperMUC Phase 2 for each run versus the number of cores used by this run. One can see that the power consumption scales nearly linear with the job size. For the largest jobs, a maximum of more than 1 MW was measured. To minimize the power consumption of large jobs, it is crucial to optimize the node level performance of an application. More than a relative factor of 10 in the variation of power consumption was observed for the largest jobs with 86,016 cores, which is not only important for the layout of the electrical infrastructure, but ultimately, also determines the power bill. The most power hungry compute intensive simulation code SeisSol (more than 1 PFlops sustained, cyan circles) used several times more energy in comparison to the memory bound Lattice Boltzmann Code ILBDC (red circles).

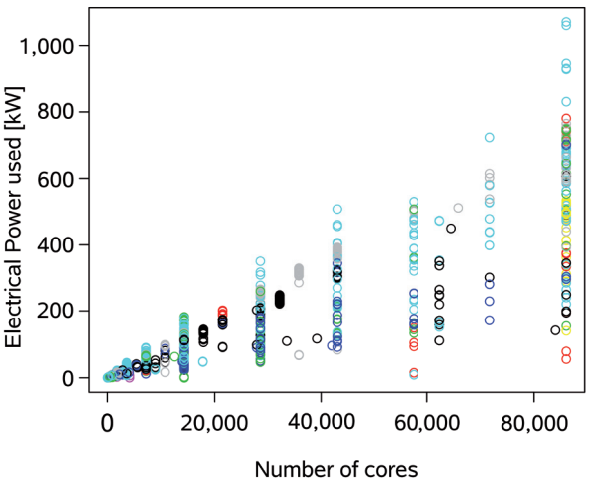

Figure 1: Scaling of the total power consumption of the system versus job size, measured by the number of cores used for each job. Different colors decode the different software packages used.

In the following, we give a brief description of some of the applications that participated in the workshop.

FLASH (Astrophysics, C. Federrath)

FLASH is a publicly available, modular, grid-based hydrodynamical code for the simulation of astrophysical flows. In the framework of the SuperMUC Phase 2 scale-out

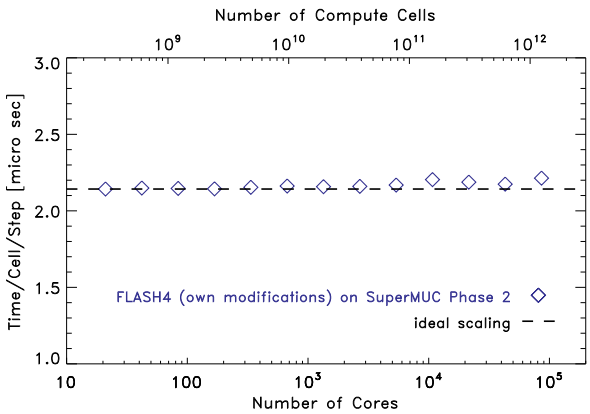

Figure 2: Weak scaling of the FLASH code.

workshop, the current version (Flash4) was optimized to reduce the memory and MPI communication requirements. In particular, non-critical operations were performed in single precision, without causing any significant impact on the accuracy of the results. This lead to a memory reduction of 75%. Also, the code was 3.6 times faster than the version used for the previous large-scale project at LRZ, and scaled remarkably well up to the whole Phase 2 system (Figure 2).

The FLASH team used the scale-out workshop to prepare a simulation of supersonic, isothermal turbulences with an unprecedented resolution exceeding $10{,}000^3$ grid elements. The simulation required several million core-hours, used about 155 TB of memory, and every data dump was 19 TB in size. This simulation was devoted to studying the complex flow pattern in the interstellar medium, with the aim of fully resolving the sonic length. It marked the transition scale between compressible and incompressible turbulence, where the turbulent cascade crosses from the supersonic into the subsonic regime.

Seven-League Hydro (Astrophysics, P. Edelmann)

The Seven-League Hydro (SLH) code is an astrophysical hydrodynamics code that solves the compressible Euler equations using a finite-volume method including

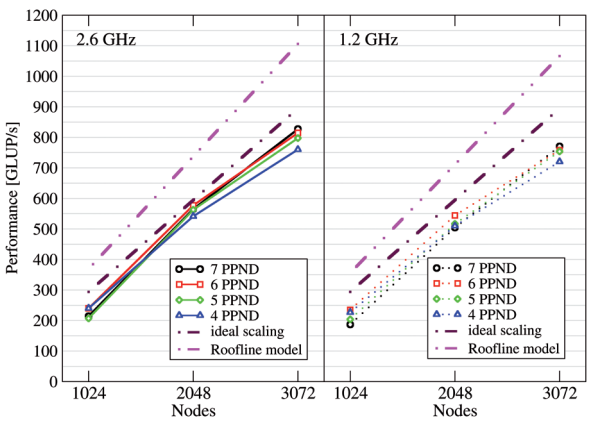

Figure 3: Strong scaling of the ILBDC code with a fixed bed reactor geometry with 17,500 x 3,500 x 3,500 nodes at two clock frequencies (ca. 1,011 fluid nodes ≈ 19, 6 TiB). PPND denotes the processes per NUMA domain.

source terms for gravity, nuclear reactions and thermal conduction, focusing on the simulation of the interior of stars. As these generally involve flows at very low Mach numbers, SLH includes a special discretization that reduces the excessive artificial dissipation which usually occurs in this regime. It also uses implicit time-stepping to overcome the CFL criterion (Courant-Friedrichs-Lewy number <1), which is exceedingly restrictive at low Mach numbers, leading to an overall more efficient scheme. The most time-consuming part of the computation is solving very large sparse Jacobian matrices using iterative linear solvers (e.g. BiCGstab, or multi- grid). These operations are mostly limited by the memory bandwidth. The code is parallelized with MPI and OpenMP, allowing pure MPI and hybrid MPI+OpenMP configurations.

For the scaling tests on SuperMUC Phase 2, the Taylor–Green vortex was chosen, a commonly used benchmark for the capability of turbulence simulation codes. It can easily be scaled to an arbitrary number of grid cells, has a very homogeneous flow pattern and does not involve any source terms. This allowed strong scaling tests of a $2{,}016^3$ grid up to the full machine. The minimum number of cores needed to store the sparse Jacobian matrix in memory is 21,504, using ~2.3 GiB/core. The code shows more than ideal scaling behavior in the step from 21,504 to 43,008 cores. This is very likely due to cache effects, since similar behavior was observed on SuperMUC Phase 1. Hybrid and pure MPI parallelization behave very similarly. Using IBM MPI, acceptable scaling was achieved even with 86,016 tasks. Intel MPI could not be tested with a high number of MPI tasks due to limitations in the Intel MPI library during the workshop.

ILBDC (Computational Fluid Dynamics, M. Wittmann)

ILBDC is a D3Q19-TRT lattice Boltzmann flow solver. Figure 3 (right panel) shows the strong scaling behavior of ILBDC from two to six islands for a CPU frequency of 2.6 GHz (left panel). Despite the 19 concurrent memory streams and indirect memory access, ILBDC achieves 87% of the memory bandwidth of the simple STREAM copy benchmark. Typically, 5 processes per NUMA domain (PPND) saturate the memory bandwidth. In the large-scale case, communication partially effects this behavior. With the Haswell-EP architecture, the sustained memory bandwidth of a node is almost independent of the core frequency. Consequently, ILBDC running at a CPU frequency of 1.2 GHz achieves 93% of the performance at a CPU frequency of 2.6 GHz (Fig. 3, right panel). For memory-bound codes, reducing the core frequency and the number of cores used per NUMA domain has only minimal performance impact. This bears the potential to drastically save energy on SuperMUC Phase 2.

Iphigenie/CPMD (Life Sciences, G. Mathias)

The MPI/OpenMP-parallel molecular dynamics package IPHIGENIE/CPMD links the publicly available molecular mechanics code IPHIGENIE [1] with the quantum–mechanical density functional theory application CPMD [2]. It is designed for highly accurate and efficient hybrid

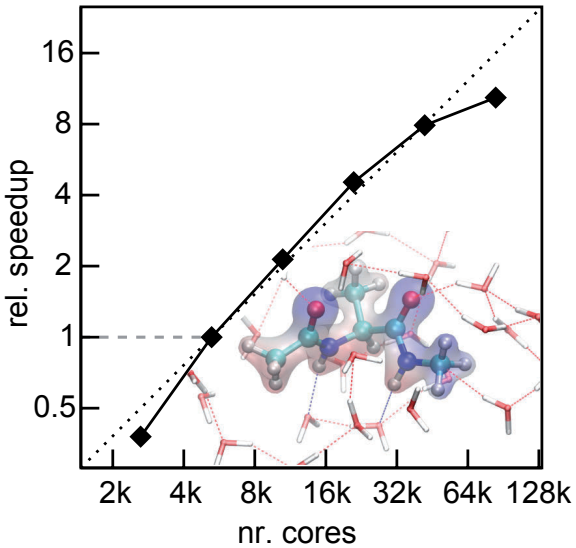

Figure 4: Strong scaling of a IPHIGENIE/CPMD generalized ensemble molecular dynamics simulation on SuperMUC Phase 2.

QM/MM–MD simulations of biomolecules in their native solvent environment. In particular, the use of polarizable molecular mechanics force fields for the aqueous solution set it apart from related approaches. Efficient sampling of bio-molecular conformations is achieved by employing generalized ensemble techniques, which jointly simulate multiple replicas of a system and require only sparse and infrequent communication between the replicas.

For the scaling tests, the Alanine dipeptide (Ac-ALA-NHMe) molecule was chosen as the DFT fragment at a plane wave cutoff of 100 Ry. It was solvated in 4,500 water molecules, described by a complex, polarizable water model. A total of 128 replicas of the system span the generalized ensemble. The best node level performance was achieved using 4 MPI tasks and 7 OpenMP threads per MPI task.

Figure 4 shows the scaling behavior on SuperMUC Phase 2, normalized to the performance at 5,376 cores. The application scales very well up to the full machine (86,016 cores). The graph even shows some super-scaling effects. Based on this excellent scaling, the team decided to use SuperMUC Phase 2 to target fundamental questions related to protein folding and vibrational spectroscopy with unprecedented accuracy.

PSC, the Plasma Simulation Code
(Plasma Physics, K. Bamberg)

The Plasma Simulation Code PSC is a general purpose framework to solve the extended Maxwell-Vlasov and extended Maxwell-Vlasov-Boltzmann equations via the Particle-In-Cell approach. Recent extensions comprise the self-field effects of radiation and electron-positron pair production in strong fields. PSC is a well-tested, widely recognized and reliable Particle-in-Cell Code that has also been used as a basis for other codes, e.g. the EPOCH code. PSC was ported to a modern modularized

C simulation framework, supporting bindings to FORTRAN and C/CUDA, featuring selectable field and particle pushers. The PSC framework allows the simulation of a large set of problems ranging from ultra-thin foils as light sources and proton-driven wake fields to quantum electrodynamic effects in extreme laser fields. PSC enables the simulation of large-scale, three-dimensional problems, e.g. ELI and AWAKE at CERN.

GASPI (Computer Science, M. Kühn)

FRTM, the framework for reverse time migration, was chosen as a showcase for GPI-2. The FRTM code is used for seismic imaging in exploration for oil and gas. It fully respects the two-way wave equation and provides high quality imaging, even in regions with complex geological structures, e.g. regions composed of steep flanks and/or high velocity contrasts. Reverse time migration has high demands on the underlying compute resources. Complex physical modeling, large target output domains, large migration aperture and/or high frequency content require efficient parallelization on the algorithmic side.

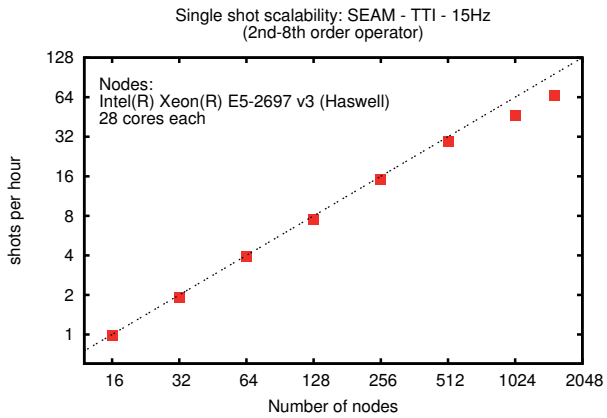

Figure 5: Scalability of FRTM using GPI-2.

In FRTM, the parallelization is based on a fine granular domain decomposition on a regular grid. It is supplemented by a fully asynchronous, data dependency driven, task based execution. Remote completion is used to take the data dependency driven execution from the process local level to the inter process level. The underlying halo exchange is implemented efficiently with GPI 2.0 which perfectly fits into this concept. GPI 2.0, the next generation Global address space Programming Interface, is a scalable communication API for one-sided, truly asynchronous communication primitives with remote completion. It is the reference implementation of the GASPI specification and is released under GPLv3 (www.gpi-site.com).

For the workshop, the strong scalability of a 15Hz single shot migration of the SEAM dataset achieved by FRTM was analyzed. The SEAM dataset is a well-established synthetic benchmark used to evaluate the quality of seismic imaging algorithms. The computational cost for a single shot migration is dominated by the finite difference modeling of the wave equation. The underlying physical model is given by a tilted transverse isotropic medium. The finite difference approximation of the temporal discretization is second order, the spatial part is an 8-th order discretization.

Up to 512 nodes, the scaling behavior on SuperMUC Phase 2 is almost perfect with a parallel efficiency of 94% at 512 nodes, 75% at 1,024 nodes, and 70% at 1,536 nodes, respectively. The declining parallel efficiency could be caused by the reduced inter-island network bandwidth or by the transition from a 2D to a full 3D domain decomposition beyond 512 nodes. Using 1,536 nodes, the absolute run time of FRTM for a single shot migration is 54.5 seconds, corresponding to a single precision floating point performance of 210 TFlop/s. The same input data set took 16 hours on a single node. This excellent strong scaling behavior of FRTM using GPI 2.0 enables interactive velocity model building based on RTM as driving engine.

Gadget (Astrophysics, K. Dolag)

GADGET is a widely used, publicly available cosmological N-body/Smoothed Particle Magnetohydrodynamics (TreePM-MHD-SPH) simulation code. It uses an explicit communication model implemented via MPI. For gravity, GADGET uses a TreePM algorithm, based on the fully MPI-parallelized FFTW library to perform the PM part of the gravity solver. Thereby, the long range gravity part is computed by sorting the particles onto a mesh and solving the Poisson equation via FFT; the short range part is determined from a Tree-walk algorithm, computing a direct sum of the forces between the particles and tree nodes (TreePM). Hydrodynamics is solved via the Smoothed Particle Hydrodynamics (SPH) method, which also make use of the Tree-walk algorithm to perform the neighbor search. Additional processes rely on various different numerical methods, for example transport processes are solved by a conjugated gradient method. These different physical modules are already optimized for mixed shared/distributed memory architectures, making use of OpenMP. In addition, most of the physics modules – for example star formation, thermal conduction, black hole treatment and on-the-fly post-processing, which are essential for modern cosmological applications – have been prepared for the latest architectures, resulting in a 30% performance improvement and efficient scalability up to 131,072 cores on SuperMUC Phase 1.

During the workshop, GADGET was used to perform one of the largest cosmological hydrodynamical simulations (in terms of resolution elements and covered cosmological volume) to date. The simulation followed 2 x $4{,}526^3$ particles and contained a detailed description of various, complex, non-gravitational, physical processes which determine the evolution of the cosmic baryons and impact their observational properties (for details, see www.magneticum.org). Amongst them are the star formation and related feedback; chemical pollution by Supernovae Type Ia, Supernovae Type II, and asymptotic giant branch (AGB) winds; transport processes like the thermal conduction, as well as the evolution of black holes, and their

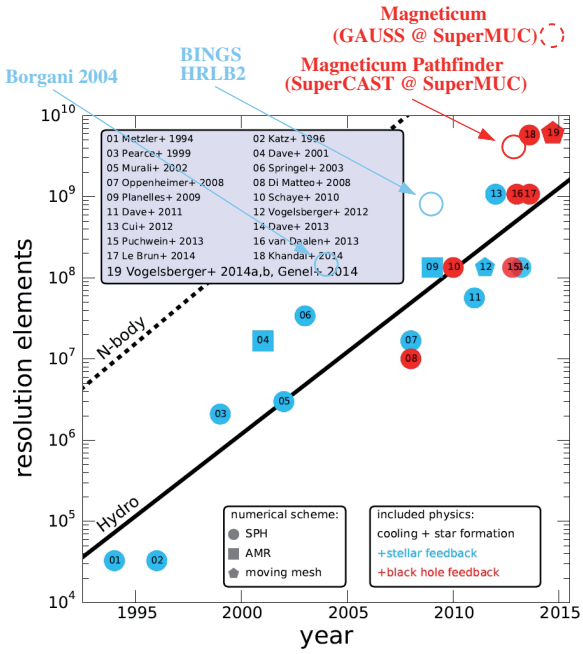

Figure 6: Magnitude of Magneticum simulation compared to previous simulations.

related active galactic nucleus (AGN) feedback. All these processes are self-consistently coupled with the underlying hydrodynamics. GADGET also uses a new optimization of the tree walk (GreenTree) as well as OpenMP lock improvements.

Including hyper-threading, GADGET was running with 172,032 threads in an OpenMP/MPI hybrid configuration on SuperMUC Phase 2. Apart from excellent scaling up to the full machine, GADGET also showed an extremely good I/O performance with 130Gbyte/s for writing and 150Gbyte/sec for reading of the regular 66 TByte checkpoint files every two hours. Overall, more than 4 Pbyte of data was produced during the workshop.

BQCD (High Energy Physics, M. Allalen)

The Berlin Quantum Chromodynamics code BQCD has implemented various communication methods: MPI, OpenMP, hybrid MPI+OpenMP, as well as single sided

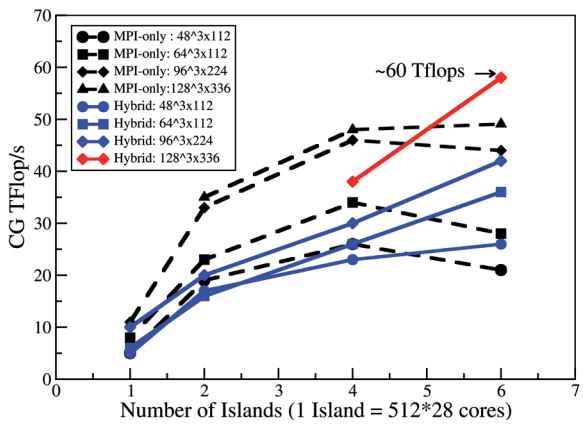

Figure 7: Strong scaling results of conjugate gradient solver for BQCD."

communication using SHMEM. The kernel is an iterative solver of a large system of linear equations, using the conjugate gradient algorithm, with more than 95% of the overall execution time spent in the solver. The dominant operation in the solver is a matrix-vector multiplication of a large and sparse hopping matrix. The entries in a row are the eight nearest neighbours of one side of the four-dimensional lattice. The conjugate gradient solver is communication intensive and represents the overall performance of the program in real world simulations. It serves as a good test to check the latency and resulting communication overhead.

Several strong and weak scaling runs were performed during the workshop. One difficulty was to find the right lattice size to fit in to the data cache, another was to describe the full system with a local lattice which fits the Infiniband network. Figure 5 shows the performance results of the conjugate gradient solver for the pure MPI versus the hybrid version, using lattice sizes of $48^3 \times 112$, $64^3 \times 112$, $96^3 \times 224$, and $128^3 \times 336$. Up to 4 islands (57,344 cores), the best performance is achieved using MPI only. On the full system, the hybrid version of the code delivers the best performance for all the lattice sizes using 7 OpenMP threads per MPI task.

CIAO (Computationial Fluid Dynamics, T. Falkenstein)

CIAO is an in-house code developed in collaboration between the Institute for Combustion Technology at RWTH Aachen University and Sogang University. Depending on the physical nature of the problem, the Navier-Stokes equations are solved for either fully compressible flows, or the low Mach limit. Various models for complex multi-physics flow configurations are available. For local mesh refinement or very stiff problems, a compressible multi-domain solver has been developed. CIAO is a structured, arbitrary order, finite difference code, which is parallelized with MPI. Using spatial and temporal staggering of flow variables increases the accuracy of stencils with a given size. The low Mach solver uses Crank-Nicolson time advancement and an iterative predictor corrector scheme. The compressible solver uses a low-storage five-stage explicit Runge-Kutta time advancement scheme. Scalar equations are discretized with a higher order WENO scheme, while momentum equations are spatially discretized with an arbitrary order central scheme. For Large-Eddy simulations, all of the subfilter stresses and scalar fluxes are modeled with dynamic Smagorinsky-type models using Lagrangian averaging along fluid particle trajectories. For the scale-out tests during the workshop, a Large-Eddy simulation of a periodic channel with 20 additional scalars was performed using the compressible Navier-Stokes solver of CIAO which does not use any third-party libraries. The code showed good strong scaling behaviour from 7,168 up to 86,0168 cores.

waLBerla

The waLBerla simulation framework is written in C++, with the main focus on CFD simulations based on the lattice Boltzmann method. Recently, the framework was

extended to support phase field models. The design goals of waLBerla are scalability to massively parallel systems as well as maximum node level efficiency. These goals are achieved by a careful performance engineering approach that includes the usage of performance models in order to determine the theoretical peak performance of the underlying algorithm running on a certain hardware platform.

The performance of the lattice Boltzmann method is best measured in MLUPS, short for "million lattice cell updates per second". For ideal weak scaling, MLUPS must scale linearly with the number of processor cores. The graph in Figure 8 also shows MLUPS per core, which is qualitatively equivalent to parallel efficiency. To evaluate the weak scaling performance of its parallel implementation of the lattice Boltzmann method, waLBerla makes us of a synthetic benchmark that simulates lid-driven cavity flow in 3D (D3Q19). The amount of work per core was kept constant by assigning 3.43 million cells to each core. Performance results are shown for the TRT collision model. The parallel implementation of the lattice Boltzmann method in waLBerla showed perfect weak scaling performance for the entire Phase 2 system. On 86,016 cores, 0.72 trillion cells could be updated each second.

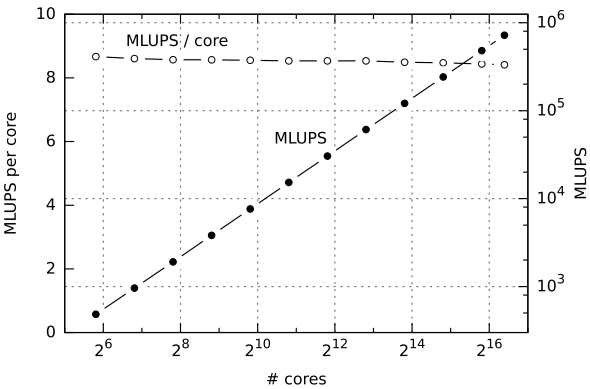

Figure 8: Scaling of waLBerla.

On SuperMUC Phase 1, waLBerla already demonstrated that compute kernels for the lattice Boltzmann method can reach the maximum performance predicted by performance models. With a memory bandwidth of 68 GB/s per socket on Phase 2 (14 cores) and 51.2 GB/s on Phase 1 (8 cores), the memory bandwidth per core on Phase 2 (4.9 GB/s) is considerably lower than on Phase 1 (6.4 GB/s). The drop in performance per core on Phase 2 corresponds to this memory bandwidth decrease. However, in strong scaling benchmarks, the parallel implementation of the lattice Boltzmann method of waLBerla reached higher throughput rates, meaning one can perform more time steps per second on Phase 2 than on Phase 1. On both machines, the best performance was achieved using a hybrid code (OpenMP+MPI): 4 threads per process on Phase 1 and 7 threads per process on Phase 2.

Results and Conclusions

The workshop at LRZ showed that preparation of a simulation campaign is crucial for the success of the project. All aspects like scaling tests, choice of OpenMP/MPI balance, interval for checkpoint and restart files, good preparation of input files, I/O strategy, and risk management have to be addressed. Under these conditions, it was possible to use a brand new system like SuperMUC Phase 2 directly after installation and obtain scientific results right from the start.

In total, there are now 25 applications that scale to the full system size of SuperMUC Phase 1 and Phase2.

One big advantage of the extreme scale-out workshop was that only one code was running at a time and this code was filling up the whole system. Hardware bugs which occur during such a "burn-in-phase" are much easier to detect and resolve. One especially hard to find bug was a combination of two timeouts and a hardware problem. During normal user operation this error would have been close to impossible to detect because of the low probability of three errors occurring simultaneously for smaller jobs.

It also became obvious that MPI is at its limits. The size of the MPI stack is growing on each node and for a system of almost 100,000 cores it occupies a significant amount of memory. The startup time can exceed the range of minutes and become a significant part of the overall run time. One way to overcome this bottleneck is the use of hybrid OpenMP/MPI programming models. However, this implies very deep system knowledge on the user side, since process pinning and the choice of the OpenMP/MPI balance has to be evaluated and decided by the user. Even for hybrid openMP/MPI set-ups with a single MPI-task per node, problems arose due to an internal limit of the MPI send/receive buffer. This limit is caused by the Integer*4 Byte implementation of the MPI index values. Such problems can be overcome by using application internal buffering.

Furthermore, I/O strategies have to be developed and tested before the complete system can be used. In the future, I/O libraries which can mediate this task become more and more important.

In summary, the Extreme Scale-Out workshop at LRZ proved that a burn-in period of a brand new HPC system can be very helpful for debugging and validating the hardware and software, especially because the whole system is filled with a single job whose characteristics are well known. Furthermore, the computing time for the burn in phase is not wasted, but can be used to provide productive scientific results.

References and Links

[1] www.sourceforge.net/projects/iphigenie
[2] www.cpmd.org

System Description

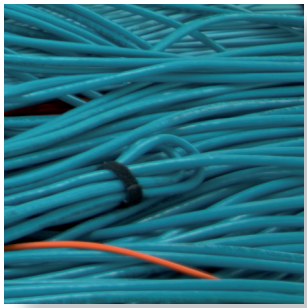

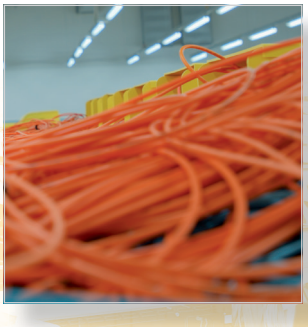

The SuperMUC Multi-Petascale System

SuperMUC is the high-end supercomputer at the Leibniz-Rechenzentrum (Leibniz Supercomputing Centre, LRZ) in Garching near Munich (the MUC suffix is borrowed from the Munich airport code). With more than 241,000 cores and a combined peak performance of the two installation phases of more than 6.8 Petaflop/s (= 6.8 x 10^{15} Floating Point Operations per second), it is one of the fastest supercomputers in the world. SuperMUC strengthens the position of Germany's Gauss Centre for Supercomputing [1] in Europe by integrating it into the European high performance computing ecosystem. With the start of operation of SuperMUC, LRZ became a European Centre for Supercomputing and a Tier-0 Centre for the Partnership for Advanced Computing in Europe (PRACE). SuperMUC is available to all German and European researchers to expand the frontiers of science and engineering.

LRZ's design goal for the architecture was a combination of a large number of thin and medium sized compute nodes with a main memory of 32 GByte (Phase 1) and 64 GByte (Phase 2), respectively, and a smaller number of fat compute nodes with a main memory of 256 GByte. The network interconnect between the nodes allows excellent scaling of parallel applications up to the level of more than 100,000 tasks. SuperMUC consists of 18 Thin Node Islands based on Intel Sandy Bridge-EP processor technology, 6 Medium Node Islands based on Intel Haswell-EP processor technology and one Fat Node Island based on Intel Westmere-EX processor technology. All compute nodes within an individual Island are connected via a fully non-blocking Infiniband network, FDR10 for the Thin nodes of Phase 1, FDR14 for the Haswell nodes of Phase 2 and QDR for the Fat Nodes of Phase 1. Above the Island level, the pruned interconnect enables a bi-directional bi-section bandwidth ratio of 4:1 (intra-Island / inter-Island). An additional system segment is called SuperMIC. It is a cluster of 32 Intel Ivy Bridge-EP nodes each having two Intel Xeon Phi accelerator cards installed (Knights Corner). See Table 1 for more details.

SuperMUC Phase 1 and Phase 2 are loosely coupled through the General Parallel File System (GPFS) and Network Attached Storage (NAS) File systems, used by both Phase 1 and Phase 2. Both phases are operated independently, but offer an identical programming environment.

SuperMUC uses a new, revolutionary form of warm water cooling developed by IBM. Active components like processors and memory are directly cooled with water that can have an inlet temperature of up to 40 degrees Celsius. This High Temperature Liquid Cooling together with very innovative system software cuts the energy consumption of the system up to 40%. In addition, LRZ buildings are heated re-using this energy.

Permanent storage for data and programs is provided by a 16-node NAS cluster from NetApp. This primary cluster has a capacity of 3.5 Petabytes and has demonstrated an aggregated throughput of more than 12 GB/s using NFSv3. Netapp's Ontap 8 Cluster-Mode provides a single namespace for several hundred project volumes on the system. Users can access multiple snapshots of data in

Figure 1: SuperMUC Phase 1 on the left and SuperMUC Phase 2 on the right, in the server room.

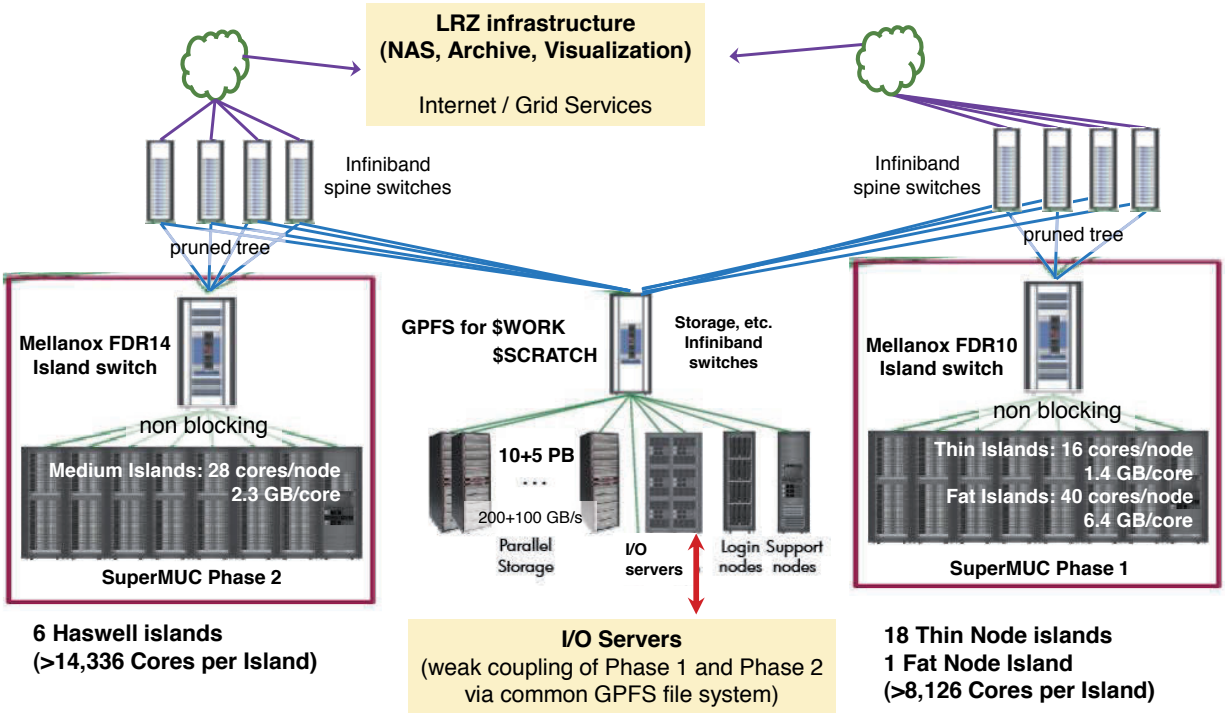

Figure 2: Schematic view of SuperMUC Phase 1 and Phase 2.

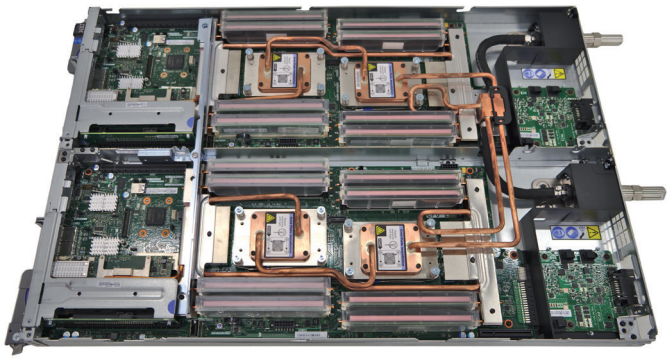

Figure 3: One blade of SuperMUC Phase 2 consists of two dual-socket compute nodes. The copper tubes distribute the warm cooling water to CPUs, memory modules and peripheral modules.

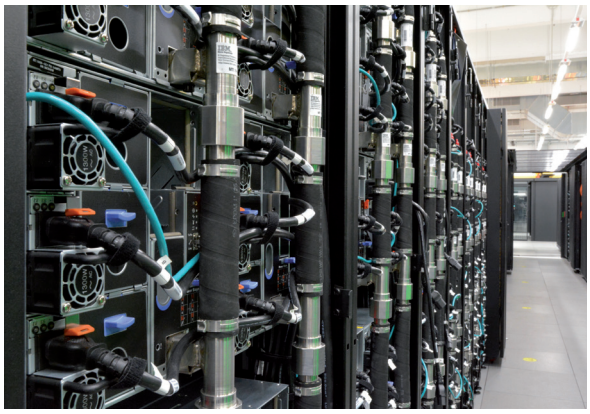

Figure 4: Rear of racks with warm water cooling.

their home directories. For additional redundancy, data is regularly replicated to a separate 4-node Netapp cluster with another 3.5 PB of storage for recovery purposes. Replication uses Snapmirror-technology and runs with up to 2 GB/s in this setup.

For high-performance I/O, IBM's GPFS with 12 PB of capacity and an aggregated throughput of 250 GB/s is available. Disk storage subsystems were built by DDN. The storage hardware consists of more than 3,400 SATA-Disks with 2 TB each, protected by double-parity RAID and integrated checksums.

LRZ's tape backup and archive systems are based on Tivoli Storage Manager (TSM) from IBM, providing more than 30 Petabytes of capacity to the users of SuperMUC. Digital long-term archives help to preserve simulation results. User archives are also transferred to a remote disaster recovery site.

Collaborations of European scientists can submit proposals to PRACE. Twice per year, the Gauss Centre for Supercomputing has a dedicated call for large scale projects that request more than 35 million core-hours. Smaller proposals by German scientists can be submitted throughout the year directly to LRZ.

References

[1] The Gauss Centre for Supercomputing (GCS) is the alliance of the three national German computing centres: Jülich Supercomputing Centre (JSC), High Performance Computing Centre Stuttgart (HLRS), and Leibniz Supercomputing Centre (LRZ).

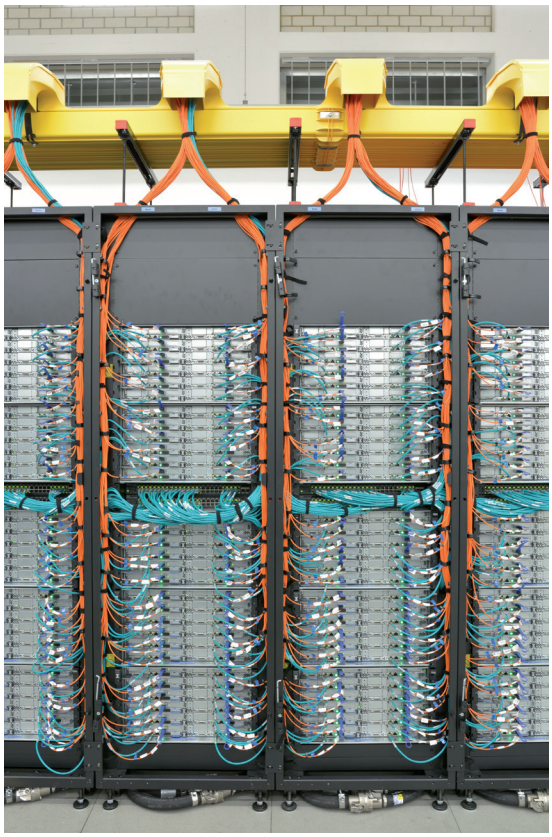

Figure 5: Several racks of SuperMUC Phase 2.

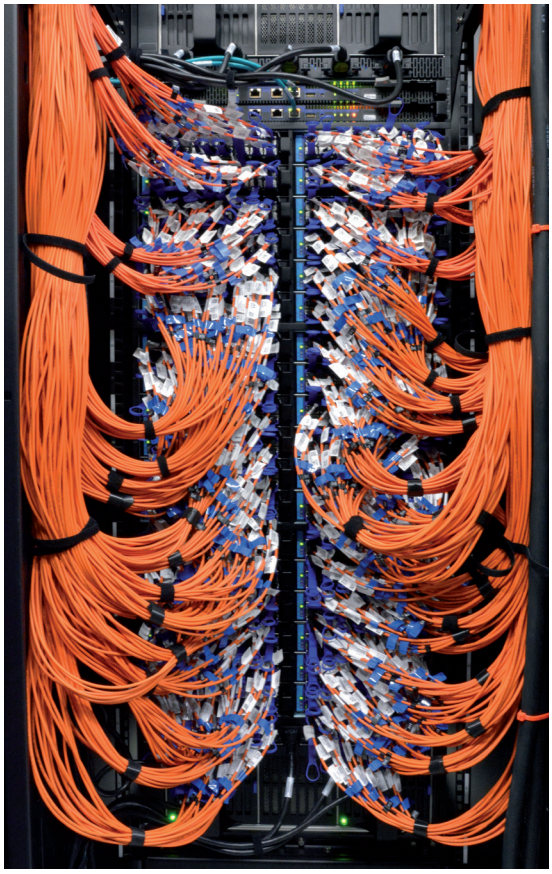

Figure 6: One of the Infiniband switches.

Technical data

Installation Phase
Installation Date
Islandtype
System
Processor Type
Nominal Frequency [GHz]
Performance per core
Total Number of nodes
Total Number of cores
Total Peak Performance [PFlop/s]
Total Linpack Performance [PFlop/s]
Total size of memory [TByte]
Total Number of Islands
Typical Power Consumption [MW]
Components
Nodes per Island
Processors per Node
Cores per Processor
Cores per Node
Logical CPUs per Node (Hyperthreading)
Memory and Caches
Memory per Core [GByte] (typically available for applications)
Size of shared Memory per node [GByte]
Bandwidth to Memory per node [Gbyte/s]
Interconnect
Technology
Intra-Island Topology
Inter-Island Topology
Bisection bandwidth of Interconnect [TByte/s]
Servers
Login Servers for users
Storage
Size of parallel storage (SCRATCH/WORK) [Pbyte]
Size of NAS storage (HOME) [PByte]
Aggregated bandwidth to/from parallel storage [GByte/s]
Aggregated bandwidth to/from NAS storage [GByte/s]
Capacity of Archive and Backup Storage [PByte]
System Software
Operating System
Batchsystem
Parallel Filesystem for SCRATCH and WORK
File System for HOME
Archive and Backup Software
System Management
Monitoring

Phase 1			Phase 2
2011	2012	2013	2015
Fat Nodes	Thin Nodes	Many Cores Nodes	Haswell Nodes
BladeCenter HX5	IBM System x iDataPlex dx360M4	IBM System x iDataPlex dx360M4	Lenovo NeXtScale nx360M5 WCT
Westmere EX Xeon E7-4870	Sandy Bridge EP Xeon E5-2680	Ivy Bridge EP and Xeon Phi 5110P	Haswell EP Xeon E5-2697 v3
2.4	2.7	1.05	2.6
4 DP Flops/cycle = 9.6 DP Flop/s 2-wide SSE2 add + 2-wide SSE2 mult	8 DP Flops/cycle =21.6 DP Flops/s 4-wide AVX add + 4-wide AVX mult	16 DP Flops/cycle =16.64 DP Flops/s 8-wide fused multi-ply-adds every cycle using 4 threads	16 DP Flops/cycle = 41.6 DP Flops/s two 4-wide fused multiply-adds
205	9,216	32	3,072
8,200	147,456	3,840 (Phi)	86,016
0.078	3.2	0.064 (Phi)	3.58
0.065	2.897	n.a.	2.814
52	288	2.56	194
1	18	1	6
< 2.3			~1.1
205	512	32	512
4	2	2 Ivy Bridge EP + 2 Phi 5110P	2
10	8	8 (Ivy Bridge EP) + 60 (Phi)	14
40	16	16 (host) + 120 (Phi)	28
80	32	32 (host) + 480 (Phi)	56
6.4 (~6.0)	2 (~1.5)	4 (host) + 2 x 0.13 (Phi)	2.3 (2.1)
256	32	64 (host) + 2 x 8 (Phi)	64 (8 nodes in job class big: 256)
136.4	102.4	Phi: 384	137
Infiniband QDR	Infiniband FDR10	Infiniband FDR10	Infiniband FDR14
non-blocking Tree			non-blocking Tree
Pruned Tree 4:1		n.a.	Pruned Tree 4:1
12.5			5.1
2	7	1	5
15			
3.5 (+ 3.5 for replication)			
250			
12			
> 30			
Suse Linux Enterprise Server (SLES)			
IBM Loadleveler			
IBM GPFS			
NetApp NAS			
IBM TSM			
xCat from IBM			
Icinga, Splunk			

In this book, the Leibniz Supercomputing Centre (LRZ), a member of the Gauss Centre for Supercomputing (GCS), reports on the results of numerical simulations, performed between 2014 and 2016 on the SuperMUC petascale system. More than 130 project reports give an impressive overview of the utilization of SuperMUC, the Tier-0 system of the Bavarian Academy of Sciences and Humanities.

SuperMUC Phase 1 began user operation in July, 2012, and **SuperMUC Phase 2** (picture above) became operational in May 2015. Each system segment has a peak performance of more than 3 PFLOP/s. Both phases are based on Intel x86 architecture and are coupled via a common parallel file system (GPFS). They are independently operated, but offer an identical programming environment. A detailed system description can be found in the appendix.

The articles provide an overview of the broad range of applications that use high performance computing to solve the most challenging scientific problems. For each project, the scientific background is described, along with the results achieved and the methodology used. References for further reading are included with each report.

ISBN 978-3-9816675-1-6

www.lrz.de