




P. Bastian • D. Kranzlmüller • H. Brüchele • M. Brehm • G. Mathias
EDITORS

High Performance Computing

in Science and Engineering
Garching/Munich 2020



KONWIHR

lrz

Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

GCS
Gauss Centre for Supercomputing

Titelbild entnommen aus dem Artikel von Christina Scharnagl und Simon Menig. Die Bilder zeigen Ergebnisse von Simulationen der Interaktionen, Desolvatisierung und langsamen Umlagerungen an Bindungsstellen. Die notwendigen langen Simulationszeiten für eine große Anzahl von Replikaten profitieren erheblich von der verbesserten Leistung von SuperMUC-NG. Dunkleres Grau zeigt häufigere und/oder länger andauernde Kontakte an. Weitere Informationen im Artikel ab Seite 252.

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: Peter Bastian, Dieter Kranzlmüller, Helmut Brühle, Matthias Brehm, Gerald Mathias
Redaktion und Layout: Helmut Brühle
Gestaltungskonzept: Tausendblauwerk, Konrad-Adenauer-Straße 22, 85221 Dachau,
www.tausendblauwerk.de
Druck und Bindung: Mayr & Abel Druck GmbH, Marktplatz 2, 87764 Legau.

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
<https://www.lrz.de/hpcbooks>

ISBN 978-3-9816675-3-0

P. Bastian • D. Kranzlmüller • H. Brüchele • M. Brehm • G. Mathias
EDITORS

High Performance Computing

in Science and Engineering
Garching/Munich 2020



Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

Table of Contents

Preface

- 10 ***SuperMUC and SuperMUC-NG: Serving Science and Society***
 PETER BASTIAN , DIETER KRANZLMÜLLER , HELMUT BRÜCHLE , MATTHIAS BREHM

Chapter 01 – Astrophysics

- 14 ***SuperSILCC – How stellar feedback regulates the interstellar medium***
 PHILIPP GIRICHIDIS
- 16 ***Cosmic structure formation on the past backwards lightcone in the MillenniumTNG simulations***
 VOLKER SPRINGEL
- 18 ***Mergers of binary neutron stars: linking simulations with multi-messenger observations***
 LUCIANO REZZOLLA
- 20 ***Unveiling the emergence of the first galaxies with state-of-the-art simulations***
 ENRICO GARALDI
- 22 ***Turbulence in astrophysical plasmas from fluid to electron scales***
 GIOVANNI LAPENTA
- 24 ***Gravitational-wave and electromagnetic signals from neutron star collisions***
 SEBASTIANO BERNUZZI, BERND BRUGMANN
- 26 ***Hunting down the cause of solar magnetism***
 MAARIT KÄPYLÄ
- 28 ***Long-term evolutions of neutron-star mergers: mass ejection, instabilities and EM counterparts***
 LUCIANO REZZOLLA
- 30 ***Testing the limits of magnetized accretion onto black holes***
 MATTEO BUGLI
- 32 ***The sonic scale revealed by the world's largest supersonic turbulence simulation***
 CHRISTOPH FEDERRATH
- 34 ***Simulating episodic outflow feedback in low-mass star formation***
 DANIEL SEIFRIED
- 36 ***Down to the dwarfs***
 STEFAN GOTTLÖBER
- 38 ***Star Formation through the Cosmic Ages: Shape and Spin of Primordial Mini-halos***
 RALF KLESSEN
- 40 ***Kinetic simulations of astrophysical and solar plasma turbulence***
 JÖRG BÜCHNER
- 42 ***Halos in Magneticum: Scaling Relations, Mass, Bias, and Concentration***
 KLAUS DOLAG
- 44 ***SuperCAST: Simulating the Universe***
 KLAUS DOLAG
- 46 ***Simulating the formation, evolution, and merging of molecular clouds***
 DANIEL SEIFRIED
- 48 ***Environmental Simulations of the Local Group***
 NOAM I LIBESKIND
- 52 ***Cracking the Convective Conundrum***
 PETRI KÄPYLÄ
- 56 ***From 3D Progenitors to 3D Supernova Explosions Including Muon Physics***
 HANS-THOMAS JANKA
- 60 ***Towards Energy Saturation in 3D Core-Collapse Supernova Simulations***
 HANS-THOMAS JANKA

Chapter 02 – Chemistry and Material Sciences

- 66 ***Role of quantum coherence in plasmon-induced hot-carrier dynamics***
EMANUELE COCCIA, STEFANO CORNI
- 68 ***Molecular vibrations reduce photovoltages in organic solar cells***
FRANK ORTMANN
- 70 ***Interface Phenomena in Chemical and Mechanical Engineering***
MAXIMILIAN KOHNS, KAI LANGENBACH, HANS HASSE, STEFFEN SECKLER, HANS-JOACHIM BUNGARTZ, MATTHIAS HEINEN, JADRAN VRABEC
- 72 ***Towards a predictive theory for unconventional superconductivity***
MARK VAN SCHILFGAARDE
- 74 ***Phase-field simulations to study eutectic colonies***
BRITTA NESTLER
- 76 ***All-electron DFT Simulations of particle-like magnetic objects***
STEFAN BLÜGEL
- 78 ***Pressure response of bulk liquid water as “seen” by THz spectroscopy***
DOMINIK MARX
- 80 ***Topology, Entanglement and Critical Phenomena in Correlated Quantum Matter***
FAKHER ASSAAD
- 82 ***Thermodynamics of the kagome lattice antiferro-magnet: tuning frustration and quantum fluctuations***
JÜRGEN SCHNACK
- 84 ***Finite-temperature Lanczos simulations of magnetic molecules***
JÜRGEN SCHNACK
- 86 ***Oxygen activation at Au/TiO₂ nanocatalysts***
DOMINIK MARX
- 88 ***Supercritical Water is not Hydrogen Bonded***
DOMINIK MARX
- 90 ***Transition metal oxide surfaces and interfaces for electronic and energy conversion applications***
ROSSITZA PENTCHEVA
- 92 ***Quantum Chemical Modeling of Actinide Interaction with Cement Phases***
ALENA KREMLEVA, SVEN KRÜGER, NOTKER RÖSCH
- 94 ***Numerical simulations of topological and correlated quantum matter***
FAKHER F. ASSAAD, EWELINA HANKIEWICZ, GIORGIO SANGIOVANNI

Chapter 03 – Computational Fluid Dynamics and Engineering

- 98 ***Simulating Combustion Systems***
ANDREAS KEMPF
- 100 ***DNS of premixed flame-wall interaction in turbulent boundary layers***
MARKUS KLEIN
- 102 ***On the dynamics of streamwise and transverse sediment patterns in turbulent open channel flows***
M. SCHERER, M. UHLMANN
- 104 ***Direct numerical simulation of partially-filled pipe flow***
MICHAEL MANHART
- 106 ***Turbulent natural convection of non-Newtonian fluids in enclosed spaces***
SAHIN YIGIT, JOSEF HASSLBERGER, MARKUS KLEIN
- 108 ***Investigation of the Flow Field and Inflow of Hovering Rotors over Inclined Ground Planes***
STEFAN PLATZER
- 110 ***Interface instability dynamics in secondary drop breakup***
NIKOLAUS A. ADAMS
- 112 ***Scale-Resolving Study of a Fluid-Structure-Acoustics Interaction Test Case***
MIRIAM MEHL
- 114 ***Thermodiffusive instabilities in lean premixed hydrogen/air flames***
HEINZ PITTSCH
- 116 ***VIVALDI—High fidelity simulation of vortex induced vibrations for flow control and energy harvesting***
ORIOLE LEHMKUHL

118	<i>Development of an integral LES model for turbulent premixed combustion at elevated pressures</i> MARKUS KLEIN
120	<i>Simulating the interaction of a strake vortex with an aircraft wing</i> SILVIA PROBST
122	<i>Aerodynamics of Innovative Rotor Configurations</i> CHRISTIAN BREITSAMTER
124	<i>Eulerian Stochastic Fields Method applied to partially premixed combustion</i> MICHAEL PFITZNER
126	<i>Bubble-collapse dynamics near biomaterial surrogates using multi-material level-set</i> NIKOLAUS A. ADAMS
128	<i>Vortical Flow Interaction on Airborne High Speed Vehicles</i> CHRISTIAN SCHNEPF
130	<i>Flashback phenomena in a Gas Turbine with hydrogen in technically premixed conditions</i> PANAGIOTIS STATHOPOULOS
132	<i>Aerodynamics and Aeroelasticity of high agility aircraft</i> CHRISTIAN BREITSAMTER
134	<i>CO Emissions in Part Load Conditions of Stationary Gas Turbines</i> ANTONIO ATTILI, HEINZ PITTSCH
136	<i>Simulating Acoustic-Flame Interactions in Rocket Engines</i> KLAUS HANNEMANN
138	<i>Numerical Investigation of turbulent heat transfer in a high aspect ratio curved duct</i> THOMAS KALLER
140	<i>Exploring Turbulence and Particle Transport Using Massively Parallel Simulations</i> MICHAEL WILCZEK
142	<i>Rocket engines under the microscope: Modeling turbulent combustion</i> OSKAR HAIDN
144	<i>Space launch vehicle aerodynamics with hot plumes</i> KLAUS HANNEMANN
146	<i>Turbulent convection at very small Prandtl numbers</i> JÖRG SCHUMACHER
148	<i>The Largest Scales in Turbulent Pipe Flow</i> CHRISTIAN BAUER
150	<i>LES of Rocket Combustion Applications Under Real-Gas Conditions</i> MICHAEL PFITZNER
152	<i>Simulations at very high Rayleigh number</i> RICHARD STEVENS, ROBERTO VERZICCO, DETLEF LOHSE
154	<i>DNS study of differential diffusion in a hydrogen jet in cross flow configuration</i> CHRISTIAN HASSE
156	<i>Development of a Stability-Based Transition Transport Modeling Framework</i> PHILIP STRÖER, ANTHONY D. GARDNER, KURT KAUFMANN
158	<i>Extreme-Scale Simulation of Turbulent Flow</i> MARTIN KRONBICHLER
160	<i>Large Eddy Simulation of Wind Turbines</i> FILIPPO CAMPAGNOLO
162	<i>Modulation of Turbulent Properties in Spray Flame Burning n-Heptane: Direct Numerical Simulation</i> DOMINIQUE THÉVENIN
164	<i>WALBERLA – A massively parallel framework for multi-physics simulations</i> HARALD KÖSTLER
166	<i>Cavitation in Injection Systems</i> THERESA TRUMMLER, STEFFEN J. SCHMIDT
168	<i>DNS study of the early flame kernel development under engine conditions</i> HEINZ PITTSCH
170	<i>Towards combustion LES of liquid alternative aviation fuels</i> PATRICK LE CLERCQ
172	<i>Buoyant-convectively driven heat and mass transfer</i> HERLINA HERLINA
174	<i>Numerical cavitation erosion prediction</i> THERESA TRUMMLER, STEFFEN J. SCHMIDT

- 176 ***Symmetry theory and turbulence***
MARTIN OBERLACK
- 178 ***Heat Exchange in Methane-based Launchers***
CHRISTIAN STEMMER
- 180 ***Enhanced Aerodynamics of Wind Turbines***
THORSTEN LUTZ

Chapter 04 – Earth, Climate and Environmental Sciences

- 184 ***Regionalization of global hydrometeorological fields***
TANJA PORTELE
- 186 ***High resolution combined global gravity field modelling***
THOMAS GRUBER
- 188 ***Advanced Simulation of Coupled Earthquake-Tsunami Events***
MICHAEL BADER, ALICE-AGNES GABRIEL
- 190 ***Towards an Exascale Hyperbolic PDE Engine***
MICHAEL BADER, ALICE-AGNES GABRIEL, TOBIAS WEINZIERL, MICHAEL DUMBSER, LUCIANO REZZOLLA
- 192 ***So2Sat – 10¹⁶ Bytes from Social Media to Earth Observation Satellites***
XIAOXIANG ZHU
- 194 ***Supercomputing of Extreme Natural Events***
DIETER KRANZLMÜLLER
- 196 ***3-D seismic wave propagation and earthquake rupture***
ALICE-AGNES GABRIEL, HEINER IGEL
- 198 ***Variable resolution meshes in climate research with the Model for Prediction Across Scales***
GERHARD SMIATEK
- 200 ***Atmospheric Chemistry and Climate***
MATTIA RIGHI
- 202 ***From New Materials to Planetary Cores***
RONALD E. COHEN
- 206 ***Deciphering Large-scale Geologic Events Through Retrodictions of Past Mantle Flow***
HANS-PETER BUNGE, BERNHARD SCHUBERTH

Chapter 05 – Elementary Particle Physics

- 212 ***The QCD phase diagram and equation of state***
ZOLTAN FODOR
- 214 ***Hadron structure observables on a fine lattice at the physical point***
SARA COLLINS
- 216 ***Isospin breaking effects in QCD***
KALMAN K. SZABO
- 218 ***Charm sea effects on heavy flavor mesons***
FRANCESCO KNECHTLI
- 220 ***TMDs and parallel transport in QCD***
ANDREAS SCHÄFER
- 222 ***How large is the color magnetic field in a proton?***
ANDREAS SCHÄFER
- 224 ***A new twist on plasma-based accelerators***
ANDREAS DÖPP
- 226 ***Study of anomalies in B decays as a probe of New Physics***
JOCHEN HEITGER
- 228 ***The QCD phase diagram in the quark mass plane***
ZOLTAN FODOR
- 230 ***Simulating bound states of gluons and gluinos***
GEORG BERGNER
- 232 ***A strong side of weak decays: How beauty and charm resonate in pions and kaons***
MARCUS PETSCHLIES

- 234 ***N=1 Supersymmetric gauge theories on the lattice***
ANDRE STERNBECK
- 236 ***Simulating Interactions for the ATLAS experiment at the LHC***
GÜNTER DUCKECK
- 238 ***Decoding the spin of the proton from quarks and gluons***
KARL JANSEN

Chapter 06 - Life Sciences

- 242 ***Structure and Dynamics on the Photosynthetic Complex I***
VILLE R. I. KAILA
- 244 ***Let there be light: the case of the major Light Harvesting Complex II***
VANGELIS DASKALAKIS
- 246 ***Simulating Blood Flow in the Virtual Human***
DIETER KRANZLMÜLLER
- 248 ***Mitigating the COVID-19 pandemic with an innovative approach to accelerate drug discovery***
DIETER KRANZLMÜLLER, PETER V. COVENEY
- 250 ***Transcription factors pioneer nucleosomes in motion***
VLAD COJOCARU
- 252 ***Understanding Substrate Recruitment by Intramembrane Proteases: Progress and Problems***
CHRISTINA SCHARNAGL
- 254 ***PIP2: Another player in biased signaling of G protein-coupled receptors?***
WOLFGANG WENZEL
- 256 ***Regulation of proton transfer in respiratory complex I***
VILLE R. I. KAILA
- 258 ***Advanced sampling of peptide and protein association and protein conformational transitions***
MARTIN ZACHARIAS
- 260 ***Multi-scale Molecular Simulations and ATPase Activity of Hsp90***
VILLE R. I. KAILA
- 262 ***Metadynamics reveals the secrets of G-protein coupled receptors***
TIMOTHY CLARK
- 264 ***Blood flow simulations***
STEPHAN GEKLE

Appendices

- 268 ***SuperMUC Phase 1 and Phase 2: System Description***
- 272 ***SuperMUC-NG: System Description***
- 274 ***More than three decades of super-computing at LRZ – peak performance and machine balance***
MATTHIAS BREHM

SuperMUC and SuperMUC-NG: Serving Science and Society

2020 was an extraordinary year: the world was, and still is, confronted with the once-in-a-century pandemic of COVID-19. This changed the way our LRZ colleagues work together and operate our HPC systems and, concomitantly, how our users conduct research on these systems. In the middle of March 2020, most employees and users had to instantly shift their work to working from home. Video conferences and chat groups became a central part of our daily work. For LRZ this helped to seamlessly keep the high-level of service, training and support provided to our users. Looking back, we can firmly say: this has changed our perspectives and practices persistently.

Starting with the first wave of the pandemic LRZ offered expedited access to SuperMUC-NG for COVID-19 related projects. This was a joint effort together with the High-Performance Computing Centre Stuttgart (HLRS) and the Jülich Supercomputing Centre (JSC) within the Gauss Centre for Supercomputing (GCS), the alliance of Germany's three national supercomputing centres. On a European level, GCS offered compute time for a Covid-19 Fast Track Call in the Partnership for Advanced Computing in Europe (PRACE). Early results of these calls are contained in the section Life Sciences of this book.

Beyond these fast track calls also long-standing efforts in fundamental science contributed to COVID-19 related research. The comprehensive computational model of the human lung established in recent years on LRZ systems provides a better understanding of the processes involved in mechanical ventilation. This can be of help to the treatment of COVID-19 patients in intensive care, as well as for other acute respiratory diseases. Extremely large computations with billions of spatial unknowns were necessary to simulate the air flow all the way from the trachea to the small-scale alveolar structures, and to assess the interaction with the lung tissue.

From SuperMUC to SuperMUC-NG

After more than seven years of operation, SuperMUC Phase 1 was shut-down by the End of 2019, followed by the shut-down of SuperMUC Phase 2 in January 2020. Ramp-up of SuperMUC-NG started in June 2019, when the system was ranked at #8 in the TOP500, with general user operation starting in August 2019.

Already with SuperMUC Phase 1 and Phase 2, LRZ established the reliable and efficient warm-water cooling that has since been adopted in computing centres around the world. This was continuously improved, and SuperMUC-NG uses water at even higher temperatures, which increases overall efficiency and allows to re-use the excess heat in adsorption chillers.

For our HPC users, we continue to provide high quality application- and user support through mentoring and dedicated High Level Support Teams. These efforts will help our users to solve the most ambitious scientific problems.



Prof. Dr. Dieter Kranzlmüller, Chairman of the Board of Directors at the Leibniz Supercomputing Centre.

Outstanding scientific research

This book covers simulations carried out on our SuperMUC systems in the time-frame from July 2018 to June 2020. The articles show how the continuously growing capabilities of supercomputers as well as the processing of rapidly growing amounts of data form the basis for more precise and highly resolved simulations in many research areas, and how they help to advance science, and to contribute to the good of society.

Astrophysical simulations on LRZ's HPC system help to understand the physics of gravitational waves and supermassive black holes, and support the breakthrough observational discoveries of recent years. Several projects bring light to the formation of stars, the universe, and the dynamics of supernova explosions. One of the largest simulations of supersonic turbulence in the universe is reported here.

Traditionally, the field of Computational fluid dynamics (CFD) consumes the largest amount of computing time on SuperMUC. Not only technical applications like simulations for the improvement of combustion, for airborne high-speed vehicles, for preventing erosion by cavitation in valve chambers, or for innovative



helicopter rotor configurations have been performed, but also the embedding of cutting-edge CFD solvers into applications of biomechanics. Other contributions in this book show how numerical simulations help to improve wind turbines, to optimize aircraft wings, or to help reduce pollutants in gas turbines.

Through global warming, mountains and coastal areas are particularly affected by extreme, weather-driven disasters. Flash floods and storms are notoriously difficult to predict in areas with complex topography. Research to improve state-of-the-art Numerical Weather Prediction models was performed on SuperMUC. Furthermore, climate research computational projects help to plan the reforestation with the Great Green Wall for the Sahara and the Sahel zone, and help to deal with sustainable water management in semi-arid regions in South America.

A better understanding of large-scale geological effects, the simulation of earthquakes and the resulting tsunamis will help to identify similar seismological and geological regions, and will help to protect sensitive areas. High resolution gravity field modelling not only gives insight into Earth's interior, it also provides an important surface reference for oceanographic applications, such as the sea level rise, or the modelling of ocean currents.

Numerical simulations in particle physics, particularly in Quantum Chromodynamics (QCD), help to understand the interaction between quarks and gluons, which make up composite nucleons like protons and neutrons, the building blocks of our visible matter. Here, SuperMUC was integrated into the LHC-ATLAS production system, to simulate proton-proton collisions.

Chemistry, Material and Life Sciences on LRZ HPC Systems continued to gain strong momentum in the last years. Calculations on SuperMUC-NG comprise the blood flow simulation in the human body, where the simulation scaled to over 300,000 cores and produced multiple terabytes of data. Several other projects simulate the roles of receptors and membranes, or the underlying mechanisms for photosynthesis. Recent simulations for an innovative approach that helps with accelerated drug design for COVID-19 are also described. Simulations in material sciences of magnetic quantum spin systems may lead to new applications in data storage, as well as quantum computing technology.

Finally, new data-driven research becomes more and more important on LRZ systems. Several petabyte of remote sensing satellite data, data from geoinformation systems and social media were systematically combined to map 3D urban infrastructures and their evolution over time. The outcome creates a unique and consistent data set of urban settlements and will help stakeholders with political decisions.

We invite you to browse through the book and to follow the many references and links provided with each article.




Acknowledgements

We gratefully acknowledge the continued support of the Free State of Bavaria and the German Federal Ministry of Education and Research (BMBF), as well as the Bavarian Competence Network for Technical and Scientific High-Performance Computing (KONWIHR), the Gauss Centre for Supercomputing (GCS), the German Research Foundation (DFG), 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 value in these projects.

Garching near Munich, April 2021

*Peter Bastian
Dieter Kranzlmüller
Helmut Brüchele
Matthias Brehm
Gerald Mathias*

Astrophysics



SuperSILCC – How stellar feedback regulates

the interstellar medium

RESEARCH INSTITUTION

¹Leibniz Institut für Astrophysik (AIP), Potsdam

PRINCIPAL INVESTIGATOR

Philipp Girichidis¹

RESEARCHERS

Tim-Eric Rathjen², Thorsten Naab², Stefanie Walch³, Frantisek Dinnbier³, Daniel Seifried³, Richard Wunsch⁴

PROJECT PARTNERS

²Max Planck Institute for Astrophysics, Garching

³Physics Institute, University of Cologne

⁴Astronomical Institute, Academy of Sciences of the Czech Republic

SuperMUC Project ID: pn34ma (Gauss Large Scale project)

Introduction

The formation and evolution of galaxies is strongly connected to the stars within the galaxy, in particular to the formation of stars and the feedback that stars provide via radiation, winds, supernovae and cosmic rays. The important agent that regulates almost all processes from the formation of molecular clouds out of the diffuse gas, to the collapse of gas into stars is the interstellar medium (ISM). Once stars have formed the coupling of stellar feedback and related to that the resulting thermal and dynamical impact is again mediated by the ISM. The ISM thus determines the transitions between different forms of energies as well as the transfer of energy between different spatial scales. Investigating the ISM together with all relevant processes is therefore crucial to understand the life cycle of gas and ultimately the evolution of galaxies (see Figure 2).

SuperSILCC is a modification based on the SILCC project [1,2] and simulates the life cycle of gas in the ISM. Compared to the original project we include more physical processes and account for more stellar feedback mechanisms.

Methods

We use the Eulerian adaptive mesh refinement code FLASH [3] to solve the magneto-hydrodynamic (MHD) equations. The thermal state of the gas is accurately computed using a chemical network that follows non-equilibrium abundances of chemical species. Besides the thermal gas, cosmic rays are included as a relativistic fluid and dynamically coupled to the MHD equations. Star formation is implemented using Lagrangean sink particles, that form at the locations of gravitational collapse. The particles that represent star clusters can accrete gas from the surroundings, which is coupled to the star formation process inside the particles as a subgrid model. We keep track of individual massive stars in the cluster sinks and compute their feedback processes according to stellar evolution tracks. Radiation, stellar winds and




Figure 1: Star formation rate as function of gas surface density [4]. Only supernova feedback (run S) results in too high star formation rates. Including stellar winds (SW & SWC) suppresses the star formation rate. Additional radiation (SWR & SWRC) shows the best result compared to Milky Way observations.

supernovae are included as time-resolved individual processes that inject mass, momentum and energy into the ISM.

We simulate a representative patch of the galaxy using a stratified box with a size of $1 \times 1 \times 10$ kiloparsec in size. The box covers multiple molecular clouds and star forming regions, which provides necessary statistics. Most of the simulations are connected to present-day star formation in the Milky Way but also more extreme conditions with larger surface density up to ten times the local solar neighborhood conditions are probed. In order to resolve the impact of individual massive stars and supernovae we use a resolution of 4 parsec at the highest level of the adaptive mesh refinement. The resolution in the star-forming disc adaptively adjust to the changes of the density. Above a height of 1 kiloparsec, where no stars form, we allow for coarser resolutions.




Figure 2: Ray-tracing image of the gas structure of the interstellar medium. Resolving the density structure.

A typical simulation uses 50-100 Million gas cells and covers 100-150 star clusters. The typical lifetime of a star cluster is 40 Million years, so we evolve the system for at least 100 Million years to form a statistically relevant sample of star clusters. Typical simulation runs use approximately 3,000 cores and require a wall clock time of three months. Data is stored in the versatile HDF5 format. We write a plot file with typically 20 field variables every 0.1 Million years in order to follow the star cluster dynamics together with the feedback processes. The typical size of a plot file is about 4 GB. Checkpoint files, from which we can restart the simulation, store more than 50 data fields and are written with a wall clock time interval of a few hours. The total amount of data sums up to approximately 150 TB so far.

Results

The individual feedback processes are key in regulating the star formation rate and outflows. Supernova feedback only occurs at the end of a massive star's life and therefore takes at least several Million years. During that time the star clusters can grow too fast and the resulting star formation rate is perceptibly larger than the observed one (see Figure 1). Radiation and stellar winds are early stellar feedback and provide an efficient mechanism to prevent too efficient gas accretion and star formation. Stellar winds alone are sufficient to reduce the star formation rate to values comparable to those observed in the Milky Way. Even more efficient is stellar radiation, which leads to the best agreement between simulations and

local observations not only in terms of the star formation rate but also concerning the clustering of stars.

Ongoing Research / Outlook

Using SuperMUC-NG allowed us to combine all the relevant feedback mechanisms into one simulation setup and simulate the full cycle of the star formation process from parsec to kiloparsec scales over the relevant time scales. The most challenging part of the simulation setup was the memory management of the numerous physical quantities that need to be included and stored like the chemical composition, radiation and cosmic ray energy, and local radiation shielding. Follow-up projects will include higher resolution in the star cluster regions.

References and Links

- [1] Walch, Stefanie et al., 2015, MNRAS, 454, 238.
- [2] Girichidis, Philipp et al., 2016, MNRAS, 456, 3432.
- [3] Dubey, Anshu et al., 2008, Phys. Scr., T132.
- [4] Rathjen, Tim-Eric et al., submitted.

Cosmic structure formation on the past backwards

1 lightcone in the MillenniumTNG simulations

RESEARCH INSTITUTION

¹Max-Planck Institute for Astrophysics, Garching

PRINCIPAL INVESTIGATOR

Volker Springel¹

RESEARCHERS

Simon White¹, Lars Hernquist²

PROJECT PARTNERS

²Harvard-Smithsonian Center for Astrophysics, Harvard University

SuperMUC Project ID: pn34mo (Gauss Large Scale project)

Introduction

Ever since the discovery of cosmic large-scale structure, the clustering of galaxies has been recognized as one of the most important observational constraints in cosmology. However, galaxies are biased tracers of the matter density field, thus they only indirectly inform about the underlying cosmic matter distribution. Moreover, a precise understanding of this galaxy bias, as a function of scale, epoch, and galaxy type, is necessary in order to make optimum use of forthcoming cosmological surveys (e.g. DESI, eBOSS, J-PAS, PFS, LSST, WFIRST, and Euclid), in particular those that target dark energy. In fact, these new surveys drive the need for high-accuracy predictions in volumes comparable to the ones sampled by the observations ($\sim 70 \text{ Gpc}^3$ for the European Euclid satellite).

Hydrodynamical simulations of galaxy formation have sufficiently matured to enable a successful prediction of the build-up of the entire galaxy population starting from cosmological initial conditions. These simulations track the non-linearly coupled evolution of both baryons and dark matter, fully accounting for their mutual influence on each other and yielding rich

predictions for galaxy properties, the diffuse gas in the circumgalactic and intergalactic media, and cosmic dark matter clustering. In addition, they yield constraints on the impact of baryonic processes on important cosmological probes, such as weak gravitational lensing, the Sunyaev-Zeldovich effects, or the mass of galaxy clusters. However, such hydrodynamic simulations cannot be pushed to sufficiently large volumes, so they are not readily available to directly predict with sufficient statistical power the non-linear observables in the upcoming cosmological surveys.

In our “MillenniumTNG” project [1], we build on the highly successful hydrodynamical simulations IllustrisTNG [2], carried out in comparatively small volumes, and combine it with the much larger volumes reached by the iconic dark matter-only Millennium simulations of a decade ago, thereby directly linking the need for ultra-large simulation volumes to the progress with hydrodynamical models.

We achieve this through a special, two-fold simulation strategy, consisting of a flagship hydrodynamical simulation in the Millennium’s original 740 Mpc volume, and a flagship dark-matter only simulation with a trillion particles in a 1.8 Gpc volume. The link between the two is established through semi-analytic galaxy formation models calibrated against the full hydro results, and by subsequently applying them to populate a seamless lightcone output produced by the large N-body calculation on the fly. In addition, we use a variance suppression technique to boost the effective volume probed by the simulations into the regime required for the upcoming surveys. This approach aims to achieve for the first time the creation of physically motivated galaxy formation models in the required volumes for the upcoming surveys.

Results and Methods

For this project, we have engaged in extensive code-development, producing the new GADGET-4 code, i.e. a new major version of our GADGET code, which is one of the most widely used cosmological codes in the field. GADGET-4 now uses a new approach to shared-




Figure 1: Power spectrum convergence of our runs carried out with GADGET-4. In each case, the shot-noise (which is subtracted) is plotted as a grey inclined line. The linear theory power spectrum at $z=0$ is shown as a thin solid line. Convergence is excellent. Also, the baryonic acoustic oscillations at $k \sim 0.1 \text{ h Mpc}^{-1}$ are represented accurately.




Figure 2: A small wedge of thickness 22 Mpc excised from the lightcone of MillenniumTNG, extending out to $z=2$, corresponding to a comoving distance of $\sim 5,410$ Mpc. The insets show successive factor of 10 zooms.

memory hybrid parallelization based on MPI-3, where all MPI ranks on a node synchronize their work through shared memory access based on the C++11 memory model, bypassing the MPI stack. One MPI task per shared memory node is set aside to serve incoming communication requests from other nodes, thereby establishing a highly efficient one-sided communication model. The new code also supports sophisticated lightcone output and new gravitational solvers that are both highly accurate and efficient. We have just made the new version publicly available [3].

The other code we use in this project is our moving-mesh code AREPO, which is particularly well adapted to cosmological hydrodynamical simulations of galaxy formation. We use it to run hydrodynamical twins of some of our GADGET-4 simulations, so that the impact of baryonic physics on cosmic structure formation can be studied in detail, and the semi-analytic galaxy formation code can be calibrated for these effects.

In Figure 1, we show measurements of the non-linear matter power spectrum today, illustrating the high dynamic range and excellent numerical convergence we achieve in this project. In Figure 2, we show a visualization of the past backwards lightcone of the dark matter distribution in a simulation with 80 billion particles carried out with GADGET-4 on a partition with 256 nodes, i.e. 12,288 cores. This simulation was aggressively squeezed into the available memory, and represents the maximum load per core that we can fit on SuperMUC-NG while still preserving all the most important information we want to retain for our comprehensive analysis. This includes, for example, detailed merger history trees, weak gravitational lensing shear maps, and the non-linear Rees-Sciama effect causing secondary fluctuations in the cosmic microwave background. The merger trees alone contain 20.6 billion objects that track galaxy growth over more than 13 billion years from high redshift to the present. The data volume for this simulation at the

end amounts to 285 TB. For the whole project, we produce several such simulations. Our final data volume will exceed 800 TB.

Our most expensive production run, which is still in progress at the time of this writing, is a hydrodynamic model in a 740 Mpc box with the AREPO code, using 2,560 SuperMUC-NG nodes with 122,880 cores. It will cost nearly 80 million core hours to completion, and the restart files alone amount to a gigantic size of 70 TB. The most difficult part for us in getting the simulation running was to stay in the memory envelope available on SuperMUC-NG. A side effect of the large memory need is that we needed to go to a very large partition size, which in turn required involved improvements in the scalability of the code, and the development of work-arounds for MPI stability issues for certain communication patterns.

Ongoing Research / Outlook

The MillenniumTNG project is still ongoing, with the most challenging simulations still underway. The ones completed are already redefining the state of the art in precision predictions for cosmic large-scale structure, and are now used in numerous scientific analysis projects. A main challenge in this project is the huge produced data volume. We can cope with it because we have inlined critical parts of the postprocessing right into the simulation code, such that the data that is output to disk can be processed in a comparatively convenient way. We expect that the simulation methodology developed here will also be instrumental for cosmological simulations of the future, which aim to cover an even larger range of scales.

References and Links

- [1] <https://wwwmpa.mpa-garching.mpg.de/mtng>
- [2] <https://www.tng-project.org>
- [3] Springel V., Pakmor R., Zier O., Reinecke M., 2020, MNRAS, submitted.
<https://wwwmpa.mpa-garching.mpg.de/gadget4>

Mergers of binary neutron stars: linking simulations

with multi-messenger observations

RESEARCH INSTITUTION

Institute for Theoretical Physics (ITP), Goethe University Frankfurt

PRINCIPAL INVESTIGATOR

Luciano Rezzolla

RESEARCHERS

Elias Most, Jens Papenfort, Lukas Weih, Matthias Hanauske, Antonios Nathanail

PROJECT PARTNERS

—

SuperMUC Project ID: pn56bi (Gauss Large Scale project)

Introduction

This is an ongoing project, with the goal to investigate the long-term evolution of a merging binary system of two neutron stars. The investigation conducted within this project is well aligned with the past research conducted by the Relastro group in Frankfurt [1]. It is motivated by the gravitational-wave detection GW170817 and its electromagnetic counterpart, the so-called kilonova. This kilonova signal is produced by the nuclear processes within the dense and neutron rich mass that is ejected during the merger. Since a lot of mass is ejected during the long-term postmerger evolution, it is crucial to investigate this part via state-of-the-art simulations in order to fully understand the observation.

Results and Methods

Such simulations require the solutions of a number of complex equations in order to account for all the physical processes that are relevant during the long-

term evolution: general relativity for describing the strong gravitational fields, magneto-hydrodynamics for modeling the fluid making up the component neutron stars and the magnetic fields connected with that fluid, and radiative transport in order to account for neutrino radiation produced by the extreme conditions during the merger. The latter has just recently been added via the implementation of the new Frankfurt Radiation Code (FRAC) [2]. It has been used for the successful simulation of spherically symmetric accretion problems in order to test and verify the implementation in a realistic high-energy astrophysical scenario (see Fig. 1).

Currently, FRAC is being coupled to the general-relativistic magneto-hydrodynamics code FIL, which is part of the Einstein Toolkit, a framework for simulations in astrophysics. This coupling will result in a code with all the necessary physics modules in order to simulate the long-term evolution of a binary neutron star merger.




Figure 1: Evolution of the temperature (top) and radiation energy density (bottom) for a spherically symmetric perturbed accretion problem using the new radiation code FRAC.




Figure 2: Results of a general-relativistic magneto-hydrodynamics simulation of a black hole-torus system, where matter falls into the black hole, develops instabilities and deploys a powerful jet. Shown is from left to right the magnetic field strength, the electric current and the magnetization.

Such a merger is also related to the production of a short gamma-ray burst. The characteristics of such a jet have also been investigated within pn56bi supplementing the above described investigation of long-term simulations. For this purpose, magneto-hydrodynamical simulations of a black hole-torus system have been performed on SuperMUC-NG. These simulations shed light on how a jet like the one observed as counterpart to the gravitational-wave detection GW170817 is formed and how it propagates [3].

Jets are not only important for neutron stars, but also for active galactic nuclei – the centers of galaxies, where powerful jets emerge from supermassive black holes. In any case, a crucial role in producing jets is played by magnetic fields. To this scope a series of simulations was performed within pn56bi in order to study the generation of current sheets through instabilities triggered by the magnetic field [4]. Figure 2 shows the results of an exemplary simulation in terms of the magnetic field strength (left), the current (middle), and the magnetization (right). It was found that so-called plasmoids produced by these current sheets are a plausible explanation for the observed variability in the jets from active galactic nuclei.

Ongoing Research / Outlook

Currently, the Relastro Group in Frankfurt works on the coupling of FIL, FRAC and the Einstein Toolkit, whose completion will allow the long-term evolution of the first general-relativistic magneto-hydrodynamics simulation with a self-consistent treatment for gravity as well as

radiation. Such a simulation is expected to consume 15 M core hours. Publication of results will then follow after the analysis of several Terabytes of data.

Additionally, to further increase the accessible parameter space of spinning binary configurations for future studies, a significant effort has been put into developing a new state-of-the-art elliptic solver based on the publicly available Kadath library to generate challenging initial conditions. The allocation pn56bi has been used to develop and test an MPI-parallized version of the binary initial conditions solver involving matrix inversions of extreme sizes. Building on top of optimal MKL implementations, the solver scales almost perfectly to ten thousand of cores in strong scaling scenarios. First applications of high spin initial data using this new framework are already under way and promise a better understanding of the importance of spin in realistic neutron star merger gravitational wave events.

References and Links

- [1] <https://relastro.uni-frankfurt.de/>
- [2] L. R. Weih et al., MNRAS, 495:2285, 2019.
- [3] A. Nathanail, MNRAS, 495:3780, 2020.
- [4] A. Nathanail, MNRAS, 495:1549, 2020.

Unveiling the emergence of the first galaxies

with state-of-the-art simulations

RESEARCH INSTITUTION

¹Max-Planck-Institut für Astrophysik, Garching bei München

PRINCIPAL INVESTIGATOR

Enrico Garaldi¹

RESEARCHERS

Rahul Kannan², Aaron Smith³, Volker Springel¹, Mark Vogelsberger³, Federico Marinacci⁴

PROJECT PARTNERS

²Center for Astrophysics | Harvard & Smithsonian, Cambridge

³Massachusetts Institute of Technology, Cambridge

⁴University of Bologna, Bologna

SuperMUC Project ID: pn56ku

Introduction

In its infancy, more than 13 billion years ago, the Universe went through a period of darkness, with no sources of light. Then, the first stars and galaxies formed and started to pour photons into their cosmic neighborhood, ionizing and heating up gas in the so-called inter-galactic medium (IGM) between galaxies. This is known as the Epoch of Reionization (EoR), and represents the last global phase transition of the Universe, radically altering the conditions in which the progenitors of contemporary galaxies formed. It also represents an evolutionary link between the smooth matter distribution produced by the Big Bang and the large-scale structure observed today.

Despite significant progress over the past decade, little is known about the sources of EoR and their effect on the IGM. However, upcoming facilities will unleash a flood of observations that will usher in a new era of cosmic reionization studies. Most notably, they will detect galaxies in the EoR, and unveil the physical conditions of the IGM at the time when the first stars were born.

To fully exploit these observations, theoretical models need to reach sufficient physical fidelity to make accurate predictions and aid the interpretation of the observational results. This requires a realistic description of galaxy formation in an evolving radiation field. However, the galaxy assembly problem is intractable analytically, and the coupled radiation-hydrodynamics (RHD) equations can be solved analytically only in highly-idealized setups. For these reasons, numerical simulations are an essential tool to investigate the EoR. Solving the coupled RHD equations is numerically challenging, but allow us to model reionization from first principles. In this project we run the most-advanced set of RHD simulations to date, designed to understand the EoR in all its different aspects: from the formation and properties of the first galaxies to their impact on the content of the Universe.




Figure 1: Evolution of selected physical quantities (reported in the top left of each panel) in our Medium RHD simulations. Time flows from left to right spanning approximately 800 million years.

Results and Methods

We have produced a set of state-of-the-art simulations of the EoR [1]. They combine for the first time: a realistic model for galaxy formation, cosmic dust evolution, a non-equilibrium primordial chemistry solver, stellar evolution including binary systems, and a self-consistent treatment of gravity, hydrodynamics, magnetic fields, and radiation transport on cosmological scales. These simulations are able to follow a region of the Universe with volume of 95.5 comoving Mpc, while resolving scales up to 20,000 times smaller (ap-

Name	RAM	Cores	CPU-h	Storage
Flagship*	92 TB	57600	30M	220 TB
Medium	11 TB	6730	3.5M	27.5 TB
Small	1.5 TB	816	400,000	3.5 TB

Table 1: Computational and storage requirements for the simulations run within this project. *The Flagship run is still ongoing, hence we provide projected values.

proximately 5 kpc). Specifically, we run a set of three simulations with different resolutions, in order to understand the numerical convergence of our results. Table 1 summarizes the computational and storage requirements for each of them, as well as the core-h invested in each of them. Science results will be taken primarily from the flagship run. Overall, our simulations created approximately 11 million files of science output.

Figure 1 presents an overview of the results from the Medium run. It shows the evolution of a selection of simulated physical quantities, with time flowing from left to right and spanning approximately 800 million years. The top panel shows the formation of dark matter structures, which triggers the collapse of gas (second panel from the top) into galaxies. Stars formed inside this galaxies produce photons that ionize the surrounding neutral gas (the third panel shows the gas ionized fraction) and, at the same time, heat it up by 20,000 K (the bottom panel shows the gas temperature). This process alters forever the stage where cosmic evolution takes place, transforming a cold and neutral Universe into a hot and ionized one. Figure 2 shows the details of one of the 1.3 million simulated galaxy in the Medium run. Not only the gas collapse (left panel) is captured in exquisite details, but the resolution is so large to allow individual galactic discs to be simulated (right panel).

Code developments

In order to achieve this result, we had to overcome a number of technical challenges. First and foremost, the data volume of these simulations required us to use an unprecedented (for the code employed) number of processors. Hence, a significant part of the code had to be optimized in order to optimally scale on almost 60,000 computing tasks, corresponding approximately to 1/6 of the entire SuperMUC-NG machine. Our code, AREPO [2], has been developed explicitly for massively-parallel simulations, and employs a pure-MPI approach. The hydrodynamical equations are solved on an unstructured mesh defined as the Voronoi tessellation of a set of mesh-generating points. This produces a naturally-adapting Galilean-invariant mesh, on which the fluid equations are solved using a second-order-accurate, un-split, finite-volume Godunov scheme employing an HLLD Riemann solver. Two important code developments were triggered by the constraints on the memory-per-core available. In order to run on 60,000 computing tasks, we optimized the memory usage developing a hybrid




Figure 2: Large-scale (left) gas distribution at the end of the Medium simulation, and zoom-in on an individual galaxy (right), showing the dynamical range covered. The zoom factor relative to the box is shown in the top right.

shared/distributed memory approach, which removes duplicated information within the same computing node, at the price of a more complex algorithm. Additionally, in order to fit our simulation in a reasonable number of computing tasks, we carefully moved double-precision variables to single-precision, implementing safety checks in the code to prevent under/overflows. Finally, our simulations required us to integrate and connect physical computing modules that were never employed together before. In particular, the self-consistent radiative transport, the dust evolution module, and the magneto-hydrodynamics solver had to be tweaked to smoothly co-operate during the simulation.

Ongoing Research / Outlook

We are still in the process of completing the simulations described above. For this, SuperMUC-NG is vital since it allows us to employ vast numerical resources that would be otherwise unavailable to us. In particular, the possibility to get access to the machine without a EU-wide competition allowed us to significantly speed up the project timeline. Additionally, the improved performance of SuperMUC-NG over its predecessor SuperMUC-phase2 allowed us to simulate a three-times larger volume than it would have been possible with the same amount of computing time, improving the statistical power of our predictions. Our simulations mark the founding stone for a large program of investigation of the early Universe, following two main branches. On the one hand, we will select interesting simulated galaxies and perform a numerical zoom-in on these objects, in order to simulate to even-greater details their functioning. On the other hand, we will explore the still-unconstrained physics of the early Universe by running simulations with additional/different physical processes at play. For both these numerical quests, large HPC systems are vital. The amount of memory and computing cores necessary to carry out such simulations is unprecedented, because so is the quality of current and forthcoming observations.

References and Links

- [1] wwwmpa.mpa-garching.mpg.de/~egaraldi/thesan.html
 [2] Springel, V. - MNRAS 401, 791 (2010).

Turbulence in astrophysical plasmas

1 from fluid to electron scales

RESEARCH INSTITUTION

¹KU Leuven Center for Plasma Astrophysics

PRINCIPAL INVESTIGATOR

Giovanni Lapenta¹

RESEARCHERS

Francesco Pucci¹, Giuseppe Arró¹, Alfredo Micera¹, Jorge Amaya¹, Francesco Pecora², Mariangela Viviani², Sergio Servidio²

PROJECT PARTNERS

²Università della Calabria

SuperMUC Project ID: pn56ye (PRACE project)

Introduction

Similar to what happens for regular fluids on Earth, like air or water, astrophysical plasmas in space are found to behave in a turbulent way, developing vorticose motions over a broad range of scales. Due to its mathematical complexity, turbulence has been studied in the last decades by means of numerical simulations, but even with the computational power available nowadays it remains difficult to be fully described in numerical experiments. The difficult part of it relies on the need of resolving a broad range of scales that all together constitute the phenomenology of turbulence. In this project we have conducted cutting edge numerical simulations of plasma turbulence using a numerical method that overcome several limitations in the description of plasma motions from large to small scales.

Methodology

The method used is called semi-implicit Particle In Cell (PIC) method and is implemented on an open source code named iPic3D (<https://github.com/CmPA/iPic3D>). In the PIC method the plasma species, ions and electrons, are represented by macro-particles that shape and interact with a pre-existing and evolving electromagnetic field. The advantage of the semi-implicit method consists in the possibility of resolving a large scale system, i. e. a system at fluid scale, with a moderate computational cost with respect to other methods, while retaining in part small scales effects, especially due to electrons.

Results

Fields and particles are the two main characters of plasma physics and turbulence is the physical phenomenon that allows energy to be transferred from one character to another, i.e. to cause the transformation of electromagnetic energy into particle kinetic energy and viceversa. The computational time awarded by PRACE on SuperMUC-NG at GCS@LRZ has been used to study two fundamental processes of energy transfer in plasmas in two kinds of numerical

experiments: particle energization and transport in decaying turbulence and magnetic field generation in plasma dynamo.

Particle energization and transport in decaying turbulence

Decaying plasma turbulence has been simulated in three dimensions with a total number of 54 billion macroparticles per species (ions and electrons). In order to study particle transport and acceleration properties, a running time diagnostic has been implemented to store data of a reduced set of test particles. Figure 1 shows a three dimensional rendering of the simulation box. Orange and blue regions are regions of intense current (current sheets), while selected trajectories of ions and electrons are drawn in green and violet, respectively. Kinetic effects due to the presence of turbulence are observed for both species in terms of acceleration and diffusion. For instance, the helicoidal ions trajectories in green change their shapes along the ion motion due to energization produced by turbulence. At the same time, couples of particles that are found to be close to each other at a certain time of the evolution are observed to depart explosively in time due to the presence of coherent structures produced by turbulence, such as current sheets. The kinetic phenomena here observed are the base of particle energization and transport due to plasma turbulence in space. Statistical analysis on particle trajectory and energy are now carried to address the open problems of particle transport not only in the solar system, but also in our galaxy and in the Universe.

Magnetic field generation in plasma dynamo

The numerical experiment has been initialized with a plasma vortex as big as the simulation domain and zero initial magnetic field. Around 67 billion particles per species have been used. Due to velocity shears, turbulence is rapidly triggered inside the simulation domain, yielding to the generation of a magnetic field. Figure 2 (left) shows the simulation domain. The yellow and red fog represent the electron velocity magnitude, where red are those regions with higher electron speed and yellow those with a weaker electron speed.




Figure 1: Three dimensional rendering of plasma decaying turbulence. Ion and electron trajectories are shown in green and violet, respectively, and current sheets with oppositely directed currents are in blue and orange, respectively.

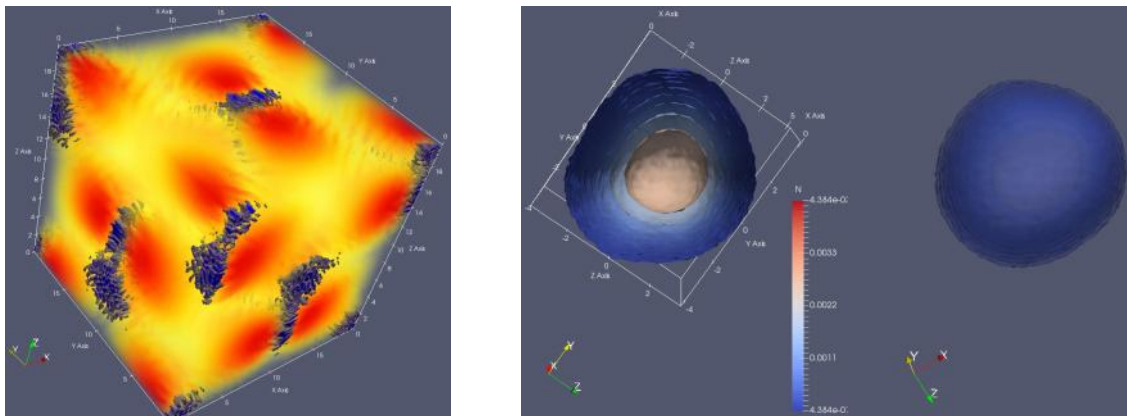


Figure 2: Three dimensional rendering of plasma dynamo simulation (left) and electron VDF (right). In the cubic box (left) the simulation domain is shown; yellow and red fog represent regions of lower and higher electron speed, respectively, and blue regions correspond to the dynamo-generated magnetic field. The electron VDFs (right) show particle counts as a function of the particle speed in the velocity space; a 3D rendering is shown on the right, while a 2D slice in the v_y - v_z plane is shown of the left.

The blue patches are regions where a magnetic field seed has been generated. Interestingly, the blue patches are located in between regions of high speed electron motions. Figure 2 (right) shows that the electron velocity distribution functions (VDFs) taken in one of the regions where a seed magnetic field is present are not spherical (Maxwellian), implying that local non thermodynamic equilibrium is one of the features of the plasma dynamo process. Further analysis on the simulation results will allow a clearer understanding of the process of plasma dynamo with an unprecedented resolution of both large and kinetic scale processes.

Ongoing Research / Outlook

New Prace Tier-0 project "AIDASpace – Artificial Intelligence Data Analysis of electron-scale turbulence in the heliosphere".

References

- [1] Goldman, M.V., Newman, D.L., Eastwood, J.P., Lapenta, G., J. Geophys. Res. Space Phys., 125 (12), 2020. doi: 10.1029/2020JA028340
- [2] Lapenta, G., Berchem, J., El Alaoui, M., Walker, R., PRL, 125 (22), 2020. doi: 10.1103/PhysRevLett.125.225101
- [3] Micera, A., Zhukov, A.N., Lopez, R.A., Innocenti, M.E., Lazar, M., Boella, E., Lapenta, G., Astrophys. J. Lett., 903 (1), 2020. doi: 10.3847/2041-8213/abc0e8
- [4] Arró, G., Califano, F., Lapenta, G., Astron. Astrophys., 642, 2020. doi: 10.1051/0004-6361/202038696
- [5] Lapenta, G., Zhukov, A., van Driel-Gesztelyi, L., Sol Phys, 295 (7), 2020. doi: 10.1007/s11207-020-01670-8
- [6] Micera, A., Boella, E., Zhukov, A.N., Shaaban, S.M., Lazar, M., Lapenta, G., Astrophys. J., 893 (2), 2020. doi: 10.3847/1538-4357/ab7faa
- [7] Lapenta, G., El Alaoui, M., Berchem, J., Walker, R., J. Geophys. Res. Space Phys., 125 (3), 2020. doi: 10.1029/2019JA027276
- [8] Divin, A., Deca, J., Eriksson, A., Henri, P., Lapenta, G., Olshevsky, V., Markidis, S., Astrophys. J. Lett., 889 (2), 2020. doi: 10.3847/2041-8213/ab6662
- [9] Dupuis, R., Goldman, M.V., Newman, D.L., Amaya, J., Lapenta, G. (2020). Characterizing magnetic reconnection regions using Gaussian mixture models on particle velocity distributions. The Astrophysical Journal: an international review of astronomy and astronomical physics, 889 (1), Art.No. 22. doi: 10.3847/1538-4357/ab5524 Open Access
- [10] Lapenta, G., Pucci, F., Goldman, M.V., Newman, D.L., Astrophys. J., 888 (2), 2020. doi: 10.3847/1538-4357/ab5a86

Gravitational-wave and electromagnetic signals from neutron star collisions

RESEARCH INSTITUTION

¹Friedrich-Schiller-Universität Jena

PRINCIPAL INVESTIGATOR

Sebastiano Bernuzzi¹, Bernd Bruggmann¹

RESEARCHERS

Matteo Breschi¹, Tim Dietrich², Vsevolod Nedora¹, Nestor Ortiz¹, Albino Perego³, Vivek Chaurasia¹, David Radice⁴

PROJECT PARTNERS

²Nikhef, Amsterdam

³Istituto Nazionale Fisica Nucleare, Milano Bicocca

⁴Princeton, USA

SuperMUC Project ID: pn56zo (Gauss Large Scale project)

Introduction

This project supported the activity of the **Computational Relativity (CoRe) collaboration** [1] whose goal is to model the gravitational waves (GWs) and the electromagnetic (EM) signals from the coalescence of binary neutron star (NS) mergers using first-principles, 4D, nonlinear simulations in general relativity. Binary neutron star mergers are unique astrophysical laboratories to investigate the **unknown matter equation of state (EOS) at extreme densities**, and they are connected to the **production of heavy elements in the Universe** (via r-process nucleosynthesis in the mass ejecta) and to some of the most energetic EM transients. Our simulations are crucial to interpret current and future observations of the multimessenger signals from this phenomenon.

This project addressed some of challenges posed by the LIGO-Virgo observations of the binary neutron star signals GW170817 and GW190425. We simulated 50 binaries to compute the gravitational waves merger signals and analyzed the strong-field dynamics of the remnant. Several binary properties and new physical processes were explored for the first time: spin, eccentricity, large mass ratios, microphysical EOS, effects of

neutrino transport and turbulent viscosity. The simulations' data product extends the CoRe database of about 30%. The latter currently represents the **largest public collection of waveform and mass ejecta** available to the gravitational-wave and astrophysics communities.

Results and Methods

The key feature of our simulations is the detailed treatment of general relativity and dense matter using adaptive mesh refinement techniques to span the multiple scales of the problem. Our largest simulations run efficiently up to 2,000-4,000 cores for more than 2,000,000 explicit timesteps in order to capture the merger dynamics from the orbital phase to the hydrodynamical and viscous timescales of the remnant. Moreover, these simulations have to be repeated at multiple resolutions to obtain error estimates on the final data, and for several parameters of the binary (masses, EOS, etc). The data generated from one simulation are snapshots of order of several TB, which are then reduced in (parallel) postprocessing steps. The simulations performed in this project used 75(+7.5)M core hours and represent **one of the latest major simulation effort in the field of numerical relativity and GW astronomy**.

The simulations conducted under pn56zo set a milestone for our research on GW modeling. We produced **high-quality gravitational waveform** at reduced eccentricity and for binaries with spins. These waveforms are used for the development of templates for LIGO-Virgo analysis. The inclusion of spin effects makes these waveforms unique data that only our codes can currently produce. We designed the **first complete gravitational-wave spectrum** model by combining numerical relativity data with semianalytical templates [2], Fig. 1. Our model will be relevant for advanced and third generation detectors observations. For example, the detection of a single GWs from a merger remnant will allow us to deliver a measurement of the NS minimal radius within the kilometer precision, thus providing unique constraints to the extreme density EOS.

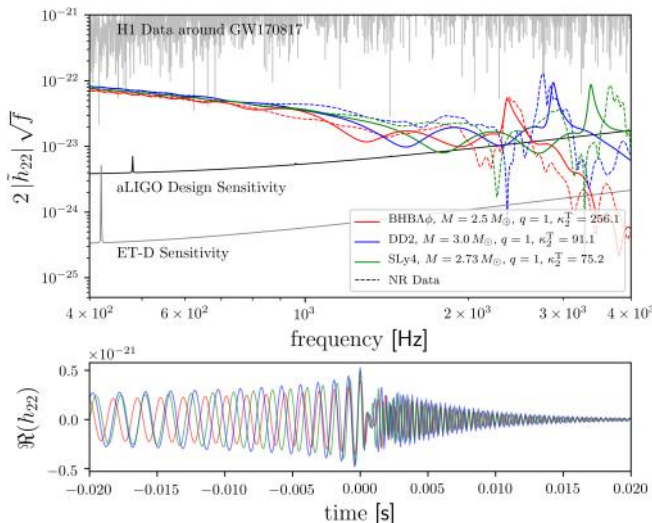


Figure 1: Waveform models for the complete spectrum of binary neutron star were developed in the projects.




Figure 2: Angular momentum transport in the merger and spiral-waves (left). Spiral-wave wind mass ejecta and kilonova lightcurves predicted from our simulations (right). Simulations are compared to observations of AT2017gfo.

Under pn56zo, we performed the first large set of BNS simulations including a general-relativistic scheme for turbulent viscosity induced by magnetic fields and a new neutrino transport scheme. We studied in detail the thermodynamics and geometrical properties of the remnant, and the winds that are emitted on **timescales of hundreds of milliseconds after the merger**.

Our simulations indicate that the bright, early-time (optical and UV) kilonova transient AT2017gfo associated to GW170817 could be explained as the signature of weak r-process nucleosynthesis in a fast expanding wind. The latter is generated by spiral density waves propagating from the remnant to the outer disk, and can be identified only with ab-initio simulations [3], Fig.2. We demonstrated that the r-process nucleosynthesis in mass ejecta with speed, temperature and composition as found in the simulations, **can robustly account for all the heavy elements** from mass number 75 to actinides and is compatible with solar abundances. By exploring five different microphysical EOS and several binary mass ratios, we concluded that this is a very general feature in binary neutron star mergers and does not require fine tuning. Indeed, the merger outcome of binaries with comparable masses and sufficiently stiff EOS (as compatible with GW170817) is a remnant NS (instead of a black hole) with lifetime of several rotational periods during which the wind can develop.

Mergers with rapid black formation, like GW190425, are particularly interesting as they provide us a connection to the maximum NS mass. Using the simulations of pn56zo we developed an inference method for LIGO-Virgo observations that, for the first time,




Figure 4: Image from the official LIGO-Virgo outreach material for the announcement of GW190425 (January 2020).




Figure 3: Mass density of the remnant disc and black hole of an accretion-induced prompt collapse merger. The orange surface is the apparent horizon of the black hole.

quantifies the probability of prompt black hole formation of the merger remnant. Applied to GW170817 and GW190425, the method supported the current astrophysical interpretation with a rigorous Bayesian inference analysis.

Using the resources of pn56zo we simulated, for the first time, mergers of binaries with mass ratio ~ 1.67 - 1.8 with microphysics. We discovered that the remnant undergoes **accretion-induced prompt collapse** [4]. In these mergers, the tidal disruption of the companion and its accretion onto the primary star determine prompt black hole formation. The tidal disruption event results in a black hole remnant with a massive and neutron-rich disc around. This is very different from the equal-masses mergers, for which there is no significant disc around the black hole, Fig.3. Challenging a common belief, our simulations showed that these accretion-induced prompt-collapse mergers can power bright electromagnetic counterparts. We predicted that the mass ejected during tidal disruption can power a red, bright and temporally extended kilonova emission. The peculiar feature of this emission might be observed in the near future and can help to constrain the binary mass ratio from future multimessenger observations.

Ongoing Research / Outlook

The HPC resources provided by LRZ covered the entire computational resources for the research conducted by the Jena group in 2018-2019 and the largest fraction of computational resources of the CoRe collaboration. We published 7 peer-review papers based on these data solely and are further developing projects using them. The project has obtained an extension in 2021 (50M core hours). We publicly release the final data product of our simulations on the CoRe website [1] or on Zenodo.org. We also produced a number of visualizations that can be found on the CoRe YouTube channel [5]. Visualization of simulated data have also been used for the **official LIGO-Virgo outreach material** for announcement of GW190425 (January 2020, see Figure 4).

References and Links

- [1] <http://www.computational-relativity.org/>
- [2] Breschi et al Phys. Rev., D100, no. 10, 104029, 2019.
- [3] Nedora et al Astrophys. J., 886, no. 2, L30, 2019.
- [4] Bernuzzi et al, MNRAS, 497, Issue 2, 2020.
- [5] <https://www.youtube.com/channel/UChmn-JGNa9mfY5H5938jnjg>

Hunting down the cause of solar magnetism

RESEARCH INSTITUTION

¹Aalto University, Department of Computer Science, Astrominformatics Group, Finland

PRINCIPAL INVESTIGATOR

Maarit Käpylä^{1,2,3}

RESEARCHERS

Lucia Duarte², Johannes Pekkilä¹, Ameya Prabhu², Matthias Rheinhardt¹, Mariangela Viviani², Jörn Warnecke²

PROJECT PARTNERS

²Max Planck Institute for Solar System Research, SOLSTAR group, Germany

³Nordita, Stockholm, Sweden

SuperMUC Project ID: pn98qu (PRACE project)

Introduction

The research of the groups of Astrominformatics in Aalto University, Finland [1], and SOLSTAR at the Max Planck Institute for Solar System Research in Göttingen, Germany [2], focuses on understanding solar and stellar dynamos. We try to achieve this by combining high-resolution, local and global numerical modelling with long-term observations. This is a challenging task: even with state of the art computational methods and resources, the stellar parameter regime remains unattainable. Therefore, we are just able to simulate “tar-stars”, namely models in which diffusivities are raised several orders of magnitude from their real values, to guarantee numerical feasibility and stability. This work is part of the PRACE project Access - Call 20 INTERDYNs, which has been granted 57M core hours on SuperMUC-NG. Our goal is to relax some approximations, in order to simulate more realistic systems, and try to connect the results with theoretical predictions and state-of-the-art observations.

Convective turbulence, together with large-scale flows, such as differential rotation (Ω -effect), plays a key role in generating and shaping the magnetic fields observed in the Sun and other stars. The most important turbulent effect is the α -effect, describing collective inductive action arising from cyclonic turbulence. The α - and Ω -effects are the two prominent generators of large scale stellar/solar magnetic fields. Turbulent effects escape observational efforts even with the largest existing and planned observational infrastructures. In theory, these effects are parametrised by transport coefficients which collectively describe the effects of turbulence without having the need to resolve smaller scales, hence the numerical determination of these coefficients which describe them is of utmost importance. For their measurement, we employ the test-field method (TFM). In its standard form, it can be used to measure the turbulent transport coefficients in the limit where a primary magnetic background turbulence vanishes (low resolution regime). If the background magnetic turbulence is present, i.e., in the high resolution regime, the non-

linear (NL) form of TFM is necessary. This regime is relevant for the Sun, which most likely exhibits vigorous small-scale dynamo action, generating magnetic background turbulence. Moreover, the TFM allows us to measure the full α tensor, therefore adding a considerable refinement to the standard theory, which models the α -effect merely by a scalar quantity.

Results and Methods

We use the Pencil Code [3], a highly modular code, to solve the fully compressible equations of magneto-hydrodynamics (MHD). The code employs a sixth-order, central finite differences scheme for spatial discretization and a 3rd-order Runge-Kutta time-integration scheme. The chosen scheme makes the code highly scalable and adaptable. To maintain the magnetic field divergence-free, the code solves for the magnetic vector potential. The output files are written out quite often in order to re-start from crashes or node-failures. The parallelization is implemented using MPI and allows the Code to scale up to 100.000 cores. The data analysis can be performed on the fly or as post-processing using Python or IDL (Interactive Data Language).

A list of disk storage, number of cores and total wall-time used for the simulations described in this report are shown in Table 1. The total short term storage on SuperMUC-NG for this part of the project will be 50TB.

Run	Grid size	# CPU	Total Walltime [h]	Total memory required
A1	256x512x256	512	1,440	17 GB
A1TF	256x512x256	512	969	83 GB
A2	512x1,024x512	2,048	1,440	137 GB
A2TFNL	512x1,024x512	4,096	969	1.52 TB
A3	1,024x2,048x1,024	4,096	1,440	1.10 TB
A3TFNL	1,024x2,048x1,024	8,192	969	12.20 TB
A4	2,048x4,096x2,048	8,192	1,632	8.90 TB

Table 1: Summary of the employed resources. Boldface: runs discussed in this report.

Even though most of the simulations are still ongoing, we have already achieved interesting results. Our new high-resolution run (A1) shows a different behaviour than an earlier low-resolution run (M0, [4]). Besides the two times higher resolution of A1, the difference between the two consists in the way heat conduction is modeled. To describe radiative heat transfer, we use in both runs the diffusion approximation. In Run M0, we prescribed a profile for the radiative heat conductivity, while in the PRACE project runs we use a Kramers-like opacity law, in which the heat conductivity depends on density and temperature. The latter, a more physically sound prescription, allows for the development of a layer which is convectively stable in the traditional sense, but in which the transport of heat is still upward-directed. This can have consequences for the properties of the flow but also for the magnetic field evolution [5]. We highlight one such difference in Figure 1. In the meridional cuts we show the component $\alpha_{\varphi\varphi}$ for run M0 (left), versus that of our new run A1 (right). The presence of negative/positive values in the northern/southern hemisphere is crucial for obtaining the equatorward propagation of the sunspot-producing zones observed during the solar cycle. In M0, the α -effect was of the wrong sign in most of the CZ, but with the improved description of heat conduction, the region where the effect is of the correct sign has grown significantly. In the new run (A1), we now observe a reversed migration direction of the magnetic field, which shows that the changes in the α -effect profile are significant, see Figure 2.

Ongoing Research / Outlook

We presented here the first test-field measurements from our higher-resolution runs with improved heat conduction description. They indicate significant changes in the profiles of the most crucial inductive effect related to solar and stellar dynamo mechanisms, which already has important consequences for understanding these dynamos. Our even higher resolution




Figure 1: A volume rendering of global convective dynamo simulations. Colors on the two spherical surfaces represent the radial velocity close to the surface and the bottom of the simulation domain. The meridional cuts show the $\alpha_{\varphi\varphi}$ component of the α tensor for a previous low-resolution run (M0) and the new high-resolution run (A1), red indicating positive, and blue negative values. The background image is from an instrument onboard the Solar Dynamics Observatory.

runs (A2-A4) currently undertaken, will bring us into an even more turbulent regime, where magnetic background fluctuations are generated by small-scale dynamo action. To measure the turbulent effects from these runs, we will apply the novel compressible test-field method. Such analysis will allow us to study, for the first time, the interaction of small- and large-scale dynamos in a quantitative way.

References and Links

- [1] <https://www.aalto.fi/en/department-of-computer-science/astroinformatics>
- [2] <https://www.mps.mpg.de/solar-stellar-magnetic-activity>
- [3] <http://pencil-code.nordita.org>
- [4] Warnecke J. et al. 2018, A&A, 609, A51.
- [5] Käpylä P. et al. 2019, GAFA, 113, 149-183.




Figure 2: Time-latitude diagram from run A1, showing the azimuthally averaged azimuthal magnetic field component near the surface as function of time. In the Sun, a prominent pattern of butterfly wings, formed by the sunspot locations migrating towards the equator during the solar cycle, is visible, but reproducing this pattern has turned out to be very challenging from direct numerical simulations. Also in our new Run (A1), such a pattern is mostly absent. The physically-motivated heat conduction was seen to strongly influence it.

Long-term evolutions of neutron-star mergers:

mass ejection, instabilities and EM counterparts

RESEARCH INSTITUTION

Institute for Theoretical Physics (ITP), Goethe University Frankfurt

PRINCIPAL INVESTIGATOR

Luciano Rezzolla

RESEARCHERS

Elias Most, Jens Papenfort, Lukas Weih, Matthias Hanauske, Antonios Nathanail

PROJECT PARTNERS

—

SuperMUC Project ID: pr27ju

Introduction

The investigation conducted within this project is well aligned with the past research conducted by the Relastro group in Frankfurt [1] and can be summarized in two categories.

One is the simulation of accreting black holes, which aided the interpretation of the world-famous image of the EHT collaboration. This image of the black hole in the center of the galaxy M87 is the first of its kind. It allows to test gravity in the strong-field regime by comparing to simulations as the ones performed within this project.

The second focus was on binary neutron stars. The extreme densities encountered inside neutron stars allow the study of so-far unexplored physics that is not accessible via experiments on earth. Two points of interest during this project were the creation of quark-gluon plasma during the merger and how to identify it from gravitational waves or the impact of spin during the binary's inspiral.

Results and Methods

Whether we simulate accreting black holes or binary neutron star systems, both require the numerical solution of the equations of (magneto-)hydrodynamics coupled to the Einstein field equations describing gravity. The correct evolution of these equations can only be achieved using a number of sophisticated numerical and of course high-performance computing methods.

In order to understand the image captured by the EHT collaboration [2] simulations of a hot plasma around a supermassive black hole have been performed (see Fig. 1). The results of these simulations were

used to develop a database of synthetic images that were computed via ray-tracing and can be compared to the observation. Such a synthetic image is shown in Fig. 2, which is obtained after post-processing the simulation snapshot shown in Fig. 1. It can be seen that this image agrees very well with the famous image obtained by the EHT collaboration. Given that such a synthetic image is produced assuming Einstein's theory of general relativity to describe gravity, the agreement with the observation is yet another strong argument in favor of this theory.

Systems of binary neutron stars, on the other hand, allow to also test the fundamental behavior of matter under extreme conditions. For doing so we performed simulations with a special prescription of the so-called equation of state. This equation describes the change in pressure for a given density, which in turn is set by the interactions between elementary particles. More precisely, we allowed the matter to undergo a phase transition at high densities. During this transition ordinary hadronic matter (neutrons and protons) disintegrate and form a soup of quarks and gluons. This is the long hypothesized quark-gluon-plasma, which is thought to have existed during the early stages of the Big Bang. Using SuperMUC-NG we performed the first ever simulation with a realistic and temperature dependent equation of state [3] that includes such a transition. In a follow-up study and with an equation of state using different parameters for




Figure 1: Simulation snapshot of a black hole-torus system (red-yellow) emitting a jet (white-blue).

this phase transition, we could even perform simulations, where the two neutron stars merge and produce a stable hybrid star, i.e. a more massive and dense neutron star with a stable core consisting of quark matter [4], which is shown in Fig. 3.

During these simulations, we also computed the gravitational waves that are emitted due to way that the massive neutron stars curve the spacetime. In this way we could show that there is a clear signature in the gravitational-wave pattern that would allow us to distinguish from this signal alone whether a phase transition to quark-gluon plasma takes place or not. It is expected that 3rd generation gravitational-wave detectors will be able to pick up such a signal in the near future and thus potentially show the presence of the quark-gluon-plasma inside the merger remnant of a neutron star binary.

The above investigation heavily relied on previous work that resulted in today's routinely performed simulations of binary neutron star mergers. However, these simulations can still be improved by so-far unexplored physics and methods. During this project, we did investigate many of these venues such as higher-order methods for the solution of the general-relativistic magneto-hydrodynamics equations or the formulation of a new method for including radiative transport. We also investigated the impact of high spins on the mass that is ejected during the merger of two neutron stars [5]. Usually simulations start from non-spinning or slowly spinning neutron stars. Within this project, we found that the merger outcome, especially the dynamically ejected mass that later undergoes nucleosynthesis and yields the observable kilonova signal, is significantly different in the case of a merger of two highly spinning neutron stars.

Ongoing Research / Outlook

This project directly continues with our new allocation pn56bi (see report on page 18). Within this allocation we investigate yet another facet of binary neutron star

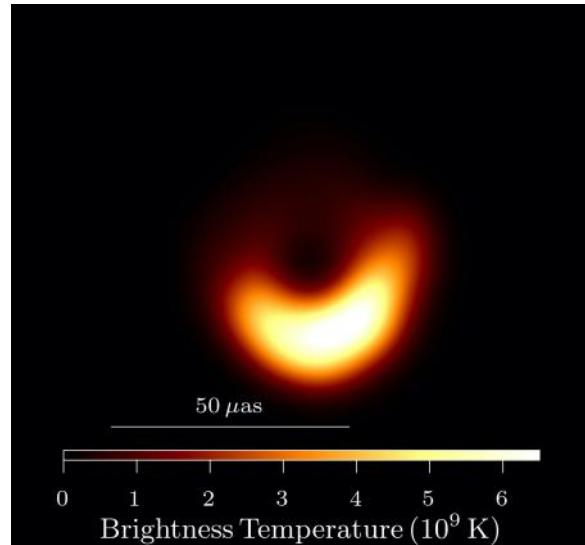


Figure 2: Example of a synthetic image of an accreting black hole produced by post-processing the results of a general-relativistic magneto-hydrodynamics simulation.

mergers that has not been explored within pr27ju. Specifically, we explore the long-term evolution of the merger remnant. During this stage radiative transport due to neutrinos has a considerable influence on the ejected mass. It will contribute to mass ejection via a neutrino-driven wind and alter the ejecta's composition, which will impact the observable kilonova signal. In interplay with the magnetic field evolution the radiation might also help to develop a short gamma-ray burst. The observations connected to the event GW170817 have proven that this phenomenon is connected to the merger of two neutron stars, but it could not yet been shown with simulations.

References and Links

- [1] <https://relastro.uni-frankfurt.de/>
- [2] EHT collaboration, *Astrophys. J. Lett.*, 875, L5 2019.
- [3] E. R. Most et. al, *Phys. Rev. Lett.*, 122:061101, 2019.
- [4] L. R. Weih et. al, *Phys. Rev. Lett.*, 124:171103, 2020.
- [5] E. R. Most et. al, , *Astrophys. J.*, 884:6, 2019.




Figure 3: Simulation snapshot of a stable hybrid star with a quark core shown in green.

Testing the limits of magnetized accretion

1 onto black holes

RESEARCH INSTITUTION

¹CEA-Saclay, Paris (France)

²Max Planck Institute for Astrophysics, Garching

PRINCIPAL INVESTIGATOR

Matteo Bugli^{1,2}

RESEARCHERS

Ewald Müller², Luca Del Zanna³, Niccolò Tomei³

PROJECT PARTNERS

³Università degli Studi di Firenze, Italy

SuperMUC Project ID: pr27qu

Introduction

Numerical simulations are an invaluable tool when it comes to investigate the complex dynamics that rule the behavior of **astrophysical plasmas**. This is even more the case if we consider the fundamental role that **General Relativity** has in the description of compact objects such as black holes and the magnetized hot accretion disks that are commonly assumed to orbit around them.

In the past decades our understanding of the physics behind accretion processes has vastly increased thanks to an always increasing number of numerical studies, but there are still several uncertainties that make challenging to draw quantitative (and sometimes qualitative) conclusions from them. This fact represents an important incentive for the improvement of numerical models, considering also the extraordinary observational discoveries made in the recent years by the **detection of gravitational waves** and the **first direct imaging of a supermassive black hole**.

Our project aims at assessing the degree of accuracy of state-of-the-art general relativistic magnetohydrodynamic (GRMHD) simulations of accretion disks around black holes using the astrophysical code ECHO [1]. We explored the impact of employing different algorithms and numerical implementations by participating to the **EHT code comparison project** [2] and investigated the limits of an ideal description of the astrophysical plasma by taking into account the **turbulent dissipation of the magnetic field** which is usually neglected.

Results and Methods

A significant part of the project focused on the impact of specific choices of numerical implementation, which we investigated within a first code **comparison study led by the Event Horizon Telescope collaboration** [2]. The 9 different codes

employed show a **remarkable agreement** between the different participating codes, with higher resolutions leading to significant decrease in deviations. The overall structure of the accretion flow is also quite similar among different results, and the same goes for the time variability of the computed light curves, displaying very similar spectra. This agreement was found between codes using different grid geometries (cartesian or spherical coordinates), programming languages (Fortran or C++) and interpolation/integration algorithms.

The role of **numerical dissipation** and the limits of an ideal description of the astrophysical plasma can also be very important, as a more efficient diffusion of the magnetic field can affect the global structure of the disk and its accretion properties. By investigating the consequences of assuming a **finite magnetic resistivity** in the plasma, we show that a weaker MHD turbulence can lead to the formation of **non-axisymmetric large-scale structures** in the disk which, for an ideal plasma, would otherwise quickly washed out [3] (see Figure 3). The necessary numerical models




Figure 1: Meridional slices showing rest mass density (left panel), magnetic to thermal pressure ratio (center) and magnetization (right) in logarithmic scale and code units [2].

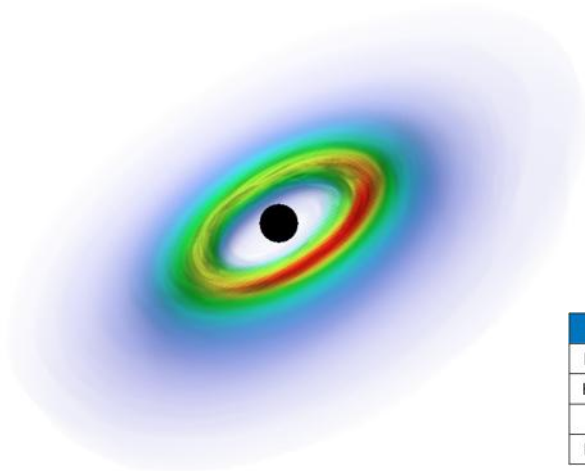


Figure 2: Volume rendering of a thick accretion disk (southern half) with finite resistivity.

	Grid points	Cores	Runtime	Core-hours	Output size
Low resolution EHT	192x96x96	560	85 h	0.05 M	0.32 GB
High resolution EHT	382x192x192	1,344	490 h	0.66 M	2.6 GB
Standard disk	256x256x256	6,840	80 h	0.55 M	1.4 GB
High resolution disk	512x512x512	51,840	170 h	8.81 M	47 GB

Table 1: Typical models run on SuperMUC Phase 2 and SuperMUC-NG.




Figure 3: Equatorial slices of the accretion disk for a hydrodynamic model (left panel) and a magnetized one (right) showing rest mass density in logarithmic scale normalized by its maximum [3].

make use of a generalized prescription for Ohm's law that accounts for the turbulent dissipation of the magnetic field due to unresolved small-scale motions.

All our three-dimensional simulations of thick magnetized disks were produced using **ECHO**, a finite-differences Fortran code solving the equations of magnetohydrodynamics in General Relativity. The **hybrid MPI-OMP parallelization scheme** employed in the code [4] allows for an efficient use of a **large number of cores** (ranging from 560 to 51,840 as shown in Table 1, depending on the model) on SuperMUC Phase 2 and SuperMUC-NG. Each simulation produced between 100 and 1000 HDF5 outputs (generated with a parallel HDF5-MPI scheme), requiring overall 20 TB of storage volume.

Ongoing Research / Outlook

Thanks to an efficient use of the HPC resources offered by LRZ we explored some important aspects of numerical models of accretion disks, such as the dependency of the results on the specific algorithm choices of the code employed and the effects of

magnetic diffusion. A second **code comparison project** is currently ongoing, which will perform a similar study to the one already published, but will mainly focus on more highly magnetized disk, which are known to trigger the formation of **strong polar jets**. This work will greatly help assessing to what extent numerical simulations can quantitatively reproduce several important dynamic properties of accreting astrophysical plasmas. On the other hand, we will continue our work on non-ideal plasmas using SuperMUC-NG, improving the description of the turbulent resistivity and testing the effects of a **turbulent dynamo** in the amplification of the magnetic field, rather than just its dissipation.

References and Links

- [1] M. Bugli, Proceedings of the 26th Euromicro International Conference on PDP, pages 674-681, 2018.
- [2] O. Porth et al., APJ Supplement Series, 243(2):26, 2019.
- [3] M. Bugli, J. Guilet, E. Müller, L. Del Zanna, N. Bucciantini and P. J. Montero, MNRAS, 475(1):108-120, 2018.
- [4] M. Bugli, L. Iapichino and F. Baruffa, Intel Parallel Universe, 34, 2018.

The sonic scale revealed by the world's largest supersonic turbulence simulation

RESEARCH INSTITUTION

¹Research School of Astronomy and Astrophysics, Australian National University

PRINCIPAL INVESTIGATOR

Christoph Federrath¹

RESEARCHERS

Ralf S. Klessen^{2,3}, Luigi Iapichino⁴, James R. Beattie¹

PROJECT PARTNERS

²Universität Heidelberg, Zentrum für Astronomie, Institut für Theoretische Astrophysik

³Universität Heidelberg, Interdisziplinäres Zentrum für Wissenschaftliches Rechnen


⁴Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften

SuperMUC Project ID: pr321o

Introduction

Understanding the physics of turbulence is crucial for many applications, including weather, industry, and astrophysics. In the interstellar medium (ISM), we believe that supersonic turbulence plays a crucial role in controlling the gas density and velocity structure, and ultimately the birth of stars. Here we present the world's largest simulation of supersonic turbulence to date. With a grid resolution of $10,048^3$ cells, this calculation allows us to determine the position and width of the *sonic scale* (λ_s)—the transition from supersonic to subsonic turbulence. The simulation is the first to simultaneously resolve the supersonic and subsonic cascade, $v(l) \propto l^p$, where we measure $p_{\text{sup}}=0.49$ and $p_{\text{sub}}=0.39$, respectively. Instead of a sharp transition, we find that the supersonic cascade transitions smoothly to the subsonic cascade over a factor of ~ 3 in scale.

Turbulence is ubiquitous on Earth, but it also plays a key role in structuring the ISM. However, there is a critical difference between astrophysical and terrestrial turbulence:



while the turbulent flows we encounter in everyday life are mostly incompressible and subsonic, many astrophysical flows are highly supersonic and compressible. For supersonic turbulence, energy dissipation occurs mainly in shocks rather than in vortex tubes, and the rate by which energy is transported through the scales is faster than in subsonic turbulence. The transition from one

regime into the other introduces a critical length scale, the *sonic scale* λ_s . For scales $l > \lambda_s$, the flow velocities exceed the speed of sound (Mach number >1) and are characterised by a complex network of interacting shocks that lead to highly dynamical, transient density fluctuations. In contrast, for $l < \lambda_s$, the flow is dominated by subsonic motions (Mach number <1) that set up the classical Kolmogorov cascade, including intermittency corrections. By definition, the sonic scale λ_s is located on scales where the Mach number is unity. Identifying λ_s and determining the transition region from supersonic to subsonic turbulent flows is crucial for advancing our understanding of turbulence.

The theoretical prediction for λ_s , previously used in models of ISM structure and star formation, is given by [3], $\lambda_s = \Phi_s L M^{-1/p_{\text{sup}}}$, where L is a characteristic scale, usually taken to be the driving scale of the turbulence in numerical calculations, or taken to be the diameter of a star-forming molecular cloud in observations. Theoretical calculations often make use of both interpretations and sometimes treat them as being identical. $M = \sigma_v / \Phi_s$ is the three-dimensional (3D) turbulent Mach number, i.e., the ratio of volume-weighted 3D velocity dispersion (σ_v) on scale L divided by the sound speed (c_s). The exponent $p_{\text{sup}} \approx 0.5$ appears in the velocity dispersion – size relation, $v(l) \propto l^p$, for supersonic turbulence, as found in observations of molecular clouds. The factor Φ_s encapsulates our lack of understanding of where exactly the sonic scale is located, and is determined and discussed in [4].

Results and Methods

Visualisations

In order to provide a sense of the dynamic range involved in the simulation, we show a visualisation in Fig. 1. This is a slice through the gas density ρ (scaled by the mean density ρ_0) seen in the x-y-plane at $z=0$, at the end of the simulation. We also provide an animation of the three-dimensional projected density field, which starts with a time evolution, and then rotates and zooms from the outer scale onto a dense filament formed in the simulation (see links below).

Figure 1: Slice through the gas density at the end of the supersonic turbulence simulation with $10,048^3$ grid cells, showing the full simulation domain and a zoom-in onto a dense region with multiple shock interactions. We believe that star clusters form preferentially at the intersection of such shocks. An animation of the three-dimensional projected density field highlighting the sonic scale is available (see links below).

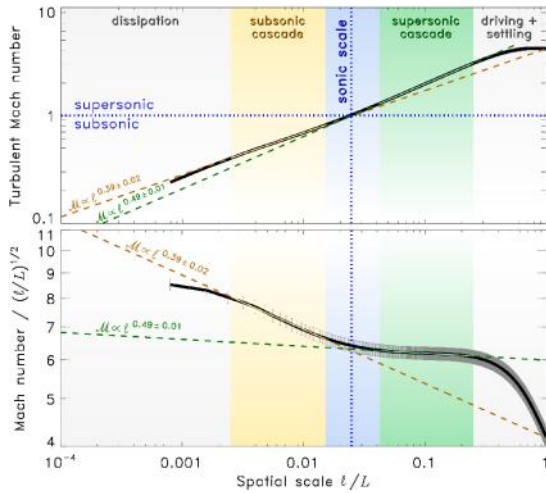


Figure 2: Top: Turbulent Mach number (M) as a function of scale l in units of the driving scale $L=L_{\text{box}}/2$, where L_{box} is the side length of the computational box. M is computed from the 2nd-order structure function (SF_2) of our simulation with $10,048^3$ grid cells. On the driving scale ($l=L$), we set the Mach number to $M=4.13$, shown by the flattening of SF_2 for $l/L \rightarrow 1$. The error bars on M and SF_2 were computed from the temporal fluctuations in the regime of fully-developed turbulence, $t=2-9t_{\text{turb}}$. The sonic scale λ_s is defined where $M(\lambda_s/L)=1$ (dotted lines). Bottom: Same as the top panel, but compensated by $(l/L)^{1/2}$ to enhance the visibility of the change in slope across λ_s . The dashed lines show power-law fits in the subsonic (gold) and supersonic (green) regimes of turbulence, with slopes of 0.39 and 0.49, respectively. This is the first time that the position and width of the sonic scale are determined. We find $\lambda_s/L=0.025$, i.e., a full-width-half-maximum (FWHM) in λ_s corresponding to a factor of 3 in scale.

Moreover, [2] present a visualisation of this simulation, which was amongst the finalists of the SC19 conference (see links below).

Basic Methods

In the framework of a Gauss Centre for Supercomputing Large-Scale Project, an allocation of ~50 million core hours has been used for this project on SuperMUC. The simulation code used for this project is FLASH, a public, modular grid-based hydrodynamical code for simulations of astrophysical flows. The parallelisation is based entirely on MPI. In the framework of the SuperMUC Phase 2 scale-out workshop, 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 (see previous Book in this series). In this way, the code runs with a factor of 4.1 less memory and 3.6 times faster than the public version of the code, and scales remarkably well up to the full size of SuperMUC [1,5]. The simulation was run on 65,536 compute cores and produced a total of 2 PB of data, split into 91 3D output files, spanning 9 turbulent crossing times (where one crossing time is defined as $t_{\text{turb}}=L/\sigma_v$). We note that it takes $2t_{\text{turb}}$ to reach stationarity (fully-developed turbulence), and thus, all results were averaged over a period of 7 turbulent crossing times, between 2–9 t_{turb} , i.e., over 71 3D snapshots of the $10,048^3$ data.

Measuring the Sonic Scale

In order to construct the Mach number M as a function of scale (l), from which we can directly determine the sonic scale (λ_s), we compute the total (transverse plus longitudinal) 2nd-order velocity structure function, and make sure that we obtain converged statistics [4].

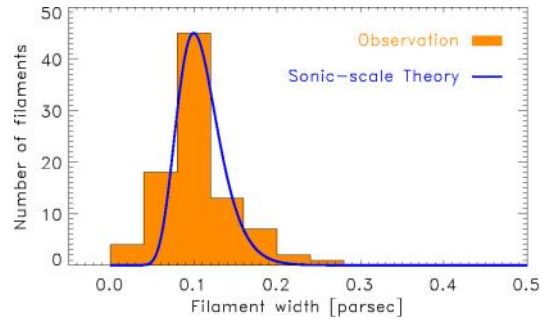


Figure 3: Filament width distribution measured from observations in the IC5146, the Aquila and Polaris molecular clouds shown as the orange histogram, together with the theoretical prediction based on the sonic scale (blue line). The theoretical prediction is not a fit; instead the peak position is set to 0.1 parsec and the lognormal standard deviation of 0.105 is determined by the width of the sonic-scale transition measured in Fig. 2.

Fig. 2 shows the time-averaged structure function (with error bars quantifying the 1σ time fluctuations). In particular, we show the Mach number $M=[SF_2/2c_s^2]^{1/2}$, constructed from SF_2 . The sonic scale is implicitly defined where $M=1$. We find $\lambda_s/L=0.025$ (vertical blue dotted line in Fig. 2). Power-law fits with $M(l) \propto l^p$ in the subsonic (gold shaded area in Fig. 2) and supersonic (green shaded area in Fig. 2) regimes yield scaling exponents of $p_{\text{sub}}=0.39$ and $p_{\text{sup}}=0.49$, respectively, consistent with the theoretical expectations in the two regimes of turbulence, i.e., $p_{\text{sub}} \sim 1/3$ for Kolmogorov turbulence and $p_{\text{sup}} \sim 1/2$ for supersonic turbulence. The subsonic slope is slightly steeper than the original Kolmogorov prediction of $p_{\text{sub}} \sim 1/3$, because of the necessary intermittency corrections for mildly compressible turbulence. We find that the transition region around the sonic scale (blue shaded area in Fig. 2) is smooth and spans a full width half maximum (FWHM) factor of 3 in l/L , which corresponds to a Gaussian standard deviation of 0.105. Thus, we measure $\lambda_s/L=0.025$.

Ongoing Research / Outlook

Many other fundamental aspects of turbulent flows can be studied with this large simulation, including probability distribution functions of key dynamic variables, the fractal dimension, etc. Fig. 3 shows the distribution of filament widths in molecular cloud observations, and how the sonic scale theory developed and tested here can provide an excellent match to the observations. Directly extracting the filament width from the simulation is one of the next steps of ongoing research. In addition, magnetic fields are ubiquitous in interstellar, star-forming clouds, so we want to extend this study to include magnetic fields. We have already started test simulations with magnetic fields, including visualisations [2]; see links below.

References and Links

- [1] Cielo, S., Iapichino, L., Baruffa, F., Bugli, M., & Federrath, C. 2020, Future Generation of Computer Systems, 12, arXiv:2002.08161.
- [2] Cielo, S., Iapichino, L., Günther, J., et al. 2019, SC19 visualization showcase finalist, arXiv:1910.07850.
- [3] Federrath, C., & Klessen, R. S. 2012, Astrophys. J., 761, 156.
- [4] Federrath, C., Klessen, R. S., Iapichino, L., & Beattie, J. R. 2020, Nature Astronomy, submitted (arXiv:2011.06238).
- [5] Hammer, N., Jamitzky, F., Satzger, H., et al. 2016, Advances in Parallel Computing, 27, 827.

Sonic-scale visualisation:
https://www.mso.anu.edu.au/~chfeder/pubs/sonic_scale/Federrath_sonic_scale_lowres.mp4

SC19 visualisation showcase finalist:
<https://youtu.be/EPe1Ho5qRuM>

Simulating episodic outflow feedback

in low-mass star formation

RESEARCH INSTITUTION

I. Physikalisches Institut, Universität zu Köln

PRINCIPAL INVESTIGATOR

Daniel Seifried

RESEARCHERS

Stefanie Walch, Paul Rohde

PROJECT PARTNERS

—

SuperMUC Project ID: pr47pi

Introduction

The star formation process is accompanied by the launching of bipolar protostellar outflows. While the young protostar accretes gas from its surrounding accretion disc, it re-ejects about 10% in bi-polar jets and disc-winds. These initially ejected jets entrain secondary core material to form massive molecular outflows. Since the entrained gas is not available for further star formation, outflow feedback lowers the stellar mass that forms within a core by about 50%. Many observations show that the accretion onto protostars is not a continuous process, but rather occurs in episodic outbursts. Since the launching of outflows is strongly coupled to the accretion, the launching of outflows is episodic, too. This process manifests itself as Herbig-Haro objects, individual outflow bullets, chained along the outflow. In this project, we simulate episodic outflow feedback in low-mass star-forming cores and study the effects on the forming stars.

Results and Methods

Sub-grid episodic outflow model

To model the accretion behaviour of observed FU Orionis type stars, we use the episodic accretion model by Stamatellos et al. (2010) [1]. Within this sub-grid model, the accretion on the protostar occurs in some ten years lasting outbursts. Between two outbursts is a quiescent phase of some thousand years with low accretion rate.

We are using this episodic accretion rate to launch bipolar outflows [3,4]. During these outbursts, 10% of the accreted gas is re-ejected. We are using the description by Matzer & McKee (1999) [2] to launch

the outflows with a velocity distribution that resembles a two-component outflow. The first component is a collimated high-velocity jet with ejection velocities that may exceed 120 km/s. The second component is a low-velocity wide-angle wind that encloses the central jet. The outflowing gas has a rotational velocity component and therefore carries away angular momentum from the central protostar. Since the accretion occurs in episodic bursts, our outflows are also episodic and resemble the observations of individual Herbig-Haro objects.

Simulation setup

We are using the highly object-orientated smooth particle hydrodynamics (SPH) code GANDALF [5] to simulate two sets of simulations. All simulations are spherically symmetric dense gas cores with a mass of $1 M_{\odot}$. We follow the collapse of the core about 500 kyr and save a snapshot file every 200 yrs. In total, we produce 2,500 snapshot files of which each is about 180 MB in disc size. Thus, one simulation adds up to about 450 GB of disc storage. Each simulation was performed on a single node, and the default runs take about 100,000 core hours.

The first simulation set covers a parameter and a resolution study. The resolution study consists of 8 simulations in which we increase the mass resolution from 25,000 to 400,000 SPH particles per M_{\odot} . In the parameter study, we use 14 simulations to evaluate the influence of varying model parameters, such as the collimation of the jet or the fraction of ejected to accreted gas [3]. We perform five additional simulations with a more realistic core setup.

In the second set of, in total, 88 simulations, we vary the initial core parameters to study the outcome of




Figure 1: Column density plot of a simulation using the episodic outflow sub-grid model at $t = 115$ kyr. One can see the outflow cavity emerging from the central protostar. Individual outflow bullets are labelled from A to D.




Figure 2: Position velocity diagram of a simulation from the first set at $t = 117$ kyr (top) and $t = 120$ kyr (bottom). The colour represents the fraction of outflowing gas with a given velocity at a certain distance from the source. The individual outflow bullets are labelled from A to F. Newly ejected outflow bullets have high velocities until they hit the leading shock front and decelerate to align the former bullets to the linear “Hubble-law” relation [2].

various initial conditions on the star formation process [4]. We perform 44 simulations with outflow feedback and compare them to the 44 reference simulations without outflow feedback.

Parameter and resolution study

The results from the resolution study show that our model is almost resolution-independent. Even with a very low mass resolution, we are able to reproduce the overall outflow properties such as outflowing gas mass and momentum. This result is especially important since the goal of this sub-grid model is to avoid the extremely high resolution requirements to compute outflow launching self-consistently.

The parameter study shows that our model is resistant against variation of all model parameter. Changing these parameters has only a small impact on the outcome of the simulation. While performing this study, we find that outflow feedback is highly self-regulated. Model parameter leading to stronger outflows do not change the outcome of the simulation significantly since these strong outflows efficiently cut off the launching protostars from fresh gas and, in turn, weaken the outflows.

Hubble-law relation and Hubble wedges

Observations of protostellar outflows show that the velocity of the outflowing gas increases linearly with distance from the source, the so-called “Hubble-law” relation. Recent observations identify individual outflow bullets of high velocity within this relation, so-called “Hubble Wedges”.

We are able to reproduce the observed “Hubble-law” relation including the “Hubble Wedges”. Since we follow the evolution of these bullets, we can explain the formation of the “Hubble-law” relation. Outflow bullets are ejected with high velocities and move through an, by former bullets, already evacuated




Figure 3: Distribution of mass ratios, mass of the secondary divided by the mass of the primary star, for all binary systems in the second set of simulations. The simulations with outflow feedback are in orange; the reference sample without outflow feedback is in blue. Outflow feedback enhances the formation of Twin-binaries.

outflow cavity. Once the outflow bullets hit the leading shock front, they overtake the former bullets (Fig. 2). Thereby the bullets decelerate and align with the former bullets to form the observed “Hubble-law” relation [3].

Twin-Binaries

Observations show an excess of almost equal mass binaries systems, called “Twin-binaries”. It is not yet clear why these Twin-binaries are so frequent. One possible formation pass is via circumbinary disc accretion. In a binary system with their main accretion channel via a surrounding accretion disc, the lower mass component will accrete faster, driving the mass ratio towards unity.

We are able to reproduce the high fraction of Twin-binaries in our simulations (Fig. 3). Interestingly, only the sample with outflow feedback shows an excess of Twin-binaries. We show that outflow feedback limits the accretion from other directions but the circumbinary disc and therefore enhances the fraction of the Twin-binaries.

Ongoing Research / Outlook

All simulations are performed on the SuperMUC Phase 2. The project is finished, and for the moment, we do not plan to extend the set of simulations. We rather continue analysing the gathered amount of data.

References and Links

- [1] Stamatellos et al., 2012, MNRAS, 427, 1182.
- [2] Matzer & McKee, 1999, ApJ, 526, L109.
- [3] Rohde et al. 2019, MNRAS, 883, 2563.
- [4] Rohde et al., 2020, in prep.
- [5] Hubber et al., 2018, MNRAS, 430, 3261.

Down to the dwarfs

RESEARCH INSTITUTION

¹Leibniz Institute for Astrophysics Potsdam

PRINCIPAL INVESTIGATOR

Stefan Gottlöber¹

RESEARCHERS

Gustavo Yepes², Peter Behroozi³, Pratika Dayal⁴, Anne Hutter⁴, Yamila Yaryura⁵

PROJECT PARTNERS

²UAM Madrid

³University of Arizona

⁴University of Groningen

⁵Universidad Nacional de Córdoba

SuperMUC Project ID: pr74no

Introduction

The observed Cosmic Microwave Background radiation tells us that the Universe was almost homogeneous some 13.8 billion years ago. At present, we observe structures on all scales from the very tiny dwarf galaxies with masses below 10^8 solar masses up to the most massive structures in the universe, the super-clusters with masses around 10^{16} solar masses. All these structures have been formed from the tiny density fluctuations at the beginning of cosmic history. One of the most demanding challenges for cosmologists is to understand how on all the different scales the observed structures have been formed. Numerical simulations are the driving force behind much of the theoretical progress in our understanding of the formation of structure in the universe.

Results and Methods

Simulations

In March 2013 we have started our MultiDark simulation project: a series of cosmological dark matter only simulations within Planck cosmology [1]. Based on the first data published by the Planck mission we assumed a spatially flat universe consisting of 30.7% of dark matter and 69.3% dark energy with a present-day expansion rate of 67.8 km/s/Mpc. Starting with the the MultiDark (MD) simulation we increased resp. decreased the box-size by a factor of 2.5. Thus in each step between the BigMultiDark (BigMD), the MultiDark (MD), the Small MultiDark (SMD), the Very Small MultiDark (VSMD) and the Extremely Small MultiDark simulations (ESMD) the volume decreases by a factor of ~ 16 and the mass resolution improves by the same factor. On top of this series we were running also the HugeMultiDark (HMD) simulation.

Simulation	box	N_p	m_p	N_{out}
HMD	4,000	$4,096^3$	7.9×10^{10}	128
BigMD	2,500	$3,840^3$	2.4×10^{10}	80
MD	1,000	$3,840^3$	8.7×10^9	129
SMD	400	$3,840^3$	9.6×10^7	117
VSMD	160	$3,840^3$	6.2×10^6	150
ESMD_2048	64	$2,048^3$	2.6×10^6	70
ESMD_4096	64	$4,096^3$	3.2×10^5	70

Table 1: The columns give the simulation identifier, the size of the simulated box in Mpc/h, the number of particles, the mass per simulation particle m_p (in M_{sun}/h) and the total number of simulation outputs stored N_{out} .

The initial conditions of these simulations are given by tiny perturbations in the distribution of the dark matter particles which follow the power spectrum predicted by the given Planck cosmology. Due to gravitational instabilities the dark matter particles clump and form dark matter halos which grow by accretion and merging events. At redshift $z = 0$ a simulation as described above contains typically 50 to 100 million halos. The halos and their sub-halos need to be identified in the simulation and their properties and formation histories must be estimated. To this end we have used two completely independent halo finders: A spherical overdensity halo finder and a hierarchical Friends-of-Friends halo finder.

In order to identify spherical halos and their sub-halos as well as their merging history we have used ROCKSTAR (Robust Overdensity Calculation using K-Space Topologically Adaptive Refinement [2]). ROCKSTAR is a massively parallel code very well designed to run over large dark matter simulations. It is designed to maximize halo consistency across time-steps and to provide reliable merger trees of the halo evolution which are the input for the generation of semi-analytical galaxy catalogs based on dark matter only simulations.

The basic idea of the Friends-of-Friends algorithm is to find halos by linking all particle pairs separated by less than a certain linking length that is defined as a fraction of the mean inter-particle separation in the simulation. This method defines a unique catalog of halos for any given linking length. By definition the halos defined at a shorter linking length are sub-halos of the parent-halos defined at larger linking length.

Galaxies

The observed stars and galaxies consist of baryonic matter which contribute only 4.8% to the total matter and energy in the Universe. Elaborated hydrodynamical numerical simulations including cooling of gas, star formation, feedback and many other effects gave new insights into the formation and evolution of galaxies. However these simulations are extremely time consuming and cannot be performed for very large volumes at very high resolution. A possible alternative to study galaxies in large cosmological volumes with the required resolution is




Figure 1: In a slice through the center of the SMD simulation box the neutral hydrogen fraction is shown 605 (left) and 739 million years after big bang (right). Black indicates the still neutral hydrogen while orange/yellow indicates almost fully ionised hydrogen.

the post-processing of dark matter only simulations with semi-analytical galaxy formation models. This has been done for the MD simulation using three different semi-analytic models, namely Galacticus, SAG, and SAGE [3] and for the SMD simulation using SAG [4]. Based on the galaxy catalogs constructed from the SMD simulations we have recently studied associations of dwarf galaxies. Studying the formation of such associations is critical to our understanding of structure formation at the low mass end, a regime known to constitute the small scale crisis of the current cosmological paradigm [4].

High redshift galaxies

The upcoming James Webb Space Telescope (JWST) will enable observations of the most distant events and objects in the Universe such as the first galaxies that have been formed more than 13 billion years ago, i.e. during the first 500 million years after big bang. JWST will provide the potential to study exotic physical systems (like population III stars and direct-collapse black holes).

The UNIVERSEMACHINE is an empirical model that links galaxy star formation rates to properties of their host halos and in particular to their formation history [5]. This results in an entire mock universe populated with galaxies. Using the UNIVERSEMACHINE, we predicted which objects JWST may observe 250 to 500 million years after big bang [6]. Such predictions are essential for proposal planning at JWST.

Reionization

The Epoch of Reionization (EoR) marks a pivotal time in the history of our Universe when the photons emitted by the first stars and galaxies permeate and gradually ionize the intergalactic medium. The VSMD simulation has been analyzed within a project funded by the European Research Council, titled "DELPHI: a framework to study dark matter and the emergence of galaxies in the epoch of reionization". We quantified the effects of radiative feedback on the underlying galaxy population as well as on the ionization topology during the EoR to answer questions including: When

and which galaxies are most affected by radiative feedback? Is the patchiness of reionization imprinted in galaxy observables? How does radiative feedback impact high-redshift observables and the 21cm signal from the neutral regions in the intergalactic medium? To this end we have used our ASTRAEUS framework that self-consistently couples a state-of-the-art N-body simulation with a semi-numerical model of galaxy evolution and reionization [7].

Ongoing Research / Outlook

Many users were interested in using the simulation data directly or analyzing the full catalogs by themselves. Therefore, we opened a possibility to download selected simulation outputs (1.7 Tb each) as well as the full ROCKSTAR or galaxy catalogs via the CosmoSim database [1] at AIP Potsdam. Besides the data, the database contains a comprehensive documentation as well as a number of images and movies for public outreach.

In addition, the CosmoSim database also provides data from some of the MultiDark simulations via SQL (Structured Query Language) queries. Since the amount of data that such simulations produce is very large, SQL is used to filter the main data products and retrieve exactly those subsets the user is interested in. Having the data directly available via SQL allows users—especially those not intimately familiar with the data format of the simulation—a far more direct path from a science question to an executable expression than a standard scripting or programming language would. More than 120 papers have been published so far using data provided by the MultiDark database.

References and Links

- [1] <http://www.cosmosim.org>.
- [2] Behroozi P. S., Wechsler R. H., Wu H.-Y., ApJ 762, 109, 2013. <https://bitbucket.org/gfcstanford/rockstar/overview>.
- [3] Knebe A. et al., MNRAS 474, 5206, 2018.
- [4] Yaryura C.Y. et al., MNRAS 499, 5932, 2020.
- [5] Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., MNRAS 488, 3143, 2019.
- [6] Behroozi P. et al., 2020, MNRAS 499,5702, 2020.
- [7] Hutter A. et al., 2020, arXiv2004.08401.

Star Formation through the Cosmic Ages:

Shape and Spin of Primordial Minihalos

RESEARCH INSTITUTION

¹Heidelberg University, Center for Astronomy

PRINCIPAL INVESTIGATOR

Ralf Klessen¹

RESEARCHERS

Anna Schauer^{2,1}, Simon Glover¹, Maik Druschke¹

PROJECT PARTNERS

²University of Texas in Austin, Department of Astronomy, USA

SuperMUC Project ID: pr74nu

Context

Project pr74nu *Star Formation through the Cosmic Ages* provides computing time to the star formation group at the Center for Astronomy at Heidelberg University, and allows the members of the team to perform small to medium-sized parameter studies, to develop and validate new numerical code modules, and to prepare for larger-scale highest-resolution numerical simulations on specific topics for which time at national and European supercomputing facilities is being requested in separate proposals.

The research activities of the group focus on identifying and characterizing the physical processes that govern the birth of stars across cosmic time, from the first stars in our universe to young protostellar objects in the solar vicinity. Scientific interests include gas dynamics and chemistry in the interstellar medium, statistical description of astrophysical turbulent flows, dynamo processes, gravitational collapse and formation of stars and stellar clusters, stellar dynamics, and the development and optimization of numerical algorithms in computational astrophysics. The use of supercomputing resources, as provided by LRZ, is instrumental for these research activities.

Here we highlight one specific example of these activities, the numerical study of a master student in the group focusing on the shape and spin of star-forming primordial minihalos in different environmental conditions and under the influence of the relative streaming velocities between baryons and dark matter. The results of this study have been published in [2,3].

Introduction

The most favored approach for describing the evolution of our universe since the big bang is the so-called Λ CDM (Lambda Cold Dark Matter) model. It constitutes the simplest theory consistent with the existing astronomical measurements and is thus considered the standard model of big bang cosmology. It assumes that most of the matter is dark and interacts with baryonic matter solely through gravity. Consequently, cosmic structure formation

occurs hierarchically with gravitational contraction setting in on the smallest scales first and with larger bound structures forming later via mergers and accretion.

Before recombination, baryons and photons were tightly coupled by Compton scattering and behaved as if they were a single fluid. On the other hand, the dark matter was not directly coupled to this baryon–photon fluid and interacted with it only via gravity. Consequently, even though fluctuations in the density of the baryon–photon fluid and the dark matter are correlated, one nevertheless expects some motion of the fluid relative to the dark matter. The imprint of this motion survives in the baryons even after recombination, in the form of a large-scale streaming motion of the baryons with respect to the dark matter. The resulting relative velocity is highly supersonic shortly after recombination, but decays with the expansion of the Universe. It therefore plays little role in the formation of galaxies at low redshifts. However, at high redshift, streaming velocities can have a significant effect on the creation and formation of dark matter minihalos and the dynamics of the gas they contain. Specifically, their presence increases the minimum mass required for efficient gas cooling and collapse, and consequently for the birth of stars in these systems.

These minihalos are the sites where the first stars in the universe, the so-called Population III (in short Pop. III) stars, are thought to form. Cooling in these systems takes place primarily via molecular hydrogen emission, with the result that the gas temperature never decreases below a few hundred Kelvin. The relatively high gas temperature inhibits fragmentation of the gas and promotes rapid gas accretion onto any protostars that form. The initial mass function of Pop. III stars is therefore expected to be dominated by massive stars, although it remains uncertain whether it also extends down to values below a solar mass due to efficient fragmentation of the accretion disk that forms due to angular momentum conservation as matter falls towards the very center of the halo. Altogether, better knowledge of the properties of the dark matter and the gas that make up the minihalos that host the first generation of stars is of great

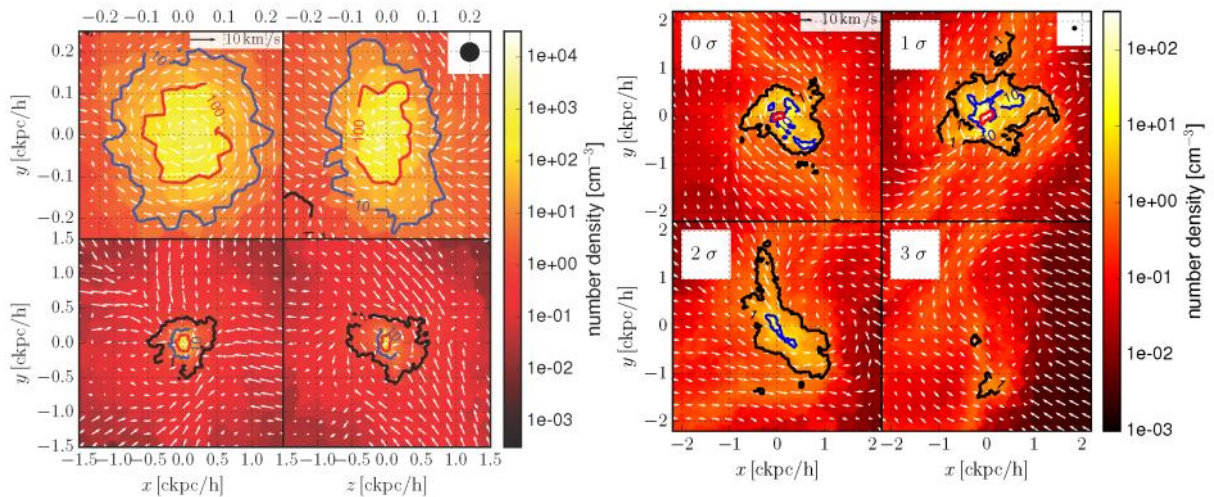


Figure 1: Example of a typical minihalo with a spin parameter of $\lambda = 0.0483$ and a mass of $M = 9.57 \times 10^6$ solar masses at a redshift of $z = 14$ in a region with a streaming velocity of 2σ . The images show the number density of gas particles in a slice centered on the density peak, with the top panel focusing on the central 500 pc of the minihalo and the lower panel depicting a larger region of 3 kpc across. The left and right columns correspond to two different viewing angles.

importance for our understanding of the onset and outcome of Pop. III star formation.

Shape and Spin of Primordial Minihalos

In the first study presented here [2], we use high resolution numerical simulations to investigate the distribution of the shape and spin of a large sample of dark matter minihalos. We find that the spin parameter of the minihalos follows a log-normal distribution with minimal dependence on redshift. Most minihalos are prolate, but those that form at the highest redshifts tend to be more prolate than those that assemble at later times. On the scale of the virial radius, there is a good correlation between the shape and spin of the gas and that of the dark matter. However, this correlation breaks down in gas which is cooling and undergoing gravitational collapse. We show that although the direction of the spin of the central dense gas correlates well with that of the halo, the magnitude of the spin of the dense gas is uncorrelated with that of the halo. Therefore, measurements of the spin on large scales tell us little about the angular momentum of the gas responsible for forming the first stars.

In a second study [3], we employ high resolution numerical simulations to investigate the impact of the relative streaming velocity between dark matter and baryons on the spin and shape distributions of high-redshift minihalos. We find that the presence of streaming velocities has a negligible effect on the spin and shape of the dark matter component of the minihalos. However, it strongly affects the dynamics of the gas in the system. The most probable spin parameter increases by a factor of five in regions of high streaming velocity compared to the case of zero streaming velocity. Consequently, the gas within the minihalos becomes increasingly less spherical and more oblate as the streaming velocity increases, with dense clumps being found at larger distances from the halo center. The impact of the streaming velocity is also mass-dependent: less massive objects are influenced more strongly, on account of their shallower gravitational potential wells. The number of halos in which gas cooling and runaway gravitational collapse

occurs decreases substantially as the streaming velocity increases. However, the spin and shape distributions of gas that does manage to cool and collapse are insensitive to the value of the streaming velocity and we therefore expect the properties of the stars that form from this collapsed gas do not depend on the value of the streaming velocity.

Numerical Methods

The numerical simulations that form the base of the study reported here were first described in [4,5]. They were carried out using the moving mesh code AREPO [6] and include both gas and dark matter. The chemical and thermal evolution of the gas were treated using a primordial chemistry network and cooling function based on the work presented in [1], but updated as described in [5]. The simulations have a particle mass of about 100 solar masses for the dark matter and a target of 20 solar masses for the gas. Cosmological parameters are taken from the current standard cosmological model. Different streaming velocities are realized as a constant offset velocity term added to the initial gas motion.

References and Links

- [1] Clark, P.C., Glover, S.C.O., Klessen, R.S., Bromm, V., *ApJ*, 727, 110 (2011).
- [2] Druschke, M., Schauer, A.T.P., Glover, S.C.O., Klessen, R.S., 481, 3266 (2018).
- [3] Druschke, M., Schauer, A.T.P., Glover, S.C.O., Klessen, R.S., *MNRAS*, 498, 4839 (2020).
- [4] Schauer, A.T.P., Regan J., Glover S.C.O., Klessen R.S., *MNRAS*, 471, 4878 (2017).
- [5] Schauer A.T.P., Glover S.C.O., Klessen R.S., Ceverino D., *MNRAS*, 484, 3510 (2019).
- [6] Springel V., *MNRAS*, 401, 791 (2010).

Kinetic simulations of astrophysical and solar plasma turbulence

RESEARCH INSTITUTION

¹Max-Planck-Institut für Sonnensystemforschung, Göttingen, Germany

PRINCIPAL INVESTIGATOR

Jörg Büchner^{1,3}

RESEARCHERS

Daniel Grošelj², Patricio Muñoz³, Frank Jenko⁴

PROJECT PARTNERS

²Columbia University, New York, USA

³Technical University Berlin, Germany

⁴Max-Planck-Institut für Plasmaphysik, Garching, Germany

SuperMUC Project ID: pr74vi

Introduction

The hot and dilute collisionless plasmas near the Earth and in the solar-wind are very turbulent. Astrophysical plasmas, e.g., in accretion disks around stars, in supernovae, the interstellar medium (ISM) or in galaxy clusters are turbulent. Despite of the omnipresence of turbulence in the plasma Universe, its properties and consequences are not well understood, yet. The reason is the complicated kinetic nature of the collective collisionless plasma phenomena at the end of the turbulence cascade, where the energy is dissipated. A kinetic, particle description is required to describe the energy dissipation in turbulence, so that non-thermal effects and wave-particle interactions are taken into account.

Turbulence is inherently linked with two fundamental processes of energy conversion in the Universe which are the focus of this SuperMUC-NG project [1]. These fundamental processes are shock waves and magnetic reconnection. In shock waves plasma flow energy is converted into thermal energy (heating) and the acceleration of particles while the plasma flow is decelerated. In the upstream region of shock waves, the plasma may become highly turbulent. Shock waves are observed in situ at space plasma boundaries, e.g. between the solar wind and magnetospheres or other heliospheric obstacles and, indirectly, between exploding supernovae and the interstellar medium (ISM). They are assumed to be behind gamma-ray bursts as well. The latter are known to occur in highly relativistic plasmas. In the first of our SuperMUC-NG sub-project we investigated the efficiency of particle acceleration by shock waves in gamma-ray bursts and the role of self-generated turbulence in this highly energetic process. The obtained particle energy spectra are used as input for models of the non-thermal emission due to, e.g., synchrotron radiation from gamma-ray bursts. The second fundamental process of energy conversion in the Universe, closely linked to turbulence, is magnetic reconnection. In the course of magnetic reconnection, magnetic energy is converted into plasma flows, thermal energy and the acceleration of highly energetic particles. The understanding of magnetic reconnection processes is crucial for the understanding of the heating of cosmic plasmas, e.g. of




Figure 1: Particle energy spectra in an unmagnetized electron-ion shock enriched with electron-positron pairs (in the case shown here one electron-positron pair per ion). From top to bottom: ion, electron, and positron spectra.

stellar coronae, the release of magnetic energy in solar and stellar flares, as well as of the dynamics and consequences of the energization of the ISM. Reconnection and turbulence are related to each other in several different ways: reconnection can proceed if turbulence allows non-ideal plasma responses, it can generate secondary turbulence itself by causing plasma instabilities, and small-scale current sheets are formed in turbulent plasmas through which reconnection happens. The latter process is the focus of our second SuperMUC-NG sub-project. A number of studies of these processes have been carried out using different plasma models. In particular we concentrated our studies on the not well understood role of the electrons in the energy dissipation, plasma heating and acceleration processes. Our studies aim at the comparison of the consequences of different ways to reduce the plasma models. The results of our investigations are important to better understand the energy dissipation and heating in astrophysical plasmas from the Earth's magnetosheath to the solar wind and the interstellar medium.

Results and Methods

Methods

To tackle the outstanding astrophysical problems of collisionless shocks and reconnection in turbulence kinetically, the method of choice are fully kinetic plasma simulations in which the full sets of Vlasov-Maxwell equations are self consistently solved together with the action of the electromagnetic fields on the charged plasma particles. For this sake we carried out Particle-in-Cell (PIC) simulations which describe the particle motion by a Lagrangian approach and solves by an Eulerian description for the electromagnetic fields that are allocated on a mesh. The particle distribution functions are the main variables to be obtained by solving Vlasov equations. Hence, the PIC-method relies on a sufficiently large number of many macro-particles, in order to statistically reliably reconstruct the particle distribution functions. Such PIC-code simulations are, therefore, computationally very demanding: they have to trace in the order of $10^9 - 10^{11}$ macro-particles on a mesh consisting about 10^8 grid points in 3D simulations. This implies that a single particle snapshot requires a memory in the order of one Terabyte. Hence, one has to make the particle output as sparse in time as possible. Grid data (e.g., electromagnetic fields) is comparably less heavy, but this data output should be carried out more frequently in order to correctly understand the evolution of the system. The resulting output field sizes are in the order of a significant fraction of a Terabyte for large production runs. In our project we utilized the optimum MPI-parallelized codes OSIRIS (<http://epp.tecnico.ulisboa.pt/osiris>) and ACRONYM (<http://plasma.nerd2nerd.org/>). These codes were proven to run efficiently on the SuperMUC-NG machine. The largest simulations that we carried out were in the order of a few million core-hours, while parametric studies are done with runs in the order of a significant fraction of a million of core-hours. A typical run requires between a few to tens of thousands of cores.

Results

With the OSIRIS code we performed the first PIC simulations of relativistic, weakly magnetized electron-ion-positron shock waves and compared them with shocks in a pure electron-ion plasma. Our simulation results apply to gamma-ray bursts, where the released radiation is thought to produce electron-positron pairs in the external medium upstream of the shock, i.e. in the unperturbed plasma flow ahead of the shock wave. In those environments, a critical parameter that controls the energization process is the magnetization parameter σ , the ratio of the upstream magnetic energy to the ion kinetic energy. We studied the dependence of the shock properties on this parameter. The shock is numerically generated by pushing against one of the reflective boundaries of the 2D simulation domain a cold stream of plasma. This kind of numerical setup is prone to a numerical instability known as Cherenkov radiation, which is taken care of by means of a modified Maxwell solver implemented in the OSIRIS code (see further details about the initial conditions in [2]). We illustrate a particular result in Figure 1, showing the downstream energy spectra in a unmagnetized shock composed of a mixture of electrons, ions, and positrons. We find that the shock




Figure 2: Current density in a 3D kinetically turbulent plasma. The (red) peak values indicate current sheets where energy is preferentially dissipated via small-scale magnetic reconnection.

can efficiently accelerate particles via scattering off the small-scale turbulence. More generally, we studied the dependence on the plasma composition (mixed electron-ion-positron versus pure electron-ion) and on σ , determining the conditions when electron-ion-positron shocks accelerate particles. With the code ACRONYM we simulated a turbulent plasma by imposing initial velocity and magnetic field fluctuations at large-scales. The fluctuations evolve in time, transferring energy to smaller and smaller scales, until it is finally dissipated at the smallest (electron-)kinetic scales. During the first period of our SuperMUC-NG project we have already found by utilizing the OSIRIS code that the turbulent fluctuations match the theoretical predictions of the so-called kinetic Alfvén wave turbulence, characterized by a specific correlation of density and perpendicular magnetic field fluctuations [3]. We then started to investigate the properties of current sheet structures that are formed in the turbulence. The current sheets release magnetic energy via magnetic reconnection. They seem to determine the energy dissipation in a kinetically turbulent collisionless plasmas (Figure 2).

Ongoing Research / Outlook

In the first sub-project (relativistic shock waves) we found the conditions under which electro-ion-positron shocks can efficiently accelerate particles. In the second sub-project (reconnection in turbulence) the question of the influence of the electron fluid with / without mass on the energy dissipation via current sheet reconnection is still open. For the sake of a better understanding of the role of the electrons, comparative simulations for different plasma models are necessary using, other than fully-kinetic PIC codes, hybrid-PIC simulations with and without electron inertia [4] which will allow to directly assess the influence of the electron dynamics.

References and Links

- [1] <https://www-astro.physik.tu-berlin.de/node/438>.
- [2] Sironi, L., Spitkovsky, A., & Arons, J. 2013, *ApJ*, 771, 54, doi: 10.1088/0004-637X/771/1/54.
- [3] Grošelj, D., Mallet, A., Loureiro, N. F., & Jenko, F. (2018). *Physical Review Letters*, 120(10), 105101. doi: 10.1103/PhysRevLett.120.105101.
- [4] Jain, N., Büchner, J., Comişel, H., & Motschmann, U. (2020). Free energy sources in current sheets in collisionless turbulence. <http://arxiv.org/abs/2008.13206>.

Halos in Magneticum: Scaling Relations,

Mass, Bias, and Concentration

RESEARCH INSTITUTION

¹University-Observatory Munich, LMU Munich

PRINCIPAL INVESTIGATOR

Klaus Dolag¹

RESEARCHERS

Antonio Ragagnin², Priyanka Singh², Tiago Castro², Alex Saro²

PROJECT PARTNERS

²INAF-Osservatorio Astronomico di Trieste

SuperMUC Project ID: pr83li, pr74do (Gauss Large Scale and KONWIHR projects)

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 centers. 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 upcoming surveys.

Scaling relations in Magneticum

The abundance of galaxy clusters as a function of mass and redshift is a well known powerful cosmological probe, which relies on underlying modelling assumptions on the mass-observable relations (MOR). Some of the MOR parameters can be constrained directly from multi-wavelength




Figure 1: Distribution of cosmological parameters for the fifteen cosmologies used in this paper. WMAP7 (i.e. our C8) is high-lighted by star symbols. The black contours represent 68%, 95% and 99% condence limits on these parameters. Taken from [1].

observations, as the normalization at some reference cosmology, the mass-slope, the redshift evolution and the intrinsic scatter. However, the cosmology dependence of MORs cannot be tested with multi-wavelength observations alone, as we observations are limited to one Universe only. We use Magneticum simulations to explore the cosmology dependence of galaxy cluster scaling relations. We run fifteen hydrodynamical cosmological simulations varying Ω_m , Ω_b , H_0 and σ_8 (around a reference cosmological model). The MORs considered are gas mass, baryonic mass, gas temperature, compton Y and velocity dispersion as a function of virial mass. We verified that the mass and redshift slopes and the intrinsic scatter of the MORs are nearly independent of cosmology with variations significantly smaller than current observational uncertain-ties. We show that the gas mass and baryonic mass sensitively depends only on the baryon fraction, velocity dispersion and gas temperature on H_0 , and Y on both baryon fraction and H_0 . We investigate the cosmological implications of our MOR




Figure 2: Mass-concentration plane of our simulations C1–C15 (black solid lines) and observations as labeled. Taken from [2].

parameterization on a mock catalog created for an idealized eROSITA-like experiment. We show that our parametrization introduces a strong degeneracy between the cosmological parameters and the normalization of the MOR. Finally, the parameter constraints derived at different overdensity ($\Delta 500c$), for X-ray bolometric gas luminosity, and for different subgrid physics prescriptions. For details see [1].

M-c and M-M relations from Magneticum

We employ a set of Magneticum cosmological hydrodynamic simulations that span 15 different cosmologies, and extract masses and concentrations of all well-resolved haloes between $z = 0-1$ for critical over-densities Δ_{vir} , Δ_{200c} , Δ_{500c} , Δ_{2500c} , and mean overdensity Δ_{200c} . We provide the first mass-concentration (M-c) relation and sparsity relation (i.e. $M\Delta 1-M\Delta 2$ mass conversion) of hydrodynamic simulations that is modelled by mass, redshift and cosmological parameters Ω_m , Ω_b , σ_8 , H_0 as a tool for observational studies. We also quantify the impact that the M-c relation scatter and the assumption of NFW density profiles have on the uncertainty of the sparsity relation. We find that converting masses with the aid of a M-c relation carries an additional fractional scatter ($\sim 4\%$) originated from deviations from the assumed NFW density profile. For this reason, we provide a direct mass-mass conversion relation fit that depends on redshift and cosmological parameters. We released the package `hydro_mc` [4], a python tool that performs all kind of conversions presented in this paper. For details see [2].

Baryon impact on halo mass and bias

Luminous matter produces very energetic events, such as active galactic nuclei and supernova explosions, that significantly affect the internal regions of galaxy clusters. Although the current uncertainty in the effect of baryonic physics on cluster statistics is subdominant as compared to other systematics, the picture is likely to change soon as the amount of high-quality data is growing fast, urging the community to keep theoretical systematic uncertainties below the ever-growing statistical precision. Therefore, we studied the effect of baryons on galaxy clusters, and their impact on the cosmological applications of clusters, using the *Magneticum* suite of cosmological hydrodynamical simulations. We show that the impact of baryons on the halo mass function $f(\sigma, z)$ as defined in Figure 3 can be recast in terms on a variation of the

$$f(\sigma, z) = A \left[(\sigma/b)^{-a} + \sigma^{-c} \right] e^{-d/\sigma^2},$$

$$A(z) = A_0 (1+z)^{A_z},$$

$$a(z) = a_0 (1+z)^{a_z},$$

$$b(z) = b_0 (1+z)^{b_z},$$

$$c(z) = c_0 (1+z)^{c_z},$$

$$d(z) = d_0 (1+z)^{d_z},$$

	A_0	A_z	a_0	a_z	b_0	b_z	c_0	c_z	d_0	d_z
200m										
Hydro	0.862	-0.475	-2.631	1.279	3.000	-1.388	1.725	0.323	1.256	0.046
DMO	0.610	-0.457	2.721	-0.113	0.846	0.213	0.731	-1.959	1.297	-0.003
200c										
Hydro	0.372	1.140	3.164	0.769	1.094	-0.629	0.528	0.927	1.710	-0.102
DMO	0.333	1.093	2.110	0.870	1.147	-0.652	0.379	0.691	1.515	-0.066




Figure 3: Confidence regions for the 200m best-fit parameters of $f(\sigma, z)$ for the DMO (blue) and Hydro (red) Magneticum simulations. The best-fit parameters for 200m and 200c are presented in the Table. Taken from [3].

mass σ of the halos simulated with pure N-body, when baryonic effects are included. The halo mass function and halo bias are only indirectly affected. Finally, we demonstrate that neglecting baryonic effects on halos mass function and bias would significantly alter the inference of cosmological parameters from high-sensitivity next-generations surveys of galaxy clusters. For more details, see [3].

References and Links

- [1] Singh, Saro, Constanzi & Dolag, MNRAS, 494 3728, 2020.
- [2] Aragagnin, Saro, Singh & Dolag, MNRAS, tmp, 3313, 2020.
- [3] Castro, Borgani, Dolag, Marra, Quartin, Saro & Sefusatti, MNRAS, 500, 2316, 2021.
- [4] http://github.com/aragagnin/hydro_mc/.

SuperCAST: Simulating the Universe

RESEARCH INSTITUTION

¹University-Observatory Munich, LMU Munich

PRINCIPAL INVESTIGATOR

Klaus Dolag¹

RESEARCHERS

Ullrich Steinwandel¹, Marcel Lotz¹, Congyao Zhang², Eugene Churazov²

PROJECT PARTNERS

²Max Planck Institute for Astrophysics

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 intergalactic 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. The focus especially lies on the interplay between feedback processes, AGN and galaxy evolution and their imprint on the intergalactic medium (IGM) or inter




Figure 1: Structure of the velocity field within the simulated galaxy for $t = 2.5$ Gyr, when the outflow is present. We see that the gas is moving in z -direction and falls back at later times in the outer part of the galactic disk. Taken from [2].

cluster medium (ICM). It is now clear that small-scale processes like the condensation of molecular clouds into stars, magnetic fields and the details of heat transport as well as 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 from the last approval period cover a link between the various scales and contribute to our understanding of crucial interplay between the formation and evolution of central black holes and galaxies [1], magnetic fields and evolution of galaxies [2] and internal processes of galaxy clusters and the ICM [3], as well as the large scale structures in the universe. It also drives the continuous effort to develop and apply new numerical methods and the next generation of multi-scale codes within the framework of numerical astrophysics within our group of students [4], PhDs, and young Post-Docs.

Rise and fall of post-starburst galaxies

Post-starburst galaxies (PSBs) belong to a short-lived transition population between star-forming and

quiescent galaxies and are found in all environments and at all redshifts. Deciphering their heavily discussed evolutionary pathways is paramount to understanding the processes that drive galaxy evolution in general. We aim to determine the dominant mechanisms governing PSB galaxy evolution in both the field and in galaxy clusters. To understand the different physical processes involved in PSB galaxy evolution, we analyse the cosmological hydrodynamical simulation suite *Magneticum Pathfinder*. At $z \sim 0$, we identify and study a set of 647 PSBs with stellar masses $M^* > 5 \cdot 10^{10} M_{\odot}$, in comparison to a quiescent and a star-forming control sample. We track the galactic evolution, merger history, and black holes of PSBs and control samples over a time-span of 3.6Gyr. Additionally, we study cluster PSBs identified at different redshifts and cluster masses. Independent of environment and redshift, we find that PSBs, similar to the star-forming control sample, have frequent mergers. At $z = 0$ 89% of PSBs have experienced at least one merger, and 65% even had a major merger event within the last 2.5Gyr, leading to episodes of strong star formation. In fact, 23% of $z = 0$ PSBs were rejuvenated during their starburst. After the mergers, field PSBs are generally shutdown via a strong increase in AGN feedback. Furthermore, we find agreement with observations for both stellar mass functions and $z = 0.9$ line-of-sight phase space distributions of galaxy cluster PSBs. Finally, we find that $z \leq 0.5$ cluster PSBs are predominantly infalling, especially in high mass clusters and show no signs of enhanced AGN activity. Thus, we conclude that the majority of cluster PSBs are shutdown via an environmental quenching mechanism such as ram-pressure stripping. For more details see [1].

Origin of magnetic driven winds

We investigate the build-up of the galactic dynamo and subsequently the origin of a magnetic driven outflow. We use a set-up of an isolated disc galaxy with a realistic circum-galactic medium (CGM). We find good agreement of the galactic dynamo with theoretical and observational predictions from the radial and toroidal components of the magnetic field as function of radius and disc scale height. We find several field reversals indicating dipole structure at early times and quadrupole structure at late times. Together with the magnetic pitch angle and the dynamo control parameters R_{α} , R_{ω} , and D , we present strong evidence for an α^2 - Ω dynamo. The formation of a bar in the centre leads to further amplification of the magnetic field via adiabatic compression which subsequently drives an outflow. Due to the Parker instability the magnetic field lines rise to the edge of the disc, break out, and expand freely in the CGM driven by the magnetic pressure. Finally, we investigate the correlation between magnetic field and star formation rate. Globally, we find that the magnetic field is increasing as function of the star formation rate surface density with a slope between 0.3 and 0.45 in good agreement with predictions from theory and observations. Locally, we find that the magnetic field can decrease while star formation increases. We find that this effect is correlated with the diffusion of magnetic field from the spiral arms to the interarm




Figure 2: Gas-entropy slice of a galaxy cluster from the cosmological simulation at redshift $z = 0$. The black solid lines and the white line compare the MA-shocks and the shock radius formed in our special 1D simulation. Taken from [3].

regions which we explicitly include by solving the induction equation and accounting for non-linear terms. For more details see [2].

Colliding shocks in galaxy clusters

Several types/classes of shocks naturally arise during formation and evolution of galaxy clusters. One such class is represented by accretion shocks, associated with deceleration of infalling baryons. Such shocks, characterized by a very high Mach number, are present even in 1D models of cluster evolution. Another class is composed of 'runaway merger shocks', which appear when a merger shock, driven by a sufficiently massive infalling subcluster, propagates away from the main-cluster centre. We argue that, when the merger shock overtakes the accretion shock, a new long-living shock is formed that propagates to large distances from the main cluster (well beyond its virial radius), affecting the cold gas around the cluster. We refer to these structures as Merger-accelerated Accretion shocks (MA-shocks) in this paper. We show examples of such MA-shocks in one-dimensional (1D) and three-dimensional (3D) simulations and discuss their characteristic properties. In particular, (1) MA-shocks shape the boundary separating the hot intracluster medium (ICM) from the unshocked gas, giving this boundary a 'flower-like' morphology. In 3D, MA-shocks occupy space between the dense accreting filaments. (2) Evolution of MA-shocks highly depends on the Mach number of the runaway merger shock and the mass accretion rate parameter of the cluster. (3) MA-shocks may lead to the misalignment of the ICM boundary and the splashback radius. For more details see [3].

References and Links

- [1] Lotz, Dolag, Remus & Burkert, MNRAS submitted, 2020arXiv201003762.
- [2] Steinwandel, Dolag, Lesch, Moster, Burkert & Pioto, MNRAS, 894, 4393, 2020.
- [3] Zhang, Curazovm Dolag, Forman & Zhuravleva, MNRAS, 494, 4539, 2020.
- [4] http://www.usm.uni-muenchen.de/CAST/student_projects.html.

Simulating the formation, evolution, and merging

of molecular clouds

RESEARCH INSTITUTION

¹1. Physikalisches Institut, Universität zu Köln

PRINCIPAL INVESTIGATOR

Daniel Seifried¹

RESEARCHERS

Stefanie Walch¹, Sebastian Haid¹, Philipp Girichidis²

PROJECT PARTNERS

²AIP Potsdam

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 or stellar radiation on the clouds. In particular we are interested in the chemical evolution of the clouds as well as the internal dynamics and structure.

Results and Methods

In our work we have so far performed a number of simulations on SuperMUC Phase 2 and SuperMUC-NG. Each of the simulations required a computational time of about 1-2 Mio. core hours with a simultaneous use of up to 1,000 cores per simulation. A few hundreds of files were produced 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 $(500 \text{ pc}) \times (500 \text{ pc}) \times (\pm 5 \text{ kpc})$ region of a galactic disk, with a gas surface density of $10 M_{\text{sun}}/\text{pc}^2$. 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 regions from these runs where we know that molecular clouds are forming and then start to zoom-in [4]. This allows us to follow the evolution of molecular clouds with a significantly higher spatial resolution of 0.06 pc.

The structure of molecular clouds

We follow the evolution of individual molecular clouds over a time span of a few Myr. In Fig. 1, we show that molecular clouds possess an extremely complicated structure, revealing a highly fragmented and filamentary shape. It appears also that the individual clouds show large morphological differences between each other. In particular, we can show that the presence of magnetic fields tends to make molecular clouds more fluffy with less dense material. Magnetic fields also tend to delay the formation of massive stars.

The impact of stellar radiation

We continue to simulate the evolution of the molecular clouds following the formation of massive stars in them in order to investigate in detail the effect of stellar radiative feedback on the clouds [5]. In Fig. 2, we show the column density of two clouds affected by this feedback process. We find that although both clouds are comparable in total mass, size, the fraction of molecular H_2 , and the velocity dispersion, the cloud in the bottom row of Fig. 2 is significantly more affected by the feedback and almost completely dispersed during the formation process of the first massive stars. We can show that this is due to the different substructure in the two molecular clouds, the cloud being affected more by the stellar radiation has less dense gas which in turn is well shielded from the radiation. This allows the radiation to penetrate deeper into the cloud and dissociated H_2 molecules leading eventually to the disruption of the cloud.




Figure 1: 3D representation of a molecular cloud modelled in the project pr94du. The reddish material represent the dense, filamentary gas, whereas the bluish material presents dilute and hot gas of nearby supernova remnants.




Figure 2: Column density of two different molecular clouds (top and bottom) comparing the effect of stellar radiative feedback (right column) to a case without stellar radiative feedback (left column). The symbols indicate the position of massive stars. Stellar feedback significantly alters the substructure of molecular clouds by blowing apart the densest structures.

Ongoing Research / Outlook

For our future research, we intend to study the evolution of further molecular clouds. In particular, we will focus on the evolution of molecular clouds under different environmental conditions like e.g. present in the early universe or in our Galactic Center close to the central black hole. This will require further large amounts of computing power in the future.

References and Links

- See also <http://www.astro.uni-koeln.de/silcc>
- [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.
 - [4] Seifried, D., Walch, S., Girichidis, P., et al., 2017, MNRAS, 472, 4797.
 - [5] Haid, S. Walch, S., Seifried, D., et al., 2019, MNRAS, 482, 4062.

Environmental Simulations of the Local Group

RESEARCH INSTITUTION

¹Leibniz Institute for Astrophysics Potsdam (AIP), Potsdam, Germany

²University of Lyon, UCB Lyon 1, France

PRINCIPAL INVESTIGATOR

Noam I Libeskind^{1,2}

RESEARCHERS

Edoardo Carlesi¹, Robert J. J. Grand³, Arman Khalatyan¹, Alexander Knebe^{4,5,6}, Ruediger Pakmor³, Sergey Piliipenko⁷, Marcel S. Pawlowski¹, Martin Sparre^{1,8}, Elmo Tempel⁹, Peng Wang¹, Helene M. Courtois², Stefan Gottloeber¹, Yehuda Hoffman¹⁰, Ivan Minchev¹, Christoph Pfrommer¹, Jenny G. Sorce^{1,11}, Volker Springel³, Matthias Steinmetz¹, R. Brent Tully¹², Mark Vogelsberger¹³, Gustavo Yepes^{4,5}

PROJECT PARTNERS

³Max-Planck-Institut für Astrophysik, Garching, Germany

⁴Departamento de Física Teórica, Modulo 15, Facultad de Ciencias, Universidad Autónoma de Madrid, Spain

⁵Centro de Investigación Avanzada en Física Fundamental (CIAFF), Facultad de Ciencias, Universidad Autónoma de Madrid, Spain

⁶International Centre for Radio Astronomy Research, University of Western Australia

⁷P.N. Lebedev Physical Institute of Russian Academy of Sciences, Moscow, Russia

⁸Potsdam University

⁹Tartu Observatory, University of Tartu, Estonia

¹⁰Racah Institute of Physics, Hebrew University, Jerusalem, Israel

¹¹Univ. Lyon, ENS de Lyon, Univ. Lyon I, CNRS, Centre de Recherche Astrophysique de Lyon, France

¹²Institute for Astronomy (IFA), University of Hawaii, USA

¹³Department of Physics, Massachusetts Institute of Technology, USA

SuperMUC Project ID: pr62ju

Introduction

The neighborhood in the immediate cosmological vicinity of the Milky Way is known as the “Near-Field”. Its main galaxies are the Milky Way-Andromeda pair, accompanied by an entourage of smaller galaxies including the Large and Small Magellanic Cloud, the Triangulum galaxy (M33) and a flock of low luminosity dwarfs. Beyond ~ 3 Mpc, there are a few similarly sized large galaxies (e.g Centaurus A, M83, M81, etc). At around 15 Mpc, the Virgo galaxy cluster dominates the scene. A dark matter filament detected by the coherent motion of galaxies stretches from us to Virgo. The entire structure described above is confined to the super-galactic plane, which is itself abutted by the Local Void. On even larger scales, the Virgo cluster and the entire local volume forms part of the local super cluster, Laniakea, which is itself embedded in the “Cosmic Web”. Our near field is the best-observed region of the universe. As such it is the most interesting to study, to understand and to try and simulate.

The issue at stake here is of a Copernican nature. To what extent the Local Group (LG) is fair representative of objects of its kind, and thereby to what extent its observed properties represent LG-like objects at large. This where constrained simulations of the LG are making the difference—by controlling both the actual local environment and by experimenting with the physics of galaxy formation, on the computer.

Cosmological structure formation is a highly non-linear process that obeys the complex (yet well understood) physics of multi-body gravitational dynamics combined

with the multi-scale and more complicated behavior of the baryons. Since these processes are essentially intractable analytically, numerical simulation, starting from initial conditions defined by the cosmological concordance model, is the preferred method of inquiry. Yet, the ability of numerical methods to simulate baryons from first principles is limited. One of the most important processes—star formation—cannot be simulated in a cosmological context for two reasons: the dynamical range is too high to handle and we still understand little of star formation since it is a fundamentally hidden process occurring deep inside dust laden giant molecular clouds that are impermeable to most wavelengths. Furthermore the life cycle of baryons is immensely complicated and involves process of heating, cooling, stellar evolution and stellar interactions, phase transitions and a variety of astrophysical winds. In cosmological simulations this process is thus modeled in a phenomenological way. Despite their limitations, numerical simulators have risen to the challenge and a number of codes and techniques have been created to cosmologically model baryon physics with great success. Simulators have succeeded in creating galaxies in a computer that match many properties to those observed in the heavens

Results and Methods

Our program seeks to use measurements of the local peculiar velocity field (obtained by comparing “standard candle” distance measures with redshifts) to obtain $z = 0$ cosmographic maps from which initial conditions for cosmological simulations can be




Figure 1: The large scale environment, including the Virgo cluster, the local Void and the local filament is accurately reproduced in these simulations (left). The gas distribution in the Local Group shows two disc galaxies (middle panel). The stellar component of these galaxies can be viewed in a mock Hubble Spac Telescope image of the Milky Way and the Andromeda galaxies (right).

constructed. The technique of such constrained realizations was first suggested by Hoffman & Ribak (1991) and has been used before to successfully reproduce LGs analogues in the correct environment (e.g. Libeskind et al 2010). Simply put, given a data set and a prior model (e.g. LCDM with Gaussian fluctuations) which specifies a covariance matrix of the field that is to be reconstructed, the Wiener Filter seeks a linear minimum variance solution. In the linear regime (~ 3 Mpc) the velocity and density fields are related by a simple convolution. Thus reconstructing the density field from the observed velocity field requires a deconvolution in the presence of noise (see Zaroubi, Hoffman & Dekel 1999). Obscured or incomplete regions can be handled, although the reconstructed Wiener Filter density field tends to zero (the mean) at the survey's limits, where the number of constraints goes to zero. The observed (linear) velocity field (reconstructed with a Wiener Filter) is translated into a linear Zeldovich displacement field which can be reversed in time. This “reversed Zeldovich approximation” (Doumler et al 2012) combined with the Wiener Filter and Constrained Realization technique has already proved its ability to produce ICs that are fit to task (see Sorce et al 2014).

Applying this method to the CosmicFlows catalogue of peculiar velocities returns initial conditions that result in a density and velocity field that looks very similar to the near field. We then embark on a process of running these initial conditions at increasing resolution in order to hone in on the best and highest resolution runs possible. The simulation set up is always the same: a periodic box is filled with particles. In the center of this box exists a spherical region filled with many more particles of higher resolution (lower mass) embedded in a low resolution background. We refer to the resolution of the simulation as the number of particles in the high resolution region, if they spanned the whole box.

The process commences with around 1000 low resolution dark matter only (N-body) simulations, set up as described in the previous paragraph with an effective resolution of 512^3 particles. The $z = 0$ snapshot is then analysed and the cosmographically successful simulations identified. The ICs are then regenerated with a higher resolution of $4,096^3$ effective particles. Small scale power is added randomly by sampling the Λ CDM power spectrum at the appropriate scales. These ICs are then run with full hydrodynamics. Gas cells are added by splitting each DM particle into a DM-gas cell pair with a mass determined by the cosmic baryon fraction. The pair is separated by a distance equal to half the mean inter-particle separation. The centre of mass and its velocity is left unchanged. In order to go to even higher resolution only three ICs are selected (owing to the high computational cost). The high resolution region is then populated with $8,192^3$ effective particles and run with gasdynamics and galaxy formation.

In Table 1 and 2 we provide some information on the data produced and computational resources used in the project. The Dark matter only simulations – used to check the fidelity of the cosmographic landscape before hydrodynamic simulation, are orders of magnitude cheaper and smaller.

The result of the procedure is a set of 13 intermediate and 3 high resolution cosmological magneto-hydrodynamical simulation of the LG within the dynamical environment as described by measurements of the peculiar velocity field. The Virgo cluster is reproduced in the correct location and at the correct distance as is a void region analogous to the Local Void. A filamentary structure, abutted by the simulated Local Void and stretched from the LG to the Virgo cluster is also reproduced. Within this environment, pairs of galaxies are found at roughly the correct distance from each other and from Virgo, with roughly the correct mass and mass ratio, negative relative velocity, and with cold gas discs.

	4,096 ³	8,192 ³
Number of runs	16	3
Core-hours (m)	7.1	21.92
Cores per job	1,024	4,000
Data produced (TByte)	49	58

Table 1: Properties of the hydrodynamic runs.

The simulated rotation curves of the LG analogs are in rough agreement with the observations. Dwarf galaxies, as well as galactic disc features (such as bars, spiral arms, etc) are resolved. The satellite galaxies have similar stellar mass distributions and radial concentration as observed. The Hestia runs also produce large Magellanic cloud and M33 like objects. The present HESTIA simulations successfully reproduce many of the observed features of the local universe on scales ranging from the $10^{14} M_{\odot}$ Virgo cluster more than 10 Mpc away, the local filament (within which the LG resides) on the scale of a few Mpc, the inter-LG sub-Mpc scale and all the way down to sub-virial scales within the MW and M31 galaxies. The large scales are strongly constrained by the CF2 velocity data and are robustly reproduced by the simulations. Going down to smaller and smaller scales, the constraining power of the data diminishes and the variance within the different realizations increases. The conclusion from this work is that the wide range of LG properties reported here—including the inner structure of the MW and M31, stellar and light distribution, distribution of satellites, environment—is reproduced within the constrained variance. Some of these properties are reproduced by some of the galaxies - not necessarily all of them.

Ongoing Research / Outlook

Our stratagem for simulating the LG opens the door to further studies on the nature of the LG and its formation. Three immediate routes are to be pursued: the increase of the ensemble of constrained simulations (thereby studying the statistical significance of our results); the increase of resolution (thereby studying the numerical robustness of the results); and experiments with the model of galaxy formation physics. These will enable a better understanding the physics of galaxy formation and to address the Copernican issue regarding our role in the universe. We intend to apply for more computing time in order to pursue higher resolution simulations.

The greatest limitation in pursuing this project was undoubtedly the extremely long queue times. We needed dozens of 48 hours periods per run. Often it took excessively long to get off the queue meaning that to be able to use the allotted time and to complete the project became a real challenge.

	512 ³	4,096 ³	8,192 ³
Number of runs	808	26	13
Core-hours (m)	0.38	0.21	0.25
Cores per job	80	320	1,024
Data produced (TByte)	3.5	6.1	13.2

Table 2: Properties of the Dark Matter only runs.

References and Links

- [1] Project website: <https://hestia.aip.de/>
 [2] N I Libeskind, E Carlesi, R Grand, A Khalatyan, A Knebe, R Pakmor, S Piliipenko, M Pawlowski, E Tempel, P Wang, H M Courtois, S Gottlober, Y Hoffman, I Minchev, C Pfrommer, J Sorce, V Springel, M Steinmetz, R B Tully, M Vogelsberger, G Yepes. The Hestia Project: High resolution Environmental Simulations of The Immediate Area, MNRAS 498 (2), 2968-2983, 2020. DOI: 10.1093/mnras/staa2541.

Cracking the Convective Conundrum

1

RESEARCH INSTITUTION

¹Institut für Astrophysik, Georg-August-Universität Göttingen

PRINCIPAL INVESTIGATOR

Petri Käpylä¹

RESEARCHERS

Axel Brandenburg², Frederick Gent³, Maarit Käpylä^{3,4}, Nigul Olsper³, Federico Spada⁴, Mariangela Viviani⁴

PROJECT PARTNERS

²NORDITA, Stockholm, Sweden

³Aalto University, Espoo, Finland

⁴Max-Planck-Institute for Solar System Research, Göttingen, Germany

SuperMUC Project ID: pr27li (Gauss Large Scale project)

Introduction

The envelopes of stars such as the Sun are convectively unstable such that fluid motions transport the energy in these regions. These motions are highly turbulent by virtue of the relatively low viscosity and the immense spatial scales involved. The stellar matter is also electrically conductive. In addition, most stars rotate and the fluid is strongly density stratified. For example, in the Sun the pressure scale height increases from 100 km near the surface to about 50,000 km at the base of the convection zone at 0.7 solar radius. The combination of turbulence, rotation, and stratification are conducive to dynamo action, such that fluid motions maintain dynamically important magnetic fields.

Recent observational and numerical results suggest that the convective velocities in simulations are consistently larger than those in the Sun. This has been found from comparisons of simulation data to helioseismic results which rely on the imprints of interior flows on sound waves in the Sun. Furthermore, global simulations with solar luminosity and rotation

rate typically produce large-scale differential rotation where the equatorial regions rotate slower than the poles or so-called anti-solar differential rotation. This is opposite to what is observed in the Sun where the equator rotates faster than the poles (solar-like differential rotation). The sense of the differential rotation depends on the rotational influence on the flow which is measured by the ratio of convective turnover time and rotation period or the Rossby number. Anti-solar (solar-like) differential rotation is expected for Rossby number roughly greater (smaller) than unity. The Sun is relatively close to this transition but simulations consistently land on the anti-solar regime, indicative of too fast convective flows.

In the current project a new dynamical recipe for radiative diffusion using Kramers opacity law was used. The Kramers opacity law leads to a smooth transition between convective and radiative layers and practically always produces an intermediate layer (Deardorff zone), where the stratification is convectively stable according to the Schwarzschild criterion but still convective with a positive convective energy flux. In such layers, a non-local non-gradient




Figure 1: Flows in a low (288^3 grid, left) and high ($1,152^3$, right) resolution Cartesian simulations of overshooting convection [2].




Figure 2: Normalized overshooting depth d_{os} as a function of normalized energy flux \mathcal{F}_n [2].

contribution, discovered by James Deardorff in the 1960s, is dominant in the convective transport. The main goal of the project was to study the effects of Kramers opacity law on the properties of convective overshooting, differential rotation, and dynamos in simulations of stellar convection.

Results and Methods

Fully compressible Cartesian and spherical wedge simulations of stellar convection were made with the Pencil Code [1]. The Pencil Code is a high-order finite-difference code written in Fortran 95/2003 for solving ordinary and partial differential equations. The code is parallelized with MPI and domain decomposition. The main application area of the code is astrophysical magnetohydrodynamics (MHD) but it has been used in a wide variety of physics applications such as turbulent combustion, planet formation, and gravitational waves.

The hydrodynamic Cartesian simulations [2] consisted of a convectively unstable layer on top of a stably stratified radiative layer. The main target was to study the dependence of convective overshooting on the

energy flux (luminosity) of the star. Using a realistic luminosity in fully compressible simulations of deep stellar convection is infeasible due to the prohibitively long thermal equilibration timescale. Instead, a range of higher normalized fluxes \mathcal{F}_n were probed. Such simulations yield a power-law dependence for the normalized overshooting depth d_{os}/H_p , where H_p is the pressure scale height. For the most realistic cases with Kramers' opacity, d_{os} is proportional to $\mathcal{F}_n^{0.08}$ (black, gray, and blue lines in Figure 2). This is much shallower than earlier numerical estimates ($\sim \mathcal{F}_n^{0.3}$). Extrapolating to the solar luminosity ($\mathcal{F}_n \approx 10^{-11}$) suggests an overshooting depth of the order of $0.1 H_p$ or about 5,000km. Furthermore, the difference to the earlier studies can be explained by a difference in the Prandtl number regimes in the respective studies. In the simulations in [2], up to 40% of the convectively mixed layer was formally stably stratified (Deardorff and overshoot zones). The bulk of the simulations in [2] used a modest resolution of 288^3 grid points with 288 cores. However, due to the long thermal saturation timescales, several simulations required up to two months wallclock time.

Furthermore, of the order of a hundred low resolution (288^3) simulations were made along with subsets a number of runs at 576^3 (576 cores) and $1,152^3$ grid points (2,304 cores) to verify the validity of the low resolution models.

Spherical wedge simulations (Figure 3) were used to probe the effects of Kramers opacity in a more realistic geometry [3,4]. These were the first (semi-)global simulations where the dynamically adapting Kramers opacity is used. The main result of [3] was that although the simulations developed a Deardorff layer, it was much thinner than in the Cartesian cases which is due to the strong rotation and less supercritical convection in the spherical cases. Furthermore, the issue regarding the convective velocity amplitudes is essentially unaffected in these simulations. That is, to achieve a solar-like rotation profile, the rotation rate of the star has to be at roughly twice the actual solar




Figure 3: Flows in representative non-rotating (left), rotating (middle) and rotating MHD (right) simulations of stellar convection in spherical wedge geometry [3].

rotation rate. The MHD cases lead to self-consistent dynamo action and large-scale magnetic fields. Solar-like cyclic dynamo solutions, where the large-scale magnetic fields migrate from mid-latitudes toward the equator in the course of the cycle, are achieved in a narrow parameter regime (Figure 4). However, the magnetic cycle period in these simulations is roughly 3.5 years instead of 22 years as in the Sun. In a subsequent publication [4], the effects of several modeling choices such as boundary conditions were studied. We found that while the magnetic boundary conditions and the centrifugal force have only a minor influence, the results are sensitive to thermal boundary conditions. These preliminary spherical wedge simulations were mostly computed at a modest resolution of $144 \times 288 \times 144$ grid points with 144 cores. However, the integration times were of the order of 30 days per run. Higher resolution spherical wedge simulations in more realistic parameter regimes will be presented in follow-up work which is currently in preparation. The project used in total 35 M core hours and 55 TB of disk space.


Ongoing Research / Outlook

The SuperMUC-NG resources were crucial for the completion of the studies discussed above. The large-scale projects in SuperMUC-NG dedicate targeted resources to endeavors tackling highly challenging physics problems. Our project did not face serious challenges. The project continues currently on SuperMUC-NG with a new allocation of 30 M core hours granted in March 2020. The current simulations continue to probe the fundamentals of turbulent convection with higher resolutions and with a particular emphasis on astro-physically more relevant low Prandtl number regime.

References and Links

- [1] <https://github.com/pencil-code>
- [2] Käpylä, P.J. (2019), *Astronomy & Astrophysics*, 631, A122. DOI: 10.1051/0004-6361/201834921
- [3] Käpylä, P.J., Viviani, M., Käpylä, M. J., Brandenburg, A. & Spada, F. (2019), *Geophys. Astrophys. Fluid Dyn.*, 113, 149. DOI: 10.1080/03091929.2019.1571584
- [4] Käpylä, P.J., Gent, F. A., Olsper, N., Käpylä, M. J. & Brandenburg, A. (2020), *Geophys. Astrophys. Fluid Dyn.*, 114, 8. DOI: 10.1080/03091929.2019.1571586

Figure 4: Six equidistant snapshots of the magnetic field in an MHD simulation with solar-like differential rotation and equatorward migrating magnetic activity. [3].



From 3D Progenitors to 3D Supernova Explosions

Including Muon Physics

RESEARCH INSTITUTION

¹Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, Garching

PRINCIPAL INVESTIGATOR

Hans-Thomas Janka¹

RESEARCHERS

Robert Bollig¹, Daniel Kresse^{1,2}, Naveen Yadav^{1,3}, Alexander Heger⁴, Bernhard Müller⁴

PROJECT PARTNERS

²Physik Department, Technische Universität München

³Excellence Cluster ORIGINS, Garching

⁴Monash University, Melbourne, Australia

SuperMUC Project ID: pr53yi (Gauss Large Scale project)

Introduction

After millions of years of stable evolution, the life of a massive star is abruptly terminated by a spectacular event that outshines the light of billions of stars for weeks and is called core-collapse supernova (CCSN). Such events are triggered when the core of the star, which consists mostly of iron, becomes gravitationally unstable and collapses to a neutron star within less than a second. The sudden halt of the implosion at the moment when the neutron star begins to form, which is called “core bounce”, creates a shock wave that ultimately expels the overlying stellar layers in the supernova explosion.

Neutron stars are among the most extraordinary objects in the Universe. With the diameter of a city like Munich they contain more mass than our Sun, compressed to densities higher than that of atomic nuclei. They start out as extremely hot objects with temperatures up to nearly 1000 billion Kelvin and emit huge numbers of neutrinos. These weakly interacting elementary particles carry away the gravitational binding energy that is released in the catastrophic collapse of the stellar core, 100 times more energy than is set free by the star’s explosion, and 10,000 times more energy than is radiated in the splendidly brilliant light of the supernova.

Enormous progress has been achieved over the past five years in our understanding of the physical mechanism that reverses the catastrophic collapse of the stellar core to the powerful supernova blast. Neutrinos have now been established as the crucial agent that transfers the power to the supernova blast, thus confirming a 50 year old hypothesis of a “neutrino-driven mechanism” first outlined in a paper by Colgate and White in 1966.

This breakthrough became possible because of decisive advances on three fronts simultaneously. First, neutrino interactions in high-density media, thermodynamic properties of hot nuclear matter, and

hydrodynamic phenomena such as convective transport, shock-wave instabilities, and turbulence, all of them playing a crucial role in the new-born neutron star and starting explosion, have been explored in great depth. Second, computationally efficient and accurate numerical tools have been developed to solve the coupled system of hydrodynamics solver and energy-dependent neutrino transport in three spatial dimensions (3D) and with the modern microphysics taken into account. And, third, massively parallel supercomputers have become available to apply these tools to full-scale, self-consistent 3D supernova simulations.

Because of access to SuperMUC at LRZ, the CCSN group at MPA Garching was worldwide first to demonstrate the viability of the neutrino-driven mechanism with modern 3D models in 2015 [1]. After this confirmation in principle, the focus has now shifted to better qualitative and quantitative insights into the role of the initial conditions in the progenitor star, the spatial resolution of the numerical models, and so far disregarded aspects of the relevant physics.

Results and Methods

In the course of this project the MPA group has taken two major steps towards more realism of CCSN modeling. First, we have started to compute initial conditions of progenitor stars in 3D, simulating the latest stages of convective oxygen-shell burning over periods of several minutes prior to iron-core collapse. This creates a highly asymmetric distribution of the chemical elements (in this case of O, Ne, Si) with large-scale fluctuations of density and velocity in the stellar layers surrounding the iron core (Figure 1), replacing the traditionally employed spherically symmetric (1D) pre-collapse conditions provided from stellar evolution modeling. Second, we have included, for the first time and uniquely, the effects of muons in the thermodynamical description of the neutron-star plasma as well as neutrino-muon interactions in the neutrino transport. Muons are about 200 times heavier

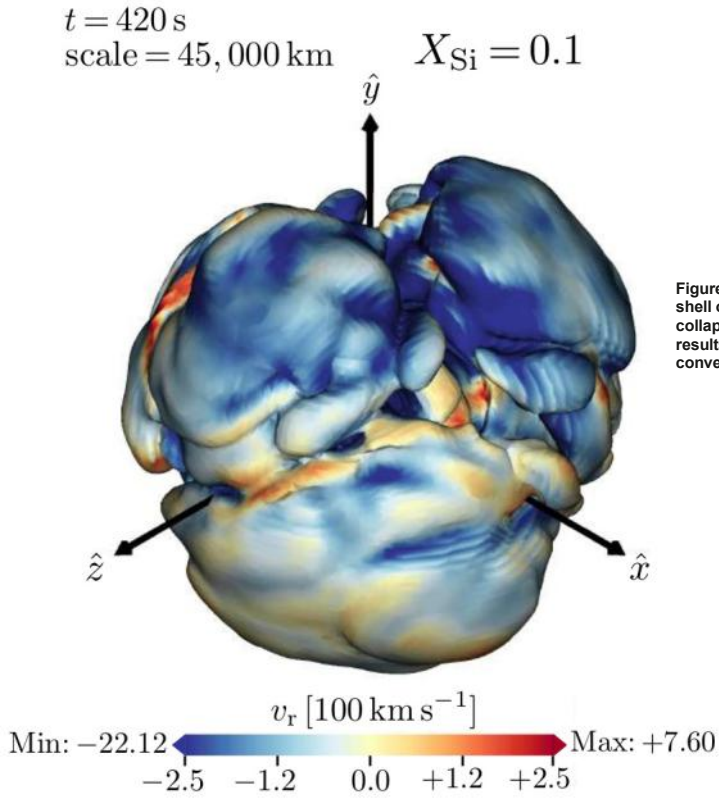


Figure 1: 3D geometry of silicon in the convectively burning oxygen-neon shell of a 19 solar-mass star, defining the initial condition for the stellar core collapse. The isosurface corresponds to a silicon mass fraction of 10% and results from 7 minutes of 3D evolution. Color coding denotes radial convective velocities [2].

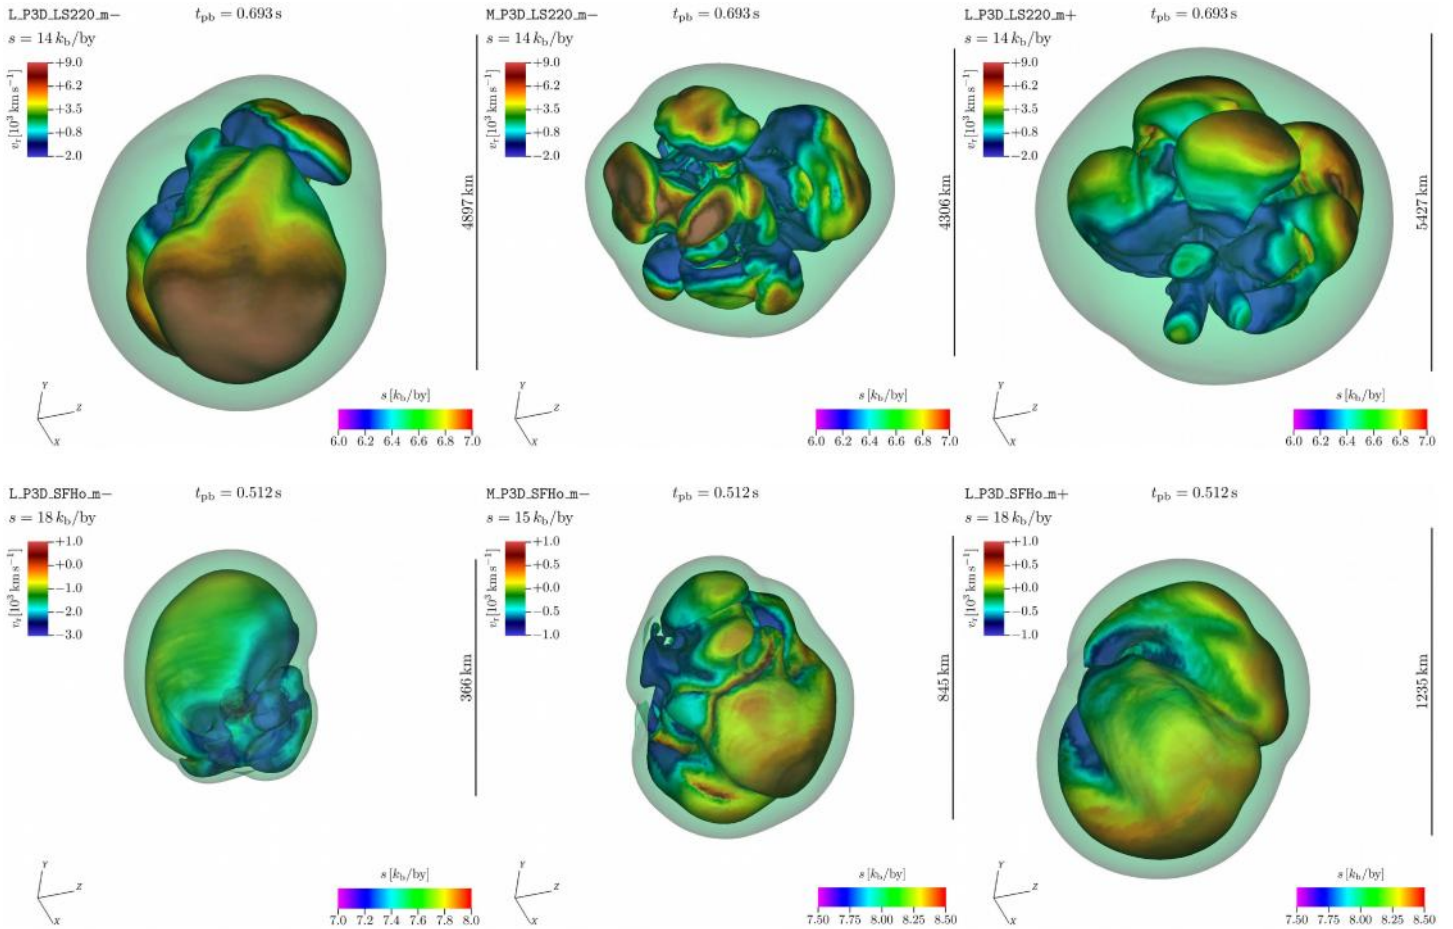


Figure 2: Early stages of neutrino-driven explosions of a 19 solar-mass star for two different theories of the nuclear plasma in the hot neutron star (top for LS220 and bottom for SFHo, each for a common evolution time). Simulations with 4-degree angular resolution (left) are compared with 2-degree simulations (middle), both without muon physics, and 4-degree simulations with muons (right) [3].




Figure 3: Development of the neutrino-driven explosions of a 19 solar-mass star (with the SFHo nuclear plasma model), visualized by the time evolution of the average shock radii (left) and diagnostic explosion energies (right). “L” denotes 4-degree and “M” 2-degree models, and suffix “m-” means without muons, “m+” with muons.

than electrons (105.66 MeV), for which reason they had been thought to be created in insufficient numbers to have any influence on supernovae. To take the corresponding physics into account, our Vertex neutrino transport code had to be generalized for a 6-species treatment with couplings between neutrinos and antineutrinos of the three lepton flavors across the whole energy-momentum space. This generalization increases the computational demand of the anyway expensive neutrino physics by another 30–50%.

In addition to studying these two novel aspects connected to the 3D progenitor structure and microphysics in hot, new-born neutron stars, we also performed corresponding 3D simulations with different angular resolutions for convergence tests: Besides 2-degree models as our default, we compared to 4-degree runs (Figure 2) and, without the expensive muon physics, also 1-degree runs. Our set of 16 simulations was done on 16800, 4992, and 27,072 cores, respectively, consumed about 190 million core hours over 3 years, and delivered roughly 3 PB of archived data for the science results.

Our findings contradict the traditional thinking that muons do not play a role. Temperature and electron degeneracy (scaling with the high density) in the neutron star interior become so huge that muons can be formed in neutrino-electron reactions and thermal processes. Their appearance converts internal energy to rest-mass energy, which does not contribute to the pressure. Therefore this so-called muonization process leads to a faster contraction of the hot neutron star, boosting the neutrino emission and thus the energy transfer to the supernova shock, strengthening the neutrino-driven mechanism. This accelerates the onset of the explosion and leads to a faster increase of the explosion energy (Figure 3).

We found that muon effects thus overcompensate the negative consequences of lower resolution, which tends to damp the growth of hydrodynamic instabilities (convection and turbulence) in the volume behind the shock due to enhanced numerical viscosity. Since these instabilities support the neutrino mechanism, the

damping influence of numerical viscosity can delay the onset of the explosion (Figure 3).

None of our simulations that were started from the 1D progenitor model ended in a successful supernova explosion. This clearly emphasizes that more realistic 3D initial conditions, constructed by our own 3D simulations of convective oxygen-shell burning, are a crucial ingredient, because the large-amplitude and large-scale asymmetries in the pre-collapse star stimulate stronger postshock flows when falling through the shock.

Ongoing Research / Outlook

Two papers with results have been finished [2,3], but the analysis of the wealth of data is still ongoing and further publications will follow, specifically for the muon and resolution effects. Unfortunately, the 1-degree models (both without muons) were too expensive to be continued beyond about 0.3 seconds after neutron-star formation. More computer time is needed to cover longer evolution periods with these 1-degree simulations, which also could not yet include our unique muon physics. SuperMUC-NG will be indispensable for this spearheading research.

References and Links

- [1] <https://www.mpa-garching.mpg.de/84411/Core-collapse-supernovae>.
- [2] Yadav, N., et al., ApJ 890, 94 (2020); arXiv:1905.04378.
- [3] Bollig, R., et al. Self-consistent 3D supernova models from -7 minutes to +7 seconds: a 1-bethe explosion of a ~19 solar-mass progenitor. ApJ, submitted. arXiv:2010.10506.

Towards Energy Saturation in 3D Core-Collapse

Supernova Simulations

RESEARCH INSTITUTION

¹Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, Garching

PRINCIPAL INVESTIGATOR

Hans-Thomas Janka¹

RESEARCHERS

Robert Bollig¹, Daniel Kresse^{1,2}, Georg Stockinger^{1,2}, Naveen Yadav^{1,3}, Alexander Heger⁴, Bernhard Müller⁴, Ken'ichi Nomoto⁵

PROJECT PARTNERS

²Physik Department, Technische Universität München, James-Franck-Str. 1, Garching

³Excellence Cluster ORIGINS, Boltzmannstr. 2, Garching

⁴Monash University, Melbourne, Australia

⁵University of Tokyo, Tokyo, Japan

SuperMUC Project ID: pn69ho, pr53yi (Gauss Large Scale project)

Introduction

This project explores the fundamental physical processes that cause and accompany the violent death of stars with at least nine times the mass of our Sun in so-called core-collapse supernova explosions [1]. Of particular interest are predictions of measurable signals such as neutrinos (elementary particles that interact with matter by the weak force) and gravitational waves (spacetime perturbations caused by asymmetric acceleration of masses). Both are emitted from the center of the explosion and can yield direct information of the processes there. Another aspect of great interest is the explanation of observed properties of supernovae and their gaseous and compact remnants, i.e., of neutron stars and black holes that are formed when the stellar iron cores collapse before the supernova blast expels the rest of the star into the circumstellar space. Moreover, the consequences of supernovae in our Universe are in the focus of intense research, for example connected to their ejection of life-enabling chemical elements such as carbon, oxygen, silicon, and iron, which are either forged by nuclear reactions over millions of years in those stars that finally explode, or which are freshly made by the supernova itself.

To address these important and timely questions of stellar astrophysics, it is crucial to develop a most detailed understanding of the physical mechanism that causes the stellar explosion. To this end and for the theoretical study of all of the mentioned problems, the most powerful available supercomputers are needed, because the involved processes are extremely complex, enormously diverse, and highly non-linear. These processes involve, tightly coupled, the hydrodynamics and thermodynamics of the stellar plasma, neutrino interactions and transport, nuclear reactions, as well as strong gravity to be described by Einstein's theory of general relativity.

Results and Methods

With the support by computing resources on SuperMUC and SuperMUC-NG of LRZ over the past years, the Garching team has achieved to demonstrate, for the first time with modern numerical methods and state-of-the-art input physics, that energy transfer by neutrinos triggers the onset of the explosion and powers the supernova blast. Despite their weak interactions with matter, these particles are produced in huge numbers at the extreme densities (up to several times the density of atomic nuclei) and temperatures (up to nearly 1,000 billion Kelvin) in the newly formed neutron star. While 99 percent of the neutrinos escape within a few seconds and thus carry away the gravitational binding energy that is released in the collapse of the stellar iron core, the remaining one percent gets stuck by particle interactions behind the supernova shock wave and powers the supernova explosion.

In this project the goal was to explain the observed properties of well studied, nearby, young supernovae and their remnants by neutrino-driven explosions. In particular, we focused on the supernova that blew up in the year 1054 and gave birth to the *Crab Nebula* with its rapidly spinning neutron star (a so-called "pulsar"), and, as a second target of interest, on *Supernova 1987A*, which exploded only 33 years ago and whose neutron star has been detected recently. The former object is thought to originate from the death of a low-mass star of about 9 to 10 solar masses, the latter case is connected to a star of 15 to 20 solar masses.

We therefore performed the first self-consistent 3D simulations that followed the evolution of such stars from the onset of stellar core collapse, through the neutrino-driven expansion of the supernova shock, to the time when the explosion energy reached its terminal value. In the cases of 8.8, 9.0, and 9.6 solar-

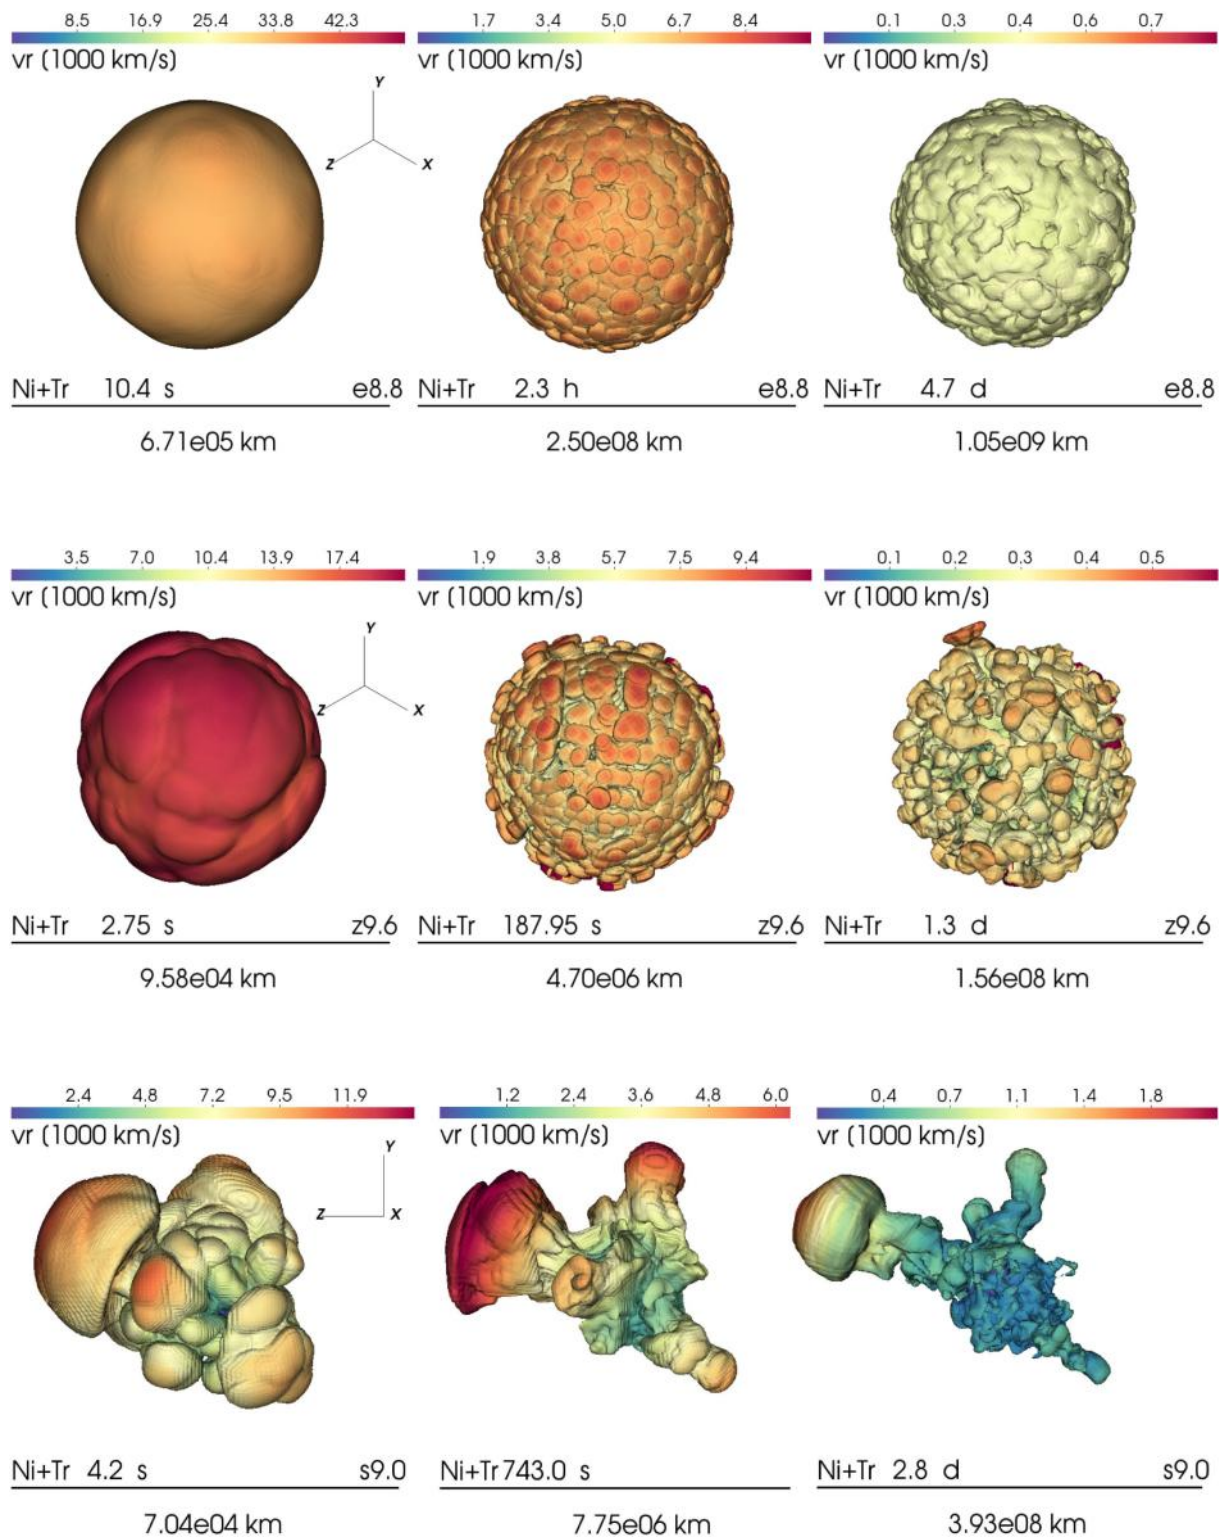


Figure 1: 3D geometry of matter ejected by neutrino-driven supernovae of low-mass progenitor stars with 8.8 (top), 9.6 (middle) and 9.0 (bottom) times the mass of the Sun. The images display the spatial distribution of radioactive nickel-56 that is produced during the explosion, at a few seconds just after the nickel production (left row), several minutes to hours later (middle row), and with the fully developed shape after some days [2].




Figure 2: Development of the neutrino-driven explosion of a 19 solar-mass star from the onset of the blast at 0.45 seconds (left), to 1.675 seconds (middle), and 7.034 seconds (right) after the stellar collapse. At 7 seconds the energy of the explosion approaches its saturation value of 10^{44} Joule, explaining observed supernovae [3].

mass stars, we continued until shock breakout from the stellar surface in order to determine the final asymmetry of the supernova and the final neutron star properties after the fallback of matter that does not become unbound in the explosion (Figure 1). In the case of a 19 solar-mass star it took already more than 7 seconds for the blast-wave energy to saturate (Figure 2). Thus we could demonstrate, for the first time by self-consistent models, that neutrino-driven explosions can explain the properties of the Crab supernova and its pulsar and of *Supernova 1987A*, i.e., their explosion energy, ejected mass of radioactive nickel, neutron star mass, kick velocity, and spin period [2,3].

Such simulations are a challenging problem, which could not be tackled before. On the one hand, because the explicit time step of the hydrodynamics solver is as small as some 10^{-7} seconds, which means that more than a million time steps are needed. On the other hand, following the growth of the energy requires the inclusion of the computationally even more demanding neutrino physics. The simulations employed our *Vertex-Prometheus* neutrino-hydrodynamics supernova code, which includes a modern shock-capturing hydrodynamics scheme based on the piecewise parabolic method (PPM) proposed by Colella and Woodward, and a two-moment ray-by-ray-plus neutrino transport solver with Boltzmann closure. They became possible for several reasons. First, the code uses a hybrid OpenMP/MPI parallelisation model that allows perfect linear scaling for up to more than 200,000 cores, as demonstrated at a dedicated workshop, and it also uses rigorously optimized vectorization on SuperMUC-NG. Second, the neutrino transport is time-implicit with recently accelerated and improved convergence, numerical stability and accuracy, which now allows for time steps that are up to 100 times bigger than the hydrodynamics steps. Third, for long-time runs over many seconds and beyond, a computationally cheaper neutrino scheme was implemented in the course of this project. It is based on data obtained from 1D neutrino-cooling simulations of the nascent neutron star and allows for a seamless, basically transient-free continuation of the 3D explosion simulations (see appendix in [2]).

A typical supernova run was done with up to 1,000 radial zones and 2-degree angular resolution on 16,800 cores (350 nodes), using several 10 million core-hours. For convergence tests and varied microphysics inputs, also runs with 4-degree and 1-degree resolution were performed on 4,992 and 54,144 cores (1,128 nodes), respectively, but because of unacceptably frequent node failures only a maximum of 27,072 cores could be used in production runs. The two projects pn69ho and pr53yi involved research of three PhD students, consumed nearly 224 million core hours over 3 years, delivered on the order of 10,000 files, needed up to 300 TB of continuous work storage, and produced about 3 PB of archived data for the science results.

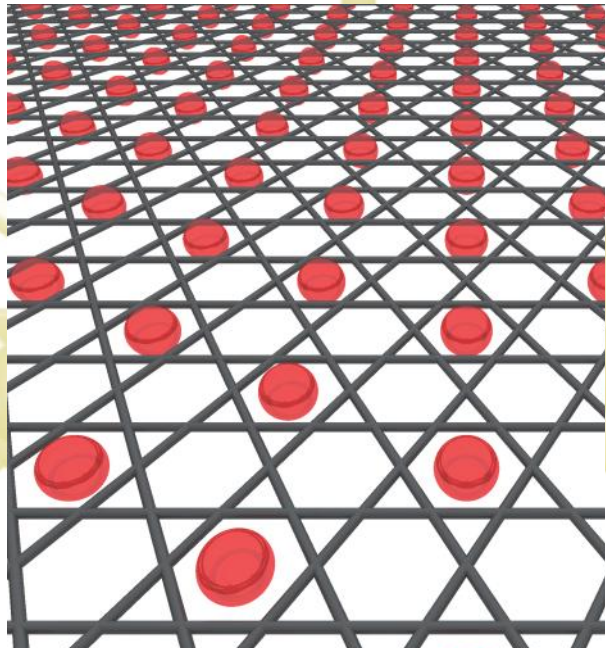


Ongoing Research / Outlook

Data analysis of the results is still going on, also with collaborators, evaluating the outputs for the predicted neutrino and gravitational-wave signals, nucleosynthesis, as well as astrophysical implications such as supernova light curves and spectra. Moreover, the 19 solar-mass model has still to be continued from 7 seconds to several days in order to track the fallback and development of the final explosion geometry. Since supernova progenitors are diverse, more simulations of the presented kind for different stars will be needed in the future for comparison with observed supernovae and remnants. SuperMUC-NG will be indispensable for this forefront research.

References

- [1] <https://www.mpa-garching.mpg.de/84411/Core-collapse-supernovae>.
- [2] Stockinger, G., et al., MNRAS 496, 2039 (2020); arXiv:2005.02420.
- [3] Bollig, R., et al. Self-consistent 3D supernova models from -7 minutes to $+7$ seconds: a 1-bethe explosion of a ~ 19 solar-mass progenitor. ApJ, submitted. arXiv:2010.10506.

Chemistry and Material Sciences



Role of quantum coherence in plasmon-induced hot-carrier dynamics

2

RESEARCH INSTITUTION

¹University of Trieste, Trieste (Italy)

²University of Padova, Padova (Italy)

³CNR-Institute of Nanoscience, Modena (Italy)

PRINCIPAL INVESTIGATORS

Emanuele Coccia¹, Stefano Corni^{2,3}

RESEARCHERS

Giulia Dall'Osto², Margherita Marsili², Mauro Stener¹, Daniele Toffoli¹, Mirko Vanzan²

PROJECT PARTNERS

—

SuperMUC Project ID: pn34qi (PRACE project)

Introduction

A growing interest in the possible role of genuine quantum effects in the dynamics of complex systems is involving chemists, physicists and biologists over the last years. We focus on molecular nanoplasmonics, which studies the effects of a metal nanostructure on the electronic and optical properties close to it.

Plasmon-enhanced catalysis is a “hot” research field which allows researchers to access new potentialities in chemical activity. The idea is to couple (large) nanostructures with well-known plasmonic activity with (smaller) transition-metal clusters which are already used as efficient catalysts in heterogeneous chemical processes. These systems, called antenna-reactor complexes (Figure 1) show favorable selectivity compared with thermal processes.

In this Project [1] we focused on rhodium nanostructures, as photocatalysts for the CO₂ hydrogenation reaction. Rhodium nanostructure plays the role of antenna and reactor at the same time. This process at atmospheric pressure presents two competing pathways: CO₂ methanation (CO₂ + 4H₂ → CH₄ + 2H₂O) and reverse water gas shift (CO₂ + H₂ → CO + H₂O). Recent experimental studies on Rh nanocubes demonstrate that mild illumination of the Rh nanostructures both reduces the activation energies of the CO₂ hydrogenation reaction below thermal activation energies and leads to strong photo-selectivity favoring CH₄ production through the methanation pathway [2]. Generation and injection of hot electrons from Rh to the molecule are considered as key steps to explain this plasmon-driven selectivity.

The fundamental question we want to answer is: does the interplay between (de)coherence and plasmonic effects play a role in ultrafast plasmon-enhanced or plasmon-induced processes? In order to realistically describe the dynamics of such a complex system, multiscale hybrid approaches are mandatory: a




Figure 1: Schematic representation of plasmon-induced hot-electron generation in antenna-reactor complexes and injection into close molecules. Dephasing and relaxation effects are also given.

subsystem of interest is treated at quantum level (QM), while a classical but reliable description of the nanostructure is usually chosen. The grand-challenge here is to provide a multiscale model where the nanostructure is treated with an approach computationally affordable (for nanostructures of tens of nm) yet accurate enough to describe electron or hole transfer to the molecule. For Rh, the idea was to define a QM vertex of the nanocube, where the process takes place, together with CHO, which is the intermediate for both possible reactions (Figure 2).

Pragmatically speaking, the approach we have chosen for this Project is based on the use, as the QM region, of the molecule plus an affordable (a few tens of atoms) portion of the metal nanostructure, to be coupled with the rest of the nanostructure described classically. The QM cluster (plus the molecule) compared with the corresponding piece of the entire nanostructure suffers, in principle, by two main artifacts. The first artifact is given by fictitious quantum size effects, related to the confinement of the electronic wavefunction within the cluster. These effects are appropriate for a real metal cluster, but they should not be there for the model cluster, as its electrons are actually delocalized all over the plasmonic nanostructure. The second one is the

creation of a time-dependent polarization at the boundary between the QM and the classical part of the nanostructure upon application of an electromagnetic field, that does not exist for the real (whole) nanostructure. How much these two effects are important, which computational strategies can be used to mitigate them and whether they can be decreased enough to obtain useful calculations are open problems we have investigated here.

Results and Methods

The theoretical framework is given by multiscale hybrid methods. In terms of the QM method for the electronic excited states of the small subsystem, we have used DFT/TDDFT and the accurate Bethe-Salpeter equation (BSE), developed within the many-body perturbation theory. It provides a correlated description of electronic excitations, giving a picture of the electronic coupling at the molecule/nanocube interface with a quality higher than that of DFT/TDDFT, and comparable to high-level quantum-chemistry methods. At the same time the BSE approach is more computationally affordable than high-level quantum-chemistry approaches.

As reported in the Introduction, the Rh vertex and the CHO intermediate species are treated at quantum level, using TDDFT or BSE approaches. Rh nanostructure is classically represented as a continuum with a frequency-dependent dielectric function.

In order to include effects of the surrounding environment, we use the stochastic Schrödinger equation (SSE) [3-4], in the Markov limit. SSE is at the core of the present project, allowing us to include (quantum) effects of the environment, as dephasing, and relaxation effects in the model in a way that is particularly suitable for HPC. In our approach we combine three main ingredients: electronic structure calculations, a time-resolved formulation of the theory of open quantum systems, and a time-resolved classical description of the nanostructure response coupled to the time-dependent description of a plasmonic nanostructure. Embarrassing parallelism is easily achieved since SSE quantum trajectories are independent and uncorrelated (Figure 3): averaging over them is a post-processing step producing the same outcomes of the density-matrix approach.

The nanostructure is represented as a continuous body characterized by its response properties to electric fields, i.e. the external field and that generated by




Figure 3: SSE computational protocol.




Figure 2: Nanocube and zoom on the vertex, described quantum-mechanically.

a neighboring QM charge distribution. A general approach to simulate the nanostructure response is to define, as we did, a set of apparent charges on the surface of the Rh nanostructure, which is discretized in terms of a mesh of elements. This choice ensures an excellent compromise between accuracy and computational cost. We used around 30 millions of core-hours. Jobs on SuperMUC-NG are hybrid MPI/OpenMP. Each simulation produces around 1,000-2,000 MPI tasks, corresponding to 1,000-2,000 SSE trajectories. This number is large enough to get a small statistical error on the quantities of interest. In the postprocessing, we are interested in the time evolution of state populations and coherences, and of descriptor useful to investigate the electron dynamics (projected density of states, i.e. PDOS, overlap PDOS, transition contribution maps etc.). We have simulated the experimental conditions [2], by using blue and UV pulses to irradiate the Rh nanocube in presence of CHO. Preliminary results show that a combination of plasmon-like collective and single-electron behaviour is found in the first femtoseconds of the dynamics. Moreover, electron injection into CHO is observed, but a quantitative evidence that the CH_4 formation is favored in these conditions is still to be found. Effects due to the cut of the Rh nanostructure have been seen to be negligible, which provides an a posteriori confirmation of the reliability of the our approach.

Ongoing Research / Outlook

The systematic analysis of the SSE trajectories for Rh nanocube/Rh corner+CHO is still ongoing. Our goals are: i) verify that the application of a multiscale approach is reliable and accurate for a plasmonic system; ii) the role of electron dephasing and of iii) the role of state-to-state relaxation in the hot-electron generation and dynamics. Using a HPC system was fundamental to produce in reasonable human time thousands of SSE trajectories. Our SSE-based multiscale hybrid computational protocol was successfully applied to study the interplay between plasmon and coherence effects in hot-electron generation dynamics in catalysts. Next step will be to systematically study other antenna-reactor systems, as recently reported in high-impact experimental works.

References and Links

- [1] <http://www.tame-plasmons.eu/>.
- [2] Zhang, X.; Li, X.; Zhang, D.; Su, N.; Yang, W.; Everitt, H. Nat. Commun. 2017, 8, 14542.
- [3] Coccia, E.; Troiani, F.; Corni, S. J. Chem. Phys. 2018, 148, 204112.
- [4] Coccia, E.; Corni, S. J. Chem. Phys. 2019, 151, 044703.

Molecular vibrations reduce photovoltages

in organic solar cells

RESEARCH INSTITUTION

¹Center for Advancing Electronics Dresden, Technische Universität Dresden

PRINCIPAL INVESTIGATOR

Frank Ortmann^{1,2}

RESEARCHER

Michel Panhans¹

PROJECT PARTNERS

²Department of Chemistry, Technische Universität München

SuperMUC Project ID: pn56hu

Introduction

Studies of exciton signatures in UV-Vis absorption of condensed matter have a long history. In semiconductors, these studies usually focus on the excitation energy, the absorption strength or the polarization direction to reveal the characteristics that are relevant for optoelectronic functionality. Here the application in solar cells or light emitting devices are a strong driving force for this research ever since. This is similar in organic semiconductor devices for which excitonic effects are much more pronounced and for which the nature of the excitons (molecular or charge-transfer (CT) excitons) and their connection to performance limitations of organic solar cells is of great interest. CT states are electronic states formed at the interface between organic electron donating and electron accepting molecules. They can be precursors for free charge carriers, or can decay to the ground-state with the emission of a photon. CT absorption tails are comparably broad but the underlying molecular and supramolecular properties responsible for the spread of energies (which causes broadness) are currently not well understood.

Results and Methods

We study the microscopic origin of exciton bands in molecular blends (see Figure 1) and investigate their role in organic solar cells (OSCs) [1,2,3]. We simulate the temperature dependence of the excitonic density of states (EDOS) and low-energy absorption features, including low-frequency molecular vibrations and multi-exciton hybridization (see Figure 2). For donor-acceptor blends with CT excitons, our numerical calculations agree very well with temperature-dependent experimental absorption spectra, i.e. we are able to predict the low-energy spectral lineshape. Based on these results, we further connect the lineshape to radiative voltage losses that reduce the solar cell's performance and should be as low as possible. In addition, we unveil that the quantum effect of zero-point vibra-




Figure 1: Generation of electron-hole pairs (excitons) after light absorption in a diluted donor-acceptor blend composed of two different organic molecular species. Different types of excitons lead to exciton bands in the optical absorption spectra.

tions, mediated by electron-phonon interaction plays a significant role in this context. It causes a large exciton bandwidth and reduces the open-circuit voltage. This is predicted from electronic and vibronic molecular parameters, which we simulated beforehand, and we find that this zero-point effect is surprisingly strong even at room temperature and can substantially reduce the OSC's efficiency.

To circumvent such radiative voltage losses, we elaborate on a better understanding of their connection to OSC parameters such as the intermolecular electronic coupling or the driving force for CT exciton formation. Radiative open circuit voltage losses are affected by exciton hybridization and delocalization and can be reduced by 100 mV or more, depending on the system.

We finally extend our theoretical predictions for a larger set of systems including non-fullerene donor-acceptor blends and different heterojunction and interface geometries of OSCs. Our proposed strategies to reduce the radiative voltage losses are supported by

excellent agreement of the present simulation approach to existing experimental data for the low-energy absorption [3].

The numerical methods and algorithms that run on SuperMUC-NG are implemented in home-made codes and no third-party code has been involved. Our code is written in Fortran90/95. It implements a domain decomposition scheme for basic matrix-vector multiplications (MVMs), where each domain is handled by an MPI process. We use pure MPI where the implementation is performed in such a way that communication is initiated and the part of the MVMs that are internal to the MPI-process are carried out subsequently. When this internal work is finished, the communication is finalized and the external part of the MVM is calculated. In this way, communication can operate in parallel to computation on the nodes. The only external library that is required is MPI. In the simulations, we calculate expectation values in an energy-resolved fashion by using an energy projection scheme [4]. This is performed by the Lanczos approach.

The key physical quantity that we calculate, is the EDOS (cf. Figs. 2c,e). All physical quantities are represented in an effective real-space basis and the supercell approach is implemented. More specifically, a tight binding parametrization for the electronic structure is used which allows to represent the Hamiltonian in a relatively sparse way.

The numerical methods and implementations have been developed by the Ortmann group over the last decade [5]. This includes the developed MPI version of the code. The numerical calculations of the EDOS are performed by a Lanczos approach [4] and a subsequent continued fraction expansion to obtain the spectral shape for the EDOS for the studied organic donor-acceptor blend models. A single EDOS calculation of a supercell with around 40,000 molecules (i.e. a total vector length of $1.6 \cdot 10^9$) with 4,800 cores per run requires about 9,000-10,000 core-hours. Besides small output log files the job output of a single EDOS calculation reduces to the file containing the exciton spectrum, which is only of MB size. For the above mentioned systems of 40,000 molecules a memory capacity of around 6-12 TB is needed during runtime that depends on the particular realization of the Hamiltonian.

Ongoing Research / Outlook

Both machines SuperMUC-Phase 2 and SuperMUC-NG are perfectly suitable for our simulations of the EDOS since they provide sufficiently large memory per core and a high speed for massive parallelization that we can exploit to run our code fast and in a highly parallel manner. Due to the high degree of parallelization, most of our jobs require a walltime of only about two hours, which reduces the queue time and enables efficient and fast data production.

In the next phase of our project, subsequent electron-transport simulations are planned in which we are using a code that is directly derived from the present one. The study of electronic transfer processes will

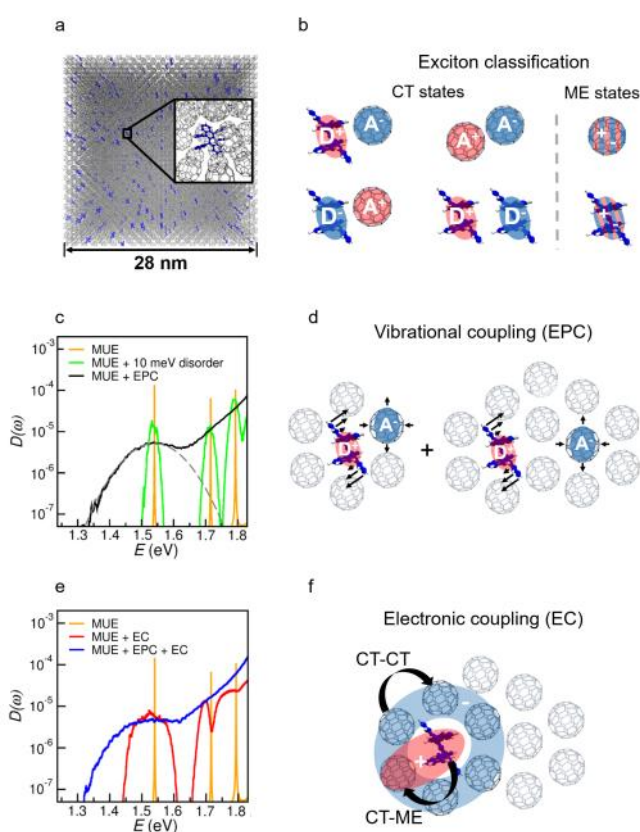


Figure 2: a) Example for a supercell used in numerical calculations. (b, d, f) Classification of excitonic states. (c, e) Low-energy absorption spectra depending on different levels of microscopic interactions considered by our model calculations. Figure adapted from Ref. [1].

require more CPU-time and memory and will produce more output data, which makes SuperMUC-NG an indispensable tool for future calculations and further progress in our research and the understanding of organic solar cells.

The only limitations we observed in the last phase were the queueing times on SuperMUC Phase 2. The only obstacles we noticed were due to the transition from SuperMUC Phase 2 to SuperMUC-NG and due to the long shutdown period during spring 2020 where many European supercomputing centers were affected and needed to be taken offline.

In future projects we plan to extend and refine our numerical calculations to a larger class of problems (exciton diffusion, charge separation, etc.) that result in valuable scientific contributions.

References and Links

- [1] Panhans, M., Hutsch, S., Benduhn, J., Schellhammer, K. S., Nikolis, V. C., Vangerven, T., Vandewal, K., Ortmann, F., *Nat Commun* 11, 1488 (2020). <https://cfaed.tu-dresden.de/ortmann-group-news>
- [2] <https://phys.org/news/2020-03-molecular-vibrations-maximum-photovoltage-solar.html>
- [3] Fan, Z., García, J. H., Cummings, A. W., Barrios-Vargas, J. E., Panhans, M., Harju, A., Ortmann, F., Roche, S., *Phys. Rep.* (2021). <https://doi.org/10.1016/j.physrep.2020.12.001>
- [4] Ortmann, F., Cresti, A., Montambaux, G., Roche, S., *EPL* 94, 47006 (2011).

Interface Phenomena in Chemical and Mechanical Engineering

2

RESEARCH INSTITUTION

¹Laboratory of Engineering Thermodynamics, TU Kaiserslautern

²Scientific Computing in Computer Science, TU München

³Chair of Thermodynamics and Thermal Separation Processes, TU Berlin

PRINCIPAL INVESTIGATORS

Maximilian Kohns¹, Kai Langenbach¹, Hans Hasse¹, Steffen Seckler², Hans-Joachim Bungartz², Matthias Heinen³, Jadran Vrabec³

RESEARCHERS

Michaela Heier¹, Aditya Kulkarni¹, Dominik Schäfer¹, Simon Stephan¹, Gabriela Guevara-Carrion³, Rene Chatwell³, Robin Fingerhut³

PROJECT PARTNERS

—

SuperMUC Project ID: pn56mo

Introduction

Interfaces are prevalent in many fields of science and engineering. In many applications, the properties of the interface between different phases – gaseous, liquid, or solid – govern the behavior of a system. Examples from engineering include machining processes (solid-solid and solid-liquid interfaces) and thermal separation processes (e.g. vapor-liquid and liquid-liquid interfaces).

However, the actual interface is only a few nanometers in width, so that most interface phenomena cannot be studied experimentally. This length scale is readily accessible with molecular dynamics (MD) simulations. For carrying out MD simulations, a model of the intermolecular interactions has to be specified. Then, a suitable starting configuration of these molecules that resembles the problem to be studied is provided, and Newton's equation of motion are solved numerically to follow the temporal evolution of the system.




Figure 1: Snapshot of a MD simulation for studying transport through vapor-liquid interfaces. Blue particles represent the heavy-boiling component. Green particles represent the light-boiling component, which move through the interface. For clarity of details, only one side of the simulation box is shown; there is another gas phase to the right of the liquid film.

Results and Methods

To study interface phenomena, we carry out massively-parallel MD simulations on SuperMUC-NG with two simulation programs: *ls1 mardyn* [1,2], which is developed jointly by the project partners, and the open source code LAMMPS [3]. Two examples are presented in the following. Both employ the Lennard-Jones potential, a simple yet physically reasonable model for intermolecular interactions.

Our research within the ERC Advanced Grant project 'ENRICO' focuses on transport through vapor-liquid interfaces. It is known that under certain circumstances, the particles of the light-boiling component (i.e. those prevalent in the gas phase) can have a favor for being directly at the interface, leading to an enrichment of that component at the interface. In terms of macroscopic theories of mass transfer, e.g. the Fickian laws, the enrichment should impose a hindrance on the flux of the light-boiling component through the interface. However, many chemical engineering separation processes, such as absorption, rely on that flux.

We carry out MD simulations in a dedicated simulation scenario to study the influence of the enrichment on mass transfer. An exemplary snapshot of that scenario is shown in Figure 1.

In these simulations, the vapor-liquid interface of a liquid film with a vapor phase is simulated. On one side of the simulation box, particles of the light-boiling component are continuously inserted into the gas phase, while particles of the light-boiling component are continuously removed from the gas phase on the other side (not shown in Figure 1). This particle insertion and removal forces a flux of the light-boiling component through the liquid film. Comparing mixtures with and without an enrichment of the light-boiling component at the interface, this allows for investigating whether the enrichment is a hindrance for




Figure 2: Snapshot of a MD simulation for studying nanoscale machining processes. A rigid indenter (grey ball) impacts a solid surface (blue), scratches over the surface and a chip forms (red particles). Particles of the lubricant are not shown in this snapshot for clarity.

transport. Our first simulation results show that there is indeed a link between the enrichment at the interface and a reduced transport. However, this needs further investigation. A first publication on this topic is currently in preparation.

Simulations such as the one shown in Figure 1 are typically carried out with approximately 20,000 to 100,000 particles. Depending on the studied conditions, they require about four to five days on SuperMUC-NG when using about 128 cores.

As part of the International Research Training Group (IRTG) 2057, we study the properties of interfaces in machining processes. A snapshot of a typical simulation scenario is shown in Figure 2.

In the MD simulations for studying nanoscale machining processes, a rigid indenter in the form of a ball is first put into contact with a solid surface by applying a force perpendicular to the solid surface. The indenter then scratches over the surface by applying an in-plane force. During the scratching process, a chip builds up from the material of the solid surface. During the process, the 3D profiles of important properties such as temperature, pressure and flow field are monitored.

The simulations show that between the contacting surfaces, there is only a small number of lubricant molecules. This means that although a lubricant is used, the direct contact between the two solids is mostly dry. Nevertheless, the lubricant has an important effect on the process in that it dissipates heat from the rigid bodies. A more detailed discussion of the results can be found in a recent publication [4].

Simulations such as the one shown in Figure 2 are typically carried out with approximately four to six million particles. Depending on the studied conditions, they require about two days on SuperMUC-NG when using about 256 cores.

Ongoing Research / Outlook

Our large scale MD simulations rely heavily on the availability of HPC resources such as those provided by SuperMUC-NG. Using the program *ls1 mardyn* on SuperMUC, the authors also hold the current world record for the largest MD simulation ever carried out: 20 trillion particles [5].

For future work, we intend to continue our cooperation in developing *ls1 mardyn*. On the application side, we aim at carrying out more systematic studies of the transport through liquid-vapor interfaces. Regarding the nanoscale scratching process, we want to study systems with molecular models for real fluids and solids, most notably for the lubricants. There, we expect to see an influence e.g. of the length of the lubricant molecules on the outcome of the process.

References and Links

- [1] C. Niethammer, S. Becker, M. Bernreuther, M. Buchholz, W. Eckhardt, A. Heinecke, S. Werth, H. J. Bungartz, C. W. Glass, H. Hasse, J. Vrabec, M. Horsch, *J. Chem. Theory. Comput.* 10 (2014) 4455-4464.
- [2] <https://www.ls1-mardyn.de/home.html>
- [3] <https://lammps.sandia.gov/>
- [4] S. Stephan, M. Dyga, H. Urbassek, H. Hasse, *Langmuir* 35 (2019) 16948-16960.
- [5] N. Tchipev, S. Steffen, M. Heinen, J. Vrabec, F. Gratl, M. T. Horsch, M. Bernreuther, C. W. Glass, C. Niethammer, N. Hammer, B. Krischok, M. Resch, D. Kranzlmüller, H. Hasse, H. J. Bungartz, P. Neumann, *Int. J. High Perform. Comput. Appl.* 33 (2019) 838-854.

Towards a predictive theory for unconventional superconductivity

2

RESEARCH INSTITUTION
King's College London, United Kingdom

PRINCIPAL INVESTIGATOR
Mark van Schilfgaarde

RESEARCHERS
Swagata Acharya, Dimitar Pashov

PROJECT PARTNERS
—

SuperMUC Project ID: pn56ya (PRACE project)

Introduction

The combination of QSGW, DMFT, and BSE for susceptibilities gives us a high fidelity general-purpose *ab initio* theory of many-body phenomena. It also contributed to the enrichment of our open source suite for electronic structure calculations (<https://questaal.gitlab.io/>) and the method paper was published in the special issue of CPC; SI: celebration of 50 years of CPC [1]. The article was also published as one of the two psi-k highlight articles for the year 2020 (<https://psi-k.net/highlights/>). In terms of its physical applications, the development resulted in some astounding insights into the interplay between charge, spin and orbital degrees of freedom in strongly correlated systems and unconventional superconductors like ruthenates, cuprates, iron based superconductors. The results were published in leading journals in physics [2,3,4]. In summary, the super-computing facility provided by PRACE, made possible the large scale application of this ability that sets the stage for a high-fidelity 'predictive *ab-initio* approach' to unconventional superconductivity, which is otherwise crowded only by the model Hamiltonian community. In that sense, the PRACE facility worked for us at the right moment when this ability is about to bring in a paradigm shift in the manner we address superconductivity in real materials.

Results and Methods

Ruthenates and disentangling degrees of freedom

One of the fundamental challenges in Sr_2RuO_4 (SRO) was to understand the nature of pairing and also to explain why T_c increases with tensile strain up to a

critical strain. *a) Spin and charge susceptibilities.* We computed *ab initio* spin and charge susceptibilities (χ^m and χ^d) in the unstrained case. χ^m has been measured, which gave us an experiment to benchmark against. We find excellent agreement with experiment (strong peaks at incommensurate q vectors ($\pm 0.3, \pm 0.3, 0$), maximum intensity at around 10 meV as observed, and weaker fluctuations extending to $q = 0$ —needed for triplet superconductivity). *b) Nature of superconducting pairing.* In a first, within our fully *ab-initio* framework we show that there are three triplet and three singlet instabilities in SRO. The analysis shows that both spin and charge contribute to the instabilities. In the charge susceptibility. We find that there are strong fluctuations close to $q = 0$, and spin and charge together can promote a triplet pairing state. However, as mentioned above, there are multiple instabilities in SRO both of singlet and triplet kind. We identify their nature, degeneracies and possible candidates that can explain the NMR observations in SRO. One primary conclusion is superconductivity can not be purely driven by spin triplet fluctuations. Most importantly we showed that the uniaxial strain that brings in more spin fluctuations to the incommensurate vector promotes a rapid enhancement in the singlet pairing instability (Figure 1). This paper is published in Communications Physics, Nature [4]. Following this, we are invited to write a review article on SRO superconductivity for the special issue of Applied-Physics Journal (Modeling transition-metal compounds: emerging developments and applications).

Iron based superconductors: LaFe_2As_2 and CaFe_2As_2 collapsed and uncollapsed phases

Recent observations of selective emergence (suppression) of superconductivity in the uncollapsed tetragonal (collapsed tetragonal) phase of LaFe_2As_2 (LFA) has rekindled the question of controlling band structure and correlations and thereby to control super-




Figure 1: Superconducting pairing functions and eigenvalues: The superconducting pairing gap symmetries for strain=0 are shown in the conventional basal plane; eigenfunctions corresponding to first three eigenvalues in singlet (s) symmetries are shown in panels (a)-(c) and triplets (t) are in the shown in panels (d)-(f). Panel (g) shows evolution of triplet and singlet eigenvalues (at 380 K) under strain. Under strain singlet eigenvalues increase and surpass the triplet eigenvalues.




Figure 2: The energy and momentum resolved spin susceptibility $ImX(q,\omega)$ (in the top panel from left to right) shown for the CT and UT phases respectively. The q -path $(H,K,L=0)$ is chosen along $(0,0)-(1/2,0)-(1/2,1/2)-(0,0)$ in the Brillouin zone corresponding to the two-Fe atom unit cell. The intensity in the CT phase is artificially multiplied by five to bring excitations for CT and UT phases to the same scale.




Figure 3: The superconducting instability is absent in the CT-LFA phase. The superconducting instability corresponding to the leading λ_1 and lagging λ_2 eigenvalues of the solutions to the Eliashberg equations, $\Delta(q,\omega=0)$ are shown for the UT-LFA phase. The evolution of the leading eigenvalue as a function of temperature is shown for CT-CFA, UT-CFA and UT-LFA in the bottom panel. In inset we zoom into the low temperature part of the curves to show the estimated T_c 's.

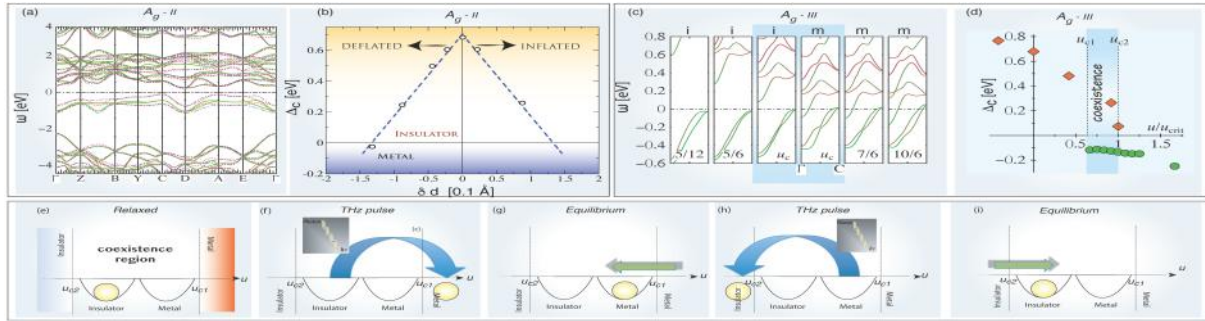


Figure 4: a) QSGW energy band calculations for the displaced structure along the Ag-II (5.47 THz) phonon mode, with both expansions and contractions of the V-V bond displacement $\sim 0.022 \text{ \AA}$ (red and blue) around the M1 phase (green). b) The charge gap shrinks for either sign of displacements and eventually closes for changes in the V-V distance as small as displacements $\sim 0.15 \text{ \AA}$. c) QSGW energy band structure for the Ag-III (e.g., 6.29 THz) phonon mode for displacements above and below the critical displacement u_c . At u_c , QSGW predicts the coexistence of two converged solutions: a metallic solution labelled m and an insulating solution labelled i. d) Charge gap as a function of displacement u : we obtain a hysteresis and purely electronic IMT at u_{c2} (diamonds) and respectively MIT at u_{c1} (circles). Bi-stability is obtained for $u_{c1} < u < u_{c2}$; if we surmise that the nuclear positions of the system can be pinned along the Ag-III mode in the coexistence region, so that the material does not relax to its pristine form, (e) bi-stability would be achieved, and (f) an applied 6.29 THz pulse would drive the system to a metal. After the system relaxed back to the middle region, (g) it would stay in the metallic solution, until (h) another pulse hits the system, after which (i) the system would relax to the insulating solution.

conducting T_c . Materials with incoherent spin fluctuations, driven away from magnetic ordering, can provide significant glue for Cooper pair formation. Controlling incoherence in strongly correlated electronic systems seem to be a future direction in engineering T_c in unconventional superconductors primarily mediated through spin fluctuations. We show that the proximity of Fe-dxy state to the Fermi energy acts as the primary source of incoherence in LaFe_2As_2 . In the collapsed tetragonal (CT) structure, the dxy state is driven far below the Fermi energy, leading to suppression of single and two-particle low-energy scattering. However, in the uncollapsed phase (UT) the dxy state is at the Fermi energy and is most strongly correlated, leading to enhanced low energy scattering in both single and two-particle channels. The emergent incoherent spin fluctuations (Figure 2) provide the desired glue for Cooper pair formation. We perform a similar analysis in CT and UT phases of CaFe_2As_2 (CFA), and show had the CaFe_2As_2 not gone through the first order structural (and therefore magnetic) transition (Figure 3), it would have the highest T_c among these four candidates. This paper is now accepted in Physical Review Letters [3].

Enigmatic metal-insulator transition in VO_2

We used our ability in conjunctions with ultra-fast spectroscopic experiments to identify the enigmatic

nature of metal-insulator transition in VO_2 and discover a bistable regime (Figure 4) that can be used for switching [4].

Ongoing Research / Outlook

In continuation to this project we got another PRACE project funded, ra5208. We are using that ability now to explore further the superconductivity in the bulk and layered variants of FeSe. In particular, we are trying to explore the the order-of-magnitude enhancement in T_c in FeSe/SrTiO₃ interface in comparison to the bulk. We are also exploring the nature of the nematicity and its role in superconductivity in FeSe. Most importantly, we are exploring if there is anything universal to the nature of the glue in iron based superconductors and how that possible differs from that of cuprates and ruthenates.

References and Links

- [1] Pashov, Dimitar, et al., Comput Phys Commun 249, 107065 (2020).
- [2] Acharya, Swagata, et al., Commun. Phys. Vol 2, 163 (2019).
- [3] Acharya, Swagata, et al., Phys. Rev. Lett. 124, 237001 (2020).
- [4] Weber, Cédric, et al., Phys. Rev. Research 2, 023076 (2020).

Phase-field simulations to study eutectic colonies

RESEARCH INSTITUTION

¹Karlsruhe Institute of Technology, IAM-CMS, Karlsruhe

PRINCIPAL INVESTIGATOR

Britta Nestler^{1,2}

RESEARCHERS

Michael Kellner¹, Marco Seiz¹

PROJECT PARTNERS

²Hochschule Karlsruhe - Technik und Wirtschaft, IDM, Karlsruhe

SuperMUC Project ID: pn69ne

Introduction

During the directional solidification of eutectic multi-component alloys the formation of eutectic colonies can be observed. The formation of these structures is driven by microscopic instabilities in a macroscopic planar solidification front, due to impurities of the components that diffuse from the solidifying phases into the melt. Within the project pn69ne [1], the formation of eutectic colonies in the high temperature material *NiAl-34Cr* is studied within phase-field simulations. For a simultaneous resolution of the fibrous structure of the eutectic *NiAl-34Cr* and the multiple times larger structures of the eutectic colonies in one simulation, large solidification fronts with a high resolution are needed. Figure 1 shows exemplarily the simulated solidification front of such a microstructure.

The focus of the conducted simulations is to investigate the interactions and growth competition between multiple colonies in two and three-dimensional large scale domains. Furthermore the transition from parallel rod growth to eutectic colonies and the conditions required for the growth of these structures are studied. Parts of the investigated results are published in Ref. [2].

Results and Methods

For the investigation of the two-phase eutectic reaction $\text{Liquid} \leftrightarrow \text{B}_{2\text{NiAl}} + \text{A}_{2\text{Cr}}$ in the ternary system *Al-Ni-Cr*, a thermodynamically consistent phase-field model based on a grand potential functional is used. The partial differential equations for the evolution of the phase-field order parameters and the chemical potentials are spatially discretized with finite differences and the temporal evolution is calculated by a forward Euler scheme. These numerical solution schemes are implemented into the massively parallel multiphysics framework Pace3D [3].

The driving force for the phase transition is described by the difference of the grand potentials. The grand potentials are derived from the Gibbs energy formulations of the thermodynamic Calphad database from Peng et al. [4].

For the evolution of new rods, which is required to ensure a stable growth of curved solidification fronts with multiple phases in 2D phase-field simulations, a concentration-driven nucleation mechanism has been added into the phase-field evolution equation of the order parameters in [2]. With each nucleation iteration, multiple nuclei are randomly set into the solid-liquid interface. This nucleus setting mechanism can be adjusted by the amplitude A , and by a concentration-driven criterion c_{dev} .

The chosen material and numerical parameters as well as the approximated thermodynamic energies for the simulation of eutectic colonies in *NiAl-34Cr* are collected in the appendix of Ref. [2]. The simulations are performed in a domain of $1,500 \times 10,000$ voxel cells, with an initial setting of random seeds, using a Voronoi tessellation. The growth direction and the velocity are controlled by an analytic temperature profile. To systematically investigate the colony formation for comparable structures, 2D large-scale simulations with different temperature gradients between 0 and 400 K/mm are performed.

Each simulation is performed with 1,200 cores for 72 hours to generate 20 million time steps. The results for each calculated field is stored in a separate file, requiring about 80 GB for one simulation.

As an example, the simulation with an applied temperature gradient of 20 K/mm is discussed in more detail. The fully evolved microstructure after 20 million time steps of this simulation is shown in Figure 2, surrounded by enlargements of selected microstructure features. In Figure 2(a), the formation of a lamellar



Figure 1: Exemplary solidification front of a simulated eutectic colony structure of NiAl-34Cr.




Figure 2: Resulting microstructure of a 2D large-scale simulation of the eutectic colony formation in NiAl-34Cr with an applied temperature gradient of 20 K/mm [2].

growth from the initially set random seeds is shown. During this formation several microstructure adjustment events such as merging, splitting, overgrowth and nucleation can be observed. Similar to these adjustments, microstructure events can also be found during the evolution of eutectic colonies. In Figure 2(b), the overgrowth of a colonies is depicted. In contrast to the events of rods, an overgrowth of a colony in 2D is not combined with a merging event. The overgrowing colonies are just forming a new contact area, instead of combining to one colony. The formation of a new colony by division of one large colony into two smaller colonies is indicated in Figure 2(c) and the evolved solidification front with multiple colonies is highlighted in Figure 2(d).

In Figure 3, the solidification fronts from the simulations with different applied temperature gradients are displayed. For all simulations, the solidification velocities converge to a similar value, as depicted in the velocity plot in the lower part of the figure. Multiple well-pronounced colonies with convex solidification fronts evolve for all simulations in this study.

In general, the rods in the middle of a colony mainly show straight growth. In the vicinity of the contact areas between two colonies, the rod growth is no longer directed in the direction of the applied temperature gradient, which results in a curved growth with repeated splitting and nucleation events. As expected, the height difference at the solidification front decreases with an increasing temperature gradient. However, independent from the temperature gradient, all

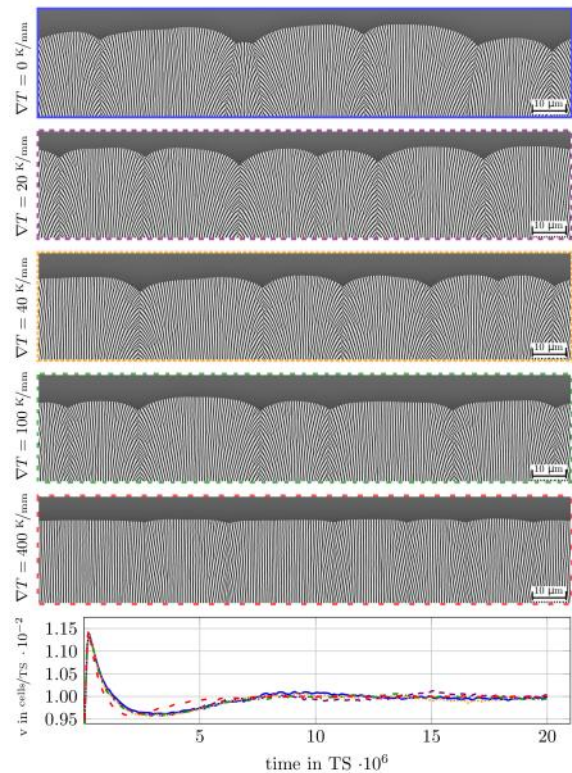


Figure 3: Comparison of solidification fronts for the simulations with different temperature gradients after 20 million time steps. The corresponding velocity profiles over the simulation time are plotted in the diagram [2].

solidification fronts consist of six colonies after simulating 20 million time steps. For the temperature gradient of zero, it can be assumed that the smallest colony is overgrown by further solidification. Quantitative analyzes of the visually described trends for the colony widths and heights as well as for the rod spacing are also given in Ref. [2].

Ongoing Research / Outlook

Due to the provided computational resources, the growth of multiple eutectic colonies and their interactions could be simulated in realistic scales. Besides the presented studies of the temperature gradient dependencies, the influence of the growth velocity on the colony formation has been investigated [5]. However, the influence of other conditions such as surface energies or melt compositions has not been fully investigated. Furthermore, the simulation of eutectic colonies in three-dimensional domains is still challenging regarding computing times. Hence, during the next period of the project pn69ne [1] the influence of different melt compositions on the eutectic colony formation in a further material system AlCu-5Ag is studied in 2D and 3D.

References and Links

- [1] <https://www.lrz.de/projekte/hlrb-projects/000000000F43B00.html>.
- [2] M. Kellner, J. Hötzer, E. Schoof, and B. Nestler, *Acta Mater.* 182, 267–277, 2020.
- [3] J. Hötzer, A. Reiter, H. Hierl, P. Steinmetz, M. Selzer, and B. Nestler, *J COMPUT SCI-NETH*, 26, 1–12, 2018.
- [4] J. Peng, P. Franke, and H. J. Seifert, *J PHASE EQUILIB DIFF*, 37, no. 5, 592–600, 2016.
- [5] M. Kellner, 2020, PhD-Thesis, KITOpen, doi: 10.5445/IR/1000119330.

All-electron DFT Simulations of particle-like

magnetic objects

RESEARCH INSTITUTION
Forschungszentrum Jülich

PRINCIPAL INVESTIGATOR
Stefan Blügel

RESEARCHERS
Daniel Wortmann, Uliana Alekseeva

PROJECT PARTNERS
—

SuperMUC Project ID: pn72qa (Gauss Large Scale project)

Introduction

Complex magnetic textures and localized particle-like structures on the nanometer scale such as chiral magnetic skyrmions (Fig. 1) with non-trivial topological properties are nowadays the most studied objects in the field of nanomagnetism. They offer the promise of new data storage and data processing technologies ranging from racetrack memories to memristive switches for neuromorphic computing.

The theoretical and computational description of nanosized magnetic objects usually follows a typical multiscale approach. Ab-initio simulations mostly based on Density Functional Theory (DFT) are utilized to determine fundamental materials properties. Based on these properties simpler purely-magnetic models like the extended Heisenberg model are constructed and spin-dynamics simulations are employed to describe the final complex magnetic structures. While this approach has demonstrated its power and efficiency, it cannot describe fundamental changes in the electronic structure induced by the complex alignment of magnetic moments. In this project we extend the realm of DFT calculations for magnetic systems to significantly larger setups to tackle these challenges from the basic description provided by our ab-initio code FLEUR [1]. Since 2015 FLEUR is one of the flagship codes of the MaX-European Center of Excellence in HPC [2]

Results and Methods

Our projects aims at providing insight into the electronic properties of large and complex non-collinear magnetic structures by applying our state-of-the-art all-electron full-potential linearized augmented plane wave (FLAPW) DFT code FLEUR [3]. While the FLAPW method is a numerical complex and computationally heavy approach it is also considered to provide the "gold standard" and reference results in DFT. Its accuracy and reliability proven in numerous applications is particularly crucial in the field of magnetism as the fundamental energy scales on which magnetic phenomena occur is usually very small. While DFT simulations of large setups containing many atoms are already possible for quite some time, the simulations we pursue extend the limits of so called all-electron DFT applied to large magnetic configurations.

A DFT simulation is an iterative process with 50-120 cycles to achieve a self-consistent solution of the fundamental quantum mechanical equations. A single of this self-consistency iteration contains many different steps like the determination of the potential, the generation of a new density or the density mixing to accelerate the self-consistency process. From the computational point of view, the most relevant task in each iteration consist of the setup of matrices for the generalized eigenvalue problem and its diagonaliza-




Figure 1: Skyrmion tubes in MnGe super cell 8x8: two Bloch skyrmions (in the left panel and in the right one) and a Néel type (in the middle panel). The direction of the magnetisation is indicated by the arrows located at each Mn site. The color alteration illustrates the change of the direction of the magnetic moments along the radii of the skyrmions. The energy of these magnetic configurations is decreasing from the left to the right (0 meV, -200 meV, -700 meV).

Structure	# atoms	Matrix size	Core-hours
Skymion	512	78k x 78k	2,048
Globule	1,024	150k x 150k	5,120

Table 1: Properties of our two magnetic setups, skymion tubes (1st row) and a globule (2nd row): number of atoms (2nd column), size of dense hermitian matrices in the generalized eigenvalue problems (3rd column) and number of core-hours needed for one self-consistence iteration (4th column).

tion. The hybrid MPI/OpenMP parallelization of the FLEUR code allows to utilize machines like SuperMUC-NG efficiently. The code has interfaces to optimized libraries (ELPA, ScaLAPACK, LAPACK, BLAS) to perform linear algebra operations as efficiently as possible.

To get insight into details of the electronic structure which are not in reach of conventional multiscale approaches, we performed simulations with two magnetic setups which are larger than any production-level FLAPW calculations reported so far: i) several skymion tubes in the MnGe 8x8 super cell (Fig. 1) and ii) a globule with two Bloch points in the MnGe 4x4x8 super cell (Fig. 2).

So far our simulations were still restrained due to some scaling limitation that might be expected if a complex code is pushed to simulations of this scale for the first time. Most of these bottlenecks have now been removed so that the simulations gained momentum. Currently, our simulations for the setups containing the Bloch points run on 256 nodes with 48 cores for about 25 minutes per iteration. Hence, a single self-consistency cycle will consume about 5000 core-hours and a full self-consistency with about 100 iterations about half a million of core-hours. As we have to compare the different magnetic states for a full analysis of the system and to gain inside knowledge of the interplay between the magnetic and electronic structure of these configurations we need to obtain self-consistency for several setups. These calculation involve the diagonalization of several complex matrices of size 150k x 150k per iteration. The corresponding data for the second system reported here are given in Table 1. Typical for a DFT calculation, our IO requirements are relatively modest and only few files are generated.

Ongoing Research / Outlook

Our simulations on skymion tubes already confirmed the fundamental predictions obtained by simple multiscale models and hence we feel confident that the basic assumptions, the setup chosen and its numerical description accurately captures the relevant physical properties. One should note here that the energy differences obtained are in the range of few hundred meV despite the significant size of the system. This again stresses the relevance of our approach. At the same time, these results are still pretty preliminary as we could not perform all simulations and all needed analysis planned due to unforeseen code problems. This also includes further simulations on thin-film setups in which a special mode of the code is employed. As a consequence, we only used about




Figure 2: A globule in MnGe super cell 4x4x8. This configuration has two Bloch points at the edges of the globule.

30% (<10Mio core hours) of our assigned computational resources.

As these problems are understood and mostly solved by now, our future research will include further simulations of similar setups, the analysis of the results, the extension of the investigation to the electronic structure of the Bloch points and the thin-film setups. As our results are still preliminary in this respect, they have not been published so far.

References and Links

- [1] <https://www.flapw.de>
- [2] "MaX-Materials Design at the Exascale, a European centre of excellence", <http://www.max-centre.eu>
- [3] U. Alekseeva, G. Michalick, D. Wortmann and S. Blügel, "Hybrid Parallelization and Performance Optimization of the FLEUR Code: New Possibilities for All-Electron Density Functional Theory", in: Springer International Publishing AG, part of Springer Nature 2018 M. Aldinucci et al. (Eds.): Euro-Par 2018, LNCS 11014, pp. 735–748, 2018.

Pressure response of bulk liquid water

as “seen” by THz spectroscopy

RESEARCH INSTITUTION

Lehrstuhl für Theoretische Chemie, Ruhr-Universität Bochum

PRINCIPAL INVESTIGATOR

Dominik Marx

RESEARCHERS

Saskia Körning

PROJECT PARTNERS

—

SuperMUC Project ID: pr23va

Introduction

The H-bond network of liquid water is characterized by a well-known tetrahedral structure offering rather large open spaces. These voids are filled to some extent with other water molecules at ambient conditions, while fully preserving the usual tetrahedral network. In contrast, at high hydrostatic pressure (HHP) conditions, where the macroscopic density of the liquid increases and the water molecules are actively pressed into the voids, they are regularly occupied by water molecules. Therefore, studying liquid water at HHP conditions is another possibility to understand its anomalies at ambient conditions. Notwithstanding this importance, the effects of pressure perturbations of liquid water up to its stability limit of about 10 kbar at room temperature are not well understood compared to temperature effects at low pressures. During the last decade, experimental and computational far-infrared (far-IR) or THz spectroscopy of aqueous solutions has been demonstrated to provide detailed insights into the changes of the H-bond network due to perturbations with respect to neat liquid water. Since THz spectroscopy directly probes the H-bond network peak close to 180 cm^{-1} , the method is expected to have huge potential to unveil the intermolecular dynamics of water and aqueous solutions not only close to ambient conditions but also at extreme thermodynamic conditions such as in the multi-kbar regime. In this project, we [1] applied instantaneous normal mode (INM) analysis to theoretically dissect the THz response of neat liquid water from 1 bar to 10 kbar at 300 K down to the level of individual environment-dependent H-bond contributions based on extensive ab initio molecular dynamics (AIMD) simulations. Although the linearization of atomic motions is certainly a drastic approximation, INM offers important advantages at the level of decomposing the total response into local contributions that depend on the individual H-bonding patterns. First, each INM can rigorously be decomposed into single-molecule hindered translations and librations as well as purely intramolecular vibrations. Second, the total dipole moment fluctuations can be decomposed within the INM approach into the effects due to the underlying atomic/molecular motions as well as the corres-

ponding dynamically induced dipole moments at the level of individual configurations and thus H-bonding patterns. These two features of the INM approach fruitfully provide novel insights into the THz spectrum of liquid water not only at ambient conditions but also deep in the kbar regime. The results of this research are published in Ref. [2,4].

Results and Methods

The INM analysis is computationally demanding. Besides the usual AIMD simulation, which is used to conduct the simulation itself, the INM analysis requires $3N$ additional electronic structure calculations, where N is the number of atoms in the system, per selected time step. Combining the INM analysis with the so-called cross correlation analysis (CCA) yields a detailed and complete dissection of the THz spectrum at ambient and high-pressure conditions as shown in Fig. 1. Here we obtain the self contribution of each individual water molecule as well as four important cross couplings between a reference molecule and (i) a H-bonded water, (ii) an interstitial water, (iii) a second-shell water, and (iv) a water molecule beyond the second solvation shell. On top of this decomposition, the INM analysis allows to dissect the spectrum further into translational-translational self contributions, $I_{t-t}(\omega)$, and translational-rotational cross couplings, $I_{t-r}(\omega)$, as shown in Fig. 1 left and right, respectively. At ambient conditions, the $t-t$ correlations $I_{t-t}(\omega)$ depicted in left panel are found to be dominated by the molecular self-correlations and the H-bonding pair correlations, whereas almost only the H-bonding pairs contribute to the $t-r$ coupling $I_{t-r}(\omega)$ in right panel. This situation does not qualitatively change upon compression to 10 kbar. Importantly, the interstitial water molecules almost contribute nothing to the THz response. This finding indicates that those structural changes upon compressing liquid water into the kilobar regime which are due to generating interstitial molecules in tetrahedral voids (while fully preserving the tetrahedral H-bond network topology) cannot be related to the observed changes of the THz lineshape.

Given the overwhelming importance of H-bonded water pairs to the THz spectrum, we further decom-




Figure 1: Partial THz spectra of bulk liquid water resulting from (left) translation–translation self-correlations $I_{t,t}(\omega)$ and (right) translation-rotation cross-correlations $I_{t,r}(\omega)$. The solid/dashed lines with filled/open symbols correspond to 1 bar/10 kbar. Red circles, blue upward triangles, green downward triangles, yellow squares, and black diamonds represent the contributions due to molecular self-correlations (self) as well as the molecular cross-correlations between H-bonded water pairs (HB), between second–neighbor pairs (2nd), between interstitial pairs (IP), and rest pairs (rest). Reproduced from Ref. [2] with permission of AIP publishing.




Figure 2: Partial THz spectra of bulk liquid water resulting from translation-translation self-correlations (left) 1 bar and (right) 10 kbar in terms of six H-bond donor pair classes. Red circles, blue triangles, and green squares represent single H-bond donor (1-d), double H-bond donor (2-d), and triple H-bond donor (3-d) pairs, respectively, where solid lines with filled symbols and dashed lines with open symbols distinguish between such short (s) and long (l) H-bond lengths, respectively. All spectra are normalized per molecule. Reproduced from Ref. [2] with permission of AIP publishing.

pose the spectrum at the molecular level by distinguishing between H-bond donor and acceptor molecules as shown in Fig. 2. This decomposition leads to the important finding that the H-bond network peak around 180 cm^{-1} cannot be obtained without considering the dynamically induced dipole moments from H-bond donor molecules (whereas acceptors contribute most of their THz intensity at much lower frequencies), which clearly highlights the outstanding importance of these H-bond donor molecules with respect to the reference molecule to understand the THz lineshape of bulk liquid water as we elaborately report in Ref. 2.

Ongoing Research / Outlook

Within this project, we have also disclosed the distinct features of hydrophilic versus hydrophobic solvation as seen by THz spectroscopy of aqueous Trimethylamine N-oxide (TMAO) solutions at ambient conditions [3]. TMAO is a well-known co-solvent used for long in biophysical chemistry which acts as an osmolyte by stabilizing the native folded state of proteins. Elevated concentrations of TMAO in deep-sea fish even counteract the perturbing effects of high pressures and actively stabilize proteins at these high

pressure conditions. Besides, there is also significant fundamental interest in TMAO due to its peculiar H-bonding properties because it has both, strongly hydrophilic and extended hydrophobic groups at its opposite ends.

Currently, we are conducting AIMD simulations of aqueous TMAO solutions at HHP conditions. We aim to decompose the THz spectra again to obtain individual contributions at the molecular level. Ultimately, we want to investigate how the solvation properties of TMAO change at HHP conditions compared to ambient conditions to understand how TMAO as a cosolvent helps to stabilize proteins at these high pressures, where the proteins would normally denature if TMAO was absent.

References and Links

- [1] <https://www.theochem.ruhr-uni-bochum.de/>
- [2] S. Imoto and D. Marx, *J. Chem. Phys.* 150, 084502 (2019).
- [3] S. Imoto, H. Forbert, and D. Marx, *Phys. Chem. Chem. Phys.* 20, 6146 (2018).
- [4] H. Vondracek, S. Imoto, L. Knake, G. Schwaab, D. Marx, and M. Havenith, *J. Phys. Chem. B* 123, 7748 (2019).
- [5] C. Hölzl, P. Kibies, S. Imoto, J. Noetzel, M. Knierbein, P. Salmen, M. Paulus, J. Nase, C. Held, G. Sadowski, D. Marx, S. M. Kast, and D. Horinek, *Biophys. Chem.* 254, 106260 (2019).

Topology, Entanglement and Critical Phenomena

in Correlated Quantum Matter

RESEARCH INSTITUTION

¹Institut für Theoretische Physik und Astrophysik, Universität Würzburg

PRINCIPAL INVESTIGATOR

Fakher Assaad¹

RESEARCHERS

Y. Liu², Z. Wang¹, T. Sato¹, B. Danu¹, M. Raczkowski¹, J. Schwab¹, F. Goth¹, A. Götz¹, F. Parisen Toldin¹, D. Hou², J. Stafusa E. Portela¹

PROJECT PARTNERS

²Beijing Normal University, China

SuperMUC Project ID: pr53ju (Gauss Large Scale project)

Introduction

The aim of our research is to carry out numerical simulations of a collection of interacting electrons or spins so as to understand experiments or to unravel the richness of collective and emergent phenomena in the solid state. Modelling experiments by overemphasizing the dominant phenomena is important for interpretation of the experimental data. Many body systems host a plethora of emergent correlated and topological phenomena, the richness of which is fascinating. Ultimately, the exotic phenomena we investigate, out of pure intellectual curiosity, in model systems complying with the basic laws of quantum mechanics, may very well be the key to new technologies beyond the silicon age. Here we will concentrate on two very different applications. The first one is a modelling of the spin system Botallackite-Cu₂(OH)₃Br [1] and the second a study of a designer model showing fascinating quantum critical phenomena [2]. The tool we use is the auxiliary field quantum Monte Carlo algorithm. Over the last years we have spent time and effort at developing a flexible open source code [3] that allows us to implement various models at a minimal programming cost. This program package is constantly being developed and we are presently preparing a new release.

Results and Methods

A) Coexistence and Interaction of Spinons and Magnons in an Antiferromagnet with Alternating Antiferromagnetic and Ferromagnetic Quantum Spin Chains [1].

Botallackite-Cu₂(OH)₃Br is a two-dimensional magnetic material that can be modeled by weakly coupled alternating spin 1/2 ferromagnetic and anti-ferromagnetic chains. The ferromagnetic chain shows well defined magnon (spin 1) excitations with quadratic dispersion at vanishing wave vector. This dispersion relation can be detected by neutron scattering experiments. On the other hand the spin 1/2 chains exhibit what is called fractionalization: a magnon decomposes into two novel entities, spinons, both carrying spin 1/2. Neutron scattering experiments can pick up the distinct signatures of this fractionalization, namely the two-spinon continuum. The interesting aspect of this material is that it contains both types of excitations and the tantalizing question is to understand how they interact. Remarkably Fig. 1 shows that we are in a position to compare unbiased numerical simulations to neutron scattering experiments. Numerical simulations and experiments show that upon coupling the chains, the spinons bind: a gap opens at the antiferromagnetic wave vector ($K=1, -1$ in Fig. 1). Furthermore, the quadratic dispersion of the ferromagnetic magnons at $K=0$ becomes linear and remarkably the two-spinon continuum is still apparent at high energies.

As mentioned previously these calculations were carried out with a fermion Monte Carlo approach [3]. To simulate spin systems with this approach we first fractionalize the spin 1/2 degree of freedom into two fermions and then impose the constraint of single occupancy using a Hubbard interaction. This approach to spin systems works very well, but is clearly cumbersome. It however has the enormous advantage that we then couple spin systems to fermions so as to tackle various aspects of mixed spin-fermion systems [4].




Figure 1: Comparison of neutron scattering results with quantum Monte Carlo data. Taken from Ref. [1]. Right hand side: a single simulation of 384 spins with required precision for the analytical continuation requires 20,000 Core hours on SuperMUC-NG.

B) Superconductivity from the Condensation of Skyrmions in a Quantum-Spin-Hall insulator [2].

In the previous section, we have used numerical tools to provide modelling of experiments. In this section, we present a model that shows, for the first time, a very exotic piece of physics. The model contains 8-component Dirac fermions akin to graphene supplemented by a very special interaction of magnitude λ . The phase diagram is shown in Fig. 2 (see Ref [2] for more details). The interaction is chosen such that a quantum spin Hall (QSH) insulator is dynamically generated. This transition can be understood in the realm of the Ginzburg-Landau-Wilson theory of phase transitions: the QSH state breaks the spin-symmetry and can be characterized by a non-vanishing value of an $O(3)$ vector. The interplay of the long wavelength fluctuations of this vector with Dirac fermions is captured by the Gross-Neveu-Yukawa model. When the vector develops a non-vanishing vacuum expectation value, the Dirac fermions acquire a gap and form a QSH state. A dynamically generated QSH state has remarkable properties. Aside from the long wave length fluctuations of the order parameter—Goldstone modes—topological skyrmion excitations turn out to carry electric charge $2e$. We can hence foresee that if we manage to condense skyrmions a superconducting state will emerge. This is precisely what happens in our model when further enhancing the interaction strength λ . The transition between the QSH and s-wave superconductor (SSC) turns out to be continuous. This is unexpected in the Ginzburg-Landau-Wilson type theories, where generically phase transitions between two different symmetry broken states are first order. Our understanding of this transition follows the idea of deconfined quantum criticality in which a novel particle emerges at the transition. This particle corresponds to a fractionalized cooper pair. The particle-hole (particle-particle) condensate of this excitation corresponds the QSH (SSC). Our model is the first explicit realization of this route to superconductivity.

C) Methods

As exemplified above, the auxiliary field quantum Monte Carlo (QMC) approach can simulate spin as well as fermion systems in thermodynamical equilibrium. It is based on the following representation of the partition function

$$Z = \int d\phi e^{-S(\phi)}$$

where Φ is a high-dimensional vector. For a given field configuration we can compute the action, S , in polynomial time. In our explicit implementation the CPU time scales as the inverse temperature times the cubed of the volume. The challenges encountered in this method include. A) The action has to be real so as to allow for Metropolis important sampling. It turns out that there is a number of non-trivial models that fall in this category. If the action is complex, we are faced with a so called negative sign problem which requires reweighting schemes and invariably renders the approach exponentially expensive in inverse temperature and volume. B) Possibly long autocorrelation times. Here we are working on a number of different approaches to improve sampling and are including




Figure 2: Phase diagram of the model introduced in Ref. 2. At half-filling, $\mu=0$, the Dirac semimetal gives way to a quantum spin Hall (QSH) insulator and then to an s-wave superconductor (SSC).

Langevin dynamics as well as hybrid Monte Carlo updates in our ALF-package. C) Reaching large volumes. This is especially important for investigations of quantum critical phenomena and requires the power SuperMUC-NG.

Our ALF-code uses an MPI-OpenMP environment. OpenMP allows us to overcome the memory requirements of our algorithm. They scale as the inverse temperature times the square of the volume. OpenMP also allows us to reduce the sweep time. This is especially important for massively parallel simulations so as to avoid using up all the CPU time in the warmup phase.

Ongoing Research / Outlook

We are presently working on many other subjects including $SU(N)$ quantum spin systems, electron-phonon interactions, spin-fermion models, new (Landau-Level based) continuum regularizations to simulate topological field theories. Access to SuperMUC-NG computational resources allows us to be very ambitious in the choice of our projects, and is an indispensable tool for our research.

References and Links

- [1] H. Zhang, Z. Zhao, D. Gautreau, M. Raczkowski, A. Saha, V. O. Garlea, H. Cao, T. Hong, H. O. Jeschke, Subhendra D. Mahanti, T. Biroli, F. F. Assaad, and X. Ke, Phys. Rev. Lett. 125 (2020), 037204.
- [2] Yuhai Liu, Zhenjiu Wang, Toshihiro Sato, Martin Hohenadler, Chong Wang, Wenan Guo, and Fakher F. Assaad, Nature Communications 10 (2019), no. 1, 2658.
- [3] Martin Berx, Florian Goth, Johannes S. Hofmann, and Fakher F. Assaad, SciPost Phys. 3 (2017), 013. <https://alf.physik.uni-wuerzburg.de>
- [4] Bimla Danu, Fakher F. Assaad, and Frédéric Mila, Phys. Rev. Lett. 123 (2019), 176601.

Thermodynamics of the kagome lattice antiferromagnet: tuning frustration and quantum fluctuations

2

RESEARCH INSTITUTION

¹Bielefeld University, Faculty of Physics

PRINCIPAL INVESTIGATOR

Jürgen Schnack¹

RESEARCHERS

Jörg Schulenburg², Johannes Richter^{2,3}

PROJECT PARTNERS

²Otto-Von-Guericke University, Magdeburg

³Max Planck Institute for the Physics of complex Systems, Dresden

SuperMUC Project ID: pr62to

Introduction

The spin- $\frac{1}{2}$ kagome Heisenberg antiferromagnet (KHAF) is one of the most prominent and at the same time challenging spin models in the field of frustrated quantum magnetism, compare Fig. 1 for a sketch of the system. The first challenge concerns the nature of the ground state on which a plethora of studies exist. The second challenge concerns the thermodynamic properties of the quantum KHAF on which far less studies are available. The reason is of course that it is in general virtually impossible to evaluate thermodynamic observables such as heat capacity for an interacting many-body quantum system. The third challenge is given by the magnetization process of the spin- $\frac{1}{2}$ KHAF, i.e., by evaluating the magnetization M as a function of the external magnetic field B .

In our project we were able to determine the thermodynamic properties of the spin- $\frac{1}{2}$ KHAF on a finite lattice of $N=42$ sites. These results were obtained by large-scale numerical calculations (5 million core hours) using the finite-temperature Lanczos method (FTLM) employed in our groups [1]. The extension to a lattice of this (record-breaking) size yields an improved insight into the low-temperature physics of the model compared to previous studies restricted to significantly smaller lattices. Our experience is that our method of choice—the Finite-Temperature Lanczos Method [2]—delivers quasi exact results for quantum spin systems [3].

Results and Methods

The investigated spin systems are modeled by a spin- $\frac{1}{2}$ Heisenberg Hamiltonian augmented with a Zeeman term,

$$\hat{H} = J \sum_{\langle i,j \rangle} \hat{s}_i \cdot \hat{s}_j + g \mu_B B \sum_i \hat{s}_i^z,$$

where J is the antiferromagnetic nearest-neighbor exchange coupling and B the external magnetic field.




Figure 1. Sketch of the structure of the kagome lattice. Spins reside at the vertices, whereas edges depict magnetic exchange interactions. The red disks represent localized magnons (of minimal size) that form a magnon crystal at small temperatures and elevated magnetic fields close to saturation.

The complete eigenvalue spectrum of a spin system composed of spins $s=\frac{1}{2}$ can be evaluated for sizes of up to about $N=24$ depending on the available symmetries. For larger systems and in particular the results presented here we approximate the partition function according to FTLM as

$$Z(T, B) \approx \sum_{\gamma=1}^{\Gamma} \frac{\dim(\mathcal{H}(\gamma))}{R} \sum_{\nu=1}^R \sum_{n=1}^{N_L} e^{-\beta \epsilon_n^{(\nu)}} |\langle n(\nu) | \nu \rangle|^2.$$

In the spirit of trace estimators, the trace is in a Monte-Carlo fashion replaced by a much smaller sum over R random vectors $|\nu\rangle$ for each symmetry-related orthogonal subspace $\mathfrak{h}(\gamma)$ of the Hilbert space, where γ labels the irreducible representations of the employed symmetries. The exponential of the Hamiltonian is then approximated by its spectral representation in a Krylov space spanned by the N_L Lanczos vectors starting from the respective random vectors $|\nu\rangle$. The vector $|n(\nu)\rangle$ is the n -th eigenvector of the Hamiltonian this Krylov space. This allows to evaluate typical observables such as magnetization and specific heat.




Figure 2. Heat capacity of small cutouts of the KHAF [4].

We could produce two major results in this project: a rather complete discussion of the thermodynamic properties of the KHAF with $N=42$ for all temperatures and fields [4] as well as the presentation of a magnon crystallization which represents a new phase transition of the KHAF [5].

The most astonishing scientific achievement of our investigation of the $N=42$ KHAF is presented in Fig. 2. While the major maximum of the heat capacity at $T \approx J$ is practically independent of system size N , the second peak at very low temperatures moves to higher temperatures with increasing N . This is counterintuitive in view of common beliefs about the KHAF. Since for the classical KHAF the ground state is macroscopically degenerate, speculations were that this would lead to a macroscopic number of singlet states below the lowest triplet state in the respective quantum KHAF. Our findings demonstrate that these singlets are unevenly distributed below the lowest triplet in so far that the density of states of these singlets tends to move to higher energies with increasing size of the lattice. We speculate that in the thermodynamic limit, i.e. for $N \rightarrow \infty$, the peak merges partially with the main peak and forms a low-temperature shoulder [4].

In the second part of the project we were able to present numerical evidence for the crystallization of magnons below the saturation field at non-zero temperatures for the KHAF. Thanks to the high magnetic field and the low temperatures at which this phase transition exist we could evaluate thermodynamic functions for much larger system sizes, compare Fig. 3.

The phenomenon of magnon crystallization can be traced back to the existence of independent localized magnons or equivalently flat-band multi-magnon states. In Ref. [5] we present a phase diagram of this transition, thus providing information for which magnetic fields and temperatures magnon crystallization can be observed experimentally. The emergence of a finite-temperature continuous transition to a magnon-crystal is expected to be generic for spin models in dimension greater than one where flat-band multi-magnon ground states break translational symmetry.




Figure 3. Specific heat for $B=0.99B_{\text{sat}}$ for various finite-size realizations of the KHAF. For $N=45, 54, 63, 72$, where too large Hilbert subspaces had to be neglected, only the low-temperature part of the specific heat is displayed; it is virtually correct for all system sizes [5].

Figure 3 shows the hall mark of the phase transition, a peak of the heat capacity at low temperatures that is expected to become sharper with increasing lattice size N . The maximum at higher temperatures is independent of N and not related to a phase transition. Our FTLM data thus confirm the very existence of a low-temperature magnon-crystal phase just below the saturation field as conjectured by the hard hexagon approximation [6] and as seen in experiment [7].

Ongoing Research / Outlook

The presented numerical investigations are the largest of their kind performed to date. They have been made possible by an adaption of the program *spinpack*, developed by Jörg Schulenburg, to the architecture first of SuperMUC Phase 2 and later SuperMUC-NG. Massive hybrid MPI/OpenMP parallelization is employed to map the sparse Hamiltonian matrix as well as the Lanczos vectors onto the nodes of SuperMUC.

Our studies demonstrate that accurate numerical investigations of highly frustrated spin systems by means of FTLM are very well possible on highly parallelized supercomputers. They pave the road for deeper investigations of other fascinating quantum magnets.

References and Links

- [1] <http://obelix.physik.uni-bielefeld.de/~schnack>, <https://wasd.urz.uni-magdeburg.de/jschulen/spin/>.
- [2] J. Jaklič and P. Prelovšek. 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; (d) J. Schnack, J. Richter, R. Steinigeweg, Phys. Rev. Research 2 (2020) 013186; (e) J. Schnack, J. Richter, T. Heitmann, J. Richter, R. Steinigeweg, Zeitschrift f. Naturforschung A 75 (2020) 465.
- [4] J. Schnack, J. Schulenburg, J. Richter, Phys. Rev. B 98 (2018) 094423.
- [5] J. Schnack, J. Schulenburg, A. Honecker, J. Richter, Phys. Rev. Lett. 125 (2020) 117207
- [6] M. E. Zhitomirsky and Hirokazu Tsunetsugu, Phys. Rev. B 70, 100403 (2004).
- [7] R. Okuma, D. Nakamura, T. Okubo, A. Miyake, A. Matsuo, K. Kindo, M. Tokunaga, N. Kawashima, S. Takeyama, and Z. Hiroi, Nat. Commun. 10, 1229 (2019).

Finite-temperature Lanczos simulations of magnetic molecules

RESEARCH INSTITUTION

Bielefeld University, Faculty of Physics

PRINCIPAL INVESTIGATOR

Jürgen Schnack

RESEARCHERS

Christian Beckmann, Michael Czopnik, Oliver Hanebaum, Christian Heesing, Henning-Timm Langwald

PROJECT PARTNERS

—

SuperMUC Project ID: pr63fa

Introduction

Magnetic molecules constitute a fascinating class of magnetic quantum spin systems with potential applications in storage and quantum computing technology as well as in magnetocalorics. The theoretical simulation of their properties often poses a massive problem since their quantum description in Hilbert space is equivalent to matrix-vector problems of very high dimension. In order to assess their thermal properties as e.g. magnetization or heat capacity approximations are needed. Our group [1] uses the finite-temperature Lanczos method (FTLM) for quantum spin systems which delivers quasi accurate results up to Hilbert space dimensions of about 10^{10} [2]. Our expertise makes us attractive for many groups working on molecular magnetism worldwide. Figure 1 shows a selection of molecules investigated by us in this project over the past years [3].

Results and Methods

In this contribution we want to discuss magnetocalorics as one prominent application in the field of molecular magnetism. Magnetocalorics denotes the property of magnetic materials to heat up or cool down while the external magnetic field is swept for instance in an adiabatic (isentropic) process, i.e. with constant entropy. This allows to construct refrigerators for sub-kelvin cooling with such materials. One goal therefore consists in finding new and better magnetic molecules for sub-kelvin cooling.

The quality of the magnetocaloric properties of a molecule can be characterized by several figures of merit: the slope of the isentropes (curves of constant entropy) corresponding to the cooling rates, the adiabatic temperature change, the isothermal entropy change, and the temperature change that can be achieved with a certain field sweep starting from an initial temperature. The latter two are displayed in Fig. 2 in the rightmost column for a molecule containing seven gadolinium ions.

There are various ways to optimize these properties for certain applications. One important concept is to




Figure 1: Structures of magnetic molecules investigated in this project over the past years [3]. The Hilbert space dimension of the system of 13 Fe^{III} spins of 5/2 for the molecules in the lower left is $6^{13}=13,060,694,016$, to provide an example of the size of such problems for linear algebra solvers.

achieve very high densities of low-lying states for certain values of the magnetic field so that in a field sweep the entropy varies considerably and large cooling rates may be obtained. Such a behavior is often observed in frustrated quantum magnets. The term “frustrated” denotes the property of the spin system to be not able to satisfy all magnetic interactions in the ground state; in other words, we observe competing interactions. Then the energy spectrum of the molecule may exhibit unusual degeneracies of ground-state energies as a function of field as is also observed in low-dimensional quantum spin systems such as the kagome lattice antiferromagnet or the sawtooth chain.

It was a lucky coincidence that a molecule with the structure of a sawtooth (or delta) chain could be synthesized by our collaboration partners a KIT (group of Annie Powell), for the structure see Fig. 3. This molecule features very unusual properties. Not only did it show the largest ever found ground-state spin at the time of its synthesis, it also exhibits a quantum

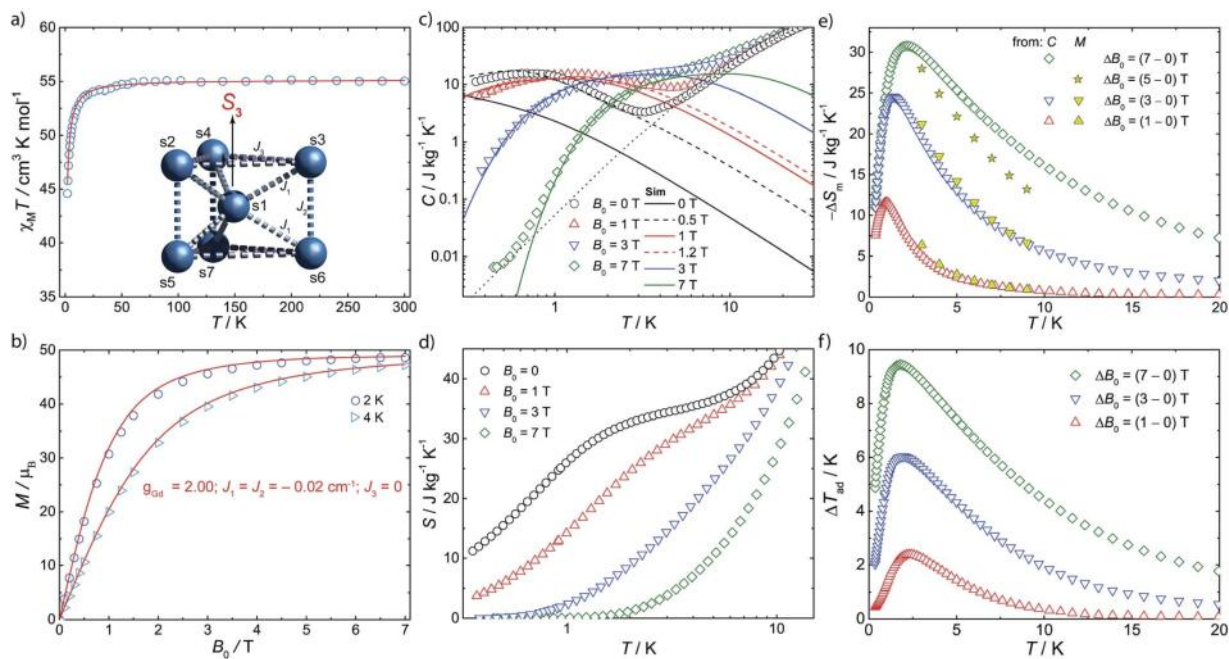


Figure 2: Example of magnetic observables usually determined for magnetic molecules. The two figures in the rightmost column display two figures of merit for magnetocalorics, the isothermal entropy change and the temperature change that can be achieved with a certain field sweep starting from an initial temperature [3].

phase transition [4]. A quantum phase transition takes place at $T=0$ as a function of another external parameter such as magnetic field or pressure and denotes a massive change of the character of the ground state. In our case it is accompanied by a huge degeneracy of ground-state energy levels at the critical point α_c . This variable denotes the ratio of the two exchange interactions present in $\text{Fe}_{10}\text{Gd}_{10}$; they are marked by red dashed and blue solid lines in Fig. 3 (left). Although the quantum phase transition happens strictly at $T=0$, it influences the magnetic properties at elevated temperatures, too. The physical properties of the quantum critical material in the cone-like region inside the two branches of the critical isentrope (red curves in Fig. 3 (right)) is drastically determined by the quantum phase transition. In this region also the magnetocaloric properties are rather unusual. For $\text{Fe}_{10}\text{Gd}_{10}$ we speculate that it should show barocaloric properties as well, since the ratio of the two exchange interactions is likely susceptible to pressure changes.

Summarizing, one could say that modern molecular magnetism teaches us a lot about quantum spin systems, see [5] for a recent review.

Ongoing Research / Outlook

Investigations such as those sketched in this contribution are only possible with the help of modern HPC architectures. The reason consists in the huge dimensions of the underlying Hilbert spaces. This dimension grows exponentially with the number of spins and yields for N spins of spin quantum number s a value of $(2s+1)^N$. Already for rather small N the dimension renders any exact linear algebra computation, such as determination of eigenvalues, impossible. Even the use of symmetries can attenuate the problem only somewhat.




Figure 3: Structure of the magnetic core of the magnetic molecule $\text{Fe}_{10}\text{Gd}_{10}$ (left) as well as phase diagram showing the critical isentrope (red curves) as well as sketches of the two magnetic phases separated by the quantum phase transition (right).

The use of trace estimators that built on imaginary time evolution in Krylov spaces has proven to deliver a very efficient and accurate approximation for the thermal observables of quantum spin problems with dimension of Hilbert spaces up to 10^{10} . Such schemes can be numerically treated as openMP, MPI or hybrid MPI/openMP programs. The core operation consists of matrix vector multiplications for which very efficient schemes and pre-compiled libraries, e.g. BLAS, can be employed.

SuperMUC-NG constitutes a perfect platform to achieve competitive, if not even world-leading, solutions in molecular magnetism.

References and Links

- [1] <http://obelix.physik.uni-bielefeld.de/~schnack>.
- [2] J. Jaklič and P. Prelovšek. 1994. Phys. Rev. B. 49, 5065–5068; J. Schnack, J. Richter, R. Steinigeweg. 2020. Phys. Rev. Research 2, 013186.
- [3] N.A.G. Bandeira, O. Sadeghi, T.J. Woods, Yuan-Zhu Zhang, J. Schnack, K.R. Dunbar, M. Nyman, C. Bo. 2017. J. Phys. Chem. A 121, 1310; T.N. Hooper, R. Inglis, G. Lorusso, J. Ujma, P.E. Barran, D. Uhrin, J. Schnack, S. Piličkos, M. Evangelisti, E.K. Brechin. 2016. Inorg. Chem. 55, 10535; E. Moreno Pineda, G. Lorusso, K. Zangana, E. Palacios, J. Schnack, M. Evangelisti, R.E.P. Winpenny, E.J. L. McInnes. 2016. Chem. Sci. 7, 4891.
- [4] A. Baniodeh, N. Magnani, Yanhua Lan, G. Buth, C.E. Anson, J. Richter, M. Affronte, J. Schnack, A.K. Powell. 2018. Nature Partner Journal Quantum Materials 3, 10.
- [5] J. Schnack. 2019. Contemporary Physics 60, 127-144.

Oxygen activation at Au/TiO₂ nanocatalysts

RESEARCH INSTITUTION

Lehrstuhl für Theoretische Chemie, Ruhr-Universität Bochum

PRINCIPAL INVESTIGATOR

Dominik Marx

RESEARCHERS

Niklas Siemer, Daniel Muñoz-Santiburcio, Johannes Frenzel

PROJECT PARTNERS

—

SuperMUC Project ID: pr74va

Introduction

The backbone of industrial chemistry is heterogeneous catalysis. Even small improvements of the used catalyst, often based on mechanistical insights, may enhance the efficiency in terms of energy consumption and/or selectivity. Such insights can be provided by computer simulations of model systems consisting of a few hundred atoms, while experiments often lack the required atomistic resolution.

Highly dispersed gold/titania catalysts are widely used for key reactions, notably including the selective oxidation of alcohols, for which we identified the activation of molecular oxygen as crucial reaction step.

Here, we take advantage of enhanced sampling ab initio molecular dynamics (AIMD) [1] simulations using a well-established model of a gold nanoparticle (AuNP) supported by titania (TiO₂) [2,3,4,5] (Figure 1) in order to elucidate the details of the mechanism and the corresponding energetic contribution of the activation of molecular oxygen at this nanocatalyst.

Results and Methods

We performed large-scale ab initio thermodynamic integration simulations with the oxygen molecule being in direct contact to the Au/TiO₂ nanoparticle and being in contact with the titania surface only, providing an internal reference case [5].




Figure 1: Close-up of the Au/TiO₂ nanocatalyst system used to simulate the activation of molecular oxygen. The TiO₂ slab is shown in dark red (O) and cyan (Ti), the AuNP in gold, and the adsorbed O₂ in red.

In summary, the direct contact between O₂ and the Au/TiO₂ interface is responsible for a 70% decrease in free activation barrier of oxygen dissociation compared to the case with O₂ adsorbed at the surface only (see Figure 2). Furthermore, the transition state (the maximum of the activation free energy curve) is located at shorter O–O bond length. This results in a much enhanced oxygen activation at this nanocatalyst compared to bare TiO₂.

This dramatic catalytic effect is explained by the charge transfer towards the oxygen molecule. In both cases most of the charge is transferred from the Ti atoms where the O₂ is bound. However, if adsorbed at the AuNP perimeter site the O₂ receives additional electronic charge density from the nearest Au atom of the AuNP. This amount of additional charge transfer increases with the elongation of the O–O bond (see Figure 3).

Further analyses allowed to explain how the enhanced charge transfer enhances the oxygen activation based on the properties of the electrons at the complicated O₂/AuNP/TiO₂ interface. The interaction of the AuNP with the oxygen molecule lowers the energy of an anti-bonding orbital of the oxygen molecule. This causes the orbital to be occupied earlier in the activation




Figure 2: Free energy profiles for oxygen dissociation at a Au/TiO₂ perimeter site (red) and far from the AuNP (green); the shaded areas illustrate the statistical error. Representative configuration snapshots (top row) show a top view of both adsorption sites. See Figure 1 for the color code.

process by the additional charge transferred, facilitating the oxygen activation.

Therefore, the direct contact between O₂ and the AuNP anchored at the TiO₂ surface is identified as the key aspect in enhanced oxygen activation at Au/TiO₂ nanocatalysts.

Performance and scaling

The AIMD simulations were carried out using CPMD [6] 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 ionic and electronic degrees of freedom of the system are time-propagated with an extended Lagrangian scheme, which is the most computationally demanding part of the calculation.

Such task is efficiently performed using 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 140 cores allowed for the highest speedup for our Au/TiO₂ nanocatalyst model. Herein we employed OpenMP parallelization with groups of seven cores per MPI task using 20 nodes (4 MPI tasks per haswell node) to allow for higher memory consumption per task. The 3D-FFT of the electronic wavefunctions is split up efficiently into two 2D-FFT per MPI task. Further levels of parallelization like computing one 2D-FFT per MPI task or further options within the CPMD code, e.g. via Kohn-Sham orbitals, resulted in lower speed-up. Thus, production runs were carried out with 2D-FFT parallelization giving the best utilization of the granted CPU time and the best time to solution. In addition, there exists another trivial level of parallelization which is the number of replicas used in the thermodynamic integration method. Each replica is independent of other replica of the system, allowing several simulations at the same time. On SuperMUC the presented results have used approximately 10 million core hours. 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 four hours and were 4.2 GiB in size.

Ongoing Research / Outlook

The new SuperMUC-NG framework allows us to include explicit water molecules in our ab initio molecular dynamics simulations. Currently, we are using the




Figure 3: Charge density differences upon adding the O₂ molecule to the surface at different O–O bond lengths as indicated by the labels in the graph of Figure 2 ($\Delta\rho = \rho(\text{O}_2/\text{Au/TiO}_2) - \rho(\text{Au/TiO}_2) - \rho(\text{O}_2)$). AuNP, Ti, O atoms and the O₂ molecule are depicted in gold, cyan, dark red, and red, respectively. Green and blue isosurfaces show charge accumulation and depletion at an isovalue of $\pm 0.01 \text{ |e}^-| \text{ \AA}^{-3}$, respectively.

Au/TiO₂ model to investigate the origin of the enhanced O₂ activation at this catalyst in the aqueous phase, with which we will be able to quantify such activation enhancement in liquid vs. gas phase in terms of free energy barriers. Such insights will be of great interest for the heterogeneous catalysis field. These simulations require about twice the computational resources as the presented thermodynamic integration simulations due to the necessary sampling of the water dynamics. We estimate an increase in computational cost by about 100 times for upcoming projects which would be required to allow for more complex reactions to be simulated.

References and Links

- [1] D. Marx and J. Hutter, "Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods", Cambridge University Press (2009).
- [2] M. Farnesi Camellone, D. Marx. J. Phys. Chem. Lett. 4, 514–518 (2013).
- [3] M. Farnesi Camellone, D. Marx. J. Phys. Chem. C 118, 20989–21000 (2014).
- [4] D. Muñoz-Santiburcio, M. Farnesi Camellone, D. Marx. Angew. Chem. Int. Ed. 57, 3327–3331 (2018).
- [5] N. Siemer, A. Lüken, M. Zalibera, J. Frenzel, D. Muñoz-Santiburcio, A. Savitsky, W. Lubitz, M. Muhler, D. Marx, and J. Strunk. J. Am. Chem. Soc. 140, 18082–18092 (2018).
- [6] CPMD, <http://www.cpmid.org/>, Copyright IBM Corp 1990-2018, Copyright MPI für Festkörperforschung Stuttgart 1997-2001.
- [7] <http://www.theochem.rub.de/research/marx>
- [8] <http://www.ruhr-uni-bochum.de/solvation>

Supercritical Water is not Hydrogen Bonded

RESEARCH INSTITUTION

Lehrstuhl für Theoretische Chemie, Ruhr-Universität Bochum

PRINCIPAL INVESTIGATOR

Dominik Marx

RESEARCHERS

Philipp Schienbein, Jan Noetzel

PROJECT PARTNERS

—

SuperMUC Project ID: pr86fo

Introduction

Hydrogen-bonding and the three-dimensional H-bond network topology certainly is the hallmark of water. Thus, thinking about water is inextricably linked to H-bonds. In our [1] recent funding period, we went out to answer the question if this common connotation remains valid in the supercritical phase of water.

The supercritical phase of matter is reached as soon as the temperature and the pressure exceed their respective critical values ($T_c = 647\text{K}$, $p_c = 221\text{bar}$, and $\rho_c = 0.332\text{kg/L}$ in case of water). Although these thermodynamic conditions might sound exotic at first glance, supercritical water (SCW) occurs also naturally, e.g. in hydrothermal formation processes and in so-called “black smokers” at the bottom of the deep sea. Besides, SCW is envisaged for or already in use in several applications as a “tunable solvent environment”.

To understand solvation processes in SCW, e.g. for chemical synthesis, a detailed knowledge about the microscopic properties of SCW is inevitable. Here it was already well known that the famous tetrahedral arrangement of the water molecules is destroyed in SCW. Moreover, H-bonds have also been investigated in the literature with the result that their extent is largely diminished, however that “some” H-bonding persists in SCW.

In this project we employed extensive ab initio molecular dynamics (AIMD) simulations to study the structure, dynamics, and electronic properties of SCW [2,3]. Furthermore, we also assessed the vibrational spectrum of SCW in the THz frequency regime [3], where the so-called intermolecular H-bond network mode is located. This mode directly monitors the intermolecular vibration between adjacent water molecules and thus

gives a direct handle to the intermolecular dynamics in water. Using this spectral response, we show that SCW should not be considered a “H-bonded” fluid [3]. Instead, we show that SCW can be regarded as an isotropic van der Waals fluid [3].

Results and Methods

Correlation lengths and times are known to be large in SCW. However, this is a problem for AIMD simulations, where electronic structure calculations are employed to obtain the forces at each time step. Since these computations are rather demanding, the system size and the simulation lengths are limited. For that reason, we employed a cascade of simulations which is illustrated in Fig. 1. First, we performed an inexpensive molecular dynamics simulation for 10 ns which is used to cover the large correlation lengths and times in SCW. From this simulation we started 10 independent AIMD simulations in the NVT ensemble at constant temperature for 25 ps each. In the last step we performed 40 subsequent AIMD simulations in the NVE ensemble, where the starting configurations are drawn from the preceding NVT simulations. All physical properties described in the following are drawn from the 40 independent NVE simulations [2,3,4].

For our analysis we selected in total 20 thermodynamic state points for which the described simulation protocol is carried out. Using these state points allows us to understand how the water structure changes systematically as a function of temperature from ambient towards supercritical conditions and from high density SCW to low density SCW. These thermodynamical pathways are illustrated in Fig. 2. In total, we simulated about 20 ns of AIMD trajectories using about 20 Mio core-hours on SuperMUC and (partly) SuperMUC-NG.




Figure 1: Schematic illustration of the employed simulation protocol (see text). Reproduced from Ref. [4] with permission from the PCCP Owner Societies.

From the AIMD simulations we extracted two key observables, namely the lifetime of a H-bond and the oscillation period of an intermolecular water—water contact. Since the H-bond lifetime is quantitatively dependent on the employed structural H-bond criterion, we employ different ways to obtain it in order to get a solid estimate as elaborately explained in Ref. [3]. These two purely dynamical quantities can then




Figure 2: Liquid-vapor coexistence curve of water including the experimental critical point and the critical point of our simulation model (RPBE-D3) as well as the simulated state points (green triangles). Reproduced from Ref. [4] with permission from the PCCP Owner Societies.

be compared as shown in Fig. 3. In case of ambient liquid water, we find an oscillation period of about 0.18 ps and a H-bond lifetime of 1.41 ps. This implies that an intermolecular H-bond between two water molecules vibrates roughly ten times, before the H-bond breaks apart. This dynamical behavior directly resembles the familiar picture of liquid water.

However, this picture changes dramatically when the temperature is increased towards supercritical conditions: At 750 K the oscillation period is about 224 fs and the H-bond lifetime is only about 78fs. This implies that the lifetime of a putative H-bond in SCW is shorter than the intermolecular vibration period. This observation holds true for all investigated supercritical state points irrespective of the density.

This imposes important consequences for H-bonding in SCW: H-bonds are broken while an intermolecular vibration is still ongoing which implies that the vibration is completely unaffected by the presence of a structural H-bond motive. It also suggests that the directionality in water, which is imprinted by the H-bonds at ambient conditions, is completely lost. Indeed, we elaborately show in our paper [3] that water molecules in SCW behave strikingly similar to an isotropic van der Waals fluid. In other words: The location of the H atoms in SCW is completely irrelevant for the intermolecular interactions.

Ongoing Research / Outlook

In the previous funding periods, we have in detail understood how the structure, dynamics, and electronic properties in SCW change with respect to ambient liquid water. Given the remarkable conclusion, that SCW is not H-bonded, we now want to advance our research towards solutions in SCW. An archetype




Figure 3: Oscillation period of the intermolecular water-water stretching vibration plotted against the H-bond lifetime. Reproduced from Ref. [3].

of a solute in liquid water are simple ions because they are the smallest possible chemical compounds. Moreover, their interaction with the adjacent water molecules at ambient conditions including their specific response in the THz frequency regime of the vibrational spectrum is well understood [5]. It is therefore suggestive to start to investigate solvation processes in SCW using simple aqueous ion solutions and investigate how the solute/water interactions change as a function of temperature and density, in full analogy to what has been done for pure sub and supercritical water [2,3,4].

From this follow up research project we aim to microscopically understand why solvation properties in SCW are dramatically different compared to ambient liquid water. Indeed, our work on pure water already suggests that we can truly expect fundamentally different microscopic interactions at supercritical solutions compared to what is known at ambient conditions.

References and Links

- [1] <https://www.theochem.ruhr-uni-bochum.de/>
- [2] P. Schienbein and D. Marx. Phys. Chem. Chem. Phys., 22, 10462 (2020).
- [3] P. Schienbein and D. Marx. Angew. Chem. Int. Ed. <http://doi.org/10.1002/anie.202009640> (2020).
- [4] P. Schienbein and D. Marx. J. Phys. Chem. B, 122, 3318 (2018).
- [5] P. Schienbein, G. Schwaab, H. Forbert, M. Havenith, and D. Marx. J. Phys. Chem. Lett., 8, 2373 (2017).

Transition metal oxide surfaces and interfaces for electronic and energy conversion applications

2

RESEARCH INSTITUTION

Department of Physics, University of Duisburg-Essen

PRINCIPAL INVESTIGATOR

Rossitza Pentcheva

RESEARCHERS

Okan Köksal, Hamidreza Hajiyani, Manish Verma, Benjamin Geisler

PROJECT PARTNERS

—

SuperMUC Project ID: pr87ro

Introduction

Transition metal oxide 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 (e.g. water splitting, thermoelectrics). During the reporting period our work within pr87ro has focused on three areas: the first one, funded by the DFG within CRC TRR80, comprises the search for topologically nontrivial phases within perovskite and corundum-derived oxide superlattices hosting a honeycomb pattern. A number of promising candidates for Chern insulators – the time-reversal symmetry broken analogues to the Z_2 topological insulators – were identified, besides a rich variety of orbital and magnetic reconstructions that are not available in the bulk or (001)-oriented superlattices. The second area funded by the priority program SPP1613 and currently CRC TRR 247 deals with understanding and optimization of the activity of anode materials for water splitting. In particular, detailed insight was obtained on the role of the surface orientation, termination and dopants on the performance of anodes in the oxygen evolution reaction. Specifically, active sites with some of the lowest reported overpotentials were identified and their efficiency correlated with the underlying electronic properties. Last but not least, we focused on the role of reduced dimensions in improving the thermo-

electric properties in oxide superlattices and predict encouraging power factors that are comparable with some of the best-performing oxide thermoelectrics, demonstrating that quantum confinement is a promising strategy to enhance the thermoelectric response in correlated transition-metal oxide superlattices that can be extended to a broader class of materials combinations.

Results and Methods

Confinement-driven electronic and topological phases in perovskite and corundum-derived oxide honeycomb lattices

Within this project funded by DFG, CRC TRR80 we investigated the possibility to stabilize nontrivial topological phases in oxide superlattices hosting a honeycomb pattern. The systematic DFT+ U calculations with the VASP and WIEN2k codes were extended from $3d$ to $4d$ and $5d$ cases [2,3]. Several promising candidates for Chern insulators – were identified in perovskite and corundum-derived superlattices (SLs) [2,3]. One example is the LaPtO_3 honeycomb bilayer confined by $\text{LaAlO}_3(111)$ shown in Fig. 1. Spin-orbit coupling leads to a band inversion and avoided crossing of the majority and minority bands around K in the Brillouin zone and the opening of a band gap of 38 meV. The Berry curvature exhibits a ring-like structure along the line of avoided crossing. Furthermore, a reversal of the spin orientation is observed in the spin texture. Thus $(\text{LaPtO}_3)_2/(\text{LaAlO}_3)_4(111)$ with symmetric sublattices of the honeycomb layer at the lateral lattice constant of LaAlO_3 emerges as a Chern insulator with $C = 1$. In contrast, increasing the lateral lattice constant leads to a disproportionation with two inequivalent Pt ions with magnetic moments of $1.06 \mu_B$ and $0.32 \mu_B$ [3].

Role of the surface orientation and termination on the OER performance of $\text{Co}_x\text{Ni}_{1-x}\text{Fe}_2\text{O}_4$ (funded by DFG, SPP1613 and CRC TRR247)

To shed light on the origin of improved OER performance in Ni-doped cobalt ferrite we explored the effect of cation mixing, surface orientation and termination in the $\text{Co}_x\text{Ni}_{1-x}\text{Fe}_2\text{O}_4$ spinel [4,5]. The DFT+ U calculations performed with the VASP code on typically 192 cores show that overall, the overpotential is low-

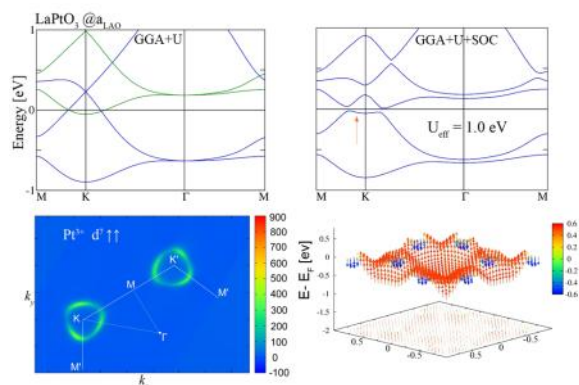


Figure 1: Band structure of $(\text{LaPtO}_3)_2/(\text{LaAlO}_3)_4(111)$ at a_{LaO} with majority (blue) and minority (green) bands crossing at E_F around K (top left). Spin-orbit coupling leads to a band inversion and avoided band crossing around K (right), visible as a ring-like structure in the Berry curvature $\Omega(\mathbf{k})$ (bottom left) and a spin reversal in the spin texture (bottom right).

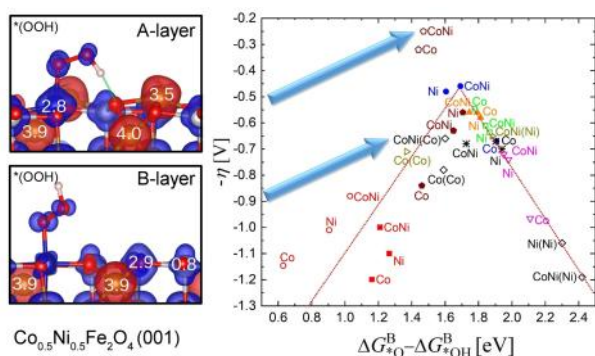


Figure 2: Overpotential versus the binding energy difference of *O and *OH for different reaction sites and terminations of the $\text{Co}_{0.5}\text{Ni}_{0.5}\text{Fe}_2\text{O}_4(001)$ and (111) surface. The left panels show a side view of the spin density of the *OOH intermediate at a Co active site for the A and B terminations [5].

ered for $x = 0.5$, compared to the end members. While the overpotential for the (111) surface is higher (0.55 V) [4] than the experimental value (0.34 V), a significant reduction is obtained for the (001)-surface.

In particular, when terminated with an additional tetrahedral Fe-layer (A-layer), $\text{Co}_{0.5}\text{Ni}_{0.5}\text{Fe}_2\text{O}_4(001)$ exhibits one of the lowest theoretically reported overpotentials of 0.26 V, that lies even above the volcano of $-\eta$ vs. the difference in binding energies of *O and *OH species (Fig. 2), due to a deviation from the standard scaling relationship. Octahedral Co cations are identified as the active sites. Analysis of the electronic properties and spin densities illustrates that the additional Fe layer stabilizes a bulk-like oxidation state of +2 for Co and Ni at the A-layer termination, while at the B-layer termination they are oxidized to +3. Moreover, the unusual relaxation pattern enables the formation of a hydrogen bond of the OOH intermediate to a neighboring surface oxygen (Fig. 2) that lowers the reaction free energy of this formerly rate-limiting step, leading to a deviation from the scaling relationship and almost equidistant reaction free-energy steps of intermediates. This renders an example of how a selective surface modification can result in a significant improvement of OER performance [4,5].

Insight and optimization of the thermoelectric properties in oxide superlattices exploiting reduced dimension (funding by DFG within CRC TR80, project G8)

We explored the effect of t_{2g} orbital occupation on the electronic, magnetic, and thermoelectric properties of $(\text{SrXO}_3)_1/(\text{SrTiO}_3)_n$ superlattices with $n=1,3$ and $X=\text{V}, \text{Cr}$, and Mn by using density functional theory calculations with the Wien2k code with an on-site Coulomb repulsion term combined with Boltzmann transport theory, using the BoltzTrap code. Disentangling the effect of quantum confinement and octahedral rotations, we find that the ground-state superlattice geometries always display finite octahedral rotations (see Fig. 3), which drive an orbital reconstruction and a concomitant metal-to-insulator transition in confined SrVO_3 and SrCrO_3 single layers with ferro- and antiferromagnetic spin alignments, respectively.

On the other hand, the confined SrMnO_3 layer exhibits electronic properties similar to the bulk. We show that confinement enhances the thermoelectric properties. Large room-temperature Seebeck coefficients are

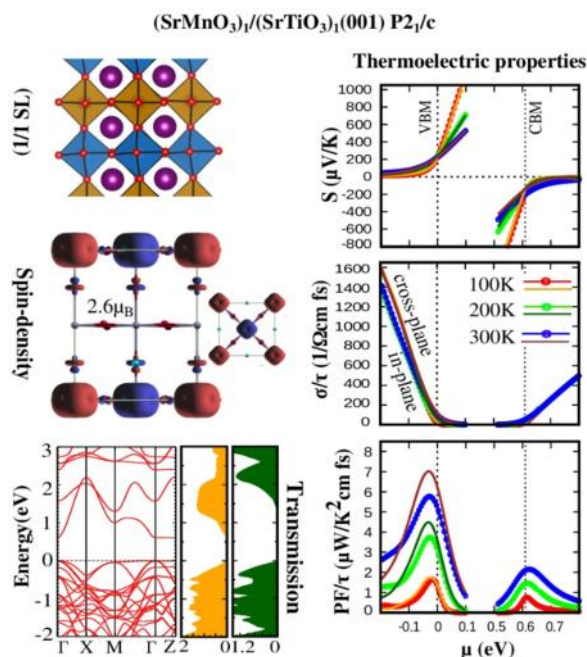


Figure 3: (left) Side view of the $(\text{SrMnO}_3)_1/(\text{SrTiO}_3)_n(001)$ SL showing the octahedral tilts, spin density and band structure with in and cross-plane transmission. (right) Seebeck coefficient, conductivity and power factor divided by the relaxation time, showing substantial values of S and PF/r at the band edges [6].

obtained for the tilted superlattices, reaching up to 600 $\mu\text{V/K}$ near the band edges (Fig. 3). The estimated attainable power factors compare favorably with some of the best performing oxide thermoelectrics, demonstrating that the idea to use quantum confinement to enhance the thermoelectric response in correlated transition-metal oxide superlattices can be applied to a broader class of materials combinations [6].

Ongoing Research / Outlook

The access to SuperMUC-NG was essential to achieve the reported results due to the large system sizes and number of studied configurations. Currently the investigations within CRC TRR80 are extended from the honeycomb to dice lattices, whereas further spinel and perovskite surfaces are studied within CRC TRR247 to relate the structural patterns and environments to their catalytic activity.

References and Links

- [1] <https://www.uni-due.de/physik/pentcheva/>
- [2] O. Köksal, S. Baidya and R. Pentcheva, Phys. Rev. B 97, 035126 (2018).
- [3] O. Köksal, and R. Pentcheva, Scientific Reports 9, 17306 (2019).
- [4] K. Chakrapani, G. Bendt, H. Hajiyani, I. Schwarzrock, T. Lunkenbein, S. Salamon, J. Landers, H. Wende, R. Schlögl, R. Pentcheva, M. Behrens, S. Schulz, ChemCatChem. 9, 2988-2995 (2017).
- [5] H. Hajiyani, R. Pentcheva, ACS Catal., 8, 11773-11782 (2018).
- [6] M. Verma, B. Geisler, and R. Pentcheva, Phys. Rev. B, 100, 165126 (2019).

Quantum Chemical Modeling of Actinide Interaction

with Cement Phases

RESEARCH INSTITUTION

Fachgebiet für Theoretische Chemie, TU München, Germany

PRINCIPAL INVESTIGATORS

Alena Kremleva, Sven Krüger, Notker Rösch

RESEARCHERS

—

PROJECT PARTNERS

—

SuperMUC Project ID: pr94je

Introduction

Selecting a site and constructing a final deep geological repository for highly radioactive waste, as well as ensuring the safety of the complex have to take into account the interaction of radioactive elements with water and rocks and also with containment and construction materials. Cement is widely used in underground constructions and for barriers against migration, as well as for solidifying radioactive waste. Thus, understanding the interaction of hazardous elements with cement is an important aspect in safety considerations and long-term modeling of a deep geological repository.

Due to nuclear power generation, uranium forms a large part of the radioactive waste that has to be safely stored. Other early actinide elements, in smaller amounts, appear as fission and decay products. Quite some experimental work has been devoted to the interaction of actinides with cement. In our project we carried out the first quantum mechanical computational studies on actinide sorption at model minerals representing essential cement phases, aiming at a mechanistic understanding at the atomic scale.

An essential product of the initial hydration of cement are calcium silicate hydrate phases (CSH), leading to the solidification of fresh cement. In aging cement CSH phases are also newly formed. During cement degradation CSH phases are rather resistant. Spectroscopic experiments have shown that the interaction and binding of actinide ions to cement is very similar to the sorption behavior on CSH phases [1]. Thus, CSH phases are regarded as the essential sorbing phase of cement for actinides. Batch and laser fluorescence spectroscopic experiments showed that actinides sorb fast at CSH phases, forming several sorbed species. These observations indicate a complex sorption system as the CSH substrate as well as the actinide sorption complexes are multifarious.

CSH phases exhibit only very small crystallites or are even gel like; thus, the structure of various CSH phases is only poorly known. For the more interesting aged cement with CSH phases with a calcium-to-silicon ratio (C/S) below 1.2, the structure is regarded

to be similar to tobermorite minerals. The latter consist of sheets formed by a mono-atomic calcium oxide layer, decorated at both sides with silicate chains (Fig. 1). These chains consist of SiO_4 tetrahedra with a periodicity of three units, where two tetrahedra are bound to the calcium oxide layer (pairing tetrahedra) and the third one bridges the pairing ones. Depending on the distance between sheets, one distinguishes several variants of tobermorite, containing a variable amount of water between the layers. For our studies we choose 14 Å tobermorite with C/S ratios of 0.67, 0.83 and 1.00 as model minerals. We also considered structures where bridging silica tetrahedra are missing, representing known defects in CSH. As the interlayer distance of CSH phases is known to decrease from about 13.5 Å for low C/S to 11 Å for high C/S, we also inspected 11 Å tobermorite as a substrate for actinide sorption.

Results and Methods

To model computationally how actinide ions interact with CSH model minerals, we carried out density functional calculations applying the plane wave approach as implemented in the MPI parallelized software VASP [2]. Electronic exchange-correlation effects were described by a functional of generalized gradient approximation type [3], including spin polarization where needed. Core electrons were represented by the projector augmented wave approach, accounting also for the relativistic effects in heavy elements [2]. Tobermorite bulk and surface models of CSH were described by periodic supercell models comprising about 450 atoms per unit cell. Water in contact with a tobermorite surface is modeled as a film of about two layers of water molecules. Preferred sorption sites and complexes in the interlayer or at surfaces of tobermorite are difficult to determine, due to a large number of soft degrees of freedom associated with the interlayer and the surface water. Facile reorganization of water arrangements leads to many minima on the potential energy surface which are close in energy. To achieve representative and comparable species and structures, which are not biased by the atomic configurations at the start, an ab initio molecular dynamics technique (AIMD) is applied [3]. Each system is equilibrated for at least 4 ps at room temper-

ature to relax the soft degrees of freedom, followed by a step where the structure of the system is optimized. Subsequently, we cycled an AIMD step of 1 ps followed by an optimization step until the total energy between subsequent optimization steps changed by less than 0.1 eV. In these calculations, we typically used 200–400 cores; determining a single sorption complex required 40–80 hours on SuperMUC or SuperMUC-NG.

We examined in some detail the sorption of uranium(VI), which forms the uranyl ion UO_2^{2+} . X-ray absorption fine structure measurements [4] show that uranyl typically forms 5 bonds to oxygen centers (coordination number $\text{CN} = 5$), 2–6 to Si and 1–4 to Ca neighbors. With our calculations for uranyl absorption in the interlayer we identified several types of sorption sites (Fig. 1). Depending on the C/S value, different sorption sites may be preferred. Uranyl is typically bound to 4 oxygen neighbors and shows 1–3 Si and up to 3 Ca neighbors. Various calculated interatomic distances satisfactorily reproduce measured values [3]. Besides bonds to tobermorite and water ligands, uranyl may also carry OH ligands, especially for higher C/S values. Compared to experiment, calculated CN values are somewhat too low. None of the sorption complexes modeled reproduces all measured interatomic distances. This may be interpreted to indicate that several sorption complexes are present in the experimental probes.

For comparison we also inspected uranyl absorption in 11 Å tobermorite. Due to its smaller distance between layers, all sorption complexes found for 11 Å tobermorite are of layer bridging character (Fig. 1). Interatomic distances and CN values were calculated to be rather similar to 14 Å tobermorite. Only the CN to Si increased on average. Thus, both tobermorite variants can be regarded as reasonable models for CSH phases.

Furthermore, we explored uranyl adsorption on tobermorite surfaces on the example of the basal (001) surface which is parallel to the tobermorite sheets. Various adsorption complexes of uranyl show similar interatomic distances as for absorption in the interlayer. CN values related to Si and Ca are lower than for species in the interlayer. Calculated CN values for complexes at the surface underestimate measured ones more strongly than for complexes in the interlayer. Thus, absorption in the interlayer should be the dominating sorption mechanism for uranyl, although estimated sorption energies are comparable for ad- and absorption.

For actinides of oxidation states III and IV, experiments suggest that sorption in CSH occurs not only in




Figure 1: Schematic representation of 14 Å tobermorite. The calcium oxide layer is indicated in green, silica chains in gray. In the central interlayer, pairing (blue) and bridging (red) silica tetrahedra of the silica chains are distinguished. Actinide ions might sorb at a surface (Adsorption), in the interlayer (Absorption) or be incorporated in the calcium oxide layer (Incorporation). Typical sites for actinide absorption in the interlayer (layer bridging, at a defect, at an edge of a bridging silica tetrahedron, or chain bridging) are indicated by blue dots.

the interlayer, but also by incorporation in the CaO layer. In ongoing studies on the examples of U(IV) and Cm(III), we are comparing both sorption modes. Calculated CN values for U(IV) in 14 Å tobermorite agree better with experiments for Np(IV) and Pu(IV) when U(IV) is incorporated in the CaO layer, in agreement with estimated energies. Experiments for the lanthanides Eu(III) and Nd(III), which are regarded as analogs to actinides, yield geometry parameters and CN values in fair agreement with our results for Cm(III) in 14 Å and 11 Å tobermorite. Only for the CN value to Ca, lower numbers than calculated for Cm(III) incorporation have been measured. This suggests that for actinide(III) absorption in the interlayer and incorporation in the CaO sheet should both occur, which is supported by similar calculated energies. Thus, our computational studies confirm sorption by incorporation for U(IV) and Cm(III), although a rather fast sorption process is reported in experimental studies for this solid-state reaction.

Ongoing Research / Outlook

The results of this project support tobermorite as a model for CSH and led to a new atomistic model of actinide sorption in CSH phases, which will be extended in ongoing studies. These results help to construct thermodynamic models of actinide sorption in cement, which are required for studying actinide migration and safety considerations. This work has been carried out in the joint project GraZ, funded by Bundesministerium für Wirtschaft und Energie under project no. 02E11415E.

References and Links

- [1] J. Tits, E. Wieland, PSI-Bericht 18-02, 2018.
- [2] www.vasp.at
- [3] A. Kremleva, S. Krüger, N. Rösch, Appl. Geochem. 113 (2020) 104463.
- [4] M. Harfouche, E. Wieland, R. Dähn, T. Fujita, J. Tits, D. Kunz, M. Tsukamoto, J. Colloid Int. Sci., 303 (2006) 195-204.

Numerical simulations of topological and correlated quantum matter

RESEARCH INSTITUTION

¹Institut für theoretische Physik und Astrophysik, Universität Würzburg

PRINCIPAL INVESTIGATORS

Fakher F. Assaad¹, Ewelina Hankiewicz¹, Giorgio Sangiovanni¹

RESEARCHERS

Florian Goth¹, Domenico Di Sante¹, Armando Consiglio¹, Severino Adler¹, Marius Fuchs¹, Philipp Eck¹, Stefan Enzner¹, Andreas Hausoel¹, Alexander Kowalski¹, Anika Götze¹, Tobias Müller¹, Yasir Iqbal²

PROJECT PARTNERS

²Indian Institute of Technology Madras, Chennai

SuperMUC Project ID: pr94vu

Introduction

This project provides computational resources to the numerically oriented PIs of the SFB 1170 on Topological and Correlated Electronics at Surfaces and Interfaces. With the discovery of graphene, Dirac electrons became a common feature in the solid state. An aim of our research, that is summarized in the first two sections, is to identify materials where correlation effects in Dirac systems may become dominant [1,2]. Spin orbit coupling is a key feature to understand spintronics as well as the surface states of strong topological insulators. In a recent paper [3], summarized in the third section, the origin of the Rashba spin-orbit coupling is discussed. The last section summarizes work studying the Heisenberg model on the pyrochlore lattice [4]. Our calculations rely on combining first principle methods, the Vienna ab initio simulation package (VASP), with the constrained random phase approximation to obtain the Coulomb interaction, and the pseudo-fermion functional renormalization group (PFFRG) to study magnetism. The success of the research hinges on access to SuperMUC-NG.

Turbulent hydrodynamics in strongly correlated Kagome metals

In an interdisciplinary effort between ab initio studies, phenomenological solid state theory, and AdS/CFT holography, we put forward correlated Kagome metals at Dirac filling as a new arena for Dirac fluids with a large effective fine structure constant α , which allows to access the regime of turbulent hydrodynamic electron flow. We propose correlated Kagome metals at Dirac filling, and the prototypical example of Sc-

substituted Herbertsmithite (Figure 1a), as a new class of materials suited to reach the strong coupling regime of hydrodynamic Dirac electron flow. Standard methods such as Fermi liquid theory are not applicable in the strongly coupled regime. We used the concept of duality, by which a strongly coupled system is equivalently described by a gravity theory involving a black hole, providing then a new extrapolation of the ratio of shear viscosity over entropy density η/s to intermediate coupling values. Analysing Sc-Herbertsmithite from this AdS/CFT angle, with a calculation of η/s as function of α , we found that a Dirac fluid with coupling strength $\alpha=2.9$ (from ab initio cRPA calculations) quickly approaches the perfect quantum fluid limit, and AdS/CFT bounds motivate a phenomenological error bar (Figure 1b). The effective electron coupling in Sc-Herbertsmithite is three times that of graphene. This renders the Reynolds number sufficiently large to bring turbulent regimes within experimental reach.

Kagome metal-organic frameworks as a platform for strongly correlated electrons

We investigated with ab initio techniques a monolayer of Copper Dicyanoanthracene (Cu-DCA). This metal-organic compound crystallises in a Kagome lattice, which crucially determines the electronic structure. A DFT simulation revealed the presence of typical three-Kagome-like bands at the Fermi level, consisting of one flat band and two bands linearly crossing at the K point of the Brillouin zone. This crossing defines the topological so-called Dirac point at the Fermi level. In its vicinity, the electrons disperse linearly with a Fermi velocity $v_F \sim 0.8$ eVÅ. There are van-Hove-singularities in the Density of States (DOS) which can stabilise the formation of new electronic phases. These Kagome bands are formed by three p_z -like orbitals originating in the DCA molecules. A subsequent cRPA calculation estimated the screened Coulomb interaction, which is an order of magnitude larger than the electronic bandwidth, indicating the importance of correlations in this system. Even though the Wannier functions spread throughout the whole DCA molecule, strong electronic correlations localise the electrons, so that they form

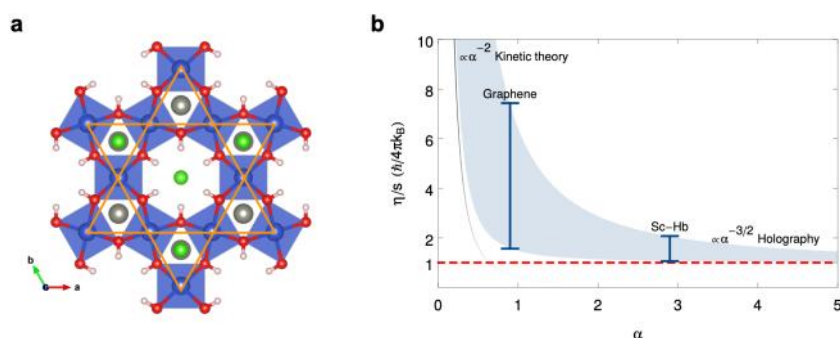


Figure 1: (a) Top view of the crystal structure of Sc-Herbertsmithite, where the CuO_4 plaquettes form a Kagome lattice (orange highlight). (b) η/s as a function of the coupling strength α . Black line: prediction in the weak coupling regime (Boltzmann kinetic theory), reliable for small values of α . The red dashed line corresponds to the universal holographic value $1/4\pi$ (in units of \hbar/k_B). The blue shaded region (η/s expanded in powers of $1/\alpha$) represents the holographic prediction.




Figure 2: (a) Atomic structure of Kagome Cu-DCA lattice. The DCA molecules form a Kagome texture. The low-energy electronic dispersion (panel (b)) is given by three molecular Wannier functions, each centred around a DCA molecule. (c) Linear extrapolations (solid lines) of the interactions U and U' as a function of the inverse out-of-plane component of the unit cell $1/c$ and cRPA calculations (circles and squares).

spin moments. The system holds promise in forming a viscous electron fluid which is difficult to study since it cannot be described by perturbation theory [1]. Studying this system would have been impossible without access to modern supercomputing facilities. The large size of the primitive cell both in the electronic degrees of freedom and volume poses great computational demands when simulating the system within DFT. The cRPA investigation adds severe memory constraints, due to the quartic scaling in both plane wave basis and Brillouin zone sampling.

Orbital driven Rashba effect in a binary honeycomb monolayer AgTe

We have studied, in collaboration with the experimental group EP-VII of the University of Würzburg, the interplay of spin-orbit coupling and inversion symmetry breaking in the 2D surface alloy AgTe on Ag(111). Te grows on Ag(111) in a triangular monolayer with an Ag vacancy in the first substrate layer (see Figure 3 panel a) forming an effective binary honeycomb lattice. This breaks the inversion symmetry in the Te lattice, as its in-plane p orbitals feel a different hybridization with the Ag s orbitals in opposite directions. This hybridization promotes orbital angular momentum polarization in bands with high radial Te p-character. The interplay of broken inversion symmetry and spin-orbit coupling gives rise to the Rashba effect, a momentum dependent spin-splitting. Calculations and experiment (see Figure 3 panel b) reveal indeed a predominant orbital angular momentum polarization and spin-splitting in the second valence and first conduction band with a high radial p-character and a suppressed splitting in the first valence band with a high tangential p-character. These findings illuminate the mechanism behind the orbital driven Rashba effect and are also of interest for other spin-polarized electron hosting material classes, e.g. in Weyl-semimetals or quantum spin Hall insulators with broken inversion symmetry. Our theoretical analysis benefited from the computational resources provided by the LRZ as the realistic simulation with density functional theory of Te on the metallic Ag(111) substrate requires a large number of substrate layers.

Quantum paramagnetism in 3D lattices

Quantum fluctuations in magnetic materials open up the possibility for a highly correlated, but disordered state of matter called quantum spin liquid. Promising candidates for this elusive state are materials with




Figure 3: (a) Top view on Te on Ag(111). The Ag vacancy in the surface substrate layer breaks the inversion symmetry in the triangular Te lattice. (b) Band character and orbital angular momentum polarization in Te on Ag(111). Bands with strong radial character carry a dominant orbital angular momentum polarization leading to a sizeable Rashba splitting.

strongly frustrated magnetic sublattices, such as Kagome in 2D. However neutron scattering experiments focused on frustrated 3D magnets and revealed magnetic fluctuation profiles not theoretically understood so far. With PFFRG We can study systems inaccessible to other methods which are unable to study parameters relevant to experiments. Using the perfectly parallelizable PFFRG equations in conjunction with the computational power of SuperMUC, we uncovered the full phase diagram of the idealized Heisenberg model on the Pyrochlore lattice, a frustrated 3D generalization of the 3D Kagome lattice [4]. We identified several magnetically ordered phases, and an extended phase space region, where quantum spins do not order. Combining PFFRG with ab-initio methods to obtain the magnetic Hamiltonian of candidates, we were able to identify $\text{Lu}_2\text{Mo}_2\text{O}_5\text{N}_2$ as a spin liquid candidate with a magnetic Pyrochlore sublattice consisting of Mo atoms. We predict the material to host a gear-wheel quantum spin liquid. Its predicted magnetic fluctuation profile shows, as measurable in neutron scattering experiments, a gearwheel like spectral distribution. In a combined effort of ab-initio calculations with PFFRG we identified the microscopic origin of a huge diversity in magnetic susceptibility profiles for a family of Spinels, all characterized by a magnetic Chromium sublattice in the form of a distorted (breathing) Pyrochlore lattice [5]. Although all materials are similar in structure, we found that the interplay of different longer range interactions up to fourth nearest neighbors leads to different ordering tendencies, that lie in proximity to spiral spin liquids with degeneracies in three, to one dimensional manifolds in ordering vector space. This shows, that the Pyrochlore lattice and its generalizations hide a rich phase diagram especially away from the (next)-nearest-neighbor model.



Ongoing Research / Outlook

The availability of a modern HPC System in the form of SuperMUC-NG was crucial for our projects and we will further pursue research into topological states of matter.

References and Links

- [1] D. Di Sante et al., Nat. Commun. 11, 3997 (2020).
- [2] M. Fuchs et al., J. of Phys.: Materials 3, 025001 (2020).
- [3] M. Ünzelmann et al., Phys. Rev. Lett. 124, 176401 (2020).
- [4] Y. Iqbal et. al, Phys. Rev. X 9, 011005 (2019).
- [5] P. Ghosh et. al, npj Quant. Mat. 4, 1-10 (2019).

Computational Fluid Dynamics and Engineering



Simulating Combustion Systems

RESEARCH INSTITUTION

Institute for Fluid Dynamics, IVG, University of Duisburg-Essen

PRINCIPAL INVESTIGATOR

Andreas Kempf

RESEARCHERS

L. Cifuentes, S. Baik, E. Inanc, J. Sellmann, T. Lipkowicz, M. Nanjaiah, P. Gruhlke

PROJECT PARTNERS

—

SuperMUC Project ID: pn34ki (GCS Large Scale project)

Introduction

This project collates individual applications from the Fluid Dynamics group at Duisburg-Essen University. The subprojects include the investigation of phenomena from the fields of nano-particle synthesis, supersonic flows, stratified burners, as well as the first simulation of staged combustion in a gas turbine combustor using LES and direct chemistry. The main objective of all subprojects is to improve the understanding of combustion systems to enable lower-emission, more efficient and explosion-proof combustion processes, partially without any CO₂ emissions if biofuels or hydrogen are used. A selection of highlights is presented in the following paragraphs.

The first subproject numerically investigated a bluff-body stabilized turbulent jet flame developed at Cambridge University that burns in stratified mode, using direct numerical simulations (DNS). The grid resolution was set to a value of $\Delta = 0.1 \text{ mm}$ such that the flame is resolved by at least 5 numerical cells. To model chemistry, an efficient premixed flamelet generated manifold (PFGM) was chosen that includes preferential diffusion without increasing the dimensions of the respective tables. The analysis focuses on dissipation rates and cross dissipation rates of mixture fraction and flame progress variable and can be used to check the applicability of common RANS (engineering) models for stratified flames.

The second subproject lead to the first simulation of staged combustion in the context of gas turbines with LES and direct chemistry. The investigated combustor consists of several jets that are separated in four

different fuel stages, namely main stages A, B and C and a pilot stage P. Three different staging concepts under similar load levels were compared, which mainly differ in terms of fuel staging and small pressure variations. A validated LES finite rate chemistry (FRC) combustion modeling




Figure 1: Iso-contour of the cross-dissipation rate super-imposed to the flame brush C=0.8.

approach was applied for a predictive study to reveal sources leading to incomplete CO burnout observed at two of the tested operating conditions, with the focus on the following goals: to demonstrate the ability to predict CO from a very complex full scale combustor and, to investigate the impact of mixture formation in the different fuel stages on CO and, to analyze the effect of local flame quenching due to secondary air on CO.

The third subproject simulated detonation-wave propagation in confined channels with- and without ozone to better understand the cellular structure and the driving mechanism towards a stable cell size. Experiments also suggest that the detonation propagation limit appears when the characteristic cells size coincides with the scale of the propagation device. Therefore, predicting the cell size will be important to allow stable propagation in small channels and to avoid unwanted detonation waves.

The fourth subproject investigated the formation of nano particles in reactive flows using DNS, while the grid resolution resolved not only the smallest flow- but also the smallest relevant length scales of the nano particle concentration field. The main objective here is to understand the physics of diffusion, coagulation, and nucleation and to use the DNS database for future modelling efforts.

Results and Methods

All simulations of subproject 1 were performed with the in-house code *PsiPhi* on a computational domain of $113 \times 120 \times 120 \text{ mm}^3$ ($1,130 \times 1,200 \times 1,200 = 1.63$ billion numerical cells) on an equidistant cartesian grid, and consumed 3.7 million core-h.

Figure 1 presents a section of the flame, showing the resolved cross-dissipation rate iso-contours of the DNS. The cross-dissipation term indicates the rate of stratification effects affecting the flame. The normalized reaction rate progress variable C denotes the progress of the reactions in the flame by marking the unburnt side of the flame (zero) and the burnt side of the flame (unity). A high cross-dissipation rate is observed downstream of the inlet as the flame interacts with the mixture stratification. This




Figure 2: Carbon monoxide mass fractions over equivalence ratio Φ in dependency of the CO reaction rate at the combustor exit, for three different operating conditions (1-3) and two computational grids (1, 1c).




Figure 4: Steady flames (top); induced vortices (bottom).

experiment reveals both the premixed (near the inlet) and stratified mode (downstream).

The open source library *OpenFOAM-6* was used for all the simulations of subproject 2. The large-eddy simulations (LES) typically used 10,000 numerical cells per core and up to 6,096 Cores. Figure 2 presents the carbon monoxide mass fractions over equivalence ratio colored by the CO source term $\dot{\omega}_{CO}$ for operating points 1 (weak stratification), 2 (strong stratification), and 3 (moderate stratification).

It can be seen that large CO levels corresponding to operating point 1 can be related to mixtures with increased equivalence ratios, whereas operating point 2 (strong stratification) reveals increased CO levels of $\approx 1,000$ ppm on the fuel-lean side. Very low CO concentrations of ≤ 50 ppm with a narrow range of equivalence ratios are identified for operating point 3. The contribution of each stage on the total mixture indicates that increased CO concentrations in the weakly stratified case 1 can be attributed to mixture inhomogeneities resulting from poor mixing, especially of the pilot stage. Since the reduction of CO and NO_x emissions is one of the major challenges in modern design of gas turbine chambers, this simulation helps to identify the mechanisms of emission production in a realistic setup.

All simulations of subproject 3 were performed with the in-house code PsiPhi. In order to resolve important




Figure 3: Numerical Schlieren of detonation wave fronts.

features of detonation waves, numerical cell widths of $\Delta = 5 \mu\text{m}$ were required. Combining these grid requirements with direct chemistry and a detailed reaction mechanism (9 species, 32 reactions) allowed simulations only in two dimensions. A simulation in a channel with a width of 24 mm was run on 27,648 Cores for 56 hours, resulting in a total cost of 1.5 M core-h. Figure 3 presents numerical Schlieren for detonation propagation in a narrow (6mm) and wide channel (24mm). The results from numerical soot foils show a significant reduction of detonation cell size, if dopant concentrations of ozone (3,000 PPM) are added to a stoichiometric hydrogen/oxygen mixture. The results further reveal an impact of the confinement geometry on cell size, as well as differences when using Navier-Stokes equations instead of Euler equations.

The simulations of the fourth subproject were performed with the in-house code PsiPhi, using direct chemistry with 32 species and 80 reactions and a sectional model for nano particles on a computational domain with $1,300 \times 1,000$ numerical cells at a grid resolution of $\Delta = 10 \mu\text{m}$. Each simulation required 7M core-h on 13,000 Cores.

Figure 4 presents temperature fields from nano-particle synthesis in a premixed and in a non-premixed case. The results help (i) to give information about the statistics, (ii) to identify the leading order of the physical mechanisms, (iii) to quantify the impact of the individual terms, and (iv) to evaluate the influence of the curvature.

References and Links

[1] <https://www.uni-due.de/ivg/fluidynamik/>.

DNS of premixed flame-wall interaction in turbulent boundary layers

RESEARCH INSTITUTION

¹Numerical methods in Aerospace Engineering, Bundeswehr University Munich, Germany

PRINCIPAL INVESTIGATOR

Markus Klein¹

RESEARCHERS

U. Ahmed², N. Chakraborty²

PROJECT PARTNERS

²Newcastle University, Newcastle upon Tyne, UK

SuperMUC Project ID: pn34xu

Introduction

The presence of walls alters the thermo-chemical and fluid-dynamical processes associated with turbulent premixed flames. However, the analysis of these phenomena poses severe challenges for both experimental and numerical investigations e.g. in terms of spatial and temporal resolutions. Hence, this aspect has not yet been sufficiently investigated in the existing literature. The increasing demands for light-weight combustors make flame-wall interaction (FWI) an inevitable event, which influences the cooling load, thermal efficiency and pollutant emission, in these applications. Therefore, a thorough physical understanding of the FWI mechanism is necessary to develop and design more energy-efficient and environmentally-friendly combustion devices.

Methods

Three-dimensional Direct Numerical Simulations (DNS) of V-flames interacting with chemically inert walls in a fully developed turbulent channel flow have been performed under adiabatic and isothermal wall boundary conditions using single-step chemistry. These simulations are representative of stoichiometric methane-air mixture at unity Lewis number under atmospheric conditions. The turbulence in the non-reacting channel is representative of the friction velocity based Reynolds number $Re_{\tau} = 110$.

A well-known three-dimensional compressible Fortran DNS code SENG+ has been used to simulate the oblique flame-wall interaction of a V-flame with inert isothermal and adiabatic walls in a fully developed turbulent channel flow. The code employs high-order finite-difference (10th order for internal points and gradually decreasing to 2nd order at the non-periodic boundaries) and Runge-Kutta (3rd order explicit) schemes for spatial differentiation and time advancement, respectively and is parallelized using MPI. The governing equations of mass, momentum, energy, and species mass fractions are solved in a non-dimensional form and a single-step irreversible reaction is used for the purpose of computational economy.

The domain size for this channel is $10.66h \times 2h \times 4h$, h being the channel half height, and has been discretized on $1,920 \times 360 \times 720$ equidistant grid points. The computational box has been decomposed into $40 \times 15 \times 8$ subdomains and run on 4,800 cores in parallel. The simulations have been performed for approximately 3 flow through times and the data has been sampled after one flow through time once the initial transience have decayed. The simulation consumed 1,500,000 CPU hours per boundary condition in total. One result file containing all variables for one time step has a size of 41 GB and altogether 180 time steps have been stored for collecting sufficient statistics.

Results

The instantaneous flame structures represented by the $c = 0.5$ isosurface of reaction progress variable for the two cases considered in this work along with the normalised vorticity magnitude Ω are shown in Fig. 1. The influence of the walls on Ω and the existence of wall ejections due to introduction of the fully developed boundary layer at the inflow are clearly visible in both cases.

It has been found that the location at which the oblique flame-wall interaction occurs is affected by the choice of the wall boundary condition. In order to investigate the aforementioned behaviour further, the mean behaviours of the Surface Density Function (SDF, $|\nabla c|$) and the strain rates affecting $|\nabla c|$ transport have been analysed. The behaviour of SDF is significantly affected by the changes in the wall boundary conditions within the viscous sub-layer region. It is found that the dilatation rate effects weaken in the viscous sub-layer region in the case with adiabatic wall boundary condition due to the constriction of the velocity gradients in the wall normal direction and also due to the existence of the cold wall in the case of isothermal wall boundary condition.

The mean behaviours of displacement speed have also been investigated and it is found that in both cases the mean displacement speed decreases in the viscous sub-layer region, but remains positive in the




Figure 1: V-flames for adiabatic (top) and isothermal (bottom) wall boundary conditions. The isosurface coloured in red represents $c = 0.5$. The instantaneous normalized vorticity magnitude is shown on the x - y plane. The grey surface denotes the bottom wall.

case of adiabatic wall boundary condition, whereas in the case of isothermal wall boundary condition the mean displacement speed has been found to be negative in some parts of the flame.

This consequently leads to differences in the normal strain rate arising from flame propagation and the curvature stretch under different wall boundary conditions.

The sensitivity of the statistics of SDF to the choice of the wall boundary conditions and the distance from the wall suggests that the sub-models of the flame surface density or scalar dissipation rate transport need to accurately capture the respective behaviours of the unclosed terms under these conditions.

Ongoing Research / Outlook

After obtaining first experience with this challenging flow configuration the next step consist of increasing the flow Reynolds number in order to obtain more representative strain values in the boundary layer.

Although it can be expected that the findings of this work will be qualitatively valid in the presence of detailed chemistry and transport, the presence of detailed chemistry, together with different possible choices of reaction progress variable, may give rise to

modified statistical behaviours of the Surface Density Function (SDF) in the vicinity of the wall. This along with the utilisation of the near wall SDF statistics to improve the FSD and SDR based reaction rate closures in the vicinity of the wall will form the basis of future investigations and follow up proposals.

References and Links

- [1] U. Ahmed, N. Chakraborty, and M. Klein. Scalar gradient and strain rate statistics in oblique premixed flame-wall interaction within turbulent channel flows. *Flow Turbulence and Combustion*, <https://doi.org/10.1007/s10494-020-00169-3>, 2020.
- [2] U. Ahmed, N. Chakraborty, and M. Klein. Effects of turbulent boundary layer on the near-wall flame dynamics in the case of adiabatic walls. In *13th International Symposium on Engineering Turbulence Modelling and Measurements*, Rhodes, May 2021.
- [3] U. Ahmed, N. Chakraborty, and M. Klein. Oblique flame-wall interaction in premixed turbulent combustion under isothermal and adiabatic wall boundary conditions. In *27th International Colloquium on the Dynamics of Explosions and Reactive Systems*, Beijing, China, August 2019.
- [4] U. Ahmed, I. Konstantinou, M. Klein, and N. Chakraborty. Near-wall behaviour of turbulence in flame-wall interaction of premixed turbulent combustion. In *9th European Combustion Meeting*, Lisboa, Portugal, April 2019.

On the dynamics of streamwise and transverse sediment patterns in turbulent open channel flows

RESEARCH INSTITUTION

Institute for Hydromechanics, Karlsruhe Institute of Technology (KIT), Karlsruhe

PRINCIPAL INVESTIGATORS

M. Scherer, M. Uhlmann

RESEARCHERS

—

PROJECT PARTNERS

—

SuperMUC Project ID: pn34yi

Introduction

The formation of subaqueous sediment bedforms in natural rivers or man-made canals has fascinated scientists and engineers likewise for over a century. Despite this long period of intensive research, the complex interaction between a turbulent stream and a sediment bed which causes various forms of sediment patterns to evolve is not fully understood up to the present day. However, an accurate description of a river's bed evolution is crucial for the hydraulic engineer in order to, for instance, correctly predict the evolution and potential risks of flooding events.

The current project aims to contribute to the fundamental understanding of the underlying physical processes and to the development of sediment transport models by providing high-fidelity datasets of sediment pattern formation. All relevant flow scales including vortices of size smaller than the particle scale are resolved by means of direct numerical simulations featuring fully-resolved particles.

In this work, we focus in particular on the different formation processes of and the interaction between transverse-oriented ripple-like patterns and streamwise-aligned sediment ridges. Although both patterns evolve from the same sediment bed, their evolution time scales, their shapes as well as the corresponding flow structures strikingly differ. A second objective of the study is to assess the influence of lateral side walls

on the two bedform classes of interest, which appear in man-made canals or, in form of side banks, also in natural rivers.

Results and Methods

The two-phase system consisting of the carrying fluid and the dispersed particle phase is simulated using an immersed boundary technique [4]. A second order finite difference scheme on a uniform staggered grid together with a standard fractional-step method and mixed explicit-implicit time integration is applied to describe the fluid motion numerically. The motion of the spherical particles is determined by integration of the Newton-Euler equations for rigid body motion, which is coupled with the fluid solver in the framework of a sub-stepping algorithm. Collision forces and torque which arise from particle-particle or particle-wall contacts are obtained by a soft-sphere collision model [1].

The reasons for the necessity to perform the current simulations on a massively-parallel computing system are twofold: on the one hand, the amount of 48 billion grid points required to resolve even the smallest relevant flow scales together with the immense number of around 1.4 million fully-resolved particles per simulation require a sufficiently high number of cores that is able to execute this work in acceptable wall-clock time. On the other hand, the slow time scales at which in particular the transverse-oriented sediment




Figure 1: Instantaneous snapshot of the sediment bed in the presence of side walls. The particle colour ranges from dark to light brown with increasing distance to the bottom wall. Mean flow and gravity are directed in positive x - and negative y -direction, respectively. Turbulent flow vortices are visualized by means of the λ_2 -criterion, while pressure fluctuations in the free surface (kept flat in the simulation) are shown as vertical deformation with low (high) values indicated by dark (light) blue colour. At the downstream end of the streamwise periodic domain, a cross-section of the flow field indicates low (high) intensity of the instantaneous streamwise fluid velocity by dark (bright) grey regions. For the sake of visualization, the second side wall in the foreground has been removed.




Figure 2: Topview of the sediment bed height variation for (a) the case without and (b) the case with lateral side walls. Details on the definition of the sediment bed surface can be found in Scherer et al. (2020). The arrow in the upper left corner indicates the mean flow direction.

patterns evolve lead to a pure computing time of several months on supercomputing systems such as SuperMUC-NG in order to capture the entire bedform evolution.

In this project, we perform two individual simulations of turbulent open channel flow over a mobile sediment bed with lateral solid side walls or with lateral periodic boundary conditions, respectively, at intermediate Reynolds number $Re_b = 3,000$ (using the bulk velocity u_b as velocity scale). In terms of the mean fluid height H_f , the streamwise periodic domains have dimensions $[12H_f \times 1.5H_f \times 16H_f]$ in the streamwise x -, wall-normal y -, and spanwise z -direction of a Cartesian coordinate system. The particles the sediment bed consists of have a relative diameter of roughly $D/H_f = 0.04$. For the current configurations, scaling tests of our MPI-parallelized simulation code with different processor arrangements have suggested that a domain decomposition strategy based on a Cartesian processor grid with $48 \times 1 \times 64 = 3,072$ cores in the streamwise, wall-normal and spanwise direction is the most economic way to use the granted computing time. Note that by choosing only one single processor in the wall-normal direction, the load imbalance due to a varying number of particles per core can be reduced to a minimum. For checkpointing and data analysis, snapshots of the flow field and the corresponding state of all particles are saved every 24 hours of simulation run time, each of which has a size of around 160 gigabytes. Until the end of the project, this procedure will lead to a total file size of roughly 45 terabytes for the two simulations.

Data analysis of the obtained simulation results reveals two different evolution stages of pattern formation: shortly after the onset of the simulation, a group of essentially parallel quasi-streamwise sediment ridges can be identified over the entire width of the sediment bed, surrounded by counter-rotating turbulent secondary flow cells of Prandtl's second kind. These flow structures naturally appear in flows with non-circular domain cross-section such as rectangular ducts. However, in the presence of mobile sediment, it has been observed that secondary flow cells can also evolve as a result of an instability of the turbulent

channel flow with respect to spanwise inhomogeneities of the sediment bed. To the best of the authors' knowledge, the current simulations are the first of their kind that are able to show that, indeed, the secondary flow cells can form independently of any side wall influence in both configurations.

With ongoing simulation time, bedforms with transverse oriented crestlines grow and eventually reach crest heights that are twice as high as the streamwise ridges.

Figure 1 provides an instantaneous snapshot of two developed transverse sediment patterns and the overlying flow field. The presence of the sediment patterns clearly modifies the global system of vortical structures in the flow and affects the pressure fluctuations along the upper free-slip plane of the channel. The observations extend the findings gained in the recent works of [2,3], which have exclusively focused on the formation of transverse patterns. The authors have studied the flow in domains significantly narrower than in the current project to intentionally suppress other sediment structures and possible spanwise heterogeneities. In the current project, however, the large channel width allows the study of pattern three-dimensionality: in the course of the simulations, several individual sediment patterns are seen to evolve and coexist for a certain period, before they merge with other bedforms to form more spanwise homogeneous structures. Interestingly, this behaviour is more pronounced in the case without lateral side walls, where the flow as well as the sediment bed are less restricted and thus show stronger spanwise heterogeneities, as can be observed in Figure 2.

Ongoing Research / Outlook

Thanks to the computing time and hardware provided on SuperMUC-NG, we have been able to create a new database of high-fidelity simulations of sediment pattern formation in turbulent open channel flows in the presence and absence of lateral side walls. The current numerical simulations at this level of detail in a similarly long and wide physical domain represent a novelty in the field and allow a unique insight into the physics of bedform evolution, propagation and interaction for a variety of arbitrarily oriented sediment patterns. An aspect which is out of the scope of the current project is the investigation of the Reynolds number dependence of the observed processes. In future studies, simulations at higher Reynolds number that allow for clear scale separation are planned to further scrutinize this point.

References and Links

- [1] A. G. Kidanemariam and M. Uhlmann. Int. J. Multiphase Flow, 67:174–188, 2014. DOI: 10.1016/j.ijmultiphaseflow.2014.08.008.
- [2] A. G. Kidanemariam and M. Uhlmann. J. Fluid Mech., 818:716–743, 2017. DOI: 10.1017/jfm.2017.147.
- [3] M. Scherer, A. Kidanemariam, and M. Uhlmann. J. Fluid Mech., 900:A1, 2020. DOI: 10.1017/jfm.2020.423.
- [4] M. Uhlmann. J. Comp. Phys., 209(2):448–476, 2005. DOI: 10.1016/j.jcp.2005.03.017.

Direct numerical simulation of partially-filled pipe flow

RESEARCH INSTITUTION

Associate Professorship of Hydromechanics, Technical University of Munich

PRINCIPAL INVESTIGATOR

Michael Manhart

RESEARCHERS

Julian Brosda, Yoshiyuki Sakai, Simon von Wenczowski

PROJECT PARTNERS

—

SuperMUC Project ID: pn56ci (KONWIHR project)

Introduction

In this project the flow in partially-filled pipes is investigated. This flow can be seen as a model flow for rivers and waste-water channels and represents a fundamental flow problem that is not yet fully understood. Nevertheless, there have neither been any high-resolution simulations nor well resolved experiments reported in literature to date for this flow configuration. In this project highly resolved 3D-simulations are performed which help further understanding narrow open duct flows. The analysis concentrates on the origin of the mean secondary flow and the role of coherent structures as well as on the time-averaged and instantaneous wall shear stress.

Results and Methods

For the direct numerical simulations within this project, the flow solver MGLET is employed. It uses a Finite Volume method to solve the incompressible Navier-Stokes equations on Cartesian grids with a staggered arrangement of the variables. A local grid refinement is implemented by adding refined grids in a hierarchical, overlapping way. An explicit third-order low-storage Runge-Kutta time step is used for time integration.

Curved surfaces are represented by an Immersed Boundary Method. MGLET is parallelized by a domain decomposition method using Message Passing Interface (MPI).

Recently, the code has been optimized for massively-parallel computing architectures within three succes-

sive KONWIHR projects in 2015, 2017 & 2019, with their outcomes being published in [1,2]. These optimisation works were done in a close collaboration with the experts from the CFDLab at LRZ. In the latest KONWIHR project, we performed a SIMD optimisation to our two pressure solvers. This was motivated by the recent trend that the modern HPC processors are equipped with ever more powerful yet more energy-efficient internal vectorisation hardware. One important example of such systems for us is SuperMUC-NG at LRZ, that is based on Intel Skylake processors being equipped with 512-bit ultrawide vector registers. By exploiting the Skylake's extensive SIMD capability, our optimised code shows up to 20% overall performance improvement proven for up to 3.2×10^4 processes and runs with reasonable efficiency up to $O(10^5)$ MPI processes (see Fig. 1).

A fully-developed turbulent flow was simulated in a straight, partially-filled pipe with various filling ratios and Reynolds numbers ranging from marginally to moderately turbulent, see Table 1.

Periodic boundary conditions were applied in stream-wise direction, no-slip for the side walls and the free-surface for low Froude number was approximated by a slip condition. To achieve converging and independent statistics all simulations were run for at least $5,000 \cdot t_{ub}/R$ within a domain length of 8π , see Fig. 2 for semi-filled pipe flow. At small Reynolds numbers mixing of the flow is not as strong as at larger Reynolds numbers, hence the simulation time was doubled and the length of the domain was enlarged by the factor 1.5.




Figure 1: Weak scaling on SuperMUC-NG. The problem size is 64,000 cells/core over a single grid level.




Figure 2: Flow setup for semi-filled pipe flow.

Pipe Filling grade [%]	25	50	75
$Re_b = u_b \cdot R/\nu$	1,500 – 6,900		
$Re_r = u_r \cdot R/\nu$	110 – 460		

Table 1: Simulation parameters.




Figure 3: Mean streamwise velocity u_1/u_b (left half); streamlines of the mean crossflow velocities u_2 and u_3 , inner vortex (red, clockwise rotation) and outer vortex (blue, anti-clockwise rotation) (right half); $Re_\tau = 230$.




Figure 4: Re_τ -dependent position of the vortex center of the inner (red) and outer (blue) secondary cell.




Figure 5: Re_b dependency of the friction factor λ . Red pluses for present study, blue crosses and diamond show measurements, all for semi-filled pipe flow. Circles are two different measurements for full pipe flow. Black dashed line is the laminar friction relation, black solid line follows Prandtl's friction law for smooth pipes.

A variety of simulations were performed—from 84 cores with $48 \cdot 10^6$ grid points and a time per cell update of about $2 \cdot 10^{-6}$ s for the smallest simulation to the biggest case including 1,056 cores with $750 \cdot 10^6$ grid points and a time per cell update of $\sim 3 \cdot 10^{-6}$ s, which lead to 80 M core-hours and 3,000 M core-hours per case, respectively.

The main results are summarised in the following. The maximum mean streamwise velocity (black cross, Fig. 3) can be found at $\sim 70\%$ of the flow depth. This has already been reported in the literature and can be found in other free surface flows as well. This so-called 'dip phenomenon' can be explained by the mean secondary flow which consists of a counter-rotating two vortex system by which at the free surface slow flow is being convected from the pipe's wall to the center. Independent of the Reynolds number the so-called 'inner secondary cell', close to the free-surface-wall juncture, rotates towards the wall at the free-surface and the 'outer secondary cell' the other way around, see Fig. 3.

The position of the two vortex centers is dependent on the Reynolds number. At higher Re the vortex centers are moving towards the free surface, see Fig. 4. The size of the inner secondary vortex cell (wall-distance of the center) scales with wall units. The distance of the center of the outer secondary cell from the free surface seems to converge for large Reynolds numbers.

Comparing the friction factor λ with measurements and full pipe flow gives surprising results, see Fig. 5. For laminar flow the friction factor for semi-filled pipe flow matches with the full pipe flow. This was expected. For




Figure 6: Instantaneous Coherent structures visualised via the Q-criterion, showing a $4R$ long section of the left channel half. $Re_\tau = 230$.

turbulent flow larger friction factors for semi-filled pipe have been reported in literature. Our results indicate, however, that semi-filled and full pipe flows have nearly identical friction factors.

Furthermore we started to investigate the link between the mean and the instantaneous flow fields, by analysing instantaneous snapshots for spatio-temporal dynamics [3], its coherent structures of different scales and their interaction. For instance, in Fig. 6 the second invariant of the velocity gradient tensor (Q-criterion) displays coherent structures, like quasi-streamwise vortices near the boundary growing towards the bulk flow and vortices being attached to the free-surface.

Currently, a publication is in preparation in which the mean flow and turbulence structures are documented in dependence of the Reynolds number. Furthermore, the generation of the secondary flow is analysed in terms of the balance equation of the mean kinetic energy.

References

- [1] Y. Sakai, S. Mendez, H. Strandenes, M. Ohlerich, I. Pasichnyk, M. Allalen, M. Manhart, Performance Optimisation of the Parallel CFD Code MGLET across Different HPC Platforms, PASC '19 Proceedings of the Platform for Advanced Scientific Computing Conference, Article No. 6, 2019.
- [2] H. Strandenes, M. Manhart, M. Allalen, I. Pasichnyk, W. Schandler. Improving scalability for the CFD software package MGLET. InSiDE - Innovatives Supercomputing in Deutschland (2), 2016, 48 – 50.
- [3] J. Brosda, M. Manhart: Dynamics of secondary currents in marginally turbulent semi-filled pipe flow. 17th European Turbulence Conference, 2019.

Turbulent natural convection of non-Newtonian fluids in enclosed spaces

RESEARCH INSTITUTION

Numerical methods in Aerospace Engineering, Bundeswehr University Munich, Germany

PRINCIPAL INVESTIGATORS

Sahin Yigit, Josef Hasslberger, Markus Klein

RESEARCHERS

—

PROJECT PARTNERS

—

SuperMUC Project ID: pn56di

Introduction

This project focuses on the modelling and physical understanding of three-dimensional turbulent natural convection of non-Newtonian fluids (i.e. where the strain rate dependence of shear stresses is non-linear in nature) in enclosures. This topic has wide relevance in many engineering applications such as preservation of canned foods, polymer and chemical processing, bio-chemical synthesis, solar and nuclear energy, thermal energy storages. Therefore, the flow and heat transfer knowledge of more complex than Newtonian fluids (fluids like water, air where viscous stress is directly proportional to strain rate) is essential from an engineering perspective since the non-Newtonian character of fluids can also be very useful for designing new adaptive thermal management systems.

Methods

Direct Numerical Simulations (DNS) of three-dimensional turbulent Rayleigh-Bénard convection of yield stress fluids obeying a Bingham model and inelastic shear thinning / thickening fluids obeying power-law model in a cubical enclosure have been performed under Dirichlet boundary conditions. The simulation configuration is schematically shown in Fig.1, which demonstrates that the differentially heated horizontal walls are subjected to constant wall temperature boundary conditions. The bottom wall is taken to be at higher temperature than the top wall (i.e. $T_H > T_C$) and all the other walls are considered to be adiabatic (i.e. the temperature gradient in the wall normal direction vanishes at the wall). No-slip and impermeability conditions are specified for all walls. The enclosure is taken to be cubic (i.e. $H = W = L$).




Figure 1: Schematic diagram of the simulation domain.

To solve the non-linear set of governing equations in a finite-volume framework, the open-source CFD package OpenFOAM has been utilised. The pressure-velocity coupling has been addressed by the use of PIMPLE algorithm. Convective and diffusive fluxes are evaluated by second-order centered difference schemes. Temporal advancement has been achieved by the second-order Crank-Nicolson scheme with constant time-stepping. It has been ensured that the Courant number is always sufficiently below unity so that the underlying physics is captured with sufficient temporal resolution.

Results

The distributions of non-dimensional temperature iso-surfaces and apparently unyielded regions (AURs are shown as grey regions) are shown in Fig. 2 for different Bingham




Figure 2: a) Distributions of iso-surfaces of non-dimensional temperature, b) three-dimensional AURs for different values of Bn at $Ra=10^8$ and $Pr=320$ [1].




Figure 3: Temporal variation of \bar{Nu} with non-dimensional simulation time for hot (i.e. Nu_{hw}) and cold (i.e. Nu_{cw}) wall for different Bn at $Ra=10^8$ and $Pr=320$. The time averaged mean Nusselt number is shown with a blue dashed line [1].

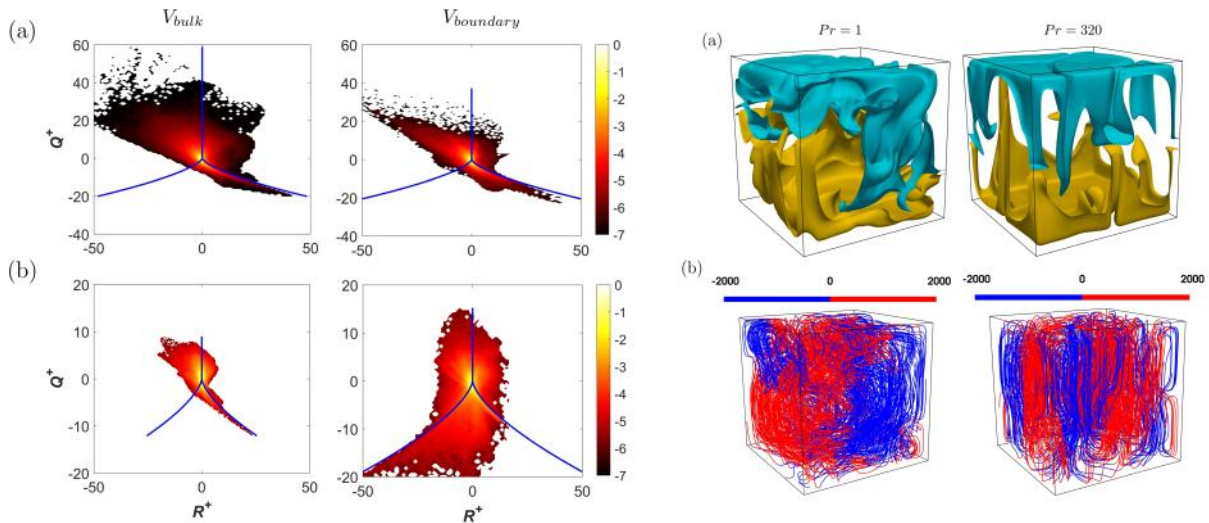


Figure 4: Left column: Contours of joint probability density functions of Q^* and R^* invariants on a logarithmic scale in the bulk region and at the heated and cooled boundaries for $Ra=10^8$ in the case of a) $Pr=1$, b) $Pr=320$ [2]. Right column: a) Non-dimensional temperature iso-surfaces $=0.4$, $=0.6$, shown in blue and yellow colour, respectively. b) Path lines coloured by non-dimensional vertical velocity magnitude for different Pr at $Ra=10^7$ [2].

number (i.e. Bn) at $Ra=10^8$, and $Pr=320$. It can be seen from Fig. 2a that the iso-surfaces become less deformed and thermal plumes from the bottom wall become less apparent as Bn increases. These are the indications of the relative weakening of buoyancy force due to the augmented flow resistance as a result of increased yield stress. This also leads to the weakening of convection with increasing Bn , which can be substantiated from Fig. 2b where the size and the probability of finding AURs increases with increasing Bn , and this gives rise to a reduction in the mean Nusselt number (see Fig. 3). For very large values of Bingham number, the boundary layer thickness becomes of the order of the enclosure size H . At that stage, the fluid flow does not influence the thermal transport and heat transfer takes place principally due to conduction, which is reflected in the unity mean Nusselt number. It is worth noting that in the context of bi-viscosity regularisation, the flow does not stop in a true sense in AURs but the flow within the islands of AURs is too weak to influence the heat transfer rate and thus the exact shape and size of AURs do not affect the heat transfer rate and the mean Nusselt number.

Ongoing Research / Outlook

After obtaining experience with yield stress fluids, the next step was the investigation of Prandtl number (Pr) effects near active walls on the velocity gradient and flow topologies. This is because liquids like non-Newtonian fluids typically have high Pr (i.e. from 10^2 to 10^3) [1]. This is important as the modelling strategy for turbulent natural convection of gaseous fluids may not be equally well suited for the simulations of turbulent natural convection of liquids with high values of Pr .

As it can be seen in Fig.4 (left column) the expected and well-known teardrop shape of the joint PDF of velocity gradient tensor invariants Q^* and R^* changes near active walls depending on Pr . This can be explained by looking at Fig.4 (right column): isolated plumes drive the convection process in the $Pr > 1$ case, whereas frequent roll ups in the $Pr = 1$ are indicative of a large-scale circulation [2].

References and Links

- [1] S. Yigit, J. Hasslberger, N. Chakraborty and M. Klein, Effects of Rayleigh-Bénard convection on spectra of viscoplastic fluids, Int. J. Heat and Mass Transfer, 147 (2020), 118947.
- [2] S. Yigit, J. Hasslberger, M. Klein and N. Chakraborty, Near wall Prandtl number effects on velocity gradient invariants and flow topologies in turbulent Rayleigh Bénard convection, Scientific Reports, 10, 14887, 2020.

Investigation of the Flow Field and Inflow of Hovering Rotors over Inclined Ground Planes

RESEARCH INSTITUTION

Institute of Helicopter Technology, Technical University of Munich

PRINCIPAL INVESTIGATOR

Stefan Platzer

RESEARCHERS

Jürgen Rauleder, Manfred Hajek

PROJECT PARTNERS

—

SuperMUC Project ID: pn56lu

Introduction

Rotorcraft regularly hover over moving surfaces (e.g., ship decks) or inclined ground planes (e.g., hillsides and mountains). However, previous research activities on hover in ground effect mostly focused on parallel ground planes. Here, performance benefits and the associated flow topology are relatively well known. In contrast, only a very limited amount of work has been done on non-parallel ground planes and the associated changes of the complex three-dimensional flow field. Furthermore, pilots require a significant amount of training to mitigate the risk in helicopter operations on ships or at hillsides/mountains. Hence, advanced flight control systems are desirable to increase safety and to reduce pilot workload in these situations. Moreover, engineers need to be able to accurately simulate the dynamic response of rotorcraft in such flight conditions when designing control systems for unmanned/autonomous aerial systems. Therefore, there is a need for accurate, computationally efficient mathematical models capable of simulating the rotor inflow in these flow environments, because of the associated

changes of the flight dynamics. Dynamic inflow-type models could fill this role. Models have been developed in the past to account for influences of ground effect and have successfully been validated for parallel ground effect. However, the lack of experimental and high-fidelity numerical data for inclined and dynamic ground effect conditions limited the ability for validation. Moreover, a prerequisite for developing new models is an understanding of the fluid dynamics of the problem, which does not fully exist to date.

The goal of the research project at the Institute of Helicopter Technology [1] was to investigate the fluid mechanics of a hovering rotor over non-parallel ground planes (e.g., hillsides). This can be regarded as a first step towards a broader understanding of the flow field of an arbitrarily moving ground plane (e.g., a moving ship deck). Results were correlated to experimental data to evaluate the accuracy of the numerical simulations [2,3]. The simulations were performed with a model scale, two-bladed teetering rotor with a rotor radius of 0.408m, at a rotational frequency of 35Hz, which resulted in a blade tip speed of 89.7m/s. The rotor plane was located at a height of one rotor radius above the ground plane. The latter had a radius of four rotor radii.




Figure 1: Sectional cut through the computational grid for in ground effect simulations with a 15° inclined ground plane. The refined grid in the tip vortex path is indicated by darker regions in the cut plane. The rotor blade is shown in blue. The chimera boundary of an additional refinement block around the blade is shown in green.

Results and Methods

The simulations were made using the computational fluid dynamics (CFD) solver TAU, developed and maintained by the German Aerospace Center (DLR). In TAU the compressible, unsteady Reynolds-averaged Navier–Stokes equations were solved on edge based dual grids. The temporal and spatial discretization was second-order accurate. Parallelization of the simulation was achieved by domain decomposition and the message passing concept using MPI.

In general, for in ground effect simulations the tip vortex needs to be resolved for multiple rotor revolutions to accurately account for the deflection of the tip vortex trajectory




Figure 2: Out-of-plane vorticity contours illustrating changes of the flow field of rotors in hover in (non-) parallel ground effect. The simulations were run with the Spalart-Allmaras turbulence model with rotation/curvature correction.

caused by the presence of the ground plane. Therefore, in the course of the project grid sizes up to 184 million cells were used on up to 2,400 cores. An illustration of the grid is shown in Figure 1.

To account for turbulence effects a variety of different turbulence models was tested, ranging from a one-equation model, over (non-linear) two-equation models, to full Reynolds stress models. No turbulence model is universally applicable and no model can be used for all types of flows to reliably predict the quantities of interest. Therefore, the project also allowed to assess turbulence model performance during the large deflection of the vortex dominated flow close to the ground plane.

In addition to the large grids, multiple rotor revolutions needed to be computed to convect the starting vortex away from the rotor plane and off the ground plane. This required 18 rotor revolutions for the parallel ground plane, with a maximum of 14 hours per revolution for the full Reynolds-stress model. To be able to extract the temporal evolution of the flow field, up to 20 TB of data needed to be stored per simulation on the LRZ hosted file systems. These extensive simula-




Figure 3: Visualization of the tip vortex trajectory at a ground plane inclination angle of 15°, by showing a Q-criterion iso-surface ($Q = 2,500 \text{ s}^{-2}$) colored by vorticity magnitude. The simulation was run with the Menter SST turbulence model including the SAS modification and rotation/curvature correction.

tions can only be performed on high performance computing systems such as SuperMUC-NG due to the required computing power, run times and storage requirements. The response of the tip vortex trajectory to (non-)parallel ground effect is illustrated in Figure 2. The overall flow field is shown in Figure 3. Upon examining the flow field beneath the rotor and the inflow distribution, the simulations revealed that the behavior of the vortex dominated flow was significantly altered when inclining the ground plane. When hovering in parallel ground effect, the rotor wake was symmetric. After inclining the ground plane, this symmetry was lost. The influence was more pronounced on the uphill side, when compared to parallel ground effect. In particular the stagnation point at the ground plane was shifted far uphill when inclining the ground plane. The inflow was mostly affected on the uphill side and at the inboard sections, whereas on the downhill side this influence was little. Moreover, the comparison to the experimental data showed good agreement of the simulations in terms of flow phenomenology, vortex trajectory, and overall vortex strength. In general the correlation between the simulation and the experiment deteriorated for the inclined ground plane. The only exception was the full Reynolds stress model, which showed comparable performance. The rotor trim was insensitive to the chosen turbulence model, despite significant differences in the ability to preserve the tip vortices.

References and Links

- [1] www.lrg.tum.de/en/ht
- [2] Milluzzo, J., Martinez, A., Drayton, S. and Davids, S.: "PIV and Performance Measurements on Rotors Hovering Above Inclined Ground Planes", American Helicopter Society 74th Annual Forum, Phoenix, AZ, 2018.
- [3] Platzer, S., Milluzzo, J., Rauleder, J.: Investigation on Hovering Rotors over Inclined Ground Planes – a Computational and Experimental Study, 44th European Rotorcraft Forum, Delft, The Netherlands, 2018.

Interface instability dynamics in secondary

drop breakup

RESEARCH INSTITUTION

Chair of Aerodynamics and Fluid, Department of Mechanical Engineering, Technical University of Munich

PRINCIPAL INVESTIGATORS

Nikolaus A. Adams

RESEARCHERS

Stefan Adami, Jakob Kaiser, Josef Winter

PROJECT PARTNERS

—

SuperMUC Project ID: pn56re (KONWIHR project)

Introduction

The breakup of a coherent liquid phase into smaller fragments is paramount to many technical applications, e.g. liquid fuel combustion or biomedicine. Two main stages are generally distinguished: primary and secondary breakup. Primary breakup thereby denotes the fragmentation of a liquid sheet or jet into drops, and secondary breakup the following fragmentation of these drops into smaller droplet.

For secondary breakup, five major regimes have classically been acknowledged: vibrational, bag, multi-mode, sheet-stripping, and catastrophic breakup. Yet, a more recent re-classification bases the occurring breakup modes on the dominating instability mechanisms at the interface of the liquid phase and the surrounding gas phase. In the Rayleigh-Taylor piercing (RTP) regime, Rayleigh-Taylor instabilities grow along the interface, resulting in the well-known bag-like deformation pattern. In the shear-induced entrainment (SIE) regime, Kelvin-Helmholtz instabilities supersede the Rayleigh-Taylor instabilities, resulting in a shearing of material from the droplet rim. This sheet later breaks up in the wake of the deformed drop. The dominating instability type varies with the ratio of inertial, viscous, and capillary forces acting on the drop. This motivates the classification of the breakup mode based on two dimensionless numbers: The Weber number We (ratio of inertial to capillary forces) and the Ohnseorge number Oh (ratio of viscous to capillary forces).

The experimental investigation of such interfacial instabilities is limited by small spatial and temporal scales. This motivates our numerical simulations. More specific, we investigate within this project the transition from RTP to SIE and the interaction with the surrounding gas flow. The work is attached to the ERC Advanced Grant No. 667483 “NANOSHOCK” [1].

Methods and Results

We have performed high-fidelity simulations of secondary drop breakup using our massively parallelized multiresolution framework ALPACA [1]. The code (written in C++20) applies a finite volume discretization

with WENO-based flux reconstruction in characteristic space. The solution is evolved in time using a strongly stable third order Runge-Kutta scheme. For multi-phase-modeling, we apply a level-set based sharp-interface method with conservative interface interaction. Explicit interface-exchange terms are proposed which include inviscid, viscous, and capillary effects at the phase interface. To improve computational efficiency, we apply a block-based multiresolution scheme for spatial adaptation. Temporal adaptivity is provided by an adaptive local time-stepping scheme, which was developed in the course of our work [2]. Our setup is based on a typical experimental setup for secondary breakup: a shock wave passes a liquid drop, which is subsequently exposed to a flow field. The breakup mode can be controlled by the physical parameter of the drop and the shock wave.

The late-stage deformation of a water cylinder is shown in Fig. 1 for three different Weber numbers. For small surface tension forces (left), the droplet breakup occurs in the shear-induced entrainment regime. Shear instabilities grow near the droplet equator which later form a sheet which is characteristic for this breakup mode. For large surface-tension forces (right), the droplet breakup occurs in the Rayleigh-Taylor piercing regime. The characteristic bag shape results from a pressure gradient between the upstream and downstream side of the droplet. The figure in the middle shows the transition between SIE and RTP. Our results allow to investigate the interplay between the dominating instability mechanisms leading to the different breakup modes.

An important aspect of our work is the analysis of the interaction between the phase interface and local flow field patterns. The pressure field in the vicinity of the droplet in the shear-induced entrainment regime is shown in Fig. 2. In our work [4], we found that during the early stages of the interface deformation, local pressure waves appear. These pressure waves are connected to local interface disturbances. These disturbances grow and later form the water sheets characteristic for this breakup mode.

For our results, we have performed multiple two- and




Figure 1: Shock-induced breakup of a water cylinder for $We=38,939$ (left), $We=16$ (middle) and $We=10$ (right). The drop is shown in white, the interface is highlighted in black. The upper half depicts numerical schlieren-type visualization and the lower half is colored by the streamwise velocity (normalized with the post-shock gas velocity). The simulation setup and further details can be found in [3].




Figure 2: Near-droplet pressure field during the early stage interface deformation [4]. The black arrows indicate pressure waves and corresponding interface disturbances.

three-dimensional simulations. Two-dimensional simulations have typically approximately 30 Mio. degrees-of-freedom and require simulation times of 4,000 Core-h on SuperMUC-NG. For three-dimensional simulations, 50 Billion degrees of freedom are necessary. This approximately requires 900,000 Core-h. Post-processing of the results requires high storage capabilities and was partly performed on SuperMUC-NG and the remote visualization cluster.

Ongoing Research / Outlook

We currently perform additional high-resolution three-dimensional simulations of RTP and SIE breakup modes. Our extended setup is aligned to experimental investigations performed at the same chair. In future, this allows us for matching experimental and numerical setups and a detailed investigation of the occurring flow field patterns.

References and Links

- [1] <https://www.nanoshock.de>
- [2] Kaiser, J. W. J. et al., JCP X 4, 1000038 (2019).
- [3] Winter, J. M. et al., Numerical investigation of 3D drop-breakup mechanisms using a sharp interface level-set method. In: Proceedings of the Eleventh International Symposium on Turbulence and Shear Flow Phenomena, TSFP 2019.
- [4] Kaiser, J. W. J. et al., Int. J. Multiph. Flow 132, 1003409 (2020).
- [2] A. Avdonin, M. Meindl, W. Polifke, 37th Int'l Symposium on Combustion (2018).

Scale-Resolving Study of a Fluid-Structure-

Acoustics Interaction Test Case

RESEARCH INSTITUTION

¹Institute for Parallel and Distributed Systems, University of Stuttgart

PRINCIPAL INVESTIGATOR

Miriam Mehl¹

RESEARCHERS

Neda Ebrahimi Pour², Amin Totounferoush¹

PROJECT PARTNERS

²University of Siegen

SuperMUC Project ID: pn56se

Introduction

The reduction of noise in early design processes of engineering devices has become more important, especially due to stricter noise regulations. The simulation of fluid-structure-acoustics is challenging from the computational perspective, as it involves different scales, which require specific numerical treatment. The target of this project is to enable these complex simulations where aero-acoustics arises in the vicinity of structures that interact with a surrounding fluid flow. To realize such simulations, a partitioned coupling approach is utilized. This strategy decomposes the large simulation domain into several subdomains, that interact at their respective coupling interfaces by exchanging boundary conditions. Each subdomain solves a different set of equations and different spatial resolution. A coupling tool is deployed, that is responsible for the communication and data-mapping between the respective subdomains. To solve the equations in the subdomains, we utilize a high-order scheme that requires an appropriate representation of the geometry as well. First investigations with a non-moving geometry have been conducted with our previous compute budget on SuperMUC Phase2 [3], where we were able to prove the accuracy of the solution and the efficiency of the partitioned coupling approach, which reduces the computational effort in comparison to a single domain that solves the same equations with the same scheme order. Further investigations with the current project are devoted to moving geometries, that are common in many engineering devices. Hence, our final goal is the 3-field coupled flow simulation with a moving object.

Results and Methods

The high-order Discontinuous Galerkin solver Ateles [2] is used for the simulations. The solver is part of the simulation framework APES [1], that provides pre- and post-processing tools. A Discontinuous Galerkin (DG) scheme is used to solve the partial differential equations. A major benefit of this method is the high locality, where information within an element is tightly coupled and the communication is limited to the element faces. It is highly parallel and can deploy two levels of parallelization: OpenMP shared-memory parallelism for operations inside the elements and MPI for the communication between the elements at partition boundaries. The moving objects in the flow domain are represented in the solver with the same discretization method as the flow domain itself. They are modelled as a porous material based on the Brinkman penalization technique (immersed boundary method). This method allows to apply a computational mesh with cubic elements and has the major advantage for moving geometries that no re-meshing of the computational mesh is necessary. Thus, the computational mesh is generated once and remains unchanged during the entire simulation. More on the method modeling the geometry can be found in [4].

Due to the decomposition of the overall problem size and the different numerical treatment of each subdomain, load imbalances have to be actively addressed. In Figure 1 we have addressed this issue using inter- and intra-subdomain load balancing for a 3-field coupled simulation, where in the innermost subdomain an airfoil is located and the compressible Navier-Stokes





Figure 1: 3-field coupled simulation strong scaling measurement with up to 640 nodes on SupermucNG. (left) Inter-subdomain load balancing (right) inter- and intra-subdomain load balancing. Reduced idling time of processing for case (right).

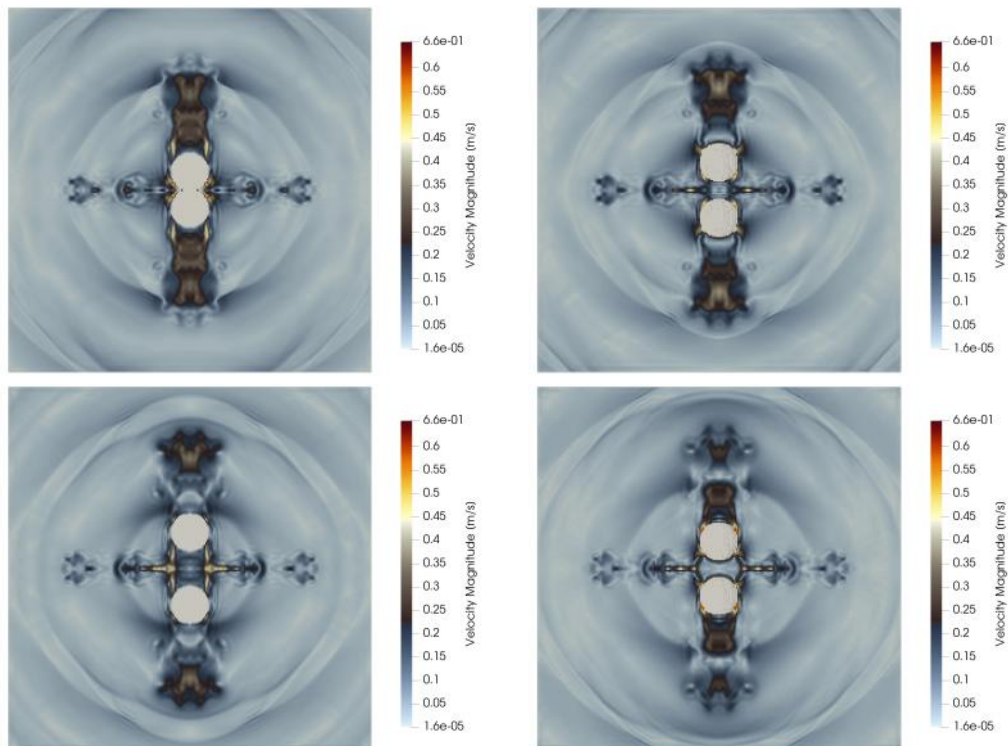


Figure 2: Simulation results of two spheres moving away and towards each other. Numerical solutions are shown after (top, left) 9.0 s, (top, right) 9.25 s, (bottom, left) 9.5 s and (bottom, right) 9.75 s simulation time.

equations solved. In the middle domain the inviscid Euler equations and in the outermost domain the linearized Euler equations are solved, respectively [3]. Inside the subdomains (intra-subdomain) the SPART algorithm is used to re-balance the workload, while the total workload per domain is computed and used to re-balance the workload between the subdomains (inter-subdomain). To capture the workload MPI timers are used, that are placed around compute intense routines. In Figure 1 (left) we only balance the load between the subdomains. In Figure 1 (right) the load is balanced inside and between the subdomains. We can clearly observe, that load balancing only between the domains (cf. Figure 1, left) involves idling time in each subdomain (vertical gap between the lines) that is an obstacle for the overall scaling. In Figure 1 (right) the idling time is reduced (vertical gap is smaller) and we can achieve an improved scaling behavior as both inter- and intra-subdomain load imbalance are addressed. Alternatively, the inter-subdomain load balancing can be addressed by using the data-based approach introduced in [5].

First results of our investigation with moving geometries are shown in Figure 2. The test case involves two moving spheres, that are located initially at their respective interfaces. During the simulation they move with a predefined cosine function ($Y(t) = 0.1 \cos 2\pi t$) away from each other. Both spheres have a diameter of 0.4 m. The simulation domain has a height and length of 4 m and a width of 1 m. The computational mesh is discretized with 65,536 elements. The compressible Euler equations are solved and discretized with a polynomial degree of seven. The pressure is set to 1 bar, the density is 1 kg/m^3 and the velocity is initially in all spatial directions 0.0 m/s. The initial position of the upper sphere is $S_1(0.0, 0.2, 0.0)$ and of the lower sphere $S_2(0.0, -0.2, 0.0)$. In Figure 2 (top, left), both spheres have reached their initial position, where both get in touch at their respective interfaces.

The fluid is compressed in that area. In Figure 2 (top, right), both spheres move again in opposite directions, to their maximum position, which they finally reach in Figure 2 (bottom, left). From that point on again both spheres move towards each other. This test case demonstrated, that the geometry can be modelled appropriately and the numerical scheme is able to handle the movement of the geometries, using a high-order scheme. Further, no stability issues were encountered, even though the geometry interfaces reach each other and the fluid is compressed in that region, resulting in very low velocity but very high pressure and density values.

Ongoing Research / Outlook

Ongoing research is devoted to further investigation of the moving geometry and its validation as well as its efficient computation. The final goal is the realization of a rigid body movement induced aero-acoustics noise problem. Where again the large simulation domain is split into subdomains and coupled together through a coupling approach. This setup allows for detailed simulations of noise generation at moving obstacles and its propagation over long distances, as needed in typical aero-acoustic problems.

References and Links

- [1] <https://apes.osdn.io>
- [2] <https://apes.osdn.io/pages/ateles>
- [3] <https://www.gauss-centre.eu/results/computational-and-scientific-engineering/article/investigating-coupled-fluid-structure-acoustic-fsa-interaction/>
- [4] N. Anand, N. Ebrahimi Pour, H. Klimach, and S. Roller. *Symmetry* 11, no. 9 (2019): 1126. <https://doi.org/10.3390/sym11091126>
- [5] A. Totounferoush, N. Ebrahimi Pour, J. Schröder, S. Roller, and M. Mehl. Submitted to *J. Comp. Sci.*
- [6] Bungartz, H. J., Lindner, F., Gatzhammer, B., Mehl, M., Scheufele, K., Shukaev, A., & Uekermann, B. (2016). *Computers & Fluids*, 141, 250-258.

Thermodiffusive instabilities in lean premixed hydrogen/air flames

RESEARCH INSTITUTION

Institute for Combustion Technology, RWTH Aachen University, Germany

PRINCIPAL INVESTIGATOR

Heinz Pitsch

RESEARCHERS

Lukas Berger

PROJECT PARTNERS

—

SuperMUC Project ID: pn56vo (Gauss Large Scale project)

Introduction

The recent rise of electricity production from renewable energy sources has increased the need for appropriate energy storage in times of excess production. Among different technologies, the concept of storing excess electricity in chemical energy is found to be of significant importance as described in the vision 'Energiekonzept 2050' of the German government [2]. In particular, the interest in using hydrogen as a clean energy carrier has increased as highlighted by the "Nationale Wasserstoffstrategie" of the German government [3]. Hydrogen can be generated by electrolysis of water and if needed, the energy can be regained by an electrification of hydrogen via thermochemical energy conversion processes like combustion. Besides the expected increased use of batteries and fuel cells, combustion will remain a critical pillar of our energy system, e.g. in the fields of heavy duty vehicles, aircrafts, the generation of heat for industry processes, and as hybrid solutions in combination with batteries and fuel cells [4].

One possibility to integrate the generated hydrogen into the existing energy infrastructure is the enrichment of conventional fuels such as natural gas by hydrogen. Such fuels are termed High Hydrogen Content (HHC) fuels. However, a variety of problems arises for combustion processes that involve hydrogen and HHC fuels, e.g., lean hydrogen/air flames are prone to combustion instabilities, which can substantially change flame dynamics and heat release rates. In particular, the instabilities lead to higher flame speeds, which is highly relevant for safety issues such as flash back. The combustion instabilities originate from the low Lewis number of hydrogen, which induces strong differential diffusion effects along the flame front, and are termed thermodiffusive instabilities. However, this instability mechanism is yet not sufficiently well understood and before HHC combustion becomes applicable in actual combustion engines a deeper understanding of its combustion behavior is required.

The development process of the next generation

combustion devices that can handle HHC fuels typically involves the use of simulations, e.g. Large Eddy Simulations (LES) to cut the high cost associated with experimental tests. However, high-fidelity, predictive simulations of hydrogen and HHC fuels combustion require LES models that accurately and reliably model thermodiffusive instabilities. Present-day combustion models developed for hydrocarbon fuels, which are unaffected by such instabilities, cannot capture these effects.

A main impediment for the development of new LES models is the lack of data. For three-dimensional thermodiffusively unstable flames, only limited data exist; the available studies only qualitatively highlight the dynamics of these flames, but do not provide enough detail for quantitative model development. The high level of detail and the availability of all desired quantities at all locations motivates the use of Direct Numerical Simulations (DNS) for model development. For example, DNS provide information about reaction rates or higher order moments, that are challenging to obtain experimentally. Recently, advances in supercomputing have enabled a range of interesting DNS studies and made DNS a powerful tool in combustion science. Within this project, a series of DNS of turbulent




Figure 1: Temperature iso-contour colored by the heat release rate of a DNS of a turbulent lean premixed hydrogen jet flame in a slot burner configuration.




Figure 2: Temperature distribution of a lean premixed hydrogen jet flame using i) non-unity Lewis numbers (left) and ii) unity-Lewis numbers (right).

premixed lean hydrogen flames is conducted to generate a unique database of thermo-diffusively unstable flames. From this database, LES models that account for the impact of thermodiffusive instabilities on turbulent flames can be rigorously developed.

Results and Methods

Numerical Framework

The governing equations of the DNS are given by the reacting Navier-Stokes equations in the low-Mach limit. For the computation, an in-house code called CIAO is employed. The code is a high-order, semi-implicit finite difference code that uses Crank-Nicolson time advancement and an iterative predictor corrector scheme. Spatial and temporal staggering is used to increase the accuracy of stencils. The Poisson equation for the pressure is solved by the multi-grid HYPRE solver. Momentum equations are spatially discretized with a second order scheme. Species and temperature equations are discretized with a fifth order WENO scheme. The temperature and species equations are advanced by utilizing an operator splitting according to Strang. The chemistry operator uses a time-implicit backward difference method, as implemented in the stiff ODE solver CVODE. For further details about the applied numerical algorithms and code verification, the reader is referred to Ref. [5]. The code uses the message passing interface (MPI) standard.

DNS of Thermodiffusively Unstable H_2 -Flames

Fig. 1 shows the DNS of a three-dimensional turbulent premixed lean hydrogen slot burner flame. In between the two blue colored walls, an unburned lean hydrogen/air mixture at $T_u = 298$ K, $p = 1$ bar, and $\phi = 0.4$ is injected in streamwise (x) direction into the simulation domain. In crosswise (y) direction, the main jet is shielded by a left and a right coflow, which consist of a burned hydrogen/air mixture at the same thermochemical conditions as the main jet. To generate a turbulent shear layer, the coflows are injected at a lower velocity compared to the main jet. The spanwise (z) direction is periodic, which allows to efficiently collect flame statistics for the analyses. The

flame in Fig. 1 is represented by a temperature iso-contour that is colored by the local heat release. Due to the interactions of turbulence and the thermodiffusive instability mechanism, the heat release rate shows strong variations along the iso-contour.

The impact of the thermodiffusive instabilities on the combustion process can be particularly seen in Fig. 2. In the left figure, an instantaneous temperature distribution of a DNS that involves realistic molecular transport such as differential diffusion is shown and is termed non-unity Lewis number case. In the right figure, another DNS, in which molecular transport is simplified such that the thermodiffusive instability mechanism is suppressed, is shown and is referred to as unity Lewis number case. Thus, a comparison of both simulations allows for a rigorous assessment of the impact of the thermodiffusive instability mechanism. In particular, the non-unity Lewis number case shows strong temperature overshoots behind the flame front which do not exist in the unity Lewis number case. Such “hot spots” are particularly relevant for NO_x pollutant formation and require further investigation. Furthermore, in contrast to the unity Lewis number case, extinction and areas of strong heat release are found in the non-unity Lewis number case (not shown here), which strongly affect the local propagation speed of the flames.

The conducted DNS clearly show the significant impact of thermodiffusive instabilities on turbulent flames. The datasets that have been generated within this project form a unique database for studying the interactions of turbulence and thermodiffusive instabilities in detail and will guide the development of combustion models for predictive LES of such flames.

References and Links

- [1] <https://www.lrz.de/projekte/hlrb-projects/000000000F43BA1.html>
- [2] <https://archiv.bundesregierung.de/resource/blob/656922/779770/794fd0c40425acd7f46afacbe62600f6/energiekonzept-final-data.pdf>
- [3] <https://www.bmbf.de/de/nationale-wasserstoffstrategie-9916.html>
- [4] <https://www.dena.de/newsroom/publikationsdetailansicht/pub/dena-factsheets-powerfuels>
- [5] O. Desjardins, G. Blanquart, G. Balarac, H. Pitsch, J. Comput. Phys., 227:7125-7159, 2008.

VIVALdI—High fidelity simulation of vortex induced vibrations for flow control and energy harvesting

RESEARCH INSTITUTION

¹Barcelona Supercomputing Center, Spain

PRINCIPAL INVESTIGATOR

Oriol Lehmkühl¹

RESEARCHERS

Juan Carlos Cajas¹, Daniel Pastrana¹, Guillaume Houzeaux¹, Ivette Rodriguez²

PROJECT PARTNERS

²Universitat Politècnica de Catalunya, Barcelona

SuperMUC Project ID: pn69fa (Prace project)

Introduction

Geophysical flows from wind to oceanic currents represent a clean source of energy widely available. Structures based on vortex-induced vibrations are one of the mechanisms for harvesting part of this energy in the range of frequencies where flow-induced vibrations originate a strong coupling between the oscillating body and the fluctuating wake. In the present project, high fidelity simulations of a cylindrical body oscillating in a free-stream from sub-critical to super-critical Reynolds numbers will be carried out for first time by means of wall-resolved LES.

The project aims at shedding light in the interactions of the cylinder with the fluid and specially with the boundary layer, but also the characterisation of the wake topology. The project focus on study the influence of the Reynolds number and the range of frequencies in which vortex induced vibrations reinforce the aerodynamic forces on the cylinder. In addition to this, for the cases corresponding to the critical and supercritical Reynolds numbers were the wake width is reduced as a consequence of the instabilities in the boundary layer, localised roughness or a trip wire is added to the cylinder surface in the location where boundary layer is forced to separate from the cylinder so as to passively control flow separation and transition to turbulence to widen the wake.

The main idea behind the use of these passive devices is to maximise the ratio of the amplitude of oscillations to the cylinder diameter (A/D). This is the first time this kind of simulations are performed at this




Figure 1: Non-dimensional crossflow and in-line maximum amplitudes of oscillation as a function of the reduced velocity. Present LES results at different Reynolds numbers compared with literature data.

level of modelisation, being a step forward in the understanding of the physics of fluid-structure interaction in the range of industrial applications.

Results and Methods

In the present project the main characteristics of the wake of 2-DOF freely vibrating cylinders have been described using results from large-eddy simulations of the flow in the range of Reynolds numbers of $Re = 3,900-7.5 \times 10^5$, and the supercritical regimes for first time.

The behaviour of 2-DOF vibrating cylinders remains almost the same along the whole subcritical regime ($Re = 3,900-10^5$) and the three-branches of response of this system is well captured by the LES. In figure 1, the maximum crossflow and in-line non-dimensional







Figure 2: Instantaneous flow in the subcritical regime in the super-upper branch at $U^* = 5.5$. $Re = 5,300$ (left) and $Re = 10^5$ (right).

amplitudes of vibration for different Reynolds numbers up to $Re = 1,000–10^5$, corresponding to the subcritical regime, are plotted. In the figure, the experimental results of Jauvtis & Williamson for a low-mass ratio cylinder (2.6) and in the range of $Re = 11,000–10^5$ are given. In addition, the numerical results by Gsell et al. at a Reynolds number of $Re = 3,900$ are also included. At these Reynolds numbers, the low-mass ratio 2-DOF cylinder response is characterised by three branches: the initial (*I*), the super-upper (*SU*) and the lower (*L*) branches. Each of these branches has different characteristics and amplitude of response.

Figure 3: The super-critical wake of a 2-DOF cylinder at $Re = 7.5 \cdot 10^5$ (right) and non-dimensional cross-flow and in-line amplitudes as a function of the reduced velocity in the super-critical regime.



The *I* branch, where the synchronisation or lock-in begins, has very small in-line vibrations (see figure 1) and the wake is characterised by the typical von Karman street vortex shedding mode or 2S vortex formation mode (following the nomenclature of Williamson & Roshko), i.e. 2 vortices shed each cycle. This wake configuration can be seen in figure 2 (top), where ensemble average vorticity contours at $U^* = 3.0$ are depicted. Actually, in the 2S mode, and in spite of the crossflow fluctuation of the cylinder, the vortices shed are in-line with the wake centreline in a similar fashion to the static cylinder as can be seen in the figure.

Of particular interest is the *SU* branch, where in-line and crossflow amplitudes of fluctuation peak. This branch is not present in 1-DOF systems and is characteristic only of 2-DOF with low-mass ratio $m^* < 6$. In this branch, the wake is broader and vortex shedding mode changes from a 2S pattern to a 2T one, i.e. 2 triplets of vortices shed each shedding cycle. This mode was first described by Williamson & Roshko. Actually, the footprint of two out of three vortices shed once the cylinder reaches its maximum elongation in the cross-flow direction is observed in our simulations. However, as the flow moves downstream, the triplets of vortices shed each cycle merge into a single vortex located off the wake centreline. Finally, the *L* branch, with a low level of fluctuation in the in-line direction and with decreasing crossflow amplitudes as the reduced velocity increases up until the de-synchronisation of the system around $U^* = 12$ occurs. In this branch the vortex shedding mode is different than in the other two branches and is characterised by two pair of vortices shed each cycle (also known as 2P mode).

Last but not the least, attention should be also paid to the fact that along the whole subcritical regime, cylinder amplitudes are only function of the reduced velocity and independent on the Reynolds number (see figure 1). In fact, the cylinder in-line and cross-flow positions at each instant of the cycle (not shown here) are almost the same regardless of the Reynolds number, which points out to the similarities with the static cylinder in which along the whole subcritical regime the wake characteristics are nearly the same. Actually, if the instantaneous flow at $Re = 5,300$ and $Re = 10^5$ for $U^* = 5.5$ are compared at the maximum elongation position, one can see the same large-scale structures of the flow and wake configuration for both Reynolds numbers (see figure 2).

Ongoing Research / Outlook

Future work will focus on the POD and DMD of the obtained data bases in order to characterize the role Gortler-like structures on the de-synchronisation of the cylinder. Additionally, effects of roughness over the super-critical behavior are also being studied.

As a final note we were awarded the best PRACE Industrial Presentation at PRACEDays19 during the EuroHPC Summit Week 2019. And the project has been highlighted in the ninth Issue of the PRACE Digest.

References and Links

- [1] <https://www.bsc.es/lehmkuhl-oriol>.
- [2] D. Pastrana, I. Rodriguez, J. C. Cajas, O. Lehmkuhl, and G. Houzeaux. On the formation of Taylor-Görtler structures in the vortex induced vibration phenomenon, *International Journal of Heat and Fluid Flow* 83, 108573, 2020.
- [3] Rodríguez, I., Lehmkühl, O., Pastrana, D., Cajas, J. C., Houzeaux, G., *Wakes and Instabilities of Static and Freely Vibrating Cylinders*, Direct and Large Eddy Simulation XII, 978-3-030-42822-8 Springer, 2020.
- [4] O. Lehmkühl, I Rodriguez, D. Pastrana, J.C. Cajas and G. Houzeaux. High fidelity simulation of vortex induced vibrations for flow control and energy harvesting, EuroHPC Summit Week 2019, PRACE Days, Poznan, Poland, 2019.
- [5] <https://prace-ri.eu/news-media/publications/prace-digest/prace-digest-2019/>.




Figure 4: Instantaneous vortical structures close to the cylinder surface. Initial branch (left), super upper branch (middle), lower branch (right).

Development of an integral LES model for turbulent premixed combustion at elevated pressures

RESEARCH INSTITUTION

¹Numerical methods in Aerospace Engineering, Bundeswehr University Munich, Germany

PRINCIPAL INVESTIGATOR

Markus Klein¹

RESEARCHERS

H. Nachtigal¹, F.B. Keil¹, R. Rasool¹, C. Kasten¹, D. Alwazzan², N. Chakraborty²

PROJECT PARTNERS

²Newcastle University, Newcastle upon Tyne, UK

SuperMUC Project ID: pn69ga

Introduction

In many engineering applications (e.g. spark ignition engines, industrial gas turbines) premixed combustion takes place at elevated pressures. However, a large part of the existing literature on numerical and theoretical analysis and model development in the context of premixed combustion has been carried out at ambient pressure conditions. Further, in high-pressure test-rigs optical access is limited, laser and signal beams must pass through several windows often resulting in signal decrease making the acquisition of reliable results more difficult and more expensive. Increased pressure affects both, flame and turbulence characteristics and in addition chemical kinetics. Typically, for hydrocarbon fuels, laminar burning velocities and flame thicknesses are reduced when pressure rises. Pressure also affects kinematic viscosity which changes the Reynolds, Damköhler and Karlovitz numbers, reaction rates and flame speed, extinction, ignition delay time or pollutant formation. The reduction of the flame thickness promotes the hydrodynamic instability which has an impact on flame morphology, flame area and turbulent flame speed. Many of these aspects have been analysed in detail by interrogating a newly established direct numerical simulation (DNS) database with the help of the project resources (i.e. 28 Mio core hours). The dataset covers a large range of parametric variations of thermodynamic pressure, inlet turbulence, nozzle diameter and heat release parameter.

Methods

The well-known DNS code SENGA [1-8] has been used for the simulations carried out in this analysis. In SENGA, the conservation equations of mass, momentum, energy and reaction progress variables are solved in non-dimensional form. The spatial discretisation is carried out using a 10th order central difference scheme for internal grid points, while the time-advancement is carried out using a low-storage 3rd order Runge-Kutta scheme. All the analysis consider a generic single-step Arrhenius type irreversible chemical mechanism for the purpose of computational

economy due to the demands of flame resolution at elevated pressures. For hydrocarbon-air combustion the unstrained laminar burning velocity and the flame thickness decrease with pressure. In the context of a simple chemical mechanism, the pre-exponential factor and kinematic viscosity have been altered to account for the desired pressure dependence.

The code has been parallelised using MPI and the results of the scaling analysis have been reported in [2]. Turbulent inflow data has been generated using a modified version of the digital filter based method. In order to overcome the efficiency problems related to the generation of synthetic turbulent inflow data on large distributed grids, several modifications have been implemented which are reported in [2]. The biggest simulation (i.e. the 25 bar Bunsen flame) consumed about 4 Mio core hours on 13,824 cores. One result file containing all variables for one time step has a size of 108GB and altogether 10 statistically independent time steps have been stored for evaluating selected statistics.

Results

The analysis of the database covered a wide range of topics from fundamental aspects of high pressure turbulent combustion processes (e.g. the role of hydrodynamic instability on flame dynamics [1], scalar gradient and strain rate statistics [2], surface density function transport statistics [3], flame curvature transport in high pressure Bunsen flames [4], the validity of Damköhler's first hypothesis [5]) to high fidelity large eddy simulation (LES) modeling, including turbulent scalar fluxes [6], sub-grid scale variance of reaction progress variable [7] as well as algebraic [8] and transport equation based [9] flame surface density modelling. The main finding in terms of LES modelling is, that most of the results indicate that for a given resolution of the flame front, the considered LES closures perform in a qualitatively similar manner irrespective of pressure variation. However, for a given computational mesh, the performance deteriorates as the flame thickness decreases with increasing pressure.



Figure 1: Turbulent premixed Bunsen flames (top) Lewis number $Le = 1$ and (bottom) $Le = 0.34$ at different pressure levels: 1 bar, 5 bar, 10 bar (left to right). The different colours represent different values of reaction progress variable.

Ongoing Research / Outlook

The unity Lewis number database considered so far has been extended recently (see Fig. 1) by variations of the ratio of thermal diffusivity to mass diffusivity. This aspect becomes increasingly important because of the presence of thermo-diffusive instability (e.g. in the case of alternative, lean high hydrogen content fuels) possibly interacting with hydrodynamic instabilities which are likely to occur under elevated pressures.

References

- [1] M. Klein, H. Nachtigal, M. Hansinger, M. Pfitzner, N. Chakraborty, *Flow Turbul Combust*, 101:1173-1187, 2018.
- [2] M. Klein, D. Alwazzan, and N. Chakraborty, *Computers and Fluids*, 173:147-156, 2018.
- [3] M. Klein, D. Alwazzan, and N. Chakraborty, *Flow Turbul Combust*, 103:439-463, 2019.
- [4] A. Alqallaf, M. Klein, C. Dopazo, and N. Chakraborty. Evolution of flame curvature in turbulent premixed Bunsen flames at different pressure levels. *Flow Turbul Combust*, 103:439-463, 2019.
- [5] N. Chakraborty, D. Alwazzan, M. Klein, R.S. Cant, *Proc. Comb. Inst.*, 37, 2019.
- [6] C. Kasten, F.B. Keil, N. Chakraborty, M. Klein, *Combustion Science and Technology*, 2020, 192:2050-2069.
- [7] F.B. Keil, N. Chakraborty, and M. Klein, *Flow Turbul Combust*, 2020, 105:869-888.
- [8] R. Rasool, N. Chakraborty, and M. Klein, Algebraic flame surface density modelling of high pressure turbulent premixed Bunsen flames, *Flow Turbul Combust*, 2020. <https://doi.org/10.1007/s10494-020-00128-y>
- [9] F.B. Keil, N. Chakraborty, and M. Klein, *Combustion Science and Technology*, 2020, 192:2070-2092.

Simulating the interaction of a strake vortex with an aircraft wing

RESEARCH INSTITUTION

¹Institute for Aerodynamics and Flow Technology, German Aerospace Center, Göttingen

PRINCIPAL INVESTIGATOR

Silvia Probst¹

RESEARCHERS

Daniela Gisele François¹, Tobias Knopp¹, Tim Landa²

PROJECT PARTNERS

²TU Braunschweig

SuperMUC Project ID: pn69mu

Introduction

On modern commercial aircraft flying at high angles of attack the flow on the main wing downstream of the area where the slat is cut out for the nacelle-wing integration is prone to massive flow separation. To counteract this effect, the engines are often equipped with small vortex generators called strakes. The longitudinal vortices generated by these strakes interact with the boundary layer of the main wing upper surface having a stabilizing effect on it. The stabilizing effect of the vortex/boundary-layer interaction is not fully understood yet. Therefore, a better understanding of the involved physics is crucial to optimize the design of the strake/slat-cut-out/main-wing configuration. The aim of the project (funded by DFG under grant numbers KN 888/2-1 and RA 595/25-1) is to establish a numerical simulation method to predict and understand the strake-vortex/boundary-layer interaction for an airfoil at a high-lift configuration and at flight Reynolds numbers. To investigate the capabilities of the assessed numerical approaches, a simplified test case setup is introduced that helps to avoid complex and unknown uncertainties. The simplified test case consists on a sharp-edged Delta wing which generates a strake-like longitudinal vortex that interacts with the boundary layer on the upper surface of an airfoil located downstream, which is the HGR-01 single element airfoil. It is well known, that state-of-the-art simulation methods based on the Reynolds Averaged Navier-Stokes (RANS) turbulence models employing the Boussinesq approximation significantly over-predict the decay of free longitudinal vortices [1]. Therefore, a full differential Reynolds Stress Model (RSM) is applied for this project. Moreover, to enhance the prediction of the longitudinal vortex and its interaction with the attached boundary layer, turbulence-resolving methods called hybrid RANS/LES methods (HRLM) are investigated.

Results and Methods

The investigations in this project were separated into two tasks. The first part comprises the simulation of the longitudinal vortex generated by the Delta wing; the second part investigates the interaction of the longitudinal vortex with the airfoil boundary layer. This




Figure 1: Split strategy for the investigation of a suitable numerical approach for the strake-like vortex/boundary layer interaction. Top: Generation and streamwise development of the strake-like vortex; Bottom: Longitudinal vortex/boundary layer interaction; Visualization with iso-surfaces of the Q-criterion colored by vorticity.

split strategy is depicted in Figure 1. For all numerical simulations performed on SuperMUC-NG, the DLR TAU-Code was applied, which is an unstructured finite-volume based, 2nd order compressible flow solver. In a first step, the focus was on establishing a simulation strategy to correctly predict the longitudinal vortex generated by the flow over the Delta wing (cf. Figure 1, top). To achieve this goal, the DLR TAU-Code was extended to introduce synthetic turbulence into off-wall, free vortex flows. These synthetic turbulent structures are injected at an interface in the flow field to reliably predict the resolved turbulence in the flow. Figure 2 shows the topology of the flow, with modeled, time-averaged turbulence on the Delta wing and the transition to resolved turbulent structures in the vortex with the interface position indicated by the red line. Three different interface positions were investigated during the project (depicted in Figure 3): At position I.1, the RANS/LES interface interferes with the vortex roll-up process impacting the bulk development of the vortex. At position I.2 the vortex is better developed, however the modeled turbulence at the vortex core has strongly decayed. The generation of synthetic turbulence uses the modeled turbulence as an input, i.e. if the modeled turbulence is not reliable at the given position also the synthetic structures are defective. This leads to a strong underestimation of the




Figure 2: Visualization of regions of modeled (on Delta wing) and resolved (longitudinal vortex) turbulence.




Figure 3: Visualization of positions for the interface between modeled and resolved turbulence.

resolved turbulence in the vortex downstream of the interface. The best agreement with experimental data that were generated within the project was obtained by the interface positioned on the Delta wing surface. A detailed discussion on these findings is given in [2]. A second part of the project focused on the interaction of the longitudinal vortex with the boundary layer of the HGR01 airfoil. To permit parameter studies resulting into a large number of simulations, the setup was further simplified: The longitudinal vortex was generated in vicinity of the HGR01 airfoil by an actuator disk boundary condition. This simplified setup is shown in Figure 1 on the right hand side. The DLR TAU-Code was extended for a tailored usage of the actuator disk boundary condition and also here, the generation of synthetic turbulence was adapted. For the assessment of the vortex/boundary layer interaction three numerical approaches were investigated.

These are:

- RANS-RANS: The vortex and the airfoil boundary layers are computed with the RANS approach hence the turbulence is fully modeled
- LES-RANS: The turbulence in the vortex is scale-resolved (LES) and the airfoil boundary layers is modeled with RANS
- LES-WMLES: The turbulence for both, the vortex and the airfoil boundary layer, is scale-resolved.

The results from the LES-RANS simulation show that fully modeling the turbulent boundary layer flow neglects the effect of the scale-resolved vortex, providing very similar results to the RANS-RANS simulation. When both, the turbulent boundary layer and the vortex are scale-resolved, the vortex/boundary layer interaction produces a stronger displacement of the boundary layer. In general, thicker boundary layers have a stronger tendency to separation. However, for the LES-WMLES, the vortex transfers additional momentum into the boundary layer which counteracts this effect and leads to a more stable boundary layer. A slice through the flow field showing the mean velocity




Figure 4: Mean streamwise velocity through the vortex core and the airfoil boundary layer: Left: RANS-RANS; Center: LES-RANS; Right: LES-WMLES.

of the vortex interaction with the underlying boundary layer is shown in Figure 4. On the right hand side, in the result for the LES-WMLES, the regions of low velocities (blue) are reduced indicating the momentum transfer from the vortex into the boundary layer.

For all computations performed within this activity, we used up to 1,680 cores for one simulation, running more than 800 hours. To ensure the progress of the simulations, restart-data were saved several times during one job running on the cluster. One data-set for restart needed 56 GB of storage. To visualize the flow in the longitudinal vortex a total amount of data of 3.75 TB was collected. The huge data generated during the simulations required a real-time script-based postprocessing and a daily transfer of the results to a local storage system.

Ongoing Research / Outlook

The granted computational resources on SuperMUC-NG were essential to achieve the goals of the project. To identify an optimal simulation strategy for a certain flow phenomenon, parameter studies are an important part. For this project, it was necessary to perform these parameter studies, applying high-quality turbulence resolving simulations, which are particularly expensive. Based on the computations performed, a deeper insight into the flow physics was obtained, and deficiencies of current simulation methods were identified. The requirement for simulations with even less modeling was outlined, leading to the successful extension of the DFG-funding. An extension of this project is planned.

References and Links

- [1] Landa T., Klug L., Radespiel R., Probst S., and Knopp T. AIAA Journal, 58(7), pp. 2857-2868, 2020. DOI: 10.2514/1.J058650.
- [2] Probst S., Knopp T., François D. G., Grabe C., Landa T., and Radespiel R. AIAA Scitech 2019 Forum, San Diego, California. DOI:10.2514/6.2019-0331.

Aerodynamics of Innovative Rotor Configurations

RESEARCH INSTITUTION

Chair of Aerodynamics and Fluid Mechanics, Technical University of Munich

PRINCIPAL INVESTIGATOR

Christian Breitsamter

RESEARCHERS

Patrick Pözlbauer, Andreas Kümmel, Michael Cerny

PROJECT PARTNERS

—

SuperMUC Project ID: pn69ni

Introduction

The present project is structured into three sub-projects and it deals with the development of innovative and environmentally friendly aircraft. The first sub-project is the Clean Sky 2 project FURADO [1], which deals with the aerodynamic design optimization of a semi-watertight full-fairing rotor head. The fairings are developed for a new compound helicopter, namely the Rapid And Cost-Effective Rotorcraft (RACER). The SuperMUC was used in the framework of a fully automated optimization tool chain allowing to simulate the investigated flow problems within a reasonable time frame. The optimized fairing geometries were examined by CPU intensive CFD simulations on the SuperMUC-NG. Within the second sub-project, AURAI, different propeller configurations are investigated. The focus is on the design optimization of the propeller blades with respect to the aerodynamic efficiency as well as the noise footprint. For this purpose, a fully automated optimization tool chain was implemented, including methods of different fidelity level. The flow simulations performed on the SuperMUC are used to validate the low-fidelity methods implemented within the toolchain and to analyze promising candidates in detail. The third sub-project, HYPROP, addresses the analysis of small-scale propeller-driven flight vehicles. The multi-copter-specific flight conditions (high inflow angles at low advance ratios) result in highly transient loads and vibrations. The focus of the project is the investigation of different propeller configurations under varying inflow conditions. The usage of the SuperMUC-NG enables a high temporal and spatial resolution of the simulations, which is required to calculate the flow field in a time-resolved manner. The results are used to derive recommendations for aircraft design purposes.




Figure 1: Surface pressure distribution on the rotor head.

Results and Methods

Regarding the sub-project FURADO, the aerodynamic design optimization of the RACER rotor head fairings was conducted in several steps [1-3]. The blade-sleeve fairing, which covers the hub region of the rotor blade, was aerodynamically optimized by means of CFD simulations. For this purpose, a genetic optimization algorithm was employed and the aerodynamic forces were evaluated. The most promising fairing shapes were examined on the isolated rotor head, which is illustrated in Fig. 1 showing the surface pressure distribution. Moreover, different full-fairing beanie shapes were investigated in combination with the optimized blade-sleeve fairing. The cyclic pitch movement of the rotor blades was realized by using mesh motion in combination with overset mesh. The best combination of fairings was investigated in more detail on the full RACER configuration, which is illustrated in Fig. 2. The numerical flow simulations were performed with ANSYS Fluent 2019.R2. The computational mesh of the full configuration consists of 67 Mio. elements. Regarding the full RACER configuration, 768 cores were used on 16 nodes. Overall, five rotor revolutions were simulated, which required 170,000 core hours for each case.

The main objective within the second sub-project AURAI is the optimization of the propeller blade shape providing an efficient and ecofriendly propulsion system. Therefore, a multi-disciplinary optimization approach is applied. For the optimization process, a fully automatic optimization tool chain is implemented. Within the tool chain, CFD simulations for a detailed




Figure 2: CFD simulation of the full RACER configuration showing near wall and flow field vortical structures.




Figure 3: Surface pressure distribution on the propeller.

aerodynamic analysis and the Ffowcs-Williams Hawkins (FW-H) method for the acoustic analysis are provided. Especially for the acoustic investigations, the computational effort is very high, which is related to the fine resolution of the computational mesh and requires the application of the SuperMUC-NG. The block-structured hexahedral mesh used for the simulation was created with ANSYS ICEMCFD and contains about 80 Mio. cells. For the unsteady flow simulations, ANSYS Fluent 2019.R2 was used. The scale resolving SBES turbulence model is applied to reduce the dissipation and to better preserve the information needed by the acoustic analysis. The simulations were conducted on 16 nodes using 768 cores. For the simulation of approximately 10 revolutions, 80,000 Core-h were needed. Figure 3 exemplarily shows the surface pressure distribution on the propeller blade. Figure 4 depicts the vortex system emitted by the propeller and illustrates the high complexity of the wake.

Regarding the sub-project HYPROP, numerical flow simulations are performed using ANSYS Fluent 2019.R2 complementary to wind tunnel measurements. ANSYS ICEMCFD was used to generate two domains comprising a fixed outer domain with 12.6 Mio. cells and a rotating inner propeller domain with 5.2 Mio. cells. The sliding mesh technique was used to realize the propeller rotation. The unsteady flow simulations are performed using the $k-\omega$ -SST turbulence model and overall, ten propeller revolutions are calculated. The usage of the SuperMUC-NG enables the calculation of the flow field with high temporal and spatial resolution allowing for a detailed analysis of the flow field characteristics. In Fig. 5, the vortex system of an 18-inch two-blade propeller is shown for an inflow condition of 10 m/s and 120° of inflow angle. The highly unsteady flow field leads to vibrations and noise. A similar setup is applied on a coaxial propeller




Figure 5: Isolated propeller. Iso-surface of the Q-criterion colored by the velocity magnitude. $\alpha=120^\circ$.




Figure 4: Iso surface of the Q-criterion colored by the axial velocity component.

configuration (see Fig. 6) to investigate the different behavior under non-axial inflow. Here, in sum 35 Mio. cells are used to resolve both propellers and the near wake region.

Ongoing Research / Outlook

SuperMUC-NG enables the complex and CPU-intensive flow simulations performed within this project. At the current stage, detailed analyses of different flight conditions are conducted for the RACER demonstrator as well as the propeller configurations shown within this article.

References and Links

- [1] www.furado.tum.de
- [2] Pözlbauer, P., Desvigne, D. and Breitsamter, C. CEAS Aeronaut. J., Vol. 10 (3), 2019, pp. 665-685. DOI: <https://doi.org/10.1007/s13272-018-0341-0>.
- [3] Pözlbauer, P., Desvigne, D. and Breitsamter, C. CEAS Aeronaut. J., Vol. 123 (1268), 2019, pp. 1456-1475. DOI: <https://doi.org/10.1017/aer.2018.172>.
- [4] Cerny, M.; Breitsamter, C. Aerospace 2020, 7, 112, <https://doi.org/10.3390/aerospace7080112>.
- [5] Cerny, M.; Breitsamter, C. Aerosp Sci Technol, 106, 2020. <https://doi.org/10.1016/j.ast.2020.106048>.




Figure 6: Coaxial propeller. Iso-surface of the Q-criterion colored by the velocity magnitude. $\alpha=30^\circ$.

Eulerian Stochastic Fields Method applied to partially premixed combustion

RESEARCH INSTITUTION

Institute for Thermodynamics, Bundeswehr University Munich

PRINCIPAL INVESTIGATOR

Michael Pfitzner

RESEARCHER

Maximilian Hansinger

PROJECT PARTNERS

—

SuperMUC Project ID: pn69we

Introduction

Numerical simulations provide a very flexible and efficient manner to investigate turbulent combustion processes. In the framework of computational fluid dynamics (CFD) the Large Eddy Simulation (LES) approach has emerged as a promising tool to reproduce accurately many important characteristics of turbulent flames. The basic idea of LES is to solve numerically the large scales of motion in turbulent reacting flows, while the small scales are modeled. Typical LES grids cannot resolve the inner structure of flames, however, combustion is a diffusion driven process, which occurs on the smallest scales and interacts with the turbulent structures of the flow field. Thus, a sub-grid combustion model is usually needed. Turbulent combustion models were originally developed for either premixed flames, where fuel and oxidizer (usually air) are perfectly mixed before they burn, or non-premixed diffusion flames, where fuel and oxidizer are initially separated and the flame stabilizes at stoichiometric conditions.

Usually, in technical combustion devices neither of the two flame types prevails. Instead, a partially premixed flame type is found. In these flames some degree of fuel-air premixing and spatio-temporal variations of fuel-air ratio can be observed. These flames may exhibit locally the character of diffusion flames or premixed flames and show also additional phenomena such as triple flames composed of two premixed flames (lean and rich) straddling a diffusion branch. In addition, localized extinction and re-ignition of the flame, as well as differential diffusion effects (if hydrogen is part of the fuel) may play a role. Standard combustion models have difficulties to cope with the effects of partially premixed combustion. However, for the design process of cleaner combustion devices (e.g. gas turbines, aircraft engines, internal combustion engines) it is highly relevant to accurately predict the combustion efficiency and pollutant formation.

The need for accurate modeling of sub-grid scale processes becomes more pressing at higher Reynolds numbers in general and elevated pressure in particular, where with decreasing viscosity and thermal




Figure 1: Comparison between different ESF formulations.

conductivity Kolmogorov length scale and flame thickness decrease, resulting in increased flame wrinkling at small scales.

Industrial applications feature complex geometries with geometrical length scales ranging from sub-mm diameter fuel injection orifices to combustor circumferences of several meters. Affordable LES of such configurations will not be able to resolve all of the flame structure even if grids are fine enough to resolve a significant fraction of turbulent kinetic energy of the flow field. LES turbulent combustion models were originally developed for either fully premixed flames or completely non-premixed diffusion flames. To simulate partially premixed / stratified combustion often a mixture fraction transported scalar is introduced to the premixed combustion model. The progress variable with the mixture fraction as a second variable to account for the equivalence ratio fluctuations are used in a chemistry look up table in terms of Flamelet / Progress-Variable models. These models perform well




Figure 2: Comparison of simulations with differential diffusion and unity Lewis analogy ($Le=1$). Neglecting differential diffusion leads to a flame lift off [2].

when applied to atmospheric, laboratory-scale flames at moderate Reynolds numbers. A more promising combustion model is the Eulerian Stochastic Fields method (ESF), which belongs to the family of transported probability density function (PDF) methods. It uses finite rate chemistry and is computationally expensive and not yet widely used in LES. In this research we investigated the suitability of the Eulerian Stochastic Fields (ESF) method to partially premixed and purely premixed combustion.

Results and Methods

All simulations were conducted with OpenFOAM 4.1. The standard implementation was extended with the ESF combustion model and to account for analytically reduced chemical schemes. Approximately 15 million Core hours were granted for project pn69we. A typical ESF simulation uses computational grids with one to four million computational cells and is carried out on up to 1,440 cores. Note also that within one production run of 48 hours the physically simulated time is typically in the range of several milliseconds. Usually, more production runs are needed to obtain at least 15-20 milliseconds of simulated time. This is sufficient to extract meaningful statistical information of the combustion process. As there exist different ESF formulations. One research question was if, and how they produce different results. Figure 1 shows a comparison of mean and RMS values between simulations with the two formulations and compares it to data from experiments and high fidelity simulations (qDNS) [1]. The considered test case is a piloted flame with inhomogeneous fuel inlet. It is found that the different formulations tend to have little effect on the




Figure 3: ESF method applied to a premixed Bunsen flame at 1, 5, and 10 bar. The flame is too long compared to the benchmark FSD model for premixed combustion.

overall simulation result. It also becomes more and more popular to enrich methane flames (Natural gas) with hydrogen, e.g. in the context of carbon capture and storage techniques. Due to its low molecular weight hydrogen diffuses faster than other, heavier species. This effect is called differential diffusion and leads to different flame structures, moreover these flames are more prone to extinction. Differential diffusion is often not included in many CFD codes. An important finding is that differential diffusion is a relevant mechanism and should not be neglected in the simulations, otherwise local extinction and flame lift-off are over-predicted (Fig. 2). The ESF method was also applied to purely premixed combustion at different pressure levels. The main finding was that this combustion model was not able to resolve the complex flame structure and is not found suitable for premixed combustion simulations, as the flame heights are not reproduced correctly (Fig. 3).

Ongoing Research / Outlook

Most of the research has been completed. An important finding is that the choice of the chemical mechanism is more relevant than the employed number of stochastic fields. Often only one or two fields are sufficient, which reduces significantly the computational cost of the simulation.

References and Links

- [1] Hansinger, M., Zirwes, T., Zips, J., Pfitzner, M., Zhang, F., Habisreuther, P., and Bockhorn, H. Flow, Turbulence and Combustion, 105.3 (2020), pp. 837-867.
- [2] Hansinger, M., Pfitzner, M., Sabelnikov, V. Proceedings of the Combustion Institute (2020).

Bubble-collapse dynamics near biomaterial

surrogates using multi-material level-set

RESEARCH INSTITUTION

Technical University of Munich, Department of Mechanical Engineering, Chair of Aerodynamics and Fluid

PRINCIPAL INVESTIGATOR

Nikolaus A. Adams

RESEARCHERS

Aleksandr Lunkov, Stefan Adami

PROJECT PARTNERS

—

SuperMUC Project ID: pn72mu

Introduction

The right dose of medication can be delivered into the respective area of the human body at the right time by means of so-called targeted drug delivery. The controlled treatment, highly localized both spatially and temporary, reduces negative side effects and increases the effectiveness of the medication. One of the promising drug delivery methods is the usage of multilayered micro emulsions. Double-emulsions, e.g., are droplets, where the water-dissolved medicine is surrounded by water-insoluble oil. Along with the surrounding water environment, this kind of emulsion is called water-in-oil-in-water (W/O/W) emulsion [2]. For such systems, the drug release can be triggered by a phase transition of the stabilizing outer oil layer by ultrasound heating. Chen et al. [3] have shown that stable triple emulsions with gas core can be generated. These gas-in-oil-in-water-in-oil (G/O/W/O) droplets are very sensitive to ultrasound, yet the details of the complex interactions and wave dynamics are not well described.

It is also known that an intense ultrasound pulse or shock in water may cause an asymmetrical collapse of a gas bubble, followed by the formation of a high-speed jet. This combination of microbubbles and ultrasound is used in medicine, for example, to destruct kidney stones (lithotripsy) or to increase cell membrane permeability (sonoporation) [4].

In this project we are investigating numerically the mechanism of G/W/O/W triple-emulsion shock interactions near fluidic biomaterial-surrogate phases. Our goal is to understand the main characteristics of the penetrating jet that is observed by a shock-induced non-symmetric bubble collapse. Eventually, this knowledge can help to improve and design non-invasive drug-delivery methods.

Results and Methods

For our investigation, we use the open-source software package ALPACA [1], which was developed at

the Chair of Aerodynamics and Fluid Mechanics (TUM) over the past five years. ALPACA is a multi-resolution compressible multiphase solver using the sharp level-set technique for phase tracking. This framework is fully MPI-parallelized and optimized for large-scale simulations. The backbone of ALPACA are latest discretization techniques combining high-resolution schemes at vanishing artificial dissipation.

We use the Finite-Volume Method (FVM) together with a generalized multi-material level-set formulation to distinguish multiple phases. Originally, the level set method can distinguish two fluids only. In this project, we use a more sophisticated interface treatment, where arbitrary numbers of phases can be tracked using a scalar marker field and a global interface level set. Locally, individual level sets are generated for each existing interface. However, globally only one geometrical interface information is tracked.

We use the SuperMUC-NG resources to perform “direct numerical simulations” of triple emulsion collapse scenarios. In the beginning, however, we limited the problem description to isolated phenomena like the shockwave propagation through different materials, the asymmetrical collapse of the bubble with formation of toroidal structure and subsequent oscillations, the high-speed jet formation, and the flow of the jet through the oil shell and its subsequent spreading outside the emulsion.

A series of qualitative snapshots for the asymmetrical bubble collapse near a gelatin interface is shown in Figure 1. By the effect of an external shockwave, the bubble collapses and initiates a highly focused liquid jet towards the gelatin phase. This liquid jet can then penetrate the gelatin phase. This scenario mimics the fluidodynamically dominating effect during sonoporation.

A representative grid for a three-dimensional simulation of the indicated setup has an effective resolution of $2,048 \times 1,024 \times 1,024$ cells. Note, the multiresolution approach allows to handle this effective resolution




Figure 1: Time series of qualitative snapshots for a non-symmetric bubble (grey) collapse near a gelatin phase (orange). Note, time increases from top left to bottom right. The remaining cavity demonstrates the liquid jet penetration.

with actually “only” about 200 million cells. For each parameter set we simulate the numerical problem setup for about two weeks using 48 nodes (48 cores each). Besides the high computational cost, also post processing these simulations is challenging since output files are quite large with sizes of 50 GB each. In total, we have generated about 170 TB of data that needs to be analyzed now in detail.

Ongoing Research / Outlook

In this project, we have seen a strong improvement in our compute capabilities with the availability of the large-memory nodes on SuperMUC-NG. Earlier, we had to sacrifice compute power since we could not load enough data into memory. This limitation basically caused cores idle and wasted compute time. Now, we can efficiently utilize all available compute units and see much improved performance and scaling of our application.

Currently, we are visualizing and analyzing the vast amount of data generated by the project and are preparing two publications on the penetration mechanism. We have seen in the data already some fascinating mechanisms that require further detailed investigations in the future.

References and Links

- [1] <https://www.nanoshock.org>
- [2] Duncanson, Wynter J., et al., *Langmuir* 30.46 (2014): 13765-13770.
- [3] Chen, Haosheng, et al., *Soft Matter* 9.1 (2013): 38-42.
- [4] Pan, Shucheng, et al., *Phys. Rev. Fluids* 3.11 (2018): 114005.

Vortical Flow Interaction on Airborne High Speed Vehicles

RESEARCH INSTITUTION

German Aerospace Center, Institute of Aerodynamic and Flow Technologies, Göttingen

PRINCIPAL INVESTIGATOR

Christian Schnepf

RESEARCHER

Michael Werner

PROJECT PARTNERS

—

SuperMUC Project ID: pn72xu

Introduction

High maneuverability is a mandatory requirement of state-of-the-art airborne high speed vehicles like missiles and fixed-wing aircrafts. This requirement demands maneuvers at high angles of attack which lead to large flow separations and the formation of multiple vortices (Figure 1 and 2). Regardless of the number of vortices, the vortices and their footprint on the surface pressure distribution of the vehicle have a strong impact on the aerodynamic forces and moments and thus on the flight stability of the vehicle. The evolution of the vortices is often affected by interactions among these vortices (vortex / vortex-interaction) and the interaction with vehicle components (vortex / vehicle component-interaction). At high speeds ($M > 0.85$), shock waves occur, which can additionally lead to a vortex / shock interaction. The interaction can range from simply influencing the vortex trajectory to a complete vortex merge, changing significantly the trajectory and strength of the vortices, and thus the footprint of the vortex on the vehicle surface and vehicle components. In the worst case, vortex burst can occur. Vortex burst, also called vortex breakdown, is characterized by an abrupt change of the vortex trajectory and topology, respectively, often including an internal stagnation point and an expansion of the vortical structure, which goes along with a pressure rise (Bubble-type breakdown [1]). On a delta wing, vortex breakdown often occurs at first on one side of the vehicle, leading to a highly asymmetric flow field and surface pressure distribution. Important factors for the different kinds of interactions are the vortex size and strength, which are changing with angle of attack, roll and sideslip angle, and Mach number. The different kinds of vortex-interaction mentioned above have been investigated for decades, but important aspects of the vortex-interaction flow physics are still not well understood. Therefore, a non-confidential NATO Science and Technology Organization (STO) working group (AVT 316) [2] has been established recently to improve the prediction of vortex-interaction. Numerical simulation can help to get a deeper insight into the flow physics. However, for this purpose the roll-up and the evolution of the vortices have to be modelled with a certain accuracy, since this affects the trajectory,

strength, and size of the vortices. In this respect, standard turbulence models are often inadequate [3]. Considering vortex-interaction with a shock wave or a vehicle component, both the development of the attached vortex as well as the evolution of the separated vortex in downstream direction towards the interaction zone is of importance. However, standard turbulence models have been proven to be too dissipative to correctly reproduce vortex evolution [2,3]. Therefore, the aim of this project is to investigate the application of state-of-the-art scale-resolving methods like DDES and IDDES to these kinds of flows. Furthermore, the performance of a seven equation Reynolds stress model will be investigated. In the final report, a comparison of the results of eddy-resolving methods with the results of Reynolds-Averaged Navier-Stokes (RANS) simulations, using standard 1 and 2 equation turbulence models, will be made.

Results and Methods

For a generic double-delta wing configuration and a generic missile configuration, steady state (RANS) and time accurate flow simulations (URANS), as well as scale resolving simulations have been carried out with the DLR-TAU Code in this study. TAU is a hybrid structured/unstructured finite volume flow solver for the compressible Reynolds-Averaged Navier-Stokes-equations. The most demanding steady state computation needed about 15,000 Core-h (1,920 cores) and the scale resolving simulations consumed on average 500,000 Core-h (1,200 cores) to converge. The latter statement demonstrates the need for SupermucNG's HPC. In total, 5 Terabyte of data was analyzed and stored external yet.

Missiles

Regarding the investigations on the missile configuration, simulations with a state-of-the-art seven equation turbulence model have been performed in a first step. The aim is to investigate the influence of the computational grid topology and the fineness of the grid on the vortex topology, and the influence of the grid properties on the numerical dissipation, respectively. In Figure 1, the vortex dominated flow that evolves at the missile at a Mach number of 0.85, an angle of attack




Figure 1: Visualization of the vortex topology using λ_2 criterion isosurface and helicity (top: fine grid ($\approx 100^6$ grid nodes), bottom: coarse grid ($\approx 12^6$ grid nodes)).

(AoA) of 17.5° and a roll angle of 45° is visualized using the Q-criterion and the helicity. In the upper illustration (fine grid), secondary vortices on the wing are visible in addition to the large scale vortices. In the lower figure, however, they are missing due to the coarseness of the grid and the resulting greater numerical dissipation compared to the fine grid. This affects all vortices, such as the leeward vortex on the port side. Figure 1 illustrates how the strength of this vortex changes significantly with the grid fineness. However, this vortex has a strong influence on the rolling moment due to its interaction with the wing. Hence a correct determination of the vortex evolution is essential for simulation based aerodynamic data.

Delta Wings

The studies on delta wings were performed on the DLR F22 planform, a generic triple delta configuration with sharp leading edges. In a first step, URANS simulations for Mach numbers between $M = 0.5$ and $M = 1.41$ and AoA between $\alpha = 8^\circ$ and $\alpha = 28^\circ$ were performed to identify suitable flow conditions for the investigation of vortex-interactions. In a second step, scale resolving IDDES simulations were carried out for selected flow conditions. For the evaluation of the results, the unsteady simulations were averaged over a period of at least 50 convective time units. In Fig. 2, the λ_2 criterion is used to visualize the vortex topology




Figure 2: Vortex topology on delta wing configuration at $M=0.85$, $AoA=16^\circ$, visualized by an isosurface of the λ_2 criterion.




Figure 3: Comparison of velocity deficit in the vortex core between URANS and IDDES.

of an averaged IDDES simulation at $M = 0.85$, $AoA = 16^\circ$. At these flow conditions, three primary vortices develop: at the forebody, the strake and the main wing, respectively. Vortex merging between the inboard and midboard vortex can be observed on the strake, and both the inboard and the midboard vortex interact with shocks forming on the wing. While the inboard vortex is barely disturbed by the shock interaction, the vortex core velocity of the midboard vortex drops significantly, leading to vortex breakdown. This can be seen in Figure 3, where the streamwise development of the velocity deficit δ is plotted for the inboard, and midboard vortex. Both URANS and IDDES simulations predict vortex breakdown after the interaction of the midboard vortex and the shock. However, the shock position strongly varies between the two approaches.

Ongoing Research / Outlook

Regarding the missile configuration, IDDES simulations will be undertaken to investigate how not only the grid influences the evolution of the vortices and their interactions, but also the treatment of turbulence. This is still addressed in the context of this project. But, it is also a possible starting point for a follow-up project. Based on the results for the delta wing configuration, in a next step, and beyond, it is planned to perform IDDES simulations of analytical vortices in order to more closely study the vortex-shock interaction itself.

References and Links

- [1] Lucca-Negro, O., and T. O'doherty, Prog. Energy Combust. Sci., 27.4 (2001): 431-481.
- [2] TAYLOR, N., et al. The Prediction of Vortex Interactions on a Generic Missile Configuration Using CFD: Current Status of Activity in NATO AVT-316. 2019.
- [3] Braun, S., et al. Numerical simulation of vortex roll-up processes using the ssg/lrr-w model. In New Results in Numerical and Experimental Fluid Mechanics X, pages 481-491. Springer, 2016.

Flashback phenomena in a Gas Turbine with hydrogen in technically premixed conditions

RESEARCH INSTITUTION

¹Hermann-Föttinger-Institut, Technische Universität Berlin

PRINCIPAL INVESTIGATOR

Panagiotis Stathopoulos¹

RESEARCHERS

Daniel Mira², Oriol Lehmkuhl², Tom Tanneberger¹

PROJECT PARTNERS

²Barcelona Supercomputing Center, Spain

SuperMUC Project ID: pr27bo

Introduction

The use of hydrogen-enriched fuels is an alternative to reduce pollutant emissions in gas turbine (GT) applications allowing high flexibility of operation and fuel flexibility. However, the presence of hydrogen in fuel mixtures can lead to unpredictable behaviour of combustion systems and undesirable phenomena such as flashback or blow-off. The existence of flashback is one of the main drawbacks on burner-stabilized combustion systems operated with hydrogen, as the H_2 content contributes to higher flammability limits and higher burning speeds when compared to traditional hydrocarbon fuels. Swirling combustors can take advantage of an axial air injection that reduces the velocity deficit along the centerline, thus providing additional resistance against flashback [1]. Axial air injection influences the location of the flame by displacing the stagnation point of the aerodynamically induced recirculation downstream in the combustion chamber, thus increasing burner stability. An example of this design is the technically-premixed swirl-stabilized presented in the experiments [1], which shows that flashback appears at low equivalence ratios, when this combustor is operated with pure hydrogen. The axial momentum ratio between the fuel jets and the air was found to be the dominant parameter controlling the flame stabilization process and flashback resistance [1] over mixing quality and equivalence ratio fluctuations. This hypothesis motivates the present study where large-eddy simula-

tions (LES) are set with the same axial momentum ratio of the experiments and the influence of fuel/air mixing is removed by a perfectly premixed assumption.

Results and Methods

Methodology

The governing equations describing the reacting flow field correspond to the low-Mach number approximation of the Navier–Stokes equations with the energy equation represented by the total enthalpy in the context of LES with an eddy-viscosity given by the closure proposed by Vreman. The flow conditions considered in this study include the oxidation of a hydrogen/air mixture at preheated conditions in a turbulent flow field at two equivalence ratios using a flamelet method based on the tabulation of laminar premixed flamelets with a presumed-shape Probability Density Function (PDF) with a beta-function [2]. The set of governing equations is integrated in time using a third-order Runge–Kutta explicit method with a second-order low-dissipation low-Mach number scheme. This computational framework is developed into the multiphysics code Alya [3], which is used to run the LES simulations.

Numerical setup

The computational domain is depicted in Fig. 1 and corresponds to a swirl-stabilized technically premixed burner. It consists of plenum, fuel injection, mixing tube and combustion chamber. The numerical simulations have been conducted on a hybrid unstructured mesh including the combustion chamber, mixing tube, plenum and fuel injection system, so the flow distribution across the mixing tube is accounted for. The mesh is composed by prisms, tetrahedrons and pyramids, and locally refined in the regions of interest. While several meshes are considered featuring different resolutions in the combustion chamber, the mesh used of the presented results include a length scale of 0.7 mm in the reacting layer, and 1 mm everywhere else.

The three cases have a Reynolds number $Re = 75,000$ with pre-heated air at $T_{air} = 453$ K and




Figure 1: Computational domain.




Figure 2: Q-vorticity visualization and power spectra density of the inert case including PIV and LES.




Figure 3: Q-vorticity visualization and power spectra density of the inert case including PIV and LES.

hydrogen coming at $T_{H_2} = 320 \text{ K}$. The first case corresponds to inert conditions without fuel injection and it is used not only for validation purposes, but also for the evaluation of the impact of heat release on the dynamics of the flow in this burner concept. The other two cases ($\phi=0.6$ and $\phi=0.4$) correspond to reacting flow calculations for which experimental data is available.

Results

The inert calculation revealed a strong characteristic frequency at approximately 1,100 Hz that did not appear in the presence of the flames, which has been confirmed by the LES calculations. The characteristic frequency corresponds to the Precessing Vortex Core (PVC), which is a well-known global flow instability arising in swirling flows underlying vortex breakdown, see Fig. 2. A flow visualization of the reacting flow field is shown in Fig. 3 by a volumetric rendering of the density gradient. The flame topology of a swirl-stabilized flame can be distinguished in this plot. The distribution of axial velocity and temperature, during stable operation at $\phi = 0.6$ is also shown in Fig. 4. The effects of heat release also influence the vortex-breakdown mechanism. The inert case exhibits a narrow central recirculation with an axial velocity deficit at the nozzle exit plane, which can be identified as bubble-type vortex-breakdown. The heat release causes a wider opening angle with a correction of the velocity deficit at the nozzle exit that changes the stabilization mechanism to cone type vortex-breakdown as shown in Fig. 4. The LES results are able to predict this transition of the vortex breakdown in the reacting case, which is a fundamental aspect for flashback safety in this burner [1]. The propagation of the vortex breakdown into the mixing tube is fundamentally induced by the density stratification caused by the flame heat release, and suggests the burner is prone to combustion-induced vortex breakdown flashback [4]. An analysis of the spectra and POD modes indicate that the PVC is attenuated due to the increase in axial momentum and is ultimately suppressed in the reacting flow field. The LES is capable to reproduce both damping effects, which are also in agreement with the experimental data [2]. The analysis of the flames has shown certain dynamics as the flashback point is approached. As the axial momentum ratio is reduced, the flashback propensity of the burner increases due to an intensification in the




Figure 4: Axial velocity and temperature snapshots of the flow field during stable operation at $\phi=0.6$.

velocity deficit of the incoming mixture. Moreover, the recirculation region is shifted upstream, the central recirculation is altered and the flame position is displaced towards the reactants. The study of instabilities confirms there is no instability at the onset of flashback as it is suppressed by the density stratification.

Computational cost

The reacting flow simulations were run using 1,280 cores for 96 hours to obtain fully-averaged flow fields and extract first and second order statistics. The inert calculations were much faster and required 36 hours. In order to compute the POD analysis, many snapshots were stored requiring about 0.8 Terabytes of data for each case.

Ongoing Research / Outlook

Thanks to SuperMUC-NG, we were able to run these simulations, and perform data analysis and visualizations in the cluster. Our next steps are to include a detailed study of the mixing process between fuel and air.

References and Links

- [1] Reichel, T.G., Goeckeler, K., Paschereit, O. J., Eng. Gas Turbines Power 137, 111513 (2015).
- [2] Mira, D., Lehmkuhl, O., Both, A. et al., Flow Turbulence Combust 104, 479–507 (2020).
- [3] Vazquez, et al., J. Comput. Sci. 14, 15–27 (2016).
- [4] Oberleithner, K., Stöhr, M., Im, S.H., Arndt, C.M., Steinberg, A.M., Combust. Flame 162(8), 3100–3114 (2015).

Aerodynamics and Aeroelasticity

of high agility aircraft

RESEARCH INSTITUTION

Chair of Aerodynamics and Fluid Mechanics, Technical University of Munich

PRINCIPAL INVESTIGATOR

Christian Breitsamter

RESEARCHERS

Christopher Reinbold, Matteo Moiola, Dominik Sedlacek

PROJECT PARTNERS

—

SuperMUC Project ID: pr27ce

Introduction

The Chair of Aerodynamics and Fluid Mechanics (AER) is active in the research on the aerodynamic characteristics of low aspect-ratio wings. A turbulence model conditioning methodology for vortex dominated flows and the controllability and stability investigation of delta wing configurations has provided significant results and play a major role in the academic activity at AER. In this project, the aerodynamics of high agility aircraft are investigated. High agility aircraft often require an operation at extreme flight conditions. The configuration is based on wings with medium to high sweep angle and leading edges small radii. Even at low angles of attack the flow separates at the leading edge or at multiple swept leading edges and forms a-vortex or vortex systems, respectively. Extensive investigations on hybrid delta wing configurations were carried out. The project can be divided into three subprojects. In the VitAM project a turbulence model conditioning for vortex dominated flows is developed. In the second subproject, "Multiple Swept Wings", funded by DFG, the vortex dominated flow on hybrid delta wing fuselage configurations is investigated. The project "Aeroservoelasticity" deals with the development of a numerical method for the investigation of aeroelastic problems including control surface deflections which are modeled with a Chimera approach.

The usage of the High-Performance Computing system has played a fundamental role for the proper development of the simulations. Numerical simulations are routinely employed in parallel to experimental data in the design or investigation of new aerodynamic solutions. Considering a complex geometry and a progressively demanding level of accuracy, the computational grids need to be refined and the numerical methodology improved; both resulting in an increasing computational effort. In this sense, a High-Performance Computing system is essential for the effectiveness of the research. The results and future expectations are reported, showing their importance for the research projects.




Figure 1: Non-dimensional Surface distribution C_P and vorticity magnitude ω obtained by standard and optimized turbulence model simulations for a 65° delta wing.

Results and Methods

In the context of the VitAM-In-ABC group of the VitAM (Advanced Aircraft Understanding via the Virtual Aircraft Model) project, a methodology to automatically calibrate a turbulence model modification based on experimental data has been developed and applied for relevant test cases, Fig. 1. The idea consists of introducing additional source terms which are exclusively active in the vortex flow field. The remaining part of the field is not affected. The turbulence model used here is the one-equation eddy-viscosity transport model of Spalart-Allmaras. The additional terms are formulated for having different sensitivities at different regions or different topologies of the vortex. A gradient-descent method is used for the calibration of the coefficients by minimizing the error between the numerical simulation and the experimental data.




Figure 2: Iso-contour plots of the Q criterion $Q^* = 50$ (blue) and the non-dimensional axial velocity $U/U_\infty = 0$ (red) for URANS and DDES computations for $Ma_\infty = 0.15$ and $\alpha = 16^\circ$.

During the optimization procedure, CFD simulations are iteratively performed by restarting the previous solution and varying the coefficients to evaluate their sensitivity for establishing the updated parameters values. In this sense, even though the conditioned turbulence model does not introduce additional computational demand to the single numerical simulation, the optimization procedure for different test cases at different flow conditions and the investigation of the terms sensitivities as well as the production of a comprehensive data-base is computationally expensive. Results have shown the good grade of accuracy improvement obtained from the methodology for different test cases [1].

In the project, "Multiple Swept Wings", the flow field at hybrid delta wing fuselage configurations are investigated. The high-performance computing system is used for highly resolved simulations with different numerical methods. Unsteady Reynolds-averaged Navier-Stokes (URANS) simulations and delayed detached eddy simulations (DDES) are performed for the considered configurations, Fig. 2.

The simulations were conducted at subsonic speed at symmetric freestream conditions and an angle of attack of $\alpha = 16^\circ$. For the NA1 W1 configuration the results of the DDES provide a better agreement to the experiments, especially, for flow field quantities and the pitching moment coefficient. For the NA1 W2 configuration, however, the URANS simulation also shows very good results, with low deviation of the aerodynamic coefficients and flow field quantities with respect to experimental data. The computational time for a fully converged DDES simulation for one configuration at $\alpha = 16^\circ$ is about 3.5 Mio Core-h. The URANS simulations with the one-equation turbulence model of Spalart and Allmaras with rotation correction only needs about 0.2 Mio Core-h [2,3].

In the Aeroservoelasticity project, a numerical tool for the computation of aircraft control surface




Figure 3: Non-dimensional surface pressure distribution C_p of the Model53 for the rigid and flexible configuration at $Ma = 0.8$, $\alpha = 20^\circ$ and a flap deflection angle of $\delta = 30^\circ$.

aerodynamics with elasticity effects is developed and investigated. The solution is based on coupled high-fidelity Computational Fluid Dynamics (CFD) and Computational Structural Mechanics (CSM) simulations with the multidisciplinary simulation environment SimServer. In SimServer, the DLR Tau Code is embedded to obtain the CFD solution by solving the Reynolds-Averaged Navier-Stokes (RANS) equations. Structural displacements are computed with a modal solver. The Chimera implementation of SimServer is applied to model the control surfaces. Numerical simulations with the flexible Chimera method have been performed for two validation test cases and for the Model53 wing configuration, Fig. 3, which is a generic delta wing with an inboard and outboard trailing edge flap. Aerodynamic and aeroelastic simulations for several flap deflection angles are conducted in order to investigate the effect of structural deformations on the flow field and control surface effectiveness. The results revealed the significance of taking the structural flexibility into account in order to obtain accurate values and trends of the forces and moments for several flap deflection angles. Compared to the rigid simulation, the computational time increased by a factor of 1.7 for the aeroelastic simulation, which is considered to be acceptable given the considerable improvements in the accuracy of the results [4].

Ongoing Research / Outlook

The SuperMUC-NG is essential to this project. Major investigations have been performed in the past and important and interesting results were achieved. Further numerical investigations are needed in all subprojects for a deeper comprehension of the vortex flows and their structural impact on delta wing configurations.

References and Links

- [1] M. Moioli, C. Breitsamter and K. Sørensen, Int J Comput Fluid Dyn, 33:4, 149-170, 2019. DOI: 10.1080/10618562.2019.1617857
- [2] S. Pfürer and C. Breitsamter, J Aircr, 56(6):2093-2107, 2019. DOI:10.2514/1.C035491.
- [3] D. Sedlacek, S. Biechele and C. Breitsamter, Numerical investigations of vortex formation on a generic multiple swept wing fuselage configuration. DLRK, Aachen, 2020 (submitted).
- [4] C. Reinbold, K. Sørensen and C. Breitsamter, Development of a Numerical Method for Aeroelastic Simulations with a Chimera Approach for Moveable Control Surfaces. DLRK, Aachen, 2020 (submitted).

CO Emissions in Part Load Conditions of Stationary Gas Turbines

RESEARCH INSTITUTION

Institute for Combustion Technology, RWTH Aachen University, Germany

PRINCIPAL INVESTIGATORS

Antonio Attili, Heinz Pitsch

RESEARCHERS

Konstantin Kleinheinz, Kai Niemi

PROJECT PARTNERS

—

SuperMUC Project ID: pr27gu

Introduction

For the development of the next generation gas turbine combustors, the assessment of pollutant formation plays a key role. Due to increasingly stringent emission limits, further improvements and optimization are required. Numerical simulations can contribute to design cost reduction and acceleration of the development process. However, a lack in the predictive quality of state-of-the-art combustion models often limits the usefulness of numerical results. In stationary gas turbines, the most relevant pollutants are carbon monoxide (CO) and nitrogen oxides (NO_x). This work focuses on the formation of these pollutants under part load conditions of stationary gas turbines. Part load conditions are characterized by low equivalence ratios of the air/fuel mixture leading to low exhaust temperatures. Under such conditions, NO_x emissions are not critical, but CO emissions can increase and limit load flexibility. One way to reduce CO emissions in low load conditions is “staging” of the fuel jets. In this mode some of the jets are run with air only. Thus, the remaining, burning jets can be operated at higher individual equivalence ratios while retaining a low total equivalence ratio in the combustion chamber. This research focuses on the staged combustion and the interaction of the burning flames and the air jets. These interactions influence the CO chemistry and are not well captured by current combustion models. With better insight into the chemistry-turbulence interaction via direct numerical simulations (DNS) improved models can be developed.

Results and Methods

Numerical Framework

The governing equations of the DNS are given by the reacting Navier-Stokes equations in the low-Mach limit. For the computation, an in-house code called CIAO is employed. CIAO solves the Navier-Stokes equations along with multi-physics effects. It is a structured, arbitrary order, finite difference code. A low-Mach solver is available within the framework, which uses Crank-Nicolson type time advancement along with a predictor-corrector update scheme. The momentum equations are discretized by fourth order

central differences resulting in low numerical dissipation. The simulations are run with finite rate chemistry and employing a skeletal mechanism for lean premixed methane/air combustion. Species and temperature equations are discretized by the WENO5 scheme ensuring bounded solutions. Spatial and temporal staggering of flow variables is used to increase the accuracy of the numerical stencils. Temperature and species equations are advanced by introducing the symmetric operator split of Strang. The chemistry operator uses a time-implicit backward difference method, as implemented in CVODE, which is a widespread software package for integrating stiff ODEs. For further details about the applied numerical algorithms and code verification, the reader is referred to [2]. The code uses the widely known message passing interface (MPI) standard for parallelization. The simulation layout is geared by a model combustor of a stationary gas turbine combustor, which has been used in several experimental studies. Multiple burners are arranged circularly around a central pilot flame. For the simulation the circular arrangement is transformed into three parallel slot jets. Figure 1 shows a sketch of the resulting simulation setup. The pilot flames are located between the jets. The inflow is located on the left-hand side and the outflow on the right-hand side of the domain. The domain is periodic in the spanwise and crosswise directions. The

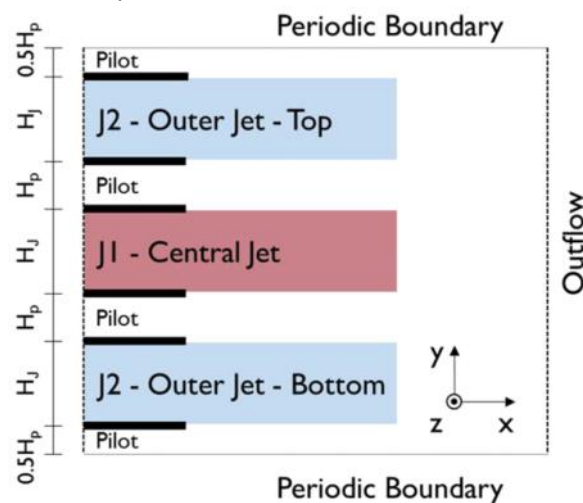


Figure 1: Sketch of the DNS domain. Inflow condition on the left hand side consists of three turbulent jets with slot width H_j and three laminar pilots with slot width H_p . The domain is periodic in y and z directions.




Figure 2: The three cases (top down: Cases A, B, C). The flame surface is shown in brown shading, the air jets in blue.

simulation domain consists of between 260 million grid points for the shorter reference case and up to 430 million grid points for the cases with air jet interaction. The production runs were performed on up to 10,752 cores, using a total of about 8.5 million core hours.

Results

One reference case with three reactive jets is compared to two cases with varying interaction between one reactive jet and two air jets. Figure 2 shows instantaneous images from the three cases. From top to bottom they are the reference case with three burning flames (Case A), a staged jet case with high jet-jet interaction (Case B) and a case where the air and fuel jets are further apart, reducing the interaction (Case C). Comparing the two staged cases shows an earlier breakup of the flame in Case B and increased mixing of flame and air jets. The effect of increased perturbations in the combustion process is seen in increased CO emissions at the outlet. The reduction in the equivalence ratio, coupled with the lower temperature from the cold air jets leads to more incomplete combustion compared to the reference case. The effect is increased by the higher interaction in Case B. Figure 3 presents the fraction of carbon bound in CO, normalized by the total carbon. While the reference Case A has the highest CO peak, it is




Figure 3: Fraction of total carbon bound in CO along stream-wise axis x .

almost completely oxidized to CO_2 downstream. The two staged cases present a less sharp peak and reduced oxidation. In order to understand the different CO peaks at $x = 20 H_j$ the joint probability density functions (PDF) of CO mass fraction and progress variable are plotted in Figure 4 for Case A and Case B. In Case A the pathway from unburnt to burnt is only along the trajectory of the laminar unstretched flame. The trajectory shows a steep increase where CO is produced in the inner reaction zone and then an equally steep decrease in the oxidation zone. This pathway is also present in Case B, though less pronounced. Additionally, the PDF of Case B presents a second, lower branch, linearly linking the burnt and the unburnt states. This pathway is formed through mixing between burnt gas and fresh air. Thus, the overall distribution of CO is lowered at this point in the combustion chamber. However, the CO formed along the lower pathway is not oxidized because of the locally lower temperatures and remains as increased CO emissions at the outlet.

Ongoing Research / Outlook

For the next step of this project a DNS setup with parallel jets and cooled walls will be investigated. This allows for further research on jet-jet interaction and additionally will provide insight into the effects of cold walls on the CO chemistry. Furthermore, it is planned to perform this DNS setup with larger Reynolds numbers to investigate the influence of increased turbulence on the combustion chemistry.

References and Links

- [1] <https://www.itv.rwth-aachen.de/en/research/ongoing-research/direct-numerical-simulation-of-the-propagation-of-a-premixed-flame-in-a-homogeneous-turbulence-field/>.
- [2] O. Desjardins, G. Blanquart, G. Balarac, H. Pitsch, J. Comput. Phys., 227:7125-7159, 2008.




Figure 4: Joint PDF of the mass fraction Y_{CO} and the progress variable for cases A (left) and B at $x/H_j = 20$. It also shows, the conditional mean (black lines) and the solution of an unstretched, laminar flame (pink).

Simulating Acoustic-Flame Interactions

in Rocket Engines

RESEARCH INSTITUTION

Institute of Aerodynamics and Flow Technology, Spacecraft Department, German Aerospace Center (DLR), Göttingen

PRINCIPAL INVESTIGATOR

Klaus Hannemann

RESEARCHERS

Tim Horchler, Stefan Fechter

PROJECT PARTNERS

—

SuperMUC Project ID: pr27ji

Introduction

Combustion instabilities are strong acoustic disturbances that often form spontaneously inside rocket combustion chambers and may lead to catastrophic failure of the launch vehicle. Although all physical processes that may contribute to the onset of such instabilities are well understood, the exact mechanisms of how these processes interact are still largely unknown. It is, however, widely accepted that the interaction between acoustic waves and the flame inside the rocket engine is very important to the formation of combustion instabilities. Even though it is very difficult to investigate the flow field in rocket combustion chambers experimentally due to the harsh flow environment, DLR researchers successfully designed a subscale combustion chamber for specifically studying the interaction between acoustics and the flame. Combustion chamber H (BKH, see Figure 1) consists of five primary shear coaxial injectors that are placed in the middle of the combustion chamber through which oxygen and hydrogen are fed into the reaction zone. Besides its large main nozzle in the axial direction, BKH contains a secondary nozzle perpendicular to the main flow direction which can be opened and closed periodically

by a siren exciter wheel. By controlling the rotation frequency of the exciter wheel, different acoustic oscillations, so-called eigenmodes, can be excited inside the combustion chamber. BKH is also equipped with two windows on its side walls allowing to observe the flame dynamics using high-speed cameras and optical measurement techniques. Even though BKH experiments provided a lot of insight into the nature of acoustic-flame interactions, many other aspects can only be observed by detailed numerical simulations. The purpose of this SuperMUC project is therefore to numerically reproduce BKH experiments for an operating condition that is vulnerable to combustion instabilities in real flight engines. Recent experiments on a different subscale rocket engine [2] showed that combustion instabilities occur when an eigenmode of the oxygen injector is in resonance with the main oscillation mode of the combustion chamber, i.e. when both modes have the same frequency. In this project we investigate for the first time such a mode coupling scenario by carefully controlling the eigenfrequency of the oxygen injector and tuning it to the combustion chamber eigenfrequency.

Results and Method

The flow field inside combustion chamber H is simulated using the DLR TAU code. In order to solve the conservation equations for mass, momentum, energy and additional turbulence quantities, the combustor volume is discretized into small non-overlapping control volumes (so-called finite volumes) in which the conserved quantities are calculated. By applying many small time steps consecutively we can accurately simulate the evolution of the complex flow field inside the combustion chamber. The simulation of rocket engines also requires sophisticated models for the thermodynamics of the propellants and the chemical reactions involved [3]. Because of the high chamber pressure in BKH (60.3 bar), oxygen is a so called supercritical fluid,




Figure 1: A schematic overview of BKH [1].




Figure 2: A visualization of the flow field inside BKH.

meaning that its density is similar to liquid water, but it behaves otherwise like a densified gas. BKH is simulated using the scale-resolving Detached Eddy Simulation (DES) model which accurately captures small turbulent features of the flow but is computationally very expensive. The DLR TAU code can be applied efficiently on supercomputing architectures like SuperMUC. Typical BKH simulations are running on 2,520 Cores in parallel. The project used 13.7 Mio. Core-h in total for the simulation of three different operating conditions of BKH and smaller preparatory studies. This project provides the first scale-resolving simulation results for combustion chamber H operating with both cryogenic oxygen and hydrogen. Figure 2 shows a typical flow field highlighting the dense oxygen cores (gray structures) which are surrounded by the hot flame. This flow field shows the undisturbed flame shape when no acoustical excitation is present. The simulated combustion chamber pressure of 61.3 bar agrees well with the experimental value. In order to determine the combustion chamber eigenmodes, which is the first step to enforce acoustic mode coupling in simulations, an artificial pressure pulse is placed inside the combustion chamber. This pulse excites all chamber eigenmodes whose shapes, frequencies and decay rates can be determined using decomposition techniques, e.g. Dynamic Mode Decomposition (DMD). The resulting mode shapes are shown in Figure 3. They consist of regions with higher pressure (red) and lower pressure (blue) oscillating at a single frequency. Chamber modes consist of longitudinal (L) and transversal (T) modes and combinations of both. The most important chamber eigenmode in terms of combustion instabilities is the $1T$ mode. Its mode frequency also agrees very well with the experimental value. The $1T$ mode is most easily excited by the siren and forces the flame to oscillate in the direction of the siren. During this motion, the dense oxygen cores are shortened significantly and flattened in a plane perpendicular to the excitation direction. Once the chamber $1T$ mode frequency is known, we can tune the oxygen injector to have its eigenfrequency matching the chamber $1T$ mode frequency.

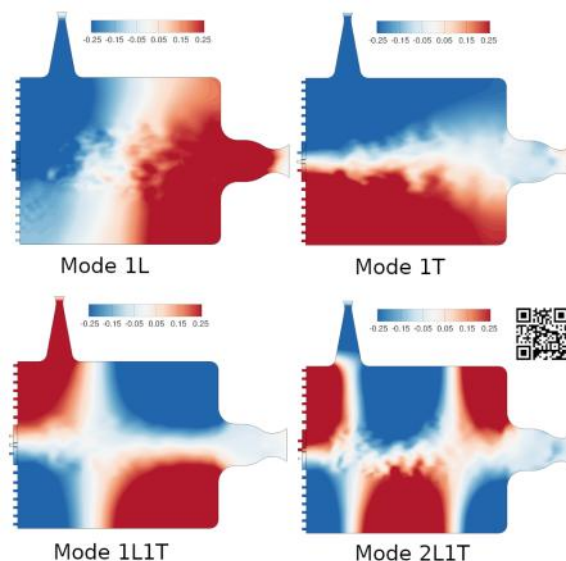


Figure 3: Combustion chamber eigenmodes of BKH. For an animated version of the plots, please follow the link in the QR code.

We then compare results from two simulations where the injector mode can couple to the chamber mode, and one without mode coupling. In a preparatory study we used computationally cheaper unsteady Reynolds-averaged Navier-Stokes (URANS) simulations to investigate the coupling scenario. Even though we successfully achieved the desired mode coupling, no flame response was observed in the simpler URANS setup. This result suggests that the simpler simulations possibly neglect a key feature for the development of instabilities, which we believe is the vortex shedding at the main injectors. This mechanism is resolved in Detached-Eddy simulations that are currently running in the last phase of this project. We expect these highly resolving simulations to clearly show if the vortex shedding plays a crucial role or if we have to focus on other possible mechanisms that might be related to the development of combustion instabilities.

Ongoing Research / Outlook

The scale-resolving simulations in this project would not have been possible without the resources from SuperMUC-NG even though we still had to limit the simulations to half of the complete combustion chamber. During the project we realized that post-processing of large amounts of time-resolved data requires a suitable framework that can process larger-than-memory datasets in reasonable time. In the future we plan on extending our simulation activities to more realistic combustion chambers and future fuels, e.g. methane, that are currently of great interest for the aerospace industry.

References and Links

- [1] S. Beinke. Dissertation. School of Mechanical Engineering. University of Adelaide, Australia, 2017.
- [2] S. Gröning, J. Hardi, D. Suslov and M. Oschwald. *Journal of Propulsion and Power*, (0):560-573, 2016.
- [3] T. Horchler, S. Fechter, S. Karl, K. Hannemann. 8th European Conference for Aeronautics and Aerospace Sciences (EUCASS), Madrid, 2019.

Numerical Investigation of turbulent heat transfer in a high aspect ratio curved duct

RESEARCH INSTITUTION

¹Chair of Aerodynamics and Fluid Mechanics, Department of Mechanical Engineering, Technical University of Munich

PRINCIPAL INVESTIGATOR

Thomas Kaller¹

RESEARCHERS

N. A. Adams¹, S. Hicel², S. J. Schmidt¹, A. Doehring¹

PROJECT PARTNERS

²Technische Universiteit Delft

SuperMUC Project ID: pr48me

Introduction

The SuperMUC-NG project pr48me is part of the subproject D4 of the SFB/Transregio 40 [1]. The main objective is the investigation of turbulent heat transfer in rocket engine cooling ducts.

A thorough understanding of cooling duct flows is required for an efficient design of structural cooling in various technical applications. Examples include ventilation systems, electrical vehicle battery cooling to rocket engines. The latter are using the carried cryogenic propellant as coolant in a supercritical state. The flow field and the turbulent heat transfer into a cooling duct is significantly influenced by secondary flows and strong non-linear thermodynamic property variations close to the so-called pseudo-boiling line. The secondary flows intensify the mixing of hot and cold fluid and therefore determine the cooling efficiency significantly. They are classified into two categories: Prandtl's flow of the first kind and the weaker turbulence-induced Prandtl's flow of the second kind. The former are reaching strengths of $\approx 20\text{-}30\%$ of the bulk flow velocity and the latter $\approx 1\text{-}3\%$. Both types exhibit a strong influence on the momentum and temperature transport.

To study the individual effects, several reference cooling ducts have been defined and investigated. To study the interaction of turbulence-induced secondary flow, a straight water high-aspect-ratio cooling duct has been investigated in cooperation with the technical university of Braunschweig conducting experiments for that setup. A curved end-section with varying curvature is added to study the interaction of the two types of secondary flow and turbulent heat transfer. Finally, to study turbulent heat transfer under transcritical conditions with strong non-linear property variations induced by intermolecular repulsive forces, a periodic channel setup has been defined and investigated using methane as working fluid.

For our simulations we use the numerically expensive large-eddy simulation (LES) method developed within our in-house LES solver. We further compare the

obtained results to the more feasible Reynolds-averaged Navier-Stokes (RANS) method using the commercial solvers ANSYS CFX and ANSYS FLUENT to show possible deficiencies of the RANS method. The main difference is, that in RANS the governing equations are solved for the averaged flow state and all scales of the turbulence cascade are modelled. In LES individual time samples are produced and the large scale turbulent structures are resolved and only the small ones modelled.

Results and Methods

For the water duct simulations [2] we solve the incompressible Boussinesq equations treating the temperature as an active scalar. Thermodynamic quantities are evaluated using the IAPWS correlations. For the pressure Poisson equation and diffusive fluxes a 2nd order central difference scheme is used. The Poisson equation is solved in each of the 3 Runge-Kutta substeps using a Krylov solver with algebraic multigrid preconditioner. For the transcritical channel simulations the compressible Navier-Stokes equations are solved. Thermodynamic and transport properties are obtained using an adaptive look-up table method based on the REFPROP database. For both we perform an implicit LES using the Adaptive Local Deconvolution Method (ALDM) providing a physically consistent subgrid-scale turbulence model.

The straight water duct simulation setup is shown in fig. 1, although with the bend added. The grid contains ≈ 280 mio cells and the $Re = 110,000$ cases have been simulated on 7,100 cores on SuperMUC Phase 2. The setup including the curvature to study the influence of different kinds of secondary flows at a lower Reynolds number of $Re = 40,000$ requires ≈ 40 mio cells and runs on $\approx 2,500$ cores of SuperMUC-NG. The equally wall-resolved transcritical channel simulations require ≈ 7.5 mio cells and the cases are simulated on SuperMUC-NG.

The results of the water cooling duct are in good agreement with the experimental PIV and PTV results provided from our project partners. The simulations have shown, that the turbulence-induced secondary




Figure 1: Setup for the high-aspect-ratio water duct simulations.

flow weakens significantly along the duct due to heating and associated viscosity decrease. Furthermore it has been shown, that the prevalent assumption of a constant turbulent Prandtl number as often used in the RANS context is not valid for asymmetrically heated high-aspect-ratio ducts. For further details we refer to the publications [2] and [3].

First RANS results for the straight-curved-duct configuration are shown in fig. 2 highlighting the significance of the different secondary flow structures on the heat transfer into the duct. In the straight part only the weak turbulence-induced secondary flow is present represented by two counter-rotating vortices in each duct corner. In the curved section the so-called Dean vortices are dominant superimposing the turbulence-induced vortices. Likewise the results for the transcritical channel showed, that the assumption of a constant turbulent Prandtl number is invalid if strong property variations are present. Furthermore various definitions of the turbulent Prandtl number for the region around the pseudo-boiling line have been investigated and a modified formulation proposed. For detailed results we refer to [3] and [4]. All investigated RANS models available within ANSYS CFX and ANSYS FLUENT showed significant deviations from our wall-resolved LES, see [3] for further details.

Ongoing Research / Outlook

In our on-going research we will continue our work on transcritical channel flows and secondary flow investigation, and finally will include and investigate the effect of wall-roughness. For the secondary flow analysis, we want to further focus on the generation mechanism of the vortices using the proper orthogonal decomposition (POD) method. For the transcritical channel flows it is planned to build a numerical database of several Reynolds numbers and temperature differences to solidify the already observed and published results. Finally, with the growing capabilities and importance of additive manufacturing also in the field of rocket engine cooling ducts, the question of the effects of an increased wall-roughness has to be investigated. The influence of rough walls on the turbulent heat transfer, boiling and secondary flows is investigated using a well-resolved LES with a setup similar as shown in fig. 1, although without the curvature at the end. The working fluid will be liquid nitrogen at a pressure of 5 bar and a temperature of 93 K. These simulations will be accompanied by experimental investigations conducted at the von Karman institute in Belgium.




Figure 2: Temperature increase and secondary flow in the lower duct region for the straight duct at (a) $x=100$ mm and (b) $x=550$ mm, and in the curved section at (c) $\alpha=45^\circ$ and (d) $\alpha=90^\circ$.

References and Links

- [1] <http://www.sfbtr40.de/>.
- [2] T. Kaller, V. Pasquariello, S. Hickel and N. A. Adams, J. Fluid Mech, 860, 258-299, 2019.
- [3] T. Kaller, A. Doehring, S. Hickel, S. J. Schmidt and N. A. Adams, 2020. Assessment of RANS turbulence models for straight cooling ducts: secondary flow and strong property variation effects. In: Future Space-Transport-System Components under High Thermal and Mechanical Loads, Notes on Numerical Fluid Mechanics and Multidisciplinary Design.
- [4] A. Doehring, S. J. Schmidt and N. A. Adams, 2020. Large-eddy simulation of turbulent channel flow at transcritical states. Int J Heat Fluid Flow (submitted).

Exploring Turbulence and Particle Transport Using Massively Parallel Simulations

RESEARCH INSTITUTION

¹Max Planck Institute for Dynamics and Self-Organization, Göttingen

PRINCIPAL INVESTIGATOR

Michael Wilczek¹

RESEARCHERS

Cristian C. Lalescu^{1,2}, Dimitar G. Vlaykov¹, B renger Bramas³, Markus Rampp²

PROJECT PARTNERS

²Max Planck Computing and Data Facility, Garching

³Inria Nancy – Grand-Est, France

SuperMUC Project ID: pr53ru

Introduction

Turbulence is virtually everywhere. It governs our atmosphere, our oceans as well as their interaction. Turbulent flows feature a broad range of dynamically active scales with a continuous transfer of energy between them. Turbulence, therefore, plays an important role for the energy budget of geophysical flows, it triggers precipitation by enhancing droplet growth in clouds and provides the main mixing mechanism for combustion processes, pollutants in the atmosphere, urban environments, and marine micro-organisms in the ocean.

Due to the inherently random character of turbulence, predictive theories of turbulence necessarily have to be of statistical nature. Despite decades of research, developing such theories, which ultimately all modeling applications rely on, remains one of the outstanding scientific challenges.

From a physics perspective, fully developed turbulence constitutes a paradigmatic problem of nonequilibrium statistical mechanics with a large number of strongly interacting degrees of freedom. The challenge of developing a statistical theory of turbulence arises from multi-scale flow structures, which introduce long-range correlations giving rise to complex, scale-dependent statistics. Numerical simulations can provide key insights into the emergence of these structures, their dynamics, and the resulting statistics.

The main goal of the project is to explore the space-time structure of fully developed turbulence with state-of-the-art simulations and, building on computational evidence, develop predictive models.

Methods and Results

For our simulation-driven theoretical investigations, we study highly turbulent flows from two complementary perspectives: The Eulerian frame, i.e. a fixed frame of reference, is well suited to investigate spatial features of turbulence. In contrast, the Lagrangian frame, i.e. following tracer particles, is particularly useful to study spatio-temporal phenomena such as turbulent mixing. We have developed the simulation framework TurTLE (Turbulence Tools: Lagrangian and Eulerian), a flexible pseudo-spectral solver for fluid and turbulence problems implemented in C++ with a hybrid MPI/OpenMP approach. TurTLE allows for efficient tracking of a large class of particles by means of a parallel programming pattern that is easy to adapt and implement. TurTLE scales well up to $O(10^4)$ computing cores for problem sizes up to $4,096^3$ grid points and up to 10^{10} particles. To make TurTLE broadly available, we plan to release it as an open source code.

For a typical simulation of a fully developed turbulent flow, we first equilibrate a simulation to a statistically stationary state at a target Reynolds number R_λ , which is limited by the grid resolution of the simulation. Snapshots of the turbulent field may then be used to




Figure 1: Turbulence is a multi-scale phenomenon. A sequence of zooms into the velocity field (two-dimensional cuts) from one of our simulations with $4,096^3$ grid points reveals significant correlations on all scales. Steep gradients in the velocity field are closely related to small-scale coherent structures such as vortex filaments (volume rendering, right panel). The close connection between coherent structures and multi-scale correlations poses a major challenge in developing a statistical theory of turbulence.




Figure 2: Summary of the main findings of our recent publication [1]: a) When a tracer particle encounters an intense vortex structure, its acceleration components oscillate strongly in time. $\langle a^2 \rangle$, the squared acceleration coarse-grained along the tracer trajectory over a few Kolmogorov time scales τ , remains approximately constant during such events (black line shows $\sqrt{\langle a^2 \rangle}$). b) Categorizing trajectories by α leads to simple statistics: The standardized probability density functions (PDFs) $f(a)$ of acceleration components conditional on α (colored lines) are all close to Gaussian. The unconditional PDF (red line) and a Gaussian distribution (black, dashed line) are plotted for comparison. c) This observation can be used to formulate a comprehensive statistical framework for single-particle statistics in turbulence: Comparison of PDFs derived from our framework (dashed) with DNS data (solid) ($R \approx 350$). From top to bottom: acceleration PDF, velocity increment (i.e. velocity differences across a time lag τ) PDFs, velocity PDF (vertically shifted for clarity).

initialize new DNS aimed at a Lagrangian analysis (e.g. [1,2]) or directly for an Eulerian analysis (e.g. [3]).

For the largest problem size ($4,096^3$) considered in this project, the simulation of an integral time (characteristic time for large-scale flow features) requires on the order of 5 M core hours on SuperMUC-NG. For such simulations, we used up to an entire 512-node island. A single flow field requires 1.65TB of disk space, such that a database spanning several integral time scales requires on the order of 100TB. These large-scale simulations are accompanied by several smaller simulations to further explore the parameter space.

Figure 1 is a visualization of the velocity field from one of our large-scale simulations. It illustrates the multi-scale nature of turbulence: while the velocity shows correlations comparable to the domain size, a considerable degree of small-scale structure in the form of coherent vortices is evident from these visualizations.

In one of the research projects related to our LRZ project, we investigated how spatial features of turbulence influence particle transport. To this end, we established so-called bridging relations, i.e. statistical connections between the Eulerian and Lagrangian frames [4]. Starting from the idea that, essentially, tracer particles sample turbulent fields in space and in time, we developed a theory that expresses Lagrangian statistics as a probabilistic mix of Eulerian statistics. Our findings have implications for a range of problems, including turbulent mixing and dispersion.

In a follow-up project, we investigated the nature of turbulent fluctuations along tracer particle paths [1]. Analyzing millions of particle tracks from large-scale simulations, we made the notable observation that the ensemble of tracer trajectories, based on their acceleration signature, can be decomposed into very simple, close-to-Gaussian sub-ensembles (see Figure 2). This led us to develop a comprehensive theoretical framework for single-particle statistics. For the first time, this work pointed toward the exciting possibility to finally unravel the complex statistics of turbulence in a data-driven theoretical approach, which inspires novel statistical field theories of turbulence.

Ongoing Research / Outlook

These are just two examples from our ongoing work, which illustrate that large-scale computing clusters such as SuperMUC-NG are an indispensable tool for pursuing simulation-driven theoretical research. We are currently continuing this line of research, focusing in particular on the statistical geometry of turbulence. By tracking continuous material lines in turbulence, we gain a better understanding of how stretching, bending, and twisting facilitate turbulent mixing. We are also using the insights of our fundamental turbulence research to address a range of important application-focused questions. For example, we investigate how turbulence enhances the encounter rates of micro-organisms in the oceans, and how it may initiate rain through enhanced droplet collision rates. Extensive large-scale simulations will continue to provide the computational foundation for such investigations.

References

- [1] L. Bentkamp, C. C. Lalescu, and M. Wilczek, Nat. Commun., 10(1):3550, 2019.
- [2] L. A. Leppin and M. Wilczek, Phys. Rev. Lett., 125:224501, 2020.
- [3] D. G. Vlaykov and M. Wilczek, J. Fluid Mech., 861:422–446, 2019.
- [4] C. C. Lalescu and M. Wilczek, New J. Phys., 20(1):013001, 2018.

Rocket engines under the microscope:

Modeling turbulent combustion

RESEARCH INSTITUTION

Department of Space Propulsion, Technical University of Munich

PRINCIPAL INVESTIGATOR

Oskar Haidn

RESEARCHERS

Nikolaos Perakis, Christian Bauer

PROJECT PARTNERS

—

SuperMUC Project ID: pr53we

Introduction

Currently almost all in-space propulsion systems depend on the propellant hydrazine or its derivatives for thrust generation. However, within the REACH framework EU has declared hydrazine a Substance of Very High Concern, due to its toxic and carcinogenic properties. Therefore this essential propellant has to be phased out in the upcoming years. The propellant combination methane / oxygen promises to be a good replacement, offering good performance, storability and handling. Nevertheless, a number of problems regarding high-pressure combustion, ignition, cooling and injection have to be solved before it can be applied in full-scale space propulsion systems. The Department of Space Propulsion (RFA) at the Technical University of Munich investigates all relevant properties of methane and oxygen for space propulsion applications by implementing various experiments. In order to improve experimental design, numerical investigations employing Computational Fluid Dynamics shall supplement the current efforts. It also aids in interpreting results and allows producing numerical models able to predict full-scale applications [1].

Results and methods

Combustion and wall heat transfer modeling

In order to obtain information about the properties of methane combustion and the interaction between the flame and the wall, simplified geometrical setups and canonical configurations are simulated in order to




Figure 1: Instantaneous density gradient and temperature in resonance igniter. The free-stream generated by a supersonic nozzle on the left creates complex flow interactions with the resonator cavity, which drive the oscillations in the cavity.

isolate the physical effects that require additional modeling. The simulation of a Piloted Premixed Jet Burner was carried out using various different fuels (methane, ethylene, n-heptane and toluene) in a 3D domain consisting of 8.1 million cells. The closure of the turbulence-chemistry interaction in the Large Eddy Simulation (LES) was carried out using a thickened flame model, while a skeletal mechanism consisting of 33 species was employed for the chemical kinetics. The difference in the oxidation process of CO for each of those fuels was examined and compared to detailed experimental mole fraction measurements. The investigation of this configuration helped shed light into the development of turbulence in the near-injector region and the influence of fuel effects on the flame length. The work was carried out in a collaboration with Stanford University, University of Southern California and University of California, Los Angeles. Figure 2 shows a 3D view of the CO₂ mass fraction in the computational domain as well as an iso-surface of the chemical heat release. Apart from the combustion dynamics, a further aspect requiring special attention in numerical simulations of rocket engines is the interaction between the flame and the wall. Chemical reactions occurring in the turbulent boundary layer, close to the wall can lead to a direct increase of the heat loads and hence precise prediction of their magnitude is required. A Direct Numerical Simulation (DNS) of a reacting boundary layer flow over an isothermal flat plate was employed in order to better understand the origin of the aforementioned exothermic reactions. An equilibrated, pre-mixed gas was injected at the domain inlet and the progress of the chemical reactions was examined. Approximately 26 million cells were used in the 3D domain and a chemical mechanism with 34 species was employed. The main reaction path responsible for the additional heat release was identified in this study: the oxidation of CO to form CO₂ at low temperatures. The instantaneous fields of CO and CO₂ mass fractions as well as the corresponding temperature field are plotted in Fig. 3, clearly indicating an increase in CO₂ concentration.

Sub-scale chamber simulations

The fact that a large number of physical phenomena takes place simultaneously within a thrust chamber, including evaporation, turbulent mixing, heat release,




Figure 2: Instantaneous CO_2 mass fraction along the radial and axial direction and iso-surface of the heat release for the piloted premixed jet burner using methane as fuel.




Figure 3: Instantaneous CO_2 and CO mass fraction as well as temperature field for the flat plate DNS.

supersonic expansion, flame/flame and flame/wall interaction renders numerical simulations of even sub-scale chambers computationally challenging. In an effort to accurately describe the chemical reactions within the chamber while keeping the required resources to a minimum, tabulated chemistry models are typically used, like the flamelet model. RANS and LES simulations of a sub-scale chamber operated with gaseous methane and gaseous oxygen consisting of 7 injector elements have been carried out using the resources of the SuperMUC-NG [3]. In both simulations, a strong interaction between the individual flames was observed, which directly influenced the circumferential distribution of wall heat loads. The resulting temperature field from the RANS as well as an instantaneous snapshot from the LES are shown in Fig. 4.

Resonance ignition

In-space propulsion systems require extreme reliability and to that end achieving reliable ignition is absolutely essential. While the combination methane/oxygen has a number of advantages compared to traditional propellants, it is a bi-propellant combination and hence requires a dedicated ignition system. Moreover, the slow chemical kinetics and the high ignition temperature of methane create challenges, that are not easy to overcome with classical ignition systems. For that reason RFA carries out research on resonance igniters, which could provide passive, simple and lightweight ignition sources for many space propulsion applications. In these devices pressurized propellants are expanded through a nozzle and excite violent oscillations in an attached resonator cavity,




Figure 4: RANS temperature field and instantaneous LES temperature in the 7-element rocket combustor.

where irreversible effects heat the propellants beyond their auto-ignition temperature, leading to self-ignition. Previous studies showed, that the commercial CFD solver ANSYS Fluent can reproduce these effects quantitatively. Current investigations focus on how to improve the preparation time of these igniters by identifying heat sources and sinks in the system. It could be shown, that in current designs most of the heat generated by irreversible effects is lost to the solid walls, which means that the heat-up time is dominated by the thermal timescales in the solid material [2]. This insight suggests new ways to optimize the configuration, which shall be developed and investigated in further works. 2D URANS simulations are performed while the low-diffusion QUICK scheme is used for spatial discretization and the bounded 2nd order implicit formulation is applied for time advancement. Fig. 1 shows the instantaneous flow-field within the resonator.

Ongoing research / Outlook

The use of the HPC resources provided by SuperMUC-NG allowed for the investigation of the considered configurations using a multitude of simulations to conduct parametric studies, helping to determine the main processes driving the design parameters. It also allowed for the simulation of more complex combustor configurations which aim to be scalable to flight hardware and therefore are a good basis for tool validation. In the topic of numerical modeling, the extension of tabulated chemistry models is planned as well as their application to fullscale hardware operating with methane/oxygen at realistic pressure and mixing conditions.

References

- [1] SFB Transregio TRR40, <https://www.sfbtr40.de/en/>, 2020.
- [2] C. Bauer et al. Numerical investigation of a resonance ignition system. In 8th EUCASS, 2019.
- [3] N. Perakis et al., Journal of Propulsion and Power, 35(6):1080–1097, 2019.

Space launch vehicle aerodynamics with hot plumes

RESEARCH INSTITUTION

Spacecraft Department, Institute of Aerodynamics and Flow Technology, German Aerospace Center (DLR)

PRINCIPAL INVESTIGATOR

Klaus Hannemann

RESEARCHERS

Jan-Erik Schumann, Volker Hannemann

PROJECT PARTNERS

—

SuperMUC Project ID: pr62po

Introduction

The aerodynamics around space launch vehicles determine the external forces that these vehicles experience. Hence, the accurate prediction of these loads is crucial for the structural design. Additionally, a better understanding of the fundamental physical phenomena responsible for the observed flow features helps to design the vehicle such that the experienced loads are reduced. Both the better understanding of the loads and the ability to reduce these loads allow for vehicles to be designed more efficiently and carry more payload mass into orbit. Of particular interest for the aerodynamic research is the region at the bottom of the vehicle, displayed in Fig. 1, where a sudden change in diameter at the end of the main body can lead to high mechanical loads. This aspect of space launch vehicles is investigated in branch B of the SFB Transregio 40 [1], funded by the German Research Foundation DFG. Due to its simpler handling both numerically and experimentally, nearly all previous investigations in the literature have investigated the associated phenomena without a propulsive plume or with one resulting from expanding air. Consequently the plume properties differ significantly from those of realistic rocket plumes. Sub-project B5 of Transregio 40, which is handled by the Spacecraft Department of the Institute of Aerodynamics and Flow Technology at the German Aerospace Center DLR, is concerned in particular with the effects that the presence of hot propulsive plumes has on the wake flow field phenomena [2]. For this purpose computational fluid

dynamics (CFD) is applied to simulate the flow around a scale-model generic space launch vehicle geometry. The advantage of CFD is the ability to analyze all regions and aspects of the flow field as required with a high level of detail. To ensure the numerical algorithm accurately predicts the real flow physics the results obtained are compared to experimental investigations of key quantities such as wall pressure measurements. In the scope of the project pr62po at LRZ, a sensitivity study based on a configuration with available comparison data in the literature was conducted first. Subsequently, it was investigated which impact changing the plume conditions and wall temperatures has on the flow field and associated vehicle loads.

Methods and Results

The CFD solver used for these simulations is the DLR TAU-Code that solves the compressible Navier-Stokes equations. TAU is a second-order accurate Finite Volume solver programmed in C and includes capabilities to handle multispecies flow and chemical reactions [3]. These capabilities are critical to simulate the flow field in the presence of hot plumes as these result from combustion processes. The flow field turbulence is modelled using a Hybrid RANS-LES method (HRLM) which allows particular regions of interest to be resolved with high accuracy Large Eddy Simulation (LES) in a time-resolved fashion whereas other regions are treated with a lower fidelity Reynolds-Averaged Navier-Stokes (RANS) approach.

The conducted investigations are the first with TAU using high-fidelity HRLM in combination with multispecies flow and finite-rate chemical reactions. In the process, several optimizations were implemented to accelerate, stabilize and improve the simulations. Additionally, optimal parameter settings in terms of computational effort and accuracy for these kinds of investigations were found. The required computational resources depend heavily on the settings used as well as the particular configuration investigated. The least resources were required for a grid in the sensitivity study




Figure 1: Visualization of the turbulent structures colored by temperature in the base region of a generic space launch vehicle.

(ca. 13 million grid points) with 200,000 core-h distributed on 1,120 cores (on SuperMUC Phase 2). The most resources were required for simulations on a fine grid (ca. 31 million grid points) with around 4 million core-h on 9,216 cores. These were more computationally expensive since multispecies flow and chemical reactions in addition to the high numerical resolution were required. Overall, about 20 different cases were simulated since the start of the project totaling about 26 million core-h. Due to the need of high time resolution the initial output of one large simulation on the SCRATCH file system is about 250TB, but this can be reduced to about 2TB per case of final data by removing quantities only required for restarts.

Scientific results

The scientific results obtained in the project include a better understanding both of the numerical methods and their requirements and sensitivities [4] as well as insights into the flow physics with different configuration changes [2]. For brevity, only the most significant of the latter are discussed below. The flow field around the investigated geometry can be found in Fig. 2. The figure shows the mean flow field with streamlines and axial velocity color contours on the top and an instantaneous view of the circumferential vorticity on the bottom for configurations with different plume conditions (top vs. center) and wall temperatures (center vs. bottom). It is visible that for a low velocity air plume the turbulent shear layer originating from the end of the main body impacts the end of the nozzle fairing, but an increase in plume velocity shifts the reattachment location further downstream onto the plume. This is attributed to a plume suction effect and increases the interaction between the plume and the external flow. Furthermore, a similar shift in reattachment location is observed if the wall temperature is increased. This is associated with the reduction in density - and hence momentum - due to the heating of the flow. Additionally, it is found that the increased wall temperatures significantly reduce the mechanical loads experienced by the nozzle.

However, it is also found that the fundamental unsteady flow features observed for configurations with an air plume are also present in those with realistic plume conditions and higher wall temperatures. This indicates that the underlying flow phenomena dominating the flow field remain unaffected by these parameters. In combination, this shows that for a general understanding of the flow field an exact replication of the full scale conditions is not necessarily required, but for the quantitative analysis of the mechanical loads accurate plume conditions and wall temperature descriptions are necessary since the detailed quantitative impact these phenomena have on the mechanical loads differs.

Ongoing Research / Outlook

Currently, the final evaluation and analysis of the results is still in progress. Additionally, remaining simulations are being conducted that tackle open

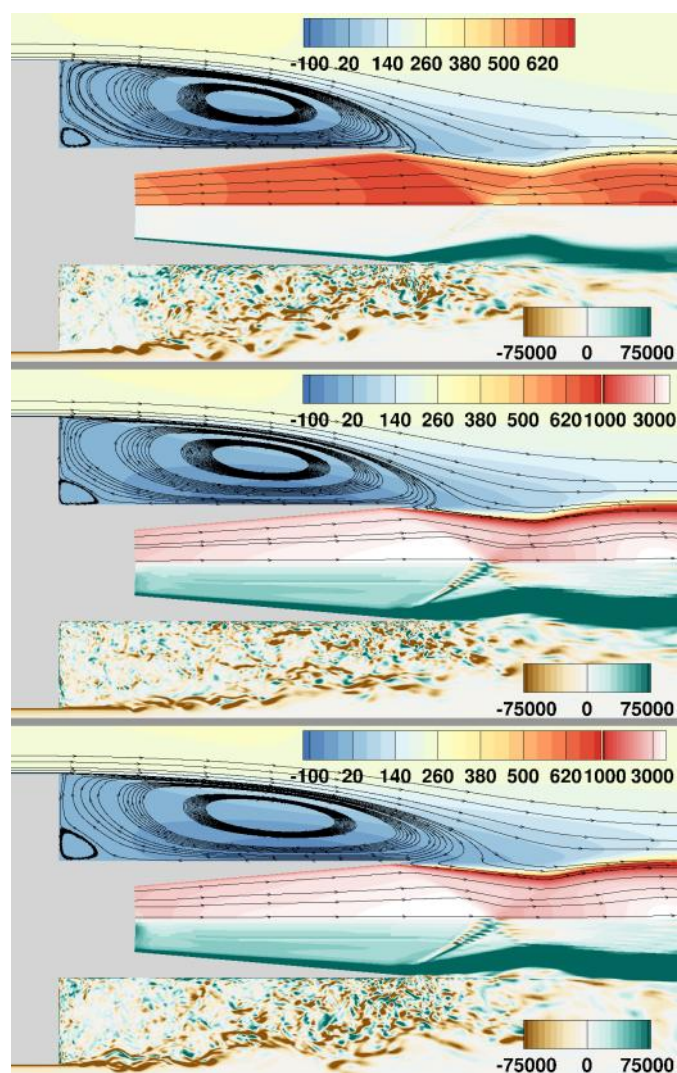


Figure 2: Comparison of the flow field with an air plume (top), a hot plume with cold walls (center) and a hot plume with hot walls (bottom). Mean axial velocity (in m/s, top) and instantaneous circumferential vorticity (in 1/s, bottom) are displayed.

questions that were raised in the process of the investigation. While the simulations could likely have been conducted on SuperMUC Phase 2 as well, the size of SuperMUC-NG helps to significantly speed up the computations by reducing queuing time. Additionally, the required number of jobs per case is reduced significantly by the improved parallel performance among other factors due to the larger number of cores per node and hence reduced cross-node-communication.

References and Links

- [1] <https://www.sfbtr40.de/>
- [2] Schumann, J.-E., Fertig, M., Hannemann, V., Eggers, T. and Hannemann, K., "Numerical Investigation of Space Launch Vehicle Base Flows with Hot Plumes", In: Future Space-Transport-System Components under High Thermal and Mechanical Loads, Springer, 2021.
- [3] Hannemann, K., Martinez-Schramm, J., Wagner, A., Karl, S., Hannemann, V., "A Closely Coupled Experimental and Numerical Approach for Hypersonic and High Enthalpy Flow Investigations Utilising the HEG Shock Tunnel and the DLR TAU Code", Technical Report, German Aerospace Center, Institute of Aerodynamics and Flow Technology, 2010.
- [4] Schumann, J.-E., Hannemann, V. and Hannemann, K., "Investigation of Structured and Unstructured Grid Topology and Resolution Dependence for Scale-Resolving Simulations of Axisymmetric Detaching-Reattaching Shear Layers", Progress in Hybrid RANS-LES Modelling, Springer, Cham, 2020, 169-179.

Turbulent convection at very small Prandtl numbers

RESEARCH INSTITUTION

¹Technische Universität Ilmenau

PRINCIPAL INVESTIGATOR

Jörg Schumacher¹

RESEARCHERS

Dmitry Krasnov¹, Amrish Pandey²

PROJECT PARTNERS

²New York University Abu Dhabi

SuperMUC Project ID: pn68ni, pr62se (Gauss Large Scale project)

Introduction

Turbulent convection is one essential process to transport heat across a fluid layer or closed domain. In many of the astrophysical or technological applications of convection the working fluid is characterized by a very low dimensionless Prandtl number $Pr = \nu/\kappa$ which relates the kinematic viscosity to temperature diffusivity. Two important cases are (i) turbulent convection in the outer shell of the Sun at $Pr \sim 10^{-6}$ in the presence of rotation, radiation, magnetic fields, and even changes of the chemical composition close to the surface [1,2] and (ii) turbulent heat transfer processes in the cooling blankets of nuclear fusion reactors at $Pr \sim 10^{-2}$ [3]. These are rectangular ducts which are exposed to very strong magnetic fields that keep the 100 million Kelvin hot plasma confined. Our understanding of the complex interaction of turbulence with the other physical processes in these two examples is still incomplete. High-resolution direct numerical simulations of the equations of turbulent fluid motion in the simplest setting of a turbulent convection flow, Rayleigh-Bénard convection in a layer or a straight duct that is uniformly heated from below and cooled from above, help to reveal some of these aspects at a reduced physical complexity and to

discuss the basic heat transfer mechanisms that have many of these applications in common. Such studies have to rely on massively parallel supercomputers.

Results and Methods

We solve the three-dimensional Boussinesq equations of thermal convection. They couple the turbulent velocity and temperature fields. The external magnetic field in the nuclear fusion application (ii) is typically very strong such that we can apply the quasi-static limit of magnetohydrodynamics [3]. The vigor of convective turbulence is quantified by a further dimensionless parameter, the Rayleigh number Ra , and the strength of the applied external magnetic field by the Hartmann number Ha . Turbulent flows at very low Prandtl numbers are known to cause highly inertial fluid turbulence which makes our numerical simulations very challenging since all turbulent vortices down to the smallest ones have to be resolved. We apply a second-order finite difference method [3]. The simulation domains are cuboid cells or ducts with no-slip boundary conditions at all walls. The sidewalls are thermally insulated. The numerical simulations at $Pr = 10^{-3}$ in domains of aspect ratio 25:25:1 require 38,400 SuperMUC-NG cores for a grid with 12,800x12,800x800 points in the non-magnetic case (i). The simulations at a Hartmann number $Ha = 10^3$ require up to 7,680 cores for a long duct with 15,360x1,280x384 grid points. All simulations are long-term runs that involved sequences of several 48-hour runs in a row. In the course of two project years, this sums up to 80 million core hours which will eventually be consumed for the Large Scale Projects pr62se and pn68ni.

Figure 1 illustrates the impact of a strong magnetic field on the turbulent mixing and thus the heat transfer properties in a rectangular duct flow (case ii). The flow enters the duct as a planar jet which is immediately rotated into the direction of the external magnetic field (along blue axis) and develops quasi-2d vortical structures downstream (along the red axis) that can mix the liquid metal coolant effectively. Also visible are so-called Shercliff layers at the front and back face that become unstable and contribute additionally to the turbulent transport.




Figure 1: Planar liquid metal jet flow inside the blanket duct in presence of homogeneous magnetic field along blue axis at $Ha=1,000$. Iso-surfaces of the streamwise velocity are shown.




Figure 2: Turbulent Rayleigh-Bénard convection at very low Prandtl number. The aspect ratio of the layer is $25H:25H:H$ with H being the height of the layer. Left: Top view on the instantaneous velocity streamlines. Right: Zooms of the corresponding temperature (top), kinetic energy dissipation rate (middle) and turbulent kinetic energy (bottom), all close to the top surface. Their size is $2H:2H:H$ taken at the black box in the center of the left panel.

Figure 2 (left) is a view from the top onto the streamlines in the whole turbulent Rayleigh-Bénard convection layer for $Ra = 10^6$ and $Pr = 10^{-3}$, one of our biggest simulation runs so far. Clearly, these parameters are by many orders of magnitude away from any realistic dynamics of solar convection. Nevertheless, they provide us already new and interesting insights into the low-Prandtl-number turbulence. We observe how the convection is organized in large patterns of circulation rolls that fill the whole layer and extend from the top to the bottom. Further turbulence fields are replotted in the right column of this figure. Characteristic for low-Prandtl-number convection is a very coarse-grained temperature field which is obvious from the blurred contours that are visible. Gradients of temperature will be washed out quickly due to the large diffusivity which implies that heat transport is very inefficient and close to the diffusive lower bound. In contrast, we observe many fine filaments for quantities that probe the small-scale structure of the corresponding fluid turbulence, such as the kinetic energy dissipation rate field or the turbulent kinetic energy. Both fields suggest a large Reynolds number flow and thus a strong momentum transfer. As the fine-scale features of the convective turbulence cannot be resolved in more complex numerical simulations of bigger domains, our present data records provide an ideal testing bed to calculate turbulent eddy viscosities and diffusivities which can

be used to close the underlying equations of motion in reduced models that describe the large-scale features of turbulent convection only. These investigations are currently underway.

Ongoing Research / Outlook

As demonstrated, our numerical investigations require access to the most powerful supercomputers. In both discussed examples, we were thus able to study turbulent convection in parameter ranges that are not accessible in laboratory experiments. For example, the smallest possible Prandtl number in a laboratory experiment is $Pr \approx 0.005$ for liquid sodium. The SuperMUC-NG computer made furthermore long-term simulations possible that resolved the evolution of the large-scale patterns. An important point of the future work will be to extend the complexity of the studies in both cases to more realistic temperature-dependent material parameters $\nu(T)$ and $\kappa(T)$ a step which has been started within the present Gauss Large Scale Project pn68ni [4].

References and Links

- [1] <https://www.tu-ilmenau.de/tsm>
- [2] J. Schumacher and K. R. Sreenivasan, Rev. Mod. Phys. 92, 041001 (2020).
- [3] I. Belyaev, D. Krasnov et al., Phys. Fluids 32, 094106 (2020).
- [4] A. Pandey, J. Schumacher and K. R. Sreenivasan, Astrophys. J., in press (2021).

The Largest Scales in Turbulent Pipe Flow

RESEARCH INSTITUTION

Institute of Aerodynamics and Flow Technology, German Aerospace Center (DLR)

PRINCIPAL INVESTIGATOR

Christian Bauer

RESEARCHERS

Christian Bauer, Claus Wagner

PROJECT PARTNERS

—

SuperMUC Project ID: pr62zu

Introduction

A large amount of the energy needed to push fluids through pipes worldwide is dissipated by viscous turbulence in the vicinity of solid walls. Therefore the study of wall-bounded turbulent flows is not only of theoretical interest but also of practical importance for many engineering applications. In wall-bounded turbulence the energy of the turbulent fluctuations is distributed among different scales. The largest energetic scales are denoted as superstructures or very-large-scale motions (VLSMs). In our project we carry out direct numerical simulations (DNSs) of turbulent pipe flow aiming at the understanding of the energy exchange between VLSMs and the small-scale coherent structures [1]. While the near-wall small-scale structures scale in viscous units, the outer flow VLSMs scale in bulk units. Hence the range of scales increases as the Reynolds number of the flow increases. In order to study the interaction between these structures, we carried out DNSs of friction Reynolds numbers up to $Re_\tau=2,880$, where $Re_\tau=u_\tau R/\nu$ is based on the friction velocity, the pipe radius and the kinematic viscosity. Besides a large Reynolds number, required for large scale separation, a sufficiently long computational domain is needed for VLSMs to settle. In a preliminary study the required computational domain length was estimated to $L=42R$ [2].

Results and Methods

Our numerical method consists of a fourth-order finite volume DNS code, which is parallelised by means of the message passing interface (MPI). The number of required cores ranges from 64 for the smallest case ($Re_\tau = 180$) up to 2,048 for the largest case ($Re_\tau = 2,880$). The heart of our Fortran90 simulation code is the Poisson solver, which solves a three-dimensional elliptic equation. Taking advantage of the homogeneity of the problem in two directions, Fast Fourier Transforms are performed in axial (z) and azimuthal (φ) direction, before $N_z N_\varphi$ one-dimensional problems are solved by a fast direct tridiagonal matrix solver. Besides flow statistics, which are accumulated

on-the-fly, a number of instantaneous flow field realisations is written out by the code in the netCDF format. With the problem sizes consisting of up to 32.5 billion finite-volume cells, the flow field snapshots consume most of the project memory of 40TB. Overall, our computations required 10 million core hours. From the five different Reynolds number simulations contributing to our analysis the two with the largest computational requirements were carried out on SuperMUC. An overview of these simulations is given in Table 1.

Resulting instantaneous streamwise velocity fluctuations for $Re_\tau=1,500$ are depicted as iso-volumes in Fig. 1. As their small-scale counterparts, the VLSMs visible in the instantaneous picture occur alternately as high- and low-speed streaky structures. In azimuthal direction three low- and high-speed structures are clearly visible. The average extension of VLSMs can be extracted from velocity correlations, as shown in Fig. 2. With a threshold of 0.1, the average streamwise length of VLSMs is measured as approximately $7R$. In addition, the velocity correlation shows an inclination towards the wall, a feature that is characteristic for both small-scale motions and VLSMs. We reported on the scaling and convergence of high-order statistical moments and contributions from VLSMs on these moments [2]. Particularly very large local wall-normal velocity fluctuations in the vicinity of the wall—so called velocity spikes—are modulated by outer flow VLSMs, which reflects in the wall-normal velocity flatness. Regarding the origin of the kinetic energy of VLSMs, we analysed the turbulent kinetic energy transport equation of the low-pass filtered velocity field, the latter basically consisting of VLSMs. By comparing our results with what is known from the small-scale near-wall cycle, we found that VLSMs are fed with energy by the mean velocity field via a turbulent production mechanism similar to their small-scale counterparts [4]. Regarding

Case	Re_τ	Grid ($N_z \times N_\varphi \times N_r$)	Cores
P1500	1,500	8,192 x 2,048 x 408	1,024
P2880	2,880	12,288 x 4,096 x 646	2,048

Table 1: Simulation cases on SuperMUC.




Figure 1: Very-large-scale motions in turbulent pipe flow at friction Reynolds number of $Re_\tau = 1,500$, visible in the shape of iso-volumes of the streamwise velocity fluctuation. Low-speed structures in blue, high-speed structures in red.

the transfer of energy between different scales, however, VLSMs behave strictly opposite to the small-scale motions. While for the last mentioned, the forward energy cascade—energy transfer from larger to smaller scales—correlates with low-speed structures and the backward transfer correlates with high-speed structures, for VLSMs the opposite is the case.

Ongoing Research / Outlook

Answering the question why the inter-scale energy transfer towards and away from VLSMs is so strikingly different from the energy transfer related to small-scale motions requires further investigations of turbulent pipe flow. Moreover, most of our analyses regarding the interactions between outer flow VLSMs and near-wall small-scale motions involved turbulent pipe flow DNSs at $Re_\tau=1,500$. In terms of real-life applications this Reynolds number is still low and VLSMs are known to become more dominant with increasing Reynolds numbers. Consequently we already started

simulations at $Re_\tau=2,880$ during the current project. Statistics obtained from these simulations are, however, not fully converged. Therefore, we plan to continue these simulations in a follow-up project. Moreover, simulations at even higher Reynolds numbers are highly desirable to obtain better comparability with real-life engineering applications. SuperMUC provides both the computational power and the memory required to carry out such large-scale DNSs of turbulent pipe flow.

References and Links

- [1] <http://scart.dlr.de/site/research-projects/direct-numerical-simulation-and-large-eddy-simulation-of-turbulent-pipe-and-channel-flows/index.htm>.
- [2] D. Feldmann, C. Bauer, and C. Wagner. J Turbul 19, 274-295 (2018).
- [3] C. Bauer, D. Feldmann, and C. Wagner. Phys Fluids 29, 125105 (2017).
- [4] C. Bauer, A. von Kameke, and C. Wagner. Phys Rev Fluids 4, 064607 (2019).




Figure 2: Iso-surfaces of the two-point velocity correlation of the streamwise velocity component R_{zz} . Red(blue) iso-surfaces exhibit values of ± 0.1 . $Re_\tau=1,500$.

LES of Rocket Combustion Applications

Under Real-Gas Conditions

RESEARCH INSTITUTION

Institute for Thermodynamics, Bundeswehr University Munich

PRINCIPAL INVESTIGATOR

Michael Pfitzner

RESEARCHERS

Julian Zips, Christoph Traxinger

PROJECT PARTNERS

—

SuperMUC Project ID: pr63ze

Introduction

Injection and mixing processes at elevated pressures play a key role in modern rocket combustion chambers. Typically, the pressure exceeds the critical pressures of fuel and oxidizer. In this high-pressure environment, molecular interactions significantly affect the fluid properties, and hence the combustion process. Consequently, the widely-applied ideal gas assumption is not suitable for the description of the fluid properties. Therefore, high fidelity equations of state have to be applied and common single-phase assumptions might not be valid due to the highly nonlinear mixing of the real fluids. In recent years, the reusability of future liquid propellant rocket engines (LREs) has become a major goal. Hereby, the most promising propellant combination is methane (CH_4) and liquid oxygen (LOx). In the context of the development process, the prediction of thermal heat loads and stresses is of major importance and interest. At the Institute for Thermodynamics at the Bundeswehr University Munich, the combustion process in LREs is investigated by means of Large-Eddy Simulations (LESs) within the research framework of the SFB TRR 40 [1].

For the conduction of the LESs a pressure-based version of the C++ toolbox OpenFOAM [2] is used with inhouse modifications concerning the thermodynamics and combustion modeling. Real-gas effects are taken into account by applying a framework based

on the cubic equations of state. For investigating multicomponent phase separation processes, a Gibbs energy minimization method together with a suitable flashing approach is employed. To model the high-pressure combustion process, different approaches are available, ranging from detailed and expensive ones like Eulerian stochastic fields (ESF) to cheaper and more common approaches like the flamelet method. The latter has been recently extended for the non-adiabatic combustion case, which enables the efficient and thorough investigation of wall heat losses in LREs.

Results

Single-Phase Instability in Non-Premixed Flames

To study the possibility of phase separation processes under initially supercritical conditions, a high-pressure LOx/ CH_4 combustion case was used. The instantaneous flow field and the comparison of the mean OH^* radiation with experimental data are shown in Fig. 1. The blue isosurface in Fig. 1 (left) indicates the region of single-phase instability in the proximity of the cryogenic oxygen core. The phase separation process is triggered by the presence of water and the moderate temperatures in this region of the flame. The comparison of the average OH^* -radiation in Fig. 1 (right) shows very good agreement with the experimental data in terms of the axial and radial extent. For more details see [3].

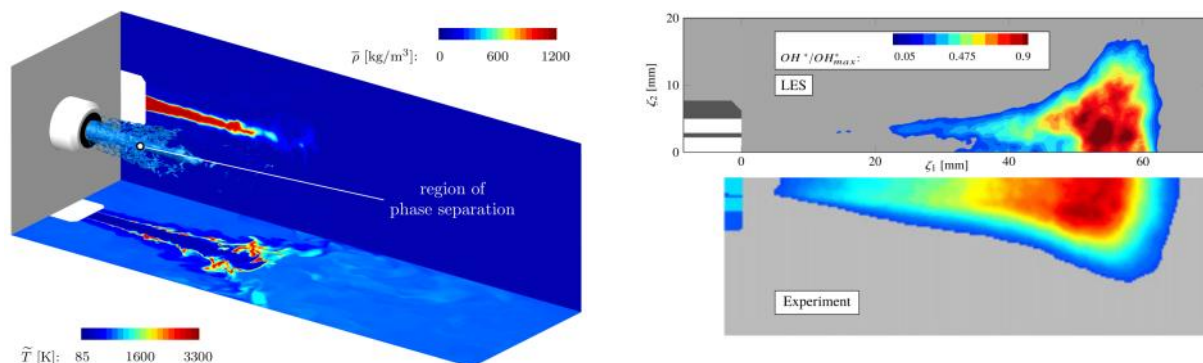


Figure 1: Large-Eddy Simulation results of the high-pressure LOx/ CH_4 combustion test case. Left: Instantaneous temperature (bottom) and density (back) fields. The isosurface shows the region of single-phase instability. Right: Comparison of the average OH^* radiation with available experimental data.

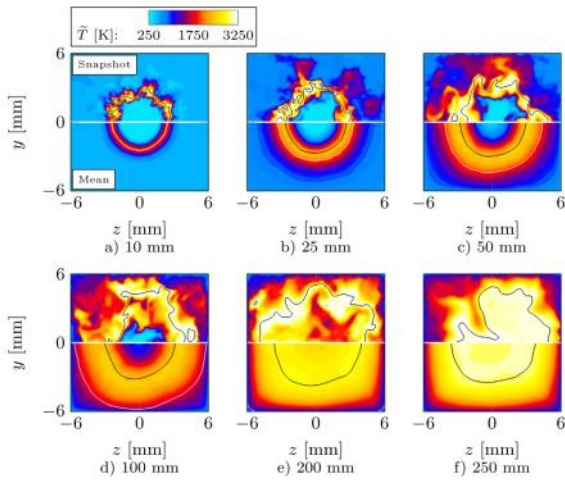


Figure 2: Large-Eddy Simulation results of the single-element combustion test case.




Figure 4: Large-Eddy Simulation results of the 7-element combustion test case.

Heat Losses in Single-Element Combustor

In the single-element case, we focused on the influence of the near-wall treatment. Three different approaches were applied, namely a wall-resolved (WR) and a wall-modeled (WM) LES and a hybrid RANS-LES model called IDDES. As a first impression, the flow field at six different axial positions is shown in Fig. 2. The quadratic chamber cross section clearly influences the axial development of the flame.

Figure 3 shows the heat flux in comparison to the experimental data taken from the literature. Overall, a good agreement between the experiment and the simulations is found. Especially the two wall-modeled approaches, WM-LES and IDDES, do a decent job. The WR-LES slightly underpredicts the wall heat flux. For more details see [4].

Heat Losses in Multi-Element Combustor

The 7-element test case was used to compare three different combustion models namely two presumed PDF approaches and one transported PDF approach. In Fig. 4, the flow field in the front part of the combustor is shown, indicating the gradual flame-flame interaction with increasing axial direction, and the highly 3D character of the flow field. Figure 5 shows the comparison of the predicted wall heat flux and the reference data from the experiment. The ESF method gives the best results, followed by the




Figure 3: Wall heat flux for the single-element case.

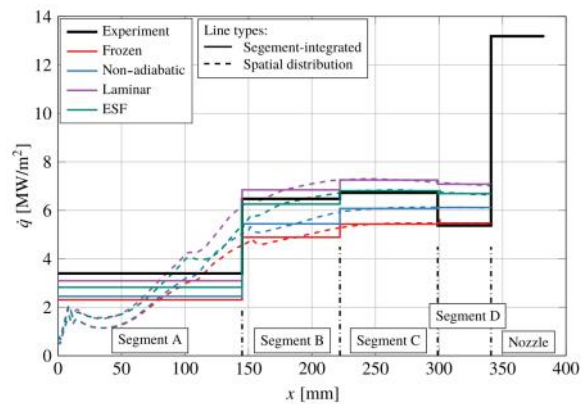


Figure 5: Wall heat flux for the 7-element case.

nonadiabatic and the frozen flamelet model which indicates the necessity of high fidelity combustion modeling for the reliable prediction of heat losses in high-pressure combustion. For more details see [5].

Ongoing Research / Outlook

Based on the successful investigation and validation of these three combustion cases, the following work packages are planned in the future: (1) Introduction of real-gas effects into the tabulated combustion models. (2) Effect of the injector geometry, e.g., recess and tapering. (3) Study of flame-flame interaction in full-scale applications. (4) Study of flame-wall interaction in full-scale applications.

References and Links

- [1] www.sfbtr40.de.
- [2] www.openfoam.org.
- [3] C. Traxinger, J. Zips, M. Pfitzner, J. Propuls. Power, 35(4), 675–689 (2019).
- [4] J. Zips, C. Traxinger, M. Pfitzner, Int. J. Heat Mass Transf., 143, 118474 (2019).
- [5] J. Zips, C. Traxinger, P. Breda, M. Pfitzner, J. Propuls. Power, 35(4), 747–764 (2019).

Simulations at very high Rayleigh number

RESEARCH INSTITUTION

¹Max Planck Institute for Dynamics and Self-Organization, Göttingen, Germany

PRINCIPAL INVESTIGATORS

Richard Stevens¹, Roberto Verzicco¹, Detlef Lohse²

RESEARCHERS

—

PROJECT PARTNERS

²Max Planck Center Twente for Complex Fluid Dynamics and Physics of Fluids Group, University of Twente, The Netherlands

SuperMUC Project ID: pr74sa (Gauss Large Scale project)

Introduction

Rayleigh-Bénard convection [2], i.e. the flow in a box heated from below and cooled from above, is one of the paradigmatic systems in fluid dynamics. The system is used to test new concepts in the field, such as instabilities, non-linear dynamics, and chaos, pattern formation, or turbulence. Rayleigh-Bénard convection is a relevant model for countless phenomena ranging from thermal convection in the atmosphere, oceans, and the outer layer of the Sun, to heating and ventilation of buildings, and convection in various industrial applications. Thus the problem is of interest in a wide range of sciences, including geology, oceanography, climatology, and astrophysics. The Rayleigh-Bénard system is ideal for studying the interaction between the boundary layer and bulk dynamics, which will also shed more light on general wall-bounded turbulent flows. Meanwhile, for 'not too strong' driving, there is a reasonable understanding of turbulent RB convection. Here, 'not too strong' means that the boundary layers still have laminar-type characteristics

and scalings. In this regime, the unifying theory describes the global transfer properties of the flow, i.e. how the Nusselt Nu number (dimensionless heat transport) and the Reynolds Re number (dimensionless flow strength) depend on the control parameters Rayleigh Ra (dimensionless temperature difference between the plates) and Prandtl Pr number (the ratio of momentum diffusivity to thermal diffusivity). However, the situation is less clear for higher Ra . For sufficiently high Ra , Kraichnan predicted in 1962 the existence of an 'ultimate' regime, where not only the bulk but also the boundary layers become turbulent. This happens when the shear Re number becomes around a few hundred and results in a strongly enhanced heat transfer.

Results and Methods

Both in experiments and simulations of Rayleigh-Bénard convection, it is an enormous challenge to reach the ultimate regime in which the boundary layers transition from laminar to turbulent. In the ultimate regime, the scaling exponent γ in the relation $Nu \sim Ra^\gamma$, increases. The critical Rayleigh number (Ra^*) for the transition to the ultimate regime has been observed around $Ra^* \approx 2 \times 10^{13}$ in the Göttingen experiments led by Prof. Bodenschatz and Prof. Ahlers. The highest Ra number obtained in direct numerical simulations (DNS) for aspect ratio $\Gamma = 0.5$ is $Ra = 2 \times 10^{12}$ [3]. Here we refer the reader to the invited general audience article by Detlef Lohse on "Turbulenz im ultimativen Regime" [2] on receiving the Max Planck Medal. In that article, some preliminary results of this project are also discussed. In this project, we perform novel direct numerical simulations in an attempt to reproduce the ultimate regime in computer simulations for the first time. We have now achieved to perform simulations up to $Ra = 10^{13}$, matching the experimental geometry of the Göttingen experiments. Figures 1 and 2 show snapshots of the temperature field from these simulations. To ensure that all turbulent length and time scales are

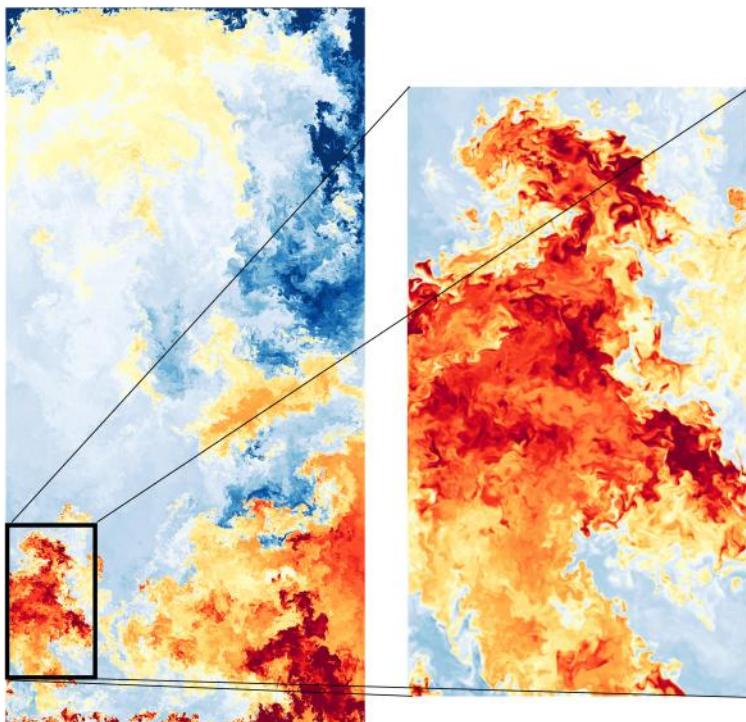


Figure 1: Visualization of Rayleigh-Bénard convection in a $\Gamma=1/2$ convection cell at $Ra=10^{13}$ for $Pr=0.7$, which is close to the onset of the ultimate regime in the corresponding Göttingen measurements. Figure 2 shows a corresponding snapshot at the boundary layer height, which is at about 0.02% of the cylinder's domain height.

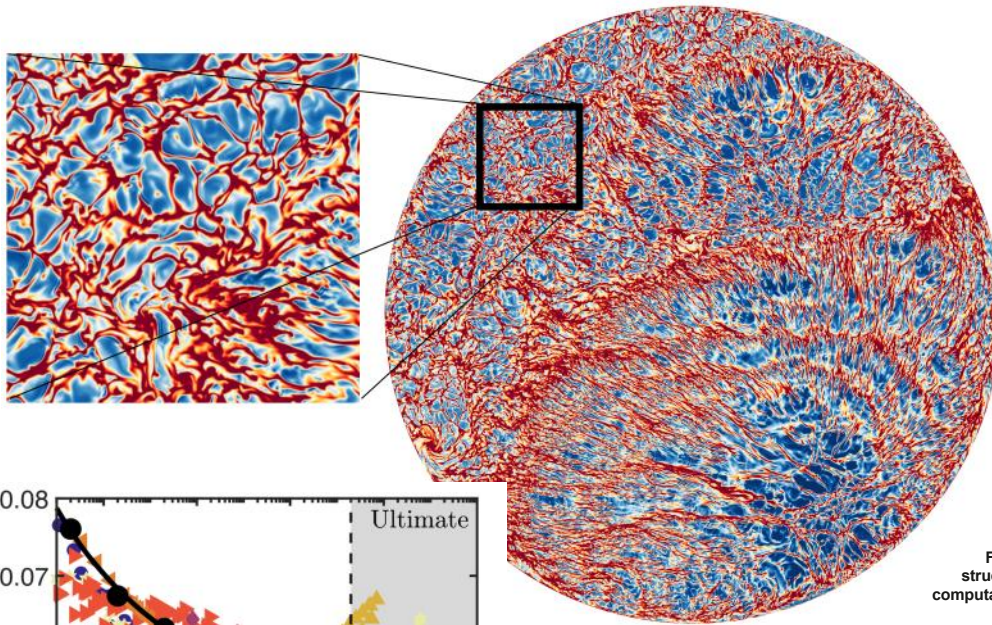


Figure 2: Snapshots of the temperature field at the boundary layer height just above the hot bottom plate for the case shown in Figure 1. One can see tiny plume structures, which require massive computational grids.




Figure 3: Compilation of experimental and numerical data for the compensated Nusselt versus Rayleigh numbers: the solid black line is the unifying theory by Grossmann and Lohse. The different color symbols indicate experimental results. The yellow down triangles are the results from the Göttingen experiments showing the transitions to the ultimate regime. Black bullets previous numerical simulations [3], black squares preliminary simulations performed in this project. For further details, see Ref [5].

properly resolved, a grid with almost 100×10^9 grid points is employed. The simulations require up to 37,000 cores on SuperMUC-NG and tens of millions of Core-h. A snapshot of the entire flow field requires up to about 5TB, while the entire database we generated is over one petabyte. The database allows us to perform advanced flow field analysis unlocking detailed flow characteristics. We find perfect agreement between the Göttingen experiments and simulations, both for the heat transfer and the mean and temperature variance profiles close to the sidewall. A comparison of the heat transfer obtained from various experiments and the simulations is shown in Figure 3. The challenge of these simulations is that they have to be performed in a fully closed domain, making the computations more time consuming and more challenging than simulations performed in periodic domains. In particular, codes that solve the flow dynamics in periodic domains rely on optimized libraries. However, unfortunately, such methods are not directly available for this problem. Therefore, to perform the large-scale simulations, we developed an in-house second-order finite-difference flow solver that is specially developed for cylindrical geometries for the simulations in this project. To perform these landmark simulations, we have completely rewritten our code to optimize its performance on large-scale computing platforms like SuperMUC-NG. Our code is written in Fortran90. Large-scale parallelization is obtained using MPI to divide the computational domain over the outer computational loops, while OpenMP is used to obtain two-

dimensional parallelization throughout. This approach is favored since it limits network communication between computational nodes by performing computations within a node in a shared memory environment. This approach is favored since it limits network communication between computational nodes by performing computations within a node in a shared memory environment. To ensure computational efficiency, great care has been taken to vectorize computationally intensive parts of the code, limit memory movement by handwriting and inlining computational operations and data movements, efficient use of different cache levels, and replace library calls by handwritten routines to perform case-specific optimizations.

Ongoing Research / Outlook

SuperMUC and SuperMUC-NG allowed us to perform unprecedented simulations, i.e. to the best of our knowledge, the largest turbulence simulations in a fully close domain. Computer simulations of such turbulent flows are notoriously computationally demanding due to the extensive range of length and time scales that needs to be resolved. Therefore, such groundbreaking simulations can only be performed on the largest supercomputers in the world, such as SuperMUC-NG. Our simulations were only made possible due to algorithmic developments that limited the communication between different computational tasks. This improved our code's computational efficiency parallel efficiency on and parallel efficiency on a vast number of processors. Long term storage and data accessibility are assured by using the open-source HDF5 data format. However, even with the massive computational and storage facilities offered by SuperMUC-NG, it is still a massive challenge to simulate the ultimate thermal convection regime observed in the Göttingen experiments. Simulating ultimate thermal convection will require immense computational resources and a new generation of supercomputers.

References and Links

- [1] <https://stevensrjam.github.io/Website/>
- [2] D. Lohse, Physik Journal 18, Nr 8/9, 39-45.
- [3] X. He et al., PRL 108, 024502 (2012).
- [4] R.J.A.M. Stevens et al., J. Fluid Mech. 688, 31-43 (2011).
- [5] R.J.A.M. Stevens et al., 12, ERCOFTAC Series 2019, volume 27, 215-224 (2020).

DNS study of differential diffusion in a hydrogen jet in cross flow configuration

RESEARCH INSTITUTION

¹Simulation of reactive Thermo-Fluid Systems, TU Darmstadt

PRINCIPAL INVESTIGATOR

Christian Hasse¹

RESEARCHERS

Sebastian Popp¹, Abouelmagd Abdelsamie², Cheng Chi²

PROJECT PARTNERS

²LSS, Otto-von-Guericke-Universität Magdeburg

SuperMUC Project ID: pr74xi

Introduction

As part of the changing energy landscape, the influence of hydrogen as fuel will further increase. Along with that, fuel flexibility of common combustion applications is a crucial issue for their development when using e.g. blends of hydrogen and fossil fuels or even pure hydrogen due to different combustion characteristics and especially the higher diffusivity of hydrogen. Thus, increasing amounts of hydrogen lead to differential diffusion effects, which alter the molecular transport mechanisms towards the reaction zone. In practical combustion applications, both molecular and turbulent transport processes play a significant role and affect the overall flame structure, flame stabilization mechanisms and pollutant formation.

Within this project, a DNS study concerning differential diffusion effects and mixing characteristics during hydrogen combustion, using a canonical jet in cross flow (JICF) flame configuration, has been performed. The investigations in the hydrogen JICF configurations are two-fold. First, a detailed analysis of the DNS data could yield a fundamental understanding of mixing characteristics in the JICF configuration and differential diffusion effects, e.g. their relation to the location of the turbulent/non-turbulent interface (T/NT) [1]. Second, commonly applied tabulated chemistry approaches and their capability of predicting differential diffusion can be validated against the DNS data. The latter, which is of highly practical interest for a related project dealing with hydrogen combustion in a novel micro gas turbine, using a JICF configuration for fuel injection [2], will be the final target of the current project.




Figure 1: Rendering of the JICF, using temperature, H₂O mass fraction and mixture fraction iso-surfaces.

Results and Methods

For the JICF simulations, the direct numerical simulation code DINO [3] was used. 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 for both periodic and non-periodic boundary conditions, but with dedicated pre- and post-processing FFT techniques in the latter case.

Although in any low-Mach number solver the time step restriction associated with acoustic waves is removed, the restriction of time step due to chemistry stiffness is still present. For that reason, an implicit time integration of the stiff chemical source terms has been implemented, relying on a semi-implicit Runge-Kutta 3rd-order. By de-




Figure 2: Location of the analysis points. Each figure shows the iso-surface of the Q-criterion applied to instantaneous flow field at $t = 2$ from different orientation; the value of Q-criterion, $Q = 2 \times 10^6 \text{ s}^{-1}$, is selected so that the iso-surface is neither too dense nor too sparse.

fault, the chemical source terms are computed using the open-source Cantera library. The transport properties are computed either with the Cantera library or with the EGLib-3.4 library. The initial turbulent field is generated by inverse Fourier transform of an analytically prescribed energy spectrum (Passot-Pouquet or Von Karman-Pao). In DINO, input/output operations rely on MPI-I/O routines provided by the 2DECOMP&FFT library. These files are used for restarting the simulations while DINO uses parallel HDF5 saving for actual post-processing data. 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.

As a first step of the study, the jet in cross flow results were compared against experimental and simulation results, published by Sandia National Laboratories [4]. The typical jet visualization of the JICF simulation is presented by volume rendering of temperature, H_2O mass fraction and mixture fraction iso-surfaces, given in Fig. 1. After the validation, the project has been divided into two sub-projects: (1) studying the mixing characteristics and (2) analyzing differential diffusion in the JICF configuration. The first one is almost finished and the publication is currently under review. In this manuscript different zones, highlighted in Fig. 2, are investigated in detail. Especially local mixing effects are assessed by spatial and temporal statistics at the corresponding locations.

During the project, typical minimum number of cores was 512, which was used for test jobs. The maximum number of cores, used during the strong scaling tests, was 65,536 cores. Two typical common numbers of cores which were often used for production cases are 2,048 and 4,096. The required resources of the project are summarized in Tab. 1.




Figure 3: 3D iso-surface plot of the CH_4 mass fraction ($Y_{\text{CH}_4} = 0.03$) at $t = 0.3 \text{ ms}$.

Total Core-h	Overall storage	Typ. no. Cores	No. of files generated
17.7 Mio	20 TB	2,048 / 4,096	6,000

Table 1: Summary of required resources.

Ongoing Research / Outlook

Even though the end of the project is almost reached, several methods are still under development. The most important modification is implementing a novel 2nd-order immersed boundary algorithm based on a ghost-cell level-set method. In this method, a local directional extrapolation scheme is employed, leading to an accurate representation of the boundaries on the DNS grid. This will help to implement the inflow condition in more accurate way, instead of using mathematical functions to implement the velocity and species profiles at the jet inlet. A first application, tested on SuperMUC, was the transient processes controlling ignition by a hot jet issued from a pre-chamber as it can be observed from Fig. 3.

All these simulations required huge computing resources, and this is where SuperMUC plays the most important role. The computational power of SuperMUC allowed to perform such simulations in relatively short time.

References and Links

- [1] F. Hunger, M. Gauding, C. Hasse, J. Fluid Mech, 802 (2016), R5.
- [2] P. Jeschke, A. Penkner. A novel gas Generator concept for jet engines using a rotating combustion chamber. In: ASME J. Turbomach.: Vol. 137, Nr. 7 (2015), 071010-1-8.
- [3] A. Abdelsamie, G. Fru, T. Oster, F. Dietzsch, J. Gábor, D. Thévenin, Comput Fluids, 131 (2016), 123-141.
- [4] S. Lyra, B. Wilde, H. Kolla, J.M. Seitzman, T.C. Lieuwen, J.H. Chen, Combust Flame, 162 (2015), 1234-1248.

Development of a Stability-Based Transition

Transport Modeling Framework

RESEARCH INSTITUTION

¹C²A²S²E, Institute of Aerodynamics and Flow Technology, German Aerospace Center (DLR), Göttingen

²Helicopter Department, Institute of Aerodynamics and Flow Technology, German Aerospace Center (DLR), Göttingen

PRINCIPAL INVESTIGATORS

Philip Ströer¹, Anthony D. Gardner², Kurt Kaufmann²

RESEARCHERS

—

PROJECT PARTNERS

—

SuperMUC Project ID: pr83su

Introduction

The physical process of a laminar boundary layer becoming turbulent is called “laminar-turbulent transition”. Since the flow physics and in particular the skin friction responsible for the drag is significantly influenced by the state of the flow, that is laminar or turbulent, transition plays an important role in external aerodynamics. For example, for rotorcraft computations it was shown that the skin friction was significantly overpredicted (depending on the geometry and the flow state) applying a fully turbulent simulation compared to a simulation that considered laminar-turbulent transition [1]. An overestimation of skin friction leads to an overestimation of power demand which needs to be avoided. However, on complex industrially relevant configurations, like a helicopter rotor, incorporating transition prediction capabilities into the simulation is not straightforward. State-of-the-art streamline-based methods require a great implementation effort, massive user input and a detailed expert knowledge. Additionally, computing helicopter rotors is inherently computationally expensive, requiring unsteady computations on a fine grid (Fig. 1), with a long computation time to establish

the flow. The addition of boundary layer transition to the computation requires even finer time-steps and additional effort to compute the transition position. While streamline-based methods are relatively effective, they are not easily parallelizable to a large number of cores. Therefore, in the last decade transition transport models have been developed that aim at simplifying the incorporation and parallelization of the laminar-turbulent transition significantly. These models are only based on information available at each grid point, using transport type partial differential equations.

Results and Methods

In this project extensive investigations of different approaches to transition modelling on rotors were undertaken, including comparison to experimental data and results of other European CFD codes, see [2]. One finding was that for flows at Reynolds numbers below 500,000 the transition transport models predict unphysically large areas of laminar flow compared to the experimental data, see Fig. 2. To solve this problem a new boundary-layer transition model was developed by Ströer et al. [3,4] to improve





Figure 1: Rotor grid for DLR-TAU, from [2].






Figure 2: Transition position on the chord over the rotor radius for a transition transport model showing that only unphysically large turbulence levels approach the measured transition position, from [2].




Figure 3: Results for a 2D airfoil extracted from the RTG helicopter rotor at a constant radial position, from [3].




Figure 4: Transition position over the rotor for the new GI transition transport model showing good agreement with experimental data, from [5].

the transition prediction for a wide range of parameters crucial to external aerodynamics. The new model was already validated for fixed-wing applications [3,4], and can be used for both, compressible and incompressible flows. A drawback of the model proposed in Refs. [3,4] is the lack of Galilean invariance (GI). For problems where the velocities can be transformed relative to an adjacent wall, the lack of Galilean-invariance can be mitigated. However, for complex test cases (like helicopter computations) the formulation of a model should be as universal as possible. Therefore, a new model is under development that is Galilean invariant [5]. The final transition modeling framework [3,4,5] includes a four, a three and a one equation model. The objective is to provide a flexible framework that can be used to meet rapidly changing industrial requirements. The new models were implemented into the DLR TAU code. TAU is the compressible finite-volume Navier-Stokes solver of the German Aerospace Center (DLR) used for all computations mentioned. The new model works on either structured or unstructured grids. The simulations performed are RANS computations using an eddy-viscosity type turbulence model for the turbulent flow region. Fig. 3 shows results for the RTG helicopter rotor. The 2D airfoil was extracted at a constant radial position following Weiss et al. [6]. The Figure shows the transition onset position for different collective pitching angles. The agreement of the proposed model with the experimental data is significantly improved compared to the results of the basic transition transport model

(γ - Re_0 model). Fig. 4 shows results for the full 3D rotor obtained by the new GI model. The transition positions as well as the surface pressure are well matched over the range of surface positions for which experimental data are available. The new transition model framework utilizes the standard MPI code parallelization of TAU, and the needs very limited expert knowledge and user interaction. The parallel speedup known from other test cases is preserved.

Ongoing Research / Outlook

The results obtained by the new model show promising improvements compared to state-of-the-art models (cf. Ströer et al. [3,4]). The development of this new model is promising, and it will now be applied to currently flying helicopters to prove its industrial relevance. The resources provided by the LRZ on SuperMUC-NG were invaluable in supporting this development.

References and Links

- [1] Dietz, M., Dieterich, O., "Towards increased industrial Application of Rotor Aeroelastic CFD", 35th European Rotorcraft Forum, Hamburg, September 22-25, 2009
- [2] Kaufmann, K., Ströer, P., Richez, F., Lienard, C., Gardarein, P., Krimmelbein, N., Gardner, A. D., "Validation of boundary-layer-transition computations for a rotor with axial inflow", Vertical Flight Society 75th Annual Forum, Philadelphia, May 13-16, 2019.
- [3] Ströer, Ph., Krimmelbein, N., Krumbain, A., Grabe, C., AIAA Journal, Vol. 85, No. 4, 2020, pp. 1506-1517.
- [4] Ströer, Ph., Krimmelbein, N., Krumbain, A., Grabe, C., AIAA SciTech 2020, Orlando, FL, USA, 2020.
- [5] Ströer, Philip, Krimmelbein, Normann, Krumbain, Andreas, Grabe, C., AIAA Journal (to be submitted).
- [6] Weiss, A., Gardner, A. D., Schwermer, T., Klein, C., and Raffel, M., AIAA Journal, Vol. 57, No. 1, 2019, pp. 252-266.

Extreme-Scale Simulation of Turbulent Flow

RESEARCH INSTITUTION

¹Institute for Computational Mechanics, Technical University of Munich

PRINCIPAL INVESTIGATOR

Martin Kronbichler¹

RESEARCHERS

Niklas Fehn¹, Peter Munch¹, Wolfgang A. Wall¹

PROJECT PARTNERS

—

SuperMUC Project ID: pr83te (KONWIHR project)

Introduction

The aim of our project is to accurately predict the fluid flow at moderate and high Reynolds numbers in order to enhance the understanding of engineering processes or biomedical settings. In order to achieve this goal, our solvers represent most or even all the relevant physics in the flow and follow them over long time intervals by direct numerical simulation (DNS) or large eddy simulation (LES). While providing high-fidelity results, the numerical solution with this approach comes along with an extremely high demand on resolution and, thus, computations at the supercomputer scale. The institute for computational mechanics [1] addresses these challenges by a unique combination of sophisticated numerical schemes, highly tuned software implementations, and supercomputing power. The interdisciplinarity of our approach is reflected in a multi-physics multi-purpose code ExaDG, a solver with modules for incompressible flow, compressible flow, transport of scalar fields, and fluid-structure interaction.

Results and Methods

The code ExaDG is implemented in the C++ programming language and parallelized with MPI and shared memory approaches. The solver has been developed with funding from the German Exascale effort SPPEXA [2] and other sources, relying heavily on the deal.II finite element library. The spatial

variables are discretized with high-order discontinuous and continuous Galerkin schemes. The computational domain is decomposed into elements, with polynomial solutions on each element. This renders the method both flexible for representing complicated computational domains and accurate with very low dispersion errors and artificial dissipation for tracking both large and small turbulent eddies. Time integration is primarily based on splitting methods, which subdivide a time step into relatively simple updates for the velocity field combined with a Poisson equation for the pressure, which is solved with multigrid methods [3]. All steps in the solver ExaDG are designed to use modern computer hardware in an optimal way [2,3]. Traditional PDE simulators relying on sparse linear algebra are a bad match for the hardware of SuperMUC-NG. This is because those algorithms are heavy on memory transfer, whereas SuperMUC-NG can perform many more arithmetic operations on cached data. The ExaDG solver resolves this conflict by matrix-free operator evaluation in iterative solvers, which only implements the action of the operator on vectors by the underlying integrals of the high-order discontinuous Galerkin scheme. This approach trades memory transfer for additional local computations [3,4], and becomes particularly efficient with AVX-512 vectorization. Furthermore, the solver contains mixed precision ingredients to leverage the higher performance of single precision, yet maintaining the required computational accuracy of double precision.




Figure 1: Strong scaling behavior of the incompressible flow solver of the ExaDG project [2] on 3D turbulence.

A large number of short computations of various sizes of our project has been devoted to develop algorithms that scale well to the full size of SuperMUC-NG, increasing both the node-level performance and avoiding scalability bottlenecks throughout the solution chain. Figure 1 reports the strong scaling of the incompressible flow solver of ExaDG for 3D turbulent flow in a box with an effective resolution of up to




Figure 2: Evolution of turbulent eddies in incompressible Euler flow in 3D Taylor-Green vortex at times $t=3$ (upper left), $t=7$ (lower left), and $t=9$ (right, enlarged) in terms of iso-contours of the Q -criterion.

2,048³ points in each solution variable. Excellent scaling down to a saturation point of around 0.03 seconds per time step has been achieved. The arithmetic performance of the geometric multigrid solver has reached almost 7 Pflop/s on the full machine [2]. While less than the LINPACK score of 19 Pflop/s, the value is very high for a PDE solver. Note that we maximize for application throughput [4] and the code is indeed limited by memory bandwidth after all optimizations. The solver has been applied to a variety of large-scale turbulence simulations, consuming almost 30 million core hours during the last two years. It has generated new insight in theoretical fluid dynamics and method development, see also the references in [2]. The most challenging simulation has been the simulation of the inviscid 3D Taylor-Green vortex described in [5]. Using the highest resolution reported in literature so far with up to 4,096³ points or 2.4×10^{11} billion spatial unknowns and 2.3×10^5 time steps run on 2,048 nodes of SuperMUC-NG, a detailed view into the development of discontinuities in the incompressible Euler model has been gained. Our results add numerical evidence for important hypotheses in fluid dynamics. The analysis of the simulation outcome has been built into the solver, rather than explicitly visualizing data, due to the sheer size of the problem with a single checkpoint snapshot of two solution vectors already generating 4 TB of data. Figure 2 shows three snapshots of the post-processing of the numerical solution with a smaller resolution of 1,024³, created on several dozen fat-memory nodes of SuperMUC-NG in collaboration with the CFDlab team at LRZ. The figure shows the turbulent eddies visualized by the Q -criterion, a common metric in the field. While the initial structures (top left) are of large-scale and smooth, the flow eventually collapses into small eddies proportional to the numerical resolution and a very rich structure.

Ongoing Research / Outlook

In ongoing research, our group is preparing for additional large-scale simulations of turbulent flows, such as the flow over periodic hills at Reynolds numbers between 19,000 and 100,000 as well as the flow




Figure 3: Geometric setup for flow in the human lung.

in a diffuser. These computations are of high interest to the community because they involve separation and other effects that are not yet fully understood. Furthermore, detailed computations can also provide training data for reduced models with machine learning concepts, helping to build the next generation of industry tools that complement highly resolved computations.

In addition, the institute for computational mechanics is currently working on embedding the cutting-edge CFD solver into an application in biomechanics. It is planned to constitute the main ingredient for a comprehensive computational model of the human lung. The main objective is to track the fluid flow all the way from the trachea to the small-scale alveolar structures, to represent the gas exchange, and to assess the interaction of the flow with the lung tissue. This research is highly relevant for a better understanding of mechanical ventilation, like for the treatment of COVID-19-related or other acute lung diseases. Figure 3 shows a sample mesh of medium scale with 11 generations of resolved airways. When approaching 16–20 generations, the large range of geometrical and flow scales implies extremely large computations with billions of spatial unknowns. The simulations will provide a view on currently unexplored effects in the field of medicine, which in turn can enhance the quality of treatments.

References and Links

- [1] www.lnm.mw.tum.de
- [2] D. Arndt, N. Fehn, G. Kanschat, K. Kormann, M. Kronbichler, P. Munch, W. A. Wall, J. Witte. 2020. ExaDG: High-Order Discontinuous Galerkin for the Exa-Scale. In *Software for Exascale Computing – SPPEXA 2016–2019* (H.-J. Bungartz et al., eds), Lecture Notes in Computational Science and Engineering 139, 189–224. Code at <https://github.com/exadg/exadg>
- [3] N. Fehn, W. A. Wall, M. Kronbichler, *Int J Numer Methods Fluids* 88(1), 32–54 (2018).
- [4] M. Kronbichler, K. Kormann, *ACM Trans Math Softw* 45(3), 29:1–29:40 (2019).
- [5] N. Fehn, M. Kronbichler, P. Munch, W. A. Wall. 2020. Numerical evidence of anomalous energy dissipation in incompressible Euler flows: Towards grid-converged results for the inviscid Taylor-Green problem. arXiv:2007.01656.

Large Eddy Simulation of Wind Turbines

RESEARCH INSTITUTION

Wind Energy Institute, Technical University of Munich

PRINCIPAL INVESTIGATOR

Filippo Campagnolo

RESEARCHERS

Chengyu Wang

PROJECT PARTNERS

—

SuperMUC Project ID: pr84be

Introduction

The goal of the Wind Energy Institute [1] is to advance wind energy science and technology towards a fully renewable future energy mix. Modern wind farms typically consist of tens or hundreds of wind turbines. It's challenging to understand the aerodynamics of each turbine, the interaction between turbines and the mutual effects between the wind farm, and the atmospheric boundary layer (ABL). Besides the major tendency of increasing the number and size of turbines, it is paramount to achieve higher efficiency with limited resources. Wind farm control is one of the most promising approaches to boost wind farm power and increase turbine life expectations by reducing fatigue loads. A wind farm controller is a supervisory controller that monitors the whole wind farm and communicates with each wind turbine. Due to the complex essence of the physical phenomena and the huge physical scale of the problem, wind tunnel testing and numerical simulations are two typical tools fostering the investigation of the problem. Scaled wind turbines with rotor diameters of about 1 meter have been developed in our group. The turbine models have pitch and yaw control capabilities, meaning that each turbine blade and the rotor can rotate following commands. The turbine models were designed by similarity laws and preserve as many non-dimensional numbers as possible to mimic the characteristics of full-scale turbines. High fidelity Computational Fluid Dynamic (CFD) simulations are employed to reveal the physical phenomena

inside a wind farm and help improve the wind turbine and farm design. Large eddy simulations coupled with actuator line methods are used to model either a wind turbine or a wind farm. The physical phenomena in a wind farm include both large-scale and small-scale problems characterized by kilometers and millimeters, respectively. The numerical solver resolves the large-scale phenomena and models the small-scale ones, and finally, it is able to investigate the interactions of different scales within a wind farm. A highlight of our investigation is the testing of various wind turbine and farm control algorithms, including individual pitch control [2], yaw misalignment control [3], and dynamic induction control. The benefits and costs of each control method are discussed and compared. During the design of control algorithms, the balance between the increases of wind farm power and wind turbine fatigue loads is always an important trade-off. Meanwhile, the mechanism of each control must be understood properly.

Results and Methods

The software used is completely open-source. The fluid solver is *foam-extend-4.0*. A wind farm simulator named SOWFA is coupled with *foam-extend-4.0*. The aeroservoelastic code FAST is used to model wind turbine blades. We have done customization of the codes to better adapt it to our desire. A cluster of three scaled wind turbines was operated in the wind tunnel. LES simulations of the scenario have been conducted and are shown in Fig. 1. This figure was generated by LES simulations. The simulations consist of two phases: the precursor simulation that was conducted with 1,120 cores for 144 hours, while the successor simulation was conducted with 560 cores for 96 hours. The storage is mainly used for the flow data sampled from the precursor simulation, which requires about 200GB. Yaw misalignment control was employed to deflect the wake of upstream turbines away from downstream turbines. The first wind turbine is yawed to be misaligned with the inflow direction, which would result in a lateral force on the flow. The lateral force deflects the turbine wake such that it does not fully impinge on the downstream turbine. The power on the first turbine is decreased because of lower wind speed proportional to the rotor plane. However, the downstream wind turbine experi-




Figure 1: Iso-vorticity surface rendered by streamwise velocity component.




Figure 2: Turbine wake measured by experiment and simulated by CFD.

ences a power boost because of higher inflow wind speed. The net effect on the total wind farm power is a power increase of about 20%. Fig. 2 shows the wakes measured by experiments and simulated by CFD for two different control scenarios: the left column corresponds to the greedy case in which every wind turbine is aligned with the inflow; the right column shows the optimal yaw misalignment control. The black lines in Fig. 2 indicate the rotor plane, while the color represents the streamwise mean wind speed. The effects of wake deflection are pronounced in the figure. The first row shows experimental measurements obtained with LiDAR, while the second row shows the results computed by the numerical solver of Technical University of Munich, which is second-order accurate in space as most finite element solvers. The third row reports the results computed by the code of Imperial College London that is sixth-order accurate in space. Both codes are able to match with the experiments well. While the aforementioned yaw misalignment control is very beneficial for wind farm power, the Individual Pitch Control (IPC) is also another promising method in terms of turbine load reduction. Wind turbines operate in the atmospheric boundary layer, which typically is governed by vertical wind shear, i.e., the wind speed increases with altitude. Under such a situation, the mean wind speed felt by the upper part of the rotor is higher than the lower part. Thus, the wind shear results in a tilting moment on the rotor, which is detrimental for the life expectancy of the turbine. An IPC controller commands a higher pitch angle when the blade is in the upper part of the rotor plane to decrease the loading. While the load-reduction characteristic of IPC is studied by many researchers, its effects on turbine wake, especially for yawed wind turbines, are not sufficiently investigated. Our project fulfills this gap by using CFD simulation to evaluate the effects of IPC on the wake accurately. Fig. 3 shows the impact of IPC on turbine wake for three different yaw misalignment angles. The red dashed circle in Fig. 3 indicate the periphery of the rotor. The velocity




Figure 3: Influence of IPC on turbine wake.

Case	1	2	3
γ_1 [°]	30	0	-30
WT1 IPC	on	on	on
WT2 IPC	off	off	off
$\Delta C_{P1}/C_{P1}$	+1.7%	-3.0%	-3.9%
$\Delta C_{P2}/C_{P2}$	+1.6%	-4.2%	-0.1%
$\Delta \bar{M}_{n1} / \bar{M}_{n1} $	-99%	-100%	-100%
$\Delta \bar{M}_{y1} / \bar{M}_{y1} $	-99%	-94%	-97%
$\Delta \bar{M}_{n2} / \bar{M}_{n2} $	-6%	-9%	-8%
$\Delta \bar{M}_{y2} / \bar{M}_{y2} $	+11%	-14%	-2%

Table 1: Effects of IPC on power and loading.

ΔU is computed as the difference between the mean velocity with IPC and the mean velocity without IPC. As expected, the value of ΔU is zero in regions far away from the rotor, and the impact of IPC on wake is clearly shown by the non-zero values. In fact, IPC results both in some local acceleration and deceleration, and the total effect on rotor effective mean wind speed, which is directly related to downstream turbine power, depends on the yaw direction of the upstream turbine. For a wind farm consisting of two aligned wind turbines, Table 1 shows the effects of IPC on power and loading. The subscript 1 indicates the upstream turbine while 2 indicates the downstream one. When the yaw angle of the upstream turbine is positive, IPC is beneficial for the power of both turbines, while it is just the opposite for a negative yaw angle of the upstream turbine. This phenomenon can be fully explained by a detailed analysis of the CFD data. With this knowledge, a better wind farm control algorithm that accounts for the asymmetrical effects of IPC can be designed.

Ongoing Research / Outlook

Our group continues the investigation with CFD for the following topics: active triggering of wake instability and enhancement of wake recovery, noise emission of wind turbines, optimal design of wind farm control algorithm. One important limitation of our research is the scalability of *foam-extend-4.0*. The parallel efficiency is not yet ideal. However, the simulation framework is able to fulfill all current requirements. Currently, our institute is involved in two projects named FLAMINGO and CompactWind II, where the potential of novel wind farm control strategies will be further investigated.

References and Links

- [1] <https://www.mw.tum.de/wind/home>
- [2] Chen, Z. J., K. A. Stol, and B. R. Mace, Renewable energy 103 (2017): 750-765.
- [3] Campagnolo F, et al., J. Phys.: Conf. Ser. 753(3) 032006 (2016).

Modulation of Turbulent Properties in Spray Flame

Burning n-Heptane: Direct Numerical Simulation

RESEARCH INSTITUTION

Laboratory of Fluid Dynamics and Technical Flows (LSS/ISUT), Otto-von-Guericke University, Magdeburg

PRINCIPAL INVESTIGATOR

Dominique Thévenin

RESEARCHERS

Abouelmagd Abdelsamie, Cheng Chi

PROJECT PARTNERS

—

SuperMUC Project ID: pr84qo (Gauss Large Scale project)

Introduction

Our main work on SuperMUC started at the end of 2014; at this time, the main topic was the fundamental investigation of spray flames in a simple geometry using direct numerical simulation (DNS). With the growing power of SuperMUC and with the further numerical optimization done concerning our code, more complex domains could later be handled: in particular, the spray flame in a spatially evolving jet has been computed. Currently, and with the help of SuperMUC, this work is further extended to include additionally models to describe nanoparticle synthesis from a spray flame in a real burner. This is a very challenging task because there are three phases that must be taken into account: a continuous gas phase, liquid droplets, and solid nanoparticles. Additionally, it is necessary to take into account detailed chemistry to describe the complexity of the kinetics. Several papers about this topic have been published by our group with the help of SuperMUC. Without the granted Core-hours, such studies would have been quite im-

possible. Even if much has been achieved, further progress is still needed in order to understand better the complex coupling processes between multiphase flow and physicochemical processes.

Algorithms and Numerical Methods

The current project relies on the in-house DNS code DINO. DINO is a third-generation, seven-year-old Fortran90 code, which has been developed in our group starting from the beginning of 2013 [1-5]. 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 for both periodic and non-periodic boundary conditions, but with dedicated pre- and post-processing FFT techniques in the latter case. The ODEs are integrated using a semi-implicit 3rd order Runge-Kutta relying on LAPACK, BLAS, and PyJac (analytical Jacobin matrix solver) libraries. By default, the chemical source terms are computed using the open-source Cantera-2.4.0 library, whereas 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 small, non-resolved droplets) or with the Immersed Boundary Method (IBM) technique for fully resolved cases. Nanoparticle synthesis leading to the development of a third solid phase in the code is simulated using the monodisperse model developed by Kruijs et al. 1993. In this model three additional equations are solved for three different parameters: nanoparticle number concentration, total surface area concentration, and total volume concentration. Regarding the wall, a novel implementation of IBM, based on directional ghost-cells with a level-set function, has been introduced into DINO in order to consider the walls in a computationally efficient manner [5]. The initial turbulent flow field is generated by inverse Fourier transform of an analytically prescribed energy spectrum (Passot-Pouquet or Von Karman-Pao). In DINO, input/output operations rely on MPI-I/O routines provided by the




Figure 1: Schematic diagram representing the complete process used in the DNS simulation to predict the nano-particle size distribution [3].




Figure 2: Instantaneous contour of instantaneous gas vorticity magnitude when considering the real injector geometry [4].

2DECOMP&FFT library. These files are used for restarting the simulations, while DINO uses parallel HDF5 saving for actual postprocessing data. The code is managed under GIT version control, which helps all users to quickly and safely carry out changes or updates, if needed. DINO uses CMake as build system; it can be compiled with both GNU and Intel Fortran compilers. Currently, DINO is used by 6 groups in 5 countries. The required resources on SuperMUC, until now, for this project are summarized in Table 1.

Scientific Results

Recently direct numerical simulations (DNS) of nanoparticle synthesis in spray flame have constituted the most ambitious part of the project. This topic itself is a part of the DFG Priority Program SPP1980; "Nanoparticle Synthesis in Spray Flames, SpraySyn: Measurement, Simulation, Processes". This project started in July 2017 and it is running for six years. The main contribution of our group in this work is to investigate the nanoparticle synthesis in the spray flame and give more details about the complete process, as it will be explained in Sub-project A. In this project, the real geometry of the burner and of the spray injector were included in the simulation. For this purpose, a novel implementation of the immersed boundary method (IBM) was developed as a side project as it will be explained in Sub-project B.

Sub-project A The main purpose of this subproject is to perform direct numerical simulations (DNS) of a burner with conditions similar to that of the SpraySyn burner. The SpraySyn burner is designed as part of a collaborative DFG project (SPP1980) entitled "Nanoparticle Synthesis in Spray Flames SpraySyn: Measurement, Simulation, Processes". This burner is de-

Total Core-h	Overall storage	Typ. no. Cores	No. of files generated
55 Mio	127 TB	4,096	12,000

Table 1: Summary of required resources (over 6 years).

signed to generate nanoparticles from spray flames. The authors are involved in this collaborative DFG project. During the last granting period of SuperMUC, several publications have taken place documenting the results of this project, in particular [2-4]. In these publications, a relevant part of domain has been simulated while taking into account the real geometry. In this manner, it was possible to estimate the resulting nanoparticle size distribution (PSD) as shown in Fig. 1. A short extension has been selected in the crosswise direction to keep the 3D effects regarding turbulent structures, relevant in particular for spray evaporation, while at the same time reducing the necessary computational time. Different parameters have been changed and investigated during this work as follows:

- (1) the solvents: ethanol or xylene;
- (2) droplets injection speed;
- (3) droplet size.

Sub-project B The impact of the wall of the SpraySyn burner on the nanoparticle synthesis is a part of the investigation in the DFG project. In order to do that, our DNS code (DINO) needed several modifications. An innovative Immersed boundary method (IBM) developed by our group [5] has been implemented in DINO to consider in a suitable manner the wall effect. During the development of this novel implementation, several cases have been tested and investigated. For example, we started with a simple, standard benchmark: the flow over a sphere. The complete validation procedure has been recently published in Ref. [5].

Later on, far more complex cases have been investigated as well: (1) transient processes controlling ignition by a hot jet issued from a pre-chamber, (2) Internal combustion engine, (3) Biomedical applications, and (4) impact of the wall in the SpraySyn burner [4] as it can be seen from Fig. 2.

Ongoing Research / Outlook

Currently, a huge analysis are performed on the running cases. The output and published data of the current project can ultimately become a reference dataset 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] Abdelsamie A., Fru G., Oster T., Dietzsch F., Janiga G., Thévenin D., *Comput. Fluids*, 131: 123-141, 2016.
- [2] Abdelsamie A. and Thévenin D., *Proc. Combust. Inst.* 37: 3373-3382, 2019.
- [3] Abdelsamie A., Kruis F.E., Wiggers H. and Thévenin D., *Flow Turbul. Combust.*, 105: 497-516, 2020.
- [4] Abdelsamie A., Chi C., Nanjiah M., Skenderovic I., Suleimann S. and Thévenin D., *Flow Turbul. Combust.*, 106: 453-469, 2021.
- [5] Chi C., Abdelsamie A. and Thévenin D., *J. Comput. Phys.* 404, 109122/1-20, 2020.

WALBERLA – A massively parallel framework for multi-physics simulations

RESEARCH INSTITUTION

¹Chair for System Simulation (Informatik 10), Friedrich-Alexander Universität Erlangen Nürnberg

PRINCIPAL INVESTIGATOR

Harald Köstler¹

RESEARCHERS

M. Bauer¹, S. Eibl¹, J. Höning¹, N. Kohl¹, C. Rettinger¹, C. Schwarzmeier¹, D. Thönnies¹, B. Vowinckel²

PROJECT PARTNERS

²TU Braunschweig

SuperMUC Project ID: pr86ma

Introduction

The open-source massively parallel software framework WALBERLA [1,2] (widely applicable lattice Boltzmann from Erlangen) provides a common basis for stencil codes on structured grids with special focus on computational fluid dynamics with the lattice Boltzmann method (LBM).

Other codes that build upon the waLBerla core are the particle dynamics module MESA-PD and the finite element framework HYTEG.

Various contributors have used waLBerla to simulate a multitude of applications, such as multiphase fluid flows, electrokinetic flows, phase-field methods and fluid-particle interaction phenomena.

The software design of WALBERLA is specifically aimed to exploit massively parallel computing architectures with highest efficiency. In order to simulate real-world scenarios, WALBERLA relies on using the immense compute power available on modern high performance computing systems such as LRZ's SuperMUC-NG.

Results and Methods

Phase-field simulations

In the SKAMPY DFG project, we extended WALBERLA with phase-field code generation [3]. We developed automatic program generation technology to create scalable phase-field methods for material science applications. To simulate the formation of microstructures in metal alloys, we employ an advanced, thermodynamically consistent phase-field method. A state-of-the-art large-scale implementation of this model requires extensive, time-consuming, manual code optimization to achieve unprecedented fine mesh resolution. Our new approach starts with an abstract description based on free-energy functionals which is formally transformed into a continuous PDE and discretized automatically to obtain a stencil-based time-stepping scheme. Subsequently, an automatized performance engineering process generates highly optimized, performance-portable code for CPUs and GPUs. We demonstrate the efficiency for real-world simulations on large-scale GPU-based (PizDaint) and CPU-based (SuperMUC-NG) supercomputers (Figure 1).

Our technique simplifies program development and optimization for a wide class of models. We further outperform existing, manually optimized implementations as our code can be generated specifically for each phase-field model and hardware configuration.

Simulation study of particulate flows

Using direct numerical simulations to study particulate flows has become a promising alternative to laboratory experiments. They allow a deeper insight into physical properties, can be controlled more easily, and are more cost efficient for parametric studies. For that reason, we are constantly developing further our fluid-particle coupling module inside the WALBERLA framework. As a result, the coupling module has been adapted to now also support WALBERLA's newly developed particle simulation module MESA-PD which now allows for much more flexibility regarding the particle-interaction algorithms. We used this gained flexibility to further improve the accuracy of the




Figure 1: Weak Scaling on SuperMUC-NG with 603 blocks per core.




Figure 2: Erosion study of a single movable layer of spheres above a fixed layer inside turbulent flow. The spheres are colored according to their velocity magnitude.

interaction algorithms. We employ these novel techniques to study particulate flow scenarios, like erosion processes of riverbeds.

To obtain statistically converged results, such systems often require a large amount of particles and long run times for which we employ the SuperMUC-NG super-computer.

Another topic of our numerical studies is to track down the source that initiates particle erosion.

In collaboration with Bernhard Vowinckel, TU Braunschweig, we revisit their recent study on the erosion of single particles in turbulent channel flow. As shown in Figure 2, a movable layer of spheres is placed on top of a fixed layer. An erosion process is characterized by a sphere leaving the top layer and traveling through the domain. It is supposedly triggered by strong fluid vortices hitting the sphere or a former collision by an already moving particle.

Besides the long run time and the fine spatial resolution, the required logging of all particle positions and velocities at all time steps is a major challenge for massively parallel execution.

This work has been presented at the PARTICLES conference 2019 in Barcelona [4].

Ongoing Research / Outlook

We are currently preparing for time-dependent runs also in the TERRANE0 project [5], and are estimating, that during the next year, these experiments may together with the riverbed experiments, consume the majority of the remaining compute time.

We are grateful that during the friendly user phase, especially the phase-field simulations could be computed without accounting of core-h on the project. Therefore we saved a lot of resources on our account.

References and Links

- [1] <https://walberla.net>
- [2] Bauer et al. "WALBERLA: A block-structured high-performance framework for multiphysics simulations". In: *Computers & Mathematics with Applications* (2020). ISSN: 0898-1221. DOI: 10.1016/j.camwa.2020.01.007.
- [3] Bauer 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*. ACM, 2015, p. 8.
- [4] <https://congress.cimne.com/particles2019/admin/files/fileabstract/a118.pdf>
- [5] <https://terraneo.fau.de>

Cavitation in Injection Systems

RESEARCH INSTITUTION

Chair of Aerodynamics and Fluid Mechanics, TUM

PRINCIPAL INVESTIGATORS

Theresa Trummler, Steffen J. Schmidt

RESEARCHERS

Felix Örley, Christian Egerer, Polina Gorkh, Stefan Hickel

PROJECT PARTNERS

—

SuperMUC Project ID: pr86ta

Introduction

Developments in direct Diesel injection systems increase rail pressures to more than 2,500 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 between cavitation and turbulence as well as flow transients due to moving geometries. We use Large-eddy simulation (LES) to understand the flow dynamics. Our simulations have been conducted on SuperMUC, SuperMUC Phase 2 and SuperMUC-NG with the massively parallelized flow solvers INCA and CATUM [1].

Numerical Method

With LES, the smallest turbulent flow scales are not resolved on the computational grid. The effects of these scales thus must be modeled. We employ an implicit LES approach based on the Adaptive Local Deconvolution (ALDM) method. In CATUM, an efficient LES method for compact stencils is implemented [2]. To consider two-phase effects, we apply the homogeneous-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 cavitation model by a non-condensable gas component [3]. Additionally, we also employ accurate full-thermodynamic tables to assess all thermodynamic effects. Complex moving bodies are considered by a conservative cut-cell method.

Results

Cavitation dynamics and effect on the jet

We have studied the cavitation dynamics and the effect of cavitation in the nozzle on the liquid jet [3,4]. The considered setup resembles a generic, scaled-up automotive fuel injector and consists of a cavitating flow in a step nozzle with subsequent injection into air. We investigated several operating conditions that lead to different cavitation characteristics. At developing cavitation, see Fig. 1 (a, b), vapor structures are present near the nozzle inlet and vapor clouds are periodically shed. The supercavitating operating point, see Fig. 1 (c, d), is characterized by a vapor sheet spanning almost over the entire nozzle length. Here,




Figure 1: Cavitation structures (purple), jet surface (grey,) and gas entrainment into the nozzle (green) for developing cavitation (a, b) and supercavitation (c, d) [3].




Figure 2: Coherent vortical structures during the main injection phase colored by velocity in the half-domain [5].

vapor clouds detach less periodically at the end of the sheet and either collapse near the nozzle outlet, where they enhance turbulent fluctuations, or are convected towards the outlet and lead to gas being sucked into the nozzle (gas entrainment). The amplified fluctuations at the outlet and gas entrainment are among the main drivers for the stronger jet break-up at supercavitation.

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 1,500 bar into air [5].

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. 2. 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. 3, and cause violent collapse events. Subsequently, the collapse of the sac hole cavity and 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.

Temperature increase in cavitating flows

The fully thermodynamic modeling of the working fluid enables the assessment of thermal effects of cavitation and temperature-dependent fluid properties. We have performed LES of a typical 8-hole diesel injector using accurate full thermodynamic tables for the equation of state.

Figure 4 depicts the time-averaged temperature distribution. The full-thermodynamic computation captures the viscous heating at the walls. Further, temperature is increased in the cavitation region. Due to




Figure 3: Cavitation structures inside the nozzle holes and sac hole shortly after full closing of the injector needle [5].




Figure 4: Time-averaged temperature distribution in a cavitating injector flow assessed with full-thermodynamic simulations.

mixing of the liquid with the heated vapor, the average temperature increases from the inlet of the hole to the outlet.

Conclusion

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 5,600 cores to compute physical time scales on the order of micro-seconds.

References and Links

- [1] <https://www.mw.tum.de/aer/>
<https://www.inca-cfd.com>
- [2] Egerer, C. P., Schmidt, S. J., Hickel, S., Adams, N. A., J. Comput. Phys., 316, 453–469, 2016. doi: 10.1016/j.jcp.2016.04.021
- [3] Trummel, T., Rahn, D., Schmidt, S. J., Adams, N. A., At. Sprays, 28(10), 931–955, 2018. doi: 10.1615/AtomizSpr.2018027386
- [4] Trummel, T., Schmidt, S. J., Adams, N. A. (2020), Int. J. Multiph. Flow (125), 103215. doi: 10.1016/j.ijmultiphaseflow.2020.103215
- [5] Örlay, F., Hickel, S., Schmidt, S. J., Adams, N. A., Int. J. Engine Res., 18 (3), 195-211, 2016. doi: 10.1177/1468087416643901

DNS study of the early flame kernel development under engine conditions

RESEARCH INSTITUTION

Institute for Combustion Technology, RWTH Aachen

PRINCIPAL INVESTIGATOR

Heinz Pitsch

RESEARCHERS

Hongchao Chu, Tobias Falkenstein

PROJECT PARTNERS

—

SuperMUC Project ID: pr89pa

Introduction

Cycle-to-cycle variations (CCVs) limit the further optimization of spark-ignited engines. Combustion instabilities related to CCVs inhibit more efficient engine operations and lower emissions. Despite the importance of CCVs, the fundamental reasons for their occurrence have not yet been fully understood or predictively described. Experiments have shown that the early period of the combustion process plays a significant role in CCVs. The flame kernel initiated by the spark is small. In the very early combustion phase, it is about the size of the turbulent scales and is not in equilibrium with the turbulence. Therefore, the evolution of the early flame kernel is dominated by local, stochastic small-scale fluctuations, which are beyond the common modeling approaches based on conventional turbulence theory and are not reflected in current models. This study aims to investigate the influence of the anisotropic turbulence and small-scale intermittency in the engine on the early flame propagation and to investigate the role of differential diffusion during early flame kernel development. For this purpose,

datasets of direct numerical simulations (DNS) for idealized engine geometries have been established and analyzed.

Results and Methods

Mixture state and composition in the simulation have been chosen to be representative for fully homogeneous, stoichiometric spark ignition (SI) engine operation at low load. Governing equations and numerical methods were extensively described in [2,3,4]. To enable a better understanding of how the flow and how differential diffusion alters the young flame kernel, Lewis numbers (Le) of all species have been artificially set to unity in some of the simulations. Three configurations with different ratio of initial flame diameter D_0 to the integral length scale l_t have been simulated and analyzed to investigate the characteristics of the early flame kernel development:

- Engine flame kernel, $D_0/l_t = 0.3$ (isochoric),
- Large flame kernel, $D_0/l_t = 2.0$ (isochoric),
- Planar flame, $D_0/l_t = \infty$ (isobaric).




Figure 1: Iso-surfaces of temperature (with arbitrary color scale) and the second invariant of the velocity gradient tensor (grey background) for the three flames at $t = 1.0 \cdot l_t$ (left: engine flame kernel, middle: large flame kernel, and right: planar flame) [2].

Visualizations of all flames at approximately one eddy turnover time after flame initialization are given in Figure 1. To illustrate early flame kernel development, a sequence of flame images has been extracted from the engine-relevant DNS dataset and is plotted in Figure 2. Comparing to the planar flame, it was found that the development of the early flame kernel differs significantly from the well-established turbulent flame configurations. The curvature distribution of the early flame kernel was found to be strongly skewed towards positive curvatures, which contradicts with the developed turbulent flames. This feature of the early flame kernel is surprising and requires more attention for two reasons. First, the observed behavior marks distinctive deviations in early flame/turbulence interactions




Figure 2: Development of the engine-relevant flame kernel: Temperature iso-contours coloured by flame curvature. The vertical lines illustrate the integral length scale l_i [3].

Case	Mio. Core-h
Large flame kernel $Le = 1$	0.37
Planar flame $Le > 1$	1.42
Planar flame $Le = 1$	2.30
Engine flame kernel $Le > 1$	1.52
Engine flame kernel $Le = 1$	1.02
4 additional realizations	2.73
Total	9.36

Table 1: Summary of the performed simulations.

of small kernels not only from the developed flames, but also from the ones that are initially as large as the energetic flow scales. Second, the occurrence of large positive curvatures and associated stretch rates may significantly change early flame propagation in situations where preferential diffusion effects are important. For non-unity Le flames, curvature influences the local heat release rate. Figure 3 shows the heat release rate on the temperature iso-surfaces corresponding to maximum heat release rate (T_{maxHR}) in a laminar unstretched flame under the same thermodynamic conditions. In the $Le = 1$ dataset, the heat release rate distribution is almost uniform. Conversely, the more engine-relevant flame with $Le > 1$ features similar high source term magnitudes only in regions of negative curvature, while large parts of the plotted iso-surface exhibit significantly reduced heat release rates. The role of differential diffusion during early flame kernel development was investigated in [4]. The in-house code called CIAO is used to perform all the simulations. It is a structured, arbitrary order, finite difference code. Spatial staggering of flow variables is used in order to increase the accuracy of the




Figure 3: Planar Flame: T_{maxHR} iso-surface colored by local heat release rate for the $Le = 1$ dataset (top) and the engine-relevant $Le > 1$ flame (bottom) [4].

numerical stencil. About 11,000 cores are used for the simulation. A summary of the performed simulations is given in Table 1.

Ongoing Research / Outlook

Current investigations focus on the run-to-run variations of early flame kernels only as a consequence of turbulence-kernel interactions. Multiple realizations of early flame kernels are computed under the same nominal conditions.

References and Links

- [1] <https://www.itv.rwth-aachen.de/forschung/aktuelle-forschung/>
- [2] Falkenstein, T., Kang, S., Cai, L., Bode, M., Pitsch, H, Combust Flame, 213, 455-466 (2020).
- [3] Falkenstein, T., Kang, S., Pitsch, H., J. Fluid Mech, 885 (2020).
- [4] Falkenstein, T., Rezchikova, A., Langer, R., Bode, M., Kang, S., Pitsch, H., Combust Flame (2020).

Towards combustion LES of liquid alternative aviation fuels

RESEARCH INSTITUTION

Institute of Combustion Technology, German Aerospace Center (DLR)

PRINCIPAL INVESTIGATOR

Patrick Le Clercq

RESEARCHERS

Georg Eckel, Benedict Enderle

PROJECT PARTNERS

—

SuperMUC Project ID: pr92ca

Introduction

Our impact on the atmosphere and the climate has been evidenced by an increasing number of critical observations. In the field of commercial aviation and air traffic, dedicated projects for instance at the DLR (SULFUR, PAZI 1&2 und CATS, ECLIF) [1] or at NASA (SNIF, SUCCESS, AAFEX, ACCESS) have contributed to the evidencing effort. It is anticipated that the introduction of alternative aviation fuels not only affects the CO₂ life cycle but also the emissions and climate. Synthetic Paraffinic Kerosene (SPK) from certified routes (F-T, HEFA) or future pathways are virtually free of aromatic compounds as well as free of sulfur-containing molecules. Blended up to 50% vol. with conventional Jet A-1 or with Synthetic Kerosene Aromatics (SKA) they form drop-in fuels, which from their initial formulation all the way to their large scale production offer unexplored optimization possibilities especially regarding emissions reductions. In this context the DLR Institute of Combustion Technology investigates the effects of the chemical composition of different multi-component alternative aviation fuels on the combustion characteristics.

One key aspect is the influence of the fuel composition on the primary atomization process which changes the quality of the liquid spray and ultimately the position of

the reaction zone and the formation of pollutants. In previous investigations it was concluded that precise and reliable spray boundary conditions (Figure 1) are required to account for this fuel effect in simulations. In order to reduce the experimental burden for their characterization, an alternative approach is investigated, consisting of an algebraic primary breakup model in combination with accurate quantification of uncertainties arising from simplifications and imperfect knowledge.

Results and Methods

A computational platform consisting of the coupling of a sophisticated evaporation model in the Lagrangian framework with a detailed chemistry solver for the chemical kinetics in the Eulerian framework for complex multi-component fuels was used for the work at hand. A lab-scale burner producing a highly turbulent swirl-stabilized spray flame served as a test case (Figure 2). Detailed validation data was available from previous experimental characterizations. The burner exhibited some of the key features of current aero-engine combustors, e.g. typical fluid dynamic instabilities and spray formation by a hybrid prefilming airblast atomizer.

To account for the highly unsteady and turbulent processes while keeping the computational costs acceptable, a hybrid URANS/LES approach was utilized. In a first step, the setup was validated using the spray boundary condition from the experiments. The simulation demonstrated a good agreement with the experiments regarding local temperature profiles and flame position (Figure 3).

In a further study, the spray boundary condition was replaced by an algebraic primary breakup model [2] which relies on few local quantities of the gas field as well as properties of the liquid fuel (Jet A-1). Inputs to the model which could not be determined precisely were interpreted as uncertain quantities and characterized as random variables from a probabilistic perspective (Table 1).




Figure 1: Liquid fuel droplets starting from the spray boundary condition (blue annulus) above the atomizer.




Figure 2: Hybrid URANS/LES of the swirl-stabilized kerosene spray flame. Red lines indicate heat release.




Figure 3: Comparison of line-of-sight integrated OH^* distribution from simulation (blue) and experiment (red). Contours show relative levels of maximum intensity.

The uncertainties from the resulting five dimensional input space were then propagated through the simulation model using Polynomial Chaos Expansion (PCE) [3]. Calculation of the expansion coefficients required 71 individual simulations with varying spray boundary conditions, involving data from the input space. PCE enabled the portrayal of the uncertain simulation results in terms of mean realizations and probability levels (Figure 4).

Ongoing Research / Outlook

As the simulation of the atomization process remains a challenge, the prediction of the combustion performance of the individual fuels by simulation still relies on accurate spray boundary conditions. A methodology has been demonstrated which reduces the required prior knowledge from experiments while enhancing simulation credibility through clear display




Figure 4: Non-deterministic simulation results of gas phase temperature. Gray areas display uncertainties.

Description	PDF	Unit
φ_D axial angle	$U(26, 29)$	[°]
ψ_D swirl angle	$U(55, 63)$	[°]
U_D absolute velocity	$U(25, 30)$	[m s ⁻¹]
U_G gas velocity	$U(70, 85)$	[m s ⁻¹]
T_{liq} liquid temperature	$U(303, 320)$	[K]

Table 1: Uncertain input parameter space.

of uncertainties. The methodology was tested for kerosene and compared to available experimental data. On-going research will focus on the application of the demonstrated methodology to new fuel candidates without prior knowledge of the spray characteristics. Besides, further boundary conditions, which are prone to uncertainties, will be investigated.

References and Links

- [1] https://www.dlr.de/vt/en/desktopdefault.aspx/tabid-10803/18874_read-43793/
- [2] Chaussonnet, G., et al., Int. J. Multiph. Flow 80 (2016): 29-42.
- [3] Najm, Habib N., Annu. Rev. Fluid Mech. 41 (2009): 35-52.

Buoyant-convectively driven heat and mass transfer

RESEARCH INSTITUTION

¹Institute for Hydromechanics, Karlsruhe Institute of Technology (KIT)

PRINCIPAL INVESTIGATOR

Herlina Herlina¹

RESEARCHERS

Jan Wissink²

PROJECT PARTNERS

²Brunel University London

SuperMUC Project ID: pr92da

Introduction

The transfer of heat and dissolved substrates, such as atmospheric gases, across the water surface is governed by the interaction between molecular diffusion and near-surface turbulence, where the latter is very effective in enhancing the transfer rate. The simulation reported here focuses on surface-cooling induced heat and mass transfer.

When a body of water is exposed to sudden surface cooling, a thin thermal boundary layer of cool water forms adjacent to the water surface. Any small perturbation, which is always present in nature, will lead to the formation of convection cells near the surface (cf. Figure 1, top panes for their footprints). The center region of these cells is predominantly occupied by warmer water rising from below, which is subsequently cooled while moving radially along the surface. At intersections of three or more convection cells cold water accumulates and eventually moves downwards in the form of mushroom-like plumes (Figure 1, bottom panes), thereby promoting the mixing of saturated cold water from the surface and unsaturated warm water from the bulk.

Results and Methods

The present simulation builds on our previous simulations, where surface-cooling induced mass transfer was accurately resolved up to a realistic Schmidt number of $Sc=500$, which is typical for oxygen in water [1, 2]. Due to the relatively small sizes of the computational domains ($5L \times 5L \times 5L$ and $10L \times 10L \times 10L$, with $L \sim 1\text{cm}$) employed in the previous simulations, falling plumes were found to reach the bottom before they naturally lost their buoyancy. Thus, the analyses were limited to the initial development of the Rayleigh instability and the early stages of transition.

Here, we aim to understand the physics of buoyancy-driven heat and mass transfer when plumes lose their buoyancy upon impinging on a thermocline separating the warm water region immediately below the surface and the cold water close to the bottom. This reflects the situation encountered in deep calm water bodies (lakes). To simulate this situation, the domain size has to be sufficiently large in order to at least incorporate both the warm water region near the surface as well as the thermocline.




Figure 1: Sequences of the normalized temperature field at one grid plane beneath the surface (top panes) and at a vertical plane crossing $x/L=10$ (bottom panes). The pattern visible in the top panes can be seen as footprints of convection cells in the water. The initial Rayleigh number based on $L=1\text{ cm}$ was 39,000.

For this purpose, a $20L \times 20L \times 20L$ computational domain was selected. As an initial condition, below the warm upper bulk, the temperature near the bottom of the domain was gradually decreased in order to obtain a stable temperature profile. Prior to this, a larger $30L \times 30L \times 30L$ domain was employed, which also included the cold water region below the thermocline. The initial results of this simulation (in terms of convection cell size development and related sinking plumes) were very similar to the present simulation using the $20L \times 20L \times 20L$ domain. Based on this, it was decided to focus our efforts on the smaller simulation, as this would allow a significantly longer running time.

A sequence of temperature distributions in the plane adjacent to the surface is shown in the top panes of Figure 1. The patterns can be seen as footprints of the convection cells in the water. On average, the convection cells were found to monotonically grow in size during the transition stage until, at $t \approx 100$ s, the plumes lose their buoyancy upon impinging on the thermocline and a quasi-steady state was reached. In the latter state, the long term average size of the convection cells remains constant, while simultaneously there is a continuous cell-renewal process, where cells merge and grow before breaking up into new cells.

To study the interaction between the falling plumes and the thermocline near the bottom of the domain, a sequence of temperature isosurfaces is shown in Figure 2. It can be seen that local peaks in the vertical temperature gradient dT/dz (corresponding to impinging plumes) result in depressions in the isosurfaces caused by the downward momentum of the plumes.

The 3D incompressible Navier-Stokes equations for the flow-field and the advection-diffusion equations for scalar fields were solved using direct numerical simulations employing a code capable of resolving all details of low-diffusivity (high Schmidt number) mass transfer which is characterized by steep concentration gradients (cf. Figure 1). The code uses 4th order discretizations of convection and diffusion to solve the incompressible Navier-Stokes equations for the fluid flow. For the scalar field, a fifth-order-accurate WENO scheme [3] for scalar convection, combined with a fourth-order central method for scalar diffusion, was implemented on a staggered and stretched mesh (cf. [4] for further details).

The flow and (two) scalar fields were resolved using 2.5088×10^9 grid points. The computation load was distributed over 14,336 cores resulting in local blocks of $50 \times 50 \times 70$ grid points. The standard Message-Passing-Interface (MPI) was employed for data exchange between processes. In addition to the temperature field (active scalar) with a Prandtl number $Pr = 7$, a concentration field (passive scalar) with a Schmidt number $Sc = 16$ was also resolved. With this arrangement, the computational time was, on average, 1.5 s / iteration. In total approximately 35 Mio Core-h were used to simulate the present quasi-steady case.




Figure 2: Dynamics and interaction of the stably-stratified layer near the bottom and cold-sinking plumes, visualized using the isosurface of the normalized temperature field $T = 0.6$, colored by the vertical temperature gradient.

The amount of data generated were 14 TB.

Ongoing Research / Outlook

SuperMUC has enabled us to perform a number of large scale numerical simulations of interfacial gas transfer at high Schmidt and/or high Reynolds numbers [2,5]. The high-fidelity data allowed a number of unbiased parametric studies of the mass transfer velocity as a function of, e.g. Schmidt number.

Apart from the occasionally long waiting times for our jobs to run, no major problems were encountered. Small updates to the run scripts were required to accommodate the running of jobs on the new SuperMUC system.

As a next step, we aim at performing large scale DNS of gas transfer in fully-developed turbulent open-channel flow in combination with additional complexities, such as buoyancy or surfactants.

References and Links

- [1] https://www.ifh.kit.edu/26_1776.php
- [2] Wissink, J.G. and Herlina, H. 2016. J. Fluid Mech. 787, 508-540.
- [3] Liu X., 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.
- [5] Herlina, H. and Wissink, J.G. 2019. J. Fluid Mech. 860, 419-440.

Numerical cavitation erosion prediction

RESEARCH INSTITUTION

Chair of Aerodynamics and Fluid Mechanics, TUM

PRINCIPAL INVESTIGATORS

Theresa Trummer, Steffen J. Schmidt

RESEARCHERS

Bruno Beban

PROJECT PARTNERS

—

SuperMUC Project ID: pr92ho

Introduction

Cavitation refers to the formation and subsequent collapse of vapor filled cavities in liquids, both of which are initiated by pressure fluctuations. At collapse, intense shock waves with post-shock pressures of up to several thousand bar are emitted. When bubbles collapse near a solid wall, an asymmetric collapse behavior occurs, which is associated with the formation of high-velocity liquid jets directed towards the material surfaces. Structure loads caused by the pressure waves and the jet impact can lead to material erosion, which may be so strong that the performance degrades severely or devices may fail. Common examples adversely affected by cavitation erosion are ship propellers, hydraulic turbines and pumps, valves and throttles in hydraulic systems and fuel injection systems. On the other hand, the energy focusing and destructive forces at collapses are also exploited for material peening, surface-cleaning or biomedical applications. The latter include urinary stone ablation and targeted drug delivery. In injection systems, cavitation erosion can be advantageous for cleaning nozzle holes and throttles from surface deposits, promoting jet break-up, or stabilizing the mass flow. In order to better control cavitation erosion, it is necessary to understand the underlying phenomena and identify the decisive mechanisms for material damage. Small characteristic length scales, short timescales and intense pressure peaks make experimental studies challenging. Computational Fluid Dynamics (CFD) can provide time-resolved information and thus have become an important tool for the understanding of cavitation erosion, complementing experiments. In this research project, we numerically investigate the cavitation erosion potential of bubble collapses, cavitating jets and in valve chambers.

Method

The simulations are conducted with the flow solver CATUM [1], developed at the Chair of

Aerodynamics and Fluid mechanics. To capture the shock waves after the collapse, full compressibility of all phases is considered in our numerical model. For the cavitating two-phase flow, a homogeneous mixture model combined with an equilibrium cavitation model is used. The actual vapor-liquid interface of cavitation structures is not reconstructed and surface tension is thus neglected. For cavitating multiphase flows, we have extended the model by an additional non-condensable gas component [2].

Results

Bubble collapse in the vicinity of a wall

We have investigated the collapse dynamics and the erosion potential of single bubble collapses in the vicinity of a wall. Figure 1 shows a time series of a collapsing wall-attached bubble exposed to atmospheric pressure (1 bar). The aspherical pressure distribution leads to an indentation of the bubble and the formation of a wall-directed jet piercing the bubble. Then, the remaining torus collapses and afterwards rebounds. Our compressible flow solver, which considers phase change, enables us to capture the emitted pressure waves at collapse and also the rebound. Further, we have studied the effect of additional non-condensable gas in the vapor bubble, see bottom of Fig. 1, which leads to an enhanced rebound and a damped pressure impact. To gain a better understanding of the damage mechanism and the decisive processes, we have conducted




Figure 1: Collapsing wall-attached vapor bubble and comparison with collapsing vapor bubble containing non-condensable gas (bottom) (PhD Thesis Trummer, 2020).




Figure 2: Cavitating jet and pressure waves of collapsing vapor structures.

simulations under high driving pressures considering varying distances to the wall [5]. For this purpose, we have also evaluated the effect of the grid resolution on the collapse dynamics and the recorded maximum pressure. The finest grid for this study contained more than 670 million cells and the conducted simulations represent the highest grid resolution ever used for 3-D simulations.

Cavitating jet eroding a target plate

Cavitating jets are typical experimental configurations for the investigation of cavitation erosion. However, these experiments do not provide detailed information about the processes on short time scales and the transient pressure pulses. To obtain more insights, we have numerically studied such a configuration at different operating conditions. In these configurations, a high-velocity liquid jet causes shear layer cavitation, as shown in Fig. 2. Close to the target plate, the vapor structures collapse emitting intense pressure waves. To obtain numerical equivalents of experimentally observed erosion pits, we have used advanced clustering algorithms of Machine Learning frameworks [4].

Cavitation in a valve chamber of a common rail injection system

Current developments in diesel direct injection systems increase the rail pressure to more than 2500 bar. High injection pressures imply a high acceleration of the fuel and thus lead to cavitation and cavitation erosion. Cavitation erosion not only affects the boreholes of injection nozzles but also other components such as the valve chamber. We have investigated the flow field, cavitation dynamics and erosion potential in different valve chamber designs and different injection pressures [5]. Figure 3 shows the flow field in a generic chamber with an injection pressure of 2,000 bar. The flow velocity exceeds 600 m/s and leads to cavitation. The subsequent collapse of formed vapor structures causes high pressure loads on the surface of the pilot valve, as visualized in Fig. 4.

Ongoing Research / Outlook

Our studies helped us to better understand the underlying mechanisms of cavitation-induced material damage. With our investigations, we have covered a broad range of relevant applications, including generic bubble collapses and the flow in a valve chamber of a




Figure 3: Instantaneous flow field (left) and formed vapor structures (right) in a valve chamber for an injection pressure of 2,000 bar (PhD Thesis Beban, 2019).




Figure 4: Recorded wall pressure loads on the surface of a pilot valve (PhD Thesis Beban, 2019).

common rail injection system. The investigations of cavitating flows impose high requirements on spatial and temporal resolution, which can only be addressed by high-performance computers. Due to the high speed of sound in liquids, the time step size usually is on the order of less than a nano-second, while cavitation structures and wall effects require a high grid resolution and thus lead to a large number of cells. Our simulations have been conducted on SuperMUC, SuperMUC Phase 2 and SuperMUC NG using our massively parallelized numerical framework CATUM. Current research topics include 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. Further, we also work on the application of deep learning and artificial intelligence techniques for numerical cavitation erosion prediction, such as the development of advanced collapse clustering algorithms [4].

References and Links

- [1] <https://www.mw.tum.de/aer/> Schmidt, S. J. (2015). PhD Thesis. TU Munich.
- [2] Trummler, T., Freytag, L., Schmidt, S. J. & Adams, N. A., Proceedings CAV2018, Chapter 125, S. 656-659 (2018). doi: 10.1115/1.861851 ch125.
- [3] Trummler, T., Schmidt, S. J., Adams, N. A., Submitted to Int. J. Multiph. Flow.
- [4] Trummler, T., Thiery, F., Schmidt, S. J., Adams, N. A., Proceedings CAV2021 (2021).
- [5] Beban, B., Schmidt, S. J., Adams, N. A., Atomization and Sprays, 27(8), 723–739 (2017). doi: 10.1615/AtomizSpr.2017020387

Symmetry theory and turbulence

RESEARCH INSTITUTION

¹TU Darmstadt

PRINCIPAL INVESTIGATOR

Martin Oberlack¹

RESEARCHERS

Sergio Hoyas², Stefanie Kraheberger¹

PROJECT PARTNERS

²Universitat Politècnica de València

SuperMUC Project ID: pr921a

Introduction

Wall turbulence is probably one of the open problems in physics with most applications in daily life. Even if the equations ruling out these flows have been known for 150 years, we still lack a complete theory. As wall-bounded turbulence is responsible for **5% of the CO₂** dumped by humankind into the atmosphere every year, this is a problem of the uttermost importance. Research of turbulent flows has been dominated by experimental techniques until the eighties of the last century, where supercomputers started to be powerful enough to solve turbulent flows. However, due to the highly non-linear behavior of wall-turbulent flows, Direct Numerical Simulation (DNS) of these flows are restricted to simplified geometries. The most successful of these idealized flows are Poiseuille turbulent channels, where the fluid is confined between two parallel plates and the flow is driven by pressure. Since the seminal paper of Kim, Moin, and Moser, the friction Reynolds number Re_τ has grown steadily. In this project, we have been able to run a simulation reaching the $Re_\tau = 10,000$ frontier. This simulation would allow the study of high Reynolds numbers effects. This Re_τ is still less than the largest realization of the flow τ obtained by experimental means. This experiment reached $Re_\tau = 20,000$. However, the main advantage is that the τ DNS allows one to compute any imaginable quantity in the whole domain.

Unfortunately, the computational cost of a DNS is very high, such that the DNS of a commercial jet-airliner is several decades away, even for the biggest supercomputers available to date. Therefore, DNS is not well-suited for design purposes. For most applications, instantaneous flow details are unnecessary to know and statistical quantities such as the mean velocity are sufficient. However, considering turbulence as a statistical process in fact leads to an infinite dimensional hierarchy of moment-equations, which are extremely difficult to solve. Hence, for most applications a truncated system is considered at the expense of introducing semi-empirical closure models (i.e. RANS methods).

An approach developed by the PI and his co-workers follows a very distinctive route as it considers the

entire infinite dimensional hierarchy of moment equations using Lie symmetry group methods. Since the work of the mathematician Lie, symmetries have experienced a striking evolution at the heart of physics. Einstein in his 1905 seminal work on special relativity contemplated the symmetry principle as the key feature of physics. In the 1920s, quantum mechanics established symmetries as an axiomatic basis of physics in general. Today, Lie symmetry group methods act as the foremost guiding principle to understand and mathematically model new physical laws. This project is the first simulation ever that allow us to test the application of the symmetry-based theory in the field of turbulence.

Results and Methods

The program code, LISO, performs a turbulence simulation using the Navier-Stokes equations for an incompressible fluid in a plane channel between two infinite parallel plates. No modelling is used. The computation is carried out in a doubly periodic domain in the two wall-parallel directions x (streamwise) and z (spanwise), which is chosen large enough to minimize artefacts due to the spurious periodicity. The $Re_\tau = 10,000$ simulation uses $2\pi h \times \pi h$ where h is the channel half-thicknesses. This is enough to avoid spurious statistics due to the finite length of the box.

Numerical methods

For time integration, the code uses a low-storage third-order semi-implicit Runge-Kutta (R-K) time-stepper. This R-K scheme allows a longer time step than purely explicit methods while at the same time, the memory requirements are as low as those of a Euler method. For the spatial derivatives, a Fourier discretization in x - and z -directions and a sixth-order Compact Finite Differences (CFD) discretization in the wall-normal y direction are used. This allows a greater flexibility in the choice of the mesh, which consists of $6,144 \times 2,101 \times 6,144$ points in space.

LISO is written in Fortran90 and for SuperMUC, it is built with the Intel ifort compiler. It has been carefully checked and optimized for SuperMUC using a test account and in a SuperMUC extreme scaling workshop. The implementation uses FFTW routines for





Figure 1: Instantaneous visualization of the streamwise velocity at the logarithmic layer (left). Details of the streaks of the flow at the viscous layer (right).




Figure 2: Indicator function. The flat region between 200 and 2,000 indicates the existence of a very long logarithmic layer. The value of 0.39 for the Karman constant can be taken as almost definitive.




Figure 3: From bottom to top, scaling of moments of the streamwise velocity from 1 to 7. Dot line, data coming from DNS. Continuous line, data from theory.

the Fourier transforms (MKL library) and MPI + OpenMP for the parallelization, where Intel MPI is our preferred environment for MPI. I/O routines use the package HDF5, more precisely its parallel implementation on SuperMUC. It took less than 15sec to save the whole field of 250 GB, distributed in as many as 4,096 cores, to disk. This must be done each 7 hours, so the time needed for the I/O routines is negligible. Due to limitations on the code, the maximum number of processors is 4,096 cores. As we need to run the code for a very large time, more than two years, to get enough statistics from the turbulent field, we ran **three independent cases** at the same time. In total LISO has been running for about 50M Core-h. This simulation has provided many interesting results about the cinematics of boundary layers. In fact, this is the first simulation where a clear separation of scales can be observed. One of the most important features is that it has allowed us to check a symmetry-theory based of Turbulence, where new scaling laws are obtained directly from theory, without any further simplification or modelling. As one example, we can cite the indicator function of Figure 2. It was predicted around 1920 that the central part of this function should be a straight line. This indicates that the flow follows a logarithmic profile in this region, which is called, for obvious reasons, logarithmic layer. This layer is the most difficult to model as it is the

region of the flow where the largest eddies of the flow, coming from the center, interact with the smaller, and short live eddies that grow from the walls.

The symmetry-based theory not only predicts this logarithmic behavior but also offers scaling laws for all the moments of the streamwise velocity, and thus characterizing completely the flow. Figure 3 shows the scaling of the moments up to 7th order, with an excellent agreement between theory and data. Our final objective is to provide a better knowledge of turbulence and improve the modelling of this thrilling phenomena.

Ongoing Research / Outlook

The simulation is still running on SuperMUC-NG, and we expect to finish it in 2020. This project was only able to run SuperMUC Phase 2, as the memory of Phase 1 was not enough. We have found, approximately, a speed up of 25 % from SuperMUC to SuperMUC-NG. This improvement is due to both communications and computations. Moreover, if we run using 24 MPI-tasks and 2 OMP tasks per node, respectively, we have been able to achieve an additional speed-up of 40%. Apart from this, we have 100TB of raw data that will be analyzed during the next years. High memory nodes should be used for this.

References and Links

- [1] Statistics from these and other simulations can be downloaded from our webpage.
 [1] https://www.fdy.tu-darmstadt.de/fdyresearch/dns/direkte_numerische_simulation.en.jsp

MPI	Memory	Memory per Task	Restart file	Tape Archive
2,048	1.8 TB	886 MB	262 GB	100TB
4,096		443 MB		

Table 1: Memory requirements in a computational box of (6,144 x 2,101 x 6,144) ~8¹⁰ points.

Heat Exchange in Methane-based Launchers

RESEARCH INSTITUTION

Technische Universität München

PRINCIPAL INVESTIGATOR

Christian Stemmer

RESEARCHER

Raffaele Olmeda

PROJECT PARTNER

Bundeswehr Universität München

SuperMUC Project ID: pr92lo

Introduction

The operating pressure and temperature of combustion chambers of rocket launchers are often well above the critical pressure and temperature of the pure injectants.

This work studies what happens inside a rocket combustion chamber during operation time and analyzes the behavior of the propellants close to walls including wall roughness and an evaluation of the heat transfer. The difficulties of simulating such scenarios come from the extreme physical condition (particularly for pressure and density) and for the complex chemical behavior. Large Eddies Simulations were performed on multiblock structured grids and the results were then compared to the available literature.

The project is part of the SFB Transregio 40, a Germany-wide research collaboration of numerous universities, research institutions (DLR) and ArianeGroup [1] funded by the Deutsche Forschungsgemeinschaft (DFG) under contract no. 26293245.

Numerical Methods

Flamelet tables have been created in a pre-processing phase in combination with the counter flow diffusion flame model approach in order to model the chemistry inside the combustion chamber. The tables are generated by solving the Flamelet equations, which consists of a mass fraction transport equation for each specie considered in the combustion process and one energy equation. The particularity of setup is that the walls of the combustion chamber act as a negative source of enthalpy for the flame, therefore the tables have been modified such that non-adiabatic processes are taken into account.

The normalized sand grain roughness is used as a measure of roughness, which is a physical quantity depending on the average roughness of the wall, the type of roughness and the friction velocity. The normalized sand-grain roughness is then used to model the effects of wall irregularities on the velocity and temperature fields which directly influences the heat transfer and the cooling requirements

The wall roughness plays its role in the Wall Model, which solves the Turbulent Boundary Layer Equations (TBLE). Three different roughness methods have been tested. The methods proposed by Cebeci et al. [2] and Feiereisen et al. [3] modify the turbulent viscosity of the TBLE, causing a downward shift of velocity and temperature profiles. The method proposed by Saito et al. [4] imposes a virtual slip velocity at the boundary between the LES grid and the TBLE grid. All three methods depend on the normalized sand-grain roughness.

Results

A correct estimation of the wall heat flux of a combustion chamber is extremely important because the design of a rocket engine needs to consider the structural limitations of the chamber walls. A methane-based combustion chamber has been simulated to study the thermodynamic behavior near the wall. The geometry of the setup has been simplified as a rectangular section.




Figure 1: Performance of the three tested roughness methods with respect to different values of Reynolds number and sand grain roughness.

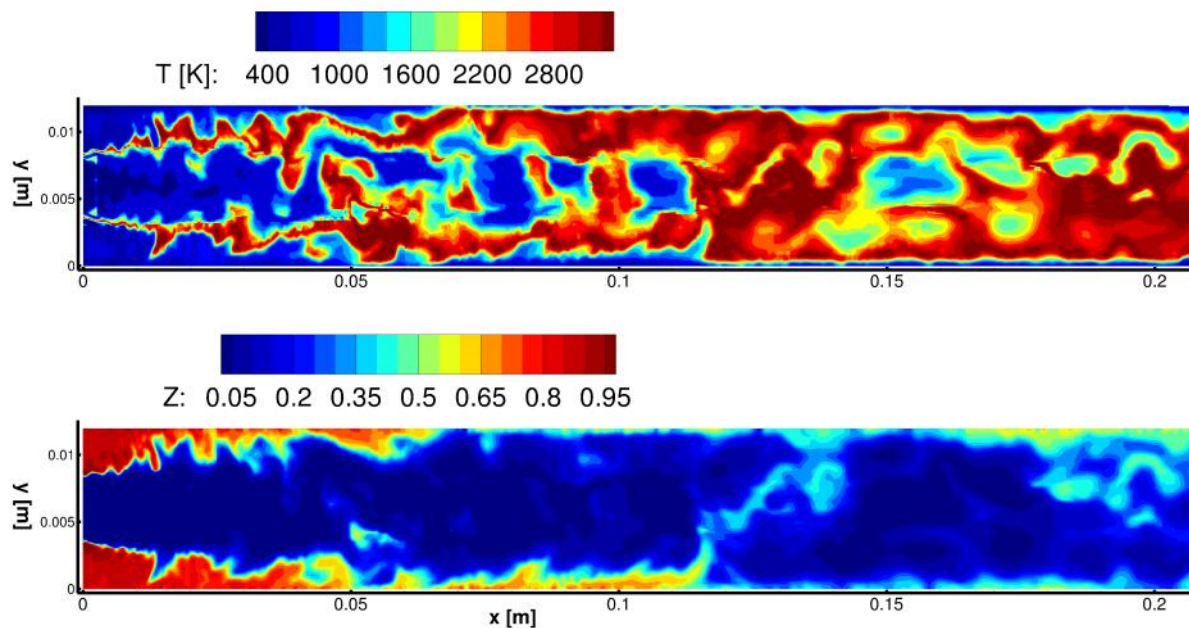


Figure 2: Snapshot of the temperature field (upper) and the mixture fraction (lower) in the methane-based combustion chamber.

Different configurations of the combustion chamber have been simulated. Simulations were performed to assess the effect of a change of operating pressure, presence of the cooling film, blowing ratio of the cooling film, and the influence of the type of the equation of state. Meshes with increasing refinements have been used to observe the convergence of the results and ensure the smaller scales were captured. In Fig. 2, a snapshot of the temperature field shows how the mixing of fuel and oxidizer is realized at half of the length of the combustion chamber.

A periodic channel has been chosen as domain for the roughness simulations, which guarantees a constant thickness of the boundary-layer during the simulation. In this case, meshes of increasing dimensions have been used once again, particularly to investigate how the wall model used to model the roughness would interact with different wall refinements. In Fig. 1, the influence of the wall roughness and of the Reynolds number on the relative error of the velocity field is shown. The Saito method has proven to be the most effective, and simulations with higher Reynolds number yielded better results.

More extensive results can be found in the final report of the project [5]. For the heat transfer simulations, 9 M Core-h were used, divided in meshes of different refinement. For the roughness simulations, 2.9 M Core-h were necessary, considering different meshes refinements and roughness methods.

Ongoing Research / Outlook

SuperMUC and its successor SuperMUC-NG have provided excellent means to efficiently perform large-scale simulations in a HPC environment. The support provided by the support team has been always fast and effective.

The methods developed so far allow to simulate the behavior of complex thermodynamic environments such as the interior of rocket-engines with a high level of accuracy. The results have particularly shown how the type of energy equation used in the Wall Model influences the quality of the results. An enthalpy equation has proved superior to a temperature equation both considering the comparison with the experiment and the convergence velocity. Furthermore, the advection term of the energy equation has proven to be important right after the injector plate, where a recirculation region is present.

The research in the future will focus on the extension of the Flamelet tables for real-gas behavior, which is of great interest for film-cooled modern combustion chambers. Cryogenic fuel is usually used as coolant, and therefore thermodynamic states close to the critical point are reached. An ideal-gas approach in this case is no longer valid, and a cubic equation of state must be implemented in the Flamelet table generation.

References and Links

- [1] <https://www.sfbtr40.de/en>.
- [2] T. Cebeci and K. C. Chang, *AIAA Journal*, pp. 730–735, 1978.
- [3] W. J. Feiereisen and M. Acharya, *AIAA Journal*, pp. 1642–1649, 1986.
- [4] N. Saito, D. I. Pullin and M. Inoue, *Physics of Fluids*, pp. 75–103, 2012.
- [5] R. Olmeda, P. Breda, C. Stemmer, and M. Pfitzner, “Large-Eddy Simulations for the Wall Heat Flux Prediction of a Film-Cooled Single-Element Combustion Chamber,” *Future Space-Transport-System Components under High Thermal and Mechanical Loads*, ISBN 978-3-030-53847-7, pp. 223–234, Springer International Publishing (2021).

Enhanced Aerodynamics of Wind Turbines

RESEARCH INSTITUTION

Institute of Aerodynamics and Gas Dynamics, University of Stuttgart

PRINCIPAL INVESTIGATOR

Thorsten Lutz

RESEARCHERS

Giorgia Guma, Florian Wenz, Ferdinand Seel, Pradip Zamre

PROJECT PARTNERS

—

SuperMUC Project ID: pr94va

Introduction

Within the present project, the aerodynamic behavior of modern wind turbines has been investigated by the use of Computational Fluid Dynamics (CFD). Three main sub-topics have been studied (results are presented for the first two topics):

- Aeroelastic effects on wind turbines exposed to turbulent inflow conditions and complex terrain.
- Low-frequency emissions from a flexible wind turbine under atmospheric conditions.
- Aeroacoustics analysis on Vertical Axis Wind Turbines (VAWT) [1].

Results and Methods

All simulations have been performed with the finite-volume CFD code FLOWer, originally developed by the German Aerospace Center (DLR) within the MEGAFLOW project [2]. It solves the Navier Stokes equations in an integral form using different turbulence RANS and hybrid RANS/LES models. The simulations requiring the highest resources were the ones regarding inflow turbulence and complex terrain, where around 450,000 Core-h and around 2,500 Cores have been used for each case. The available budget in \$WORK is around 47% of the total, while its usable space is 66% of the budget. The applied flow solver is not producing an exceeding amount of files that is why, up to now, no problem was faced in the file storage. Within the research project WINSSENT [3] and in particular in its subpart FoWEA (Forschungswindenergieanlage), the effects of aeroelasticity on wind turbines subjected to complex atmospheric inflow

conditions are investigated. A complex terrain is a terrain where topology and roughness influence the atmospheric boundary layer. In the WINSSENT project a research turbine (Figure 1, left) in complex terrain with turbulent inflow (Figure 1, right) is simulated by the use of an explicit coupling between CFD for the flow solutions and MBD/FEM for the structural response. Numerical models with different geometric complexity are used and compared. In Figure 2 (left) an only blade CFD model in uniform inflow is computed first in rigid conditions and then flexible. As it can be seen, the power increases by about 2.5%. This is due to the fact that the blade has a precone angle in rigid state, and therefore the bending of the flexible blade increases the effective rotor disk area leading to a higher power. In Figure 2 (right), the full CFD model is simulated firstly rigid and then completely flexible. Here it can be observed that the impact of blade-tower interaction is increased leading to a stronger load reduction during the tower passage. The peaks resulting from the wind fluctuations are mostly damped by the consideration of flexibility, resulting in an average lower produced power. The aim of the IEA Tasks 29 and 39 are the improvement of aerodynamic models for wind turbine design codes and assessment of noise. A high fidelity time resolved fluid-structure CFD approach was used to investigate the low-frequency emission of the DANAERO wind turbine. Firstly, the aeroelastic behavior of the turbine was investigated by means of CFD simulations coupled to the MBS solver SIMPACK for the structural response increasing the complexity from an only one blade model in uniform inflow conditions, to the full model in turbulent inflow conditions [4]. Secondly, the aeroacoustic noise propagation was computed in post-processing using a Ffowcs-William Hawking method,

following the same procedure developed during the TremAc project and prescribed in [5]. Regarding the aeroacoustic emission it was found that it is dominated by tonalities due to blade passage in front of the tower and its higher harmonics and that more broadband noise is generated in flow direction than in lateral for both observers at 100 m and 1,000 m distance. For this last observer position the Blade Passing Frequency (BPF) is not the highest peak. The tonalities are




Figure 1: Blade deformation resulting from CFD-MBS coupled simulation at uniform inflow (left). CFD-MBS coupled simulation in complex terrain (right).




Figure 2: Only blade model, rigid vs. coupled in uniform inflow (left). Full turbine in complex terrain and turbulent inflow, rigid vs. coupled (right).




Figure 4: Spectra of unweighted SPL with observer at 1,000 m distance for a turbulent inflow case where the entire turbine is flexible. Directivity of the only rotor.

dominated by the tower emission (Figure 3) because the passing blade induces pressure fluctuations on the tower surface itself. It can be observed that at an observer position of 1,000 m, the rotor emits broadband noise due to the inflow turbulence and with high intensity in the flow direction (Figure 4). The tower shows no directivity, emitting almost cylindrical at very discrete frequencies, while the rotor emits more in flow direction than in lateral, showing almost no tonality. The fluid-structure coupling has only an indirect effect on the aeroacoustic and that is because when the blades are flexible, the distance blade-tower reduces and thus the blade-tower interaction is increased.

Ongoing Research / Outlook

The project has been prolonged in order to study the additional following topics:

- Influence of operational control on low-frequency noise components.
- Analysis of the flow around thick airfoils equipped with Vortex Generators (VG).
- Influence of orography, forest and urban terrain on the flow field [6].

Indiana Wind (Interdisziplinäre Analyse und Optimierung von Windenergieanlagen und ihren Komponenten) is a joint project financed mostly by the BMWI. It deals with the analysis and optimization of an entire wind turbine including its components under consideration of the disciplines aerodynamics, aeroacoustics, structure, control and terrain. With regard to acoustic radiation, the primary objective is to outline the influence of operational control on low-frequency noise components. A further field of




Figure 3: Spectra of unweighted SPL with observer at 1000 m distance for a turbulent inflow case where the entire turbine is flexible. Influence of the different components.







Figure 5: Influence of orography and urban terrain on flow field.

application of the hybrid RANS/LES processes in this project is the flow around blade segments that are equipped with Vortex Generators (VG). At first, the detailed flow physics of the VGs for delaying flow separation is evaluated. There after the acoustic impact is deeply investigated. Within the joint graduate research Training Group program Windy Cities the economic use of small wind turbines for local distributed power generation in urban areas is investigated. Due to interaction between atmospheric boundary layer and various obstructing objects present in built environment, the inflow conditions are complex with high turbulence level, continuously changing wind direction, inclination etc. (see Figure 5). Therefore high fidelity simulations of the flow field around the building are required to assess the performance and loads of small vertical axes wind turbines.

References and Links

- [1] A. Dessoky, et al., Renewable Energy, 2019.
- [2] N. Kroll, et al., Aerosp. Sci. Technol., 2000.
- [3] WINDFORS, "WindFors Testfeld," [Online]. Available: <http://www.windfors.de/testfeld.html>.
- [4] G. Guma, et al., Wind Energy Science Discussions (2020).
- [5] L. Klein, et al., Wind Energy Science Discussions, 2019.
- [6] P. Zamre, et al., "Numerical studies of the impact of urban terrain on the loads and performance of a small vertical axis wind turbine", Torque Conference 2020, Submitted.

Earth, Climate and Environmental Sciences



Regionalization of global hydrometeorological fields

RESEARCH INSTITUTION

¹Karlsruhe Institute of Technology (KIT), Institute of Meteorology and Climate Research – Atmospheric Environmental Research (IMK-IFU), Campus Alpin, Garmisch-Partenkirchen

PRINCIPAL INVESTIGATOR

Tanja Portele¹

RESEARCHERS

Patrick Laux^{1,2}, Christof Lorenz¹, Harald Kunstmann^{1,2}

PROJECT PARTNERS

²Institute of Geography, Augsburg University of Augsburg, Augsburg

SuperMUC Project ID: pr27ko

Introduction

Increasing demands for water till 2050 will aggravate the impacts of water scarcity on agricultural production and livelihood activities [1]. To deal with this projected development, sustainable management of available water resources is required. For decision support in regional water management or agriculture, regionalized weather and climate forecasting increasingly matters. This is especially the case for semi-arid areas where more profound knowledge of variations in seasonal rainfall can help to better prepare against climate extremes. The major task of the SaWaM (Seasonal Water Resources Management in Semiarid Regions [2]) project therefore is the development of methods and tools for the practice transfer of regionalized global data for water resources management. Here, special focus is on the regionalization of global hydrometeorological fields over the semi-arid regions of Northeast Brazil and Ecuador/Peru. Hit by a multi-year drought in the last decade [3] and strongly influenced by El Niño [4], respectively, these regions are facing high climate variability especially during their rainy seasons. Water availability in those regions not only determines the production of rainfed and irrigated

agriculture, it also controls the hydroelectric energy production in their highly managed river basins.

One way to achieve regionalized hydrometeorological information is the application of regional atmospheric models, such as the Weather Research and Forecasting (WRF) model, to dynamically downscale global fields over selected areas. However, these models need to be adapted in their physical setting to the region of interest. Sub-grid scale processes such as microphysics (MP), cumulus physics (CU), planetary boundary layer physics (PBL) and radiation physics (RA) are still parameterized in these models, and their choice and combination highly influences the modeled hydrometeorological variables precipitation and temperature [5]. Therefore, within this SuperMUC-NG project, we tested the performance of different WRF parameterization combinations to find an optimized model setup for the simulation of regionalized hydro-meteorological fields in the selected semi-arid regions.

Results and Methods

As a community numerical weather prediction model, WRF is used for a wide range of both research and

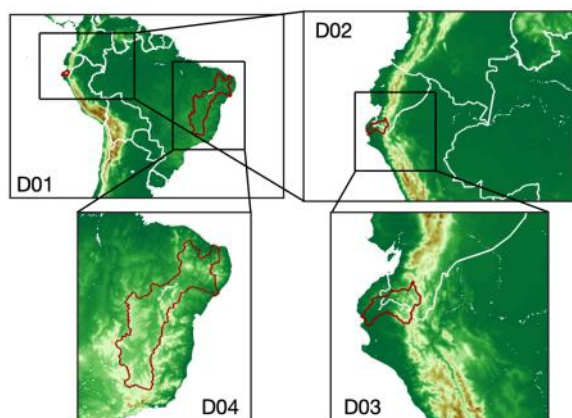


Figure 1: WRF nested domains over South America. From the input data of ERA-Interim at 80 km horizontal resolution, the domain resolution is stepwise brought to 27 km (D01), 9 km (D02 and D04), and further to 3 km (D03). Country borders are shown with white lines, basin borders of the São Francisco river in Northeast Brazil and the Catamayo-Chira river in Ecuador/Peru with red lines.




Figure 2: Median (red dot) and total range of averaged required computing time per simulated time step using 576 cores (12 nodes) at SuperMUC-NG for runs of different WRF parameterization groups. One simulation included all four WRF domains (D01-D04, Fig. 1). The cumulus parameterization groups (G3D, Tiedtke, BMJ, KF) consist of 8 members, all others consist of 16 members using the respective scheme.

	Storage
WRF output (6-hourly)	25.9 TB
WRF output nc4 (6-hourly)	5.2 TB
WRF diagnostics (daily)	1.1 TB
WRF restart files (monthly)	8.4 TB

Table 1: Required storage of different WRF output files.

operational applications. Applying a nested approach in WRF (Fig. 1), the horizontal resolution of global ERA-Interim reanalyses is step-wise brought from 80 km to 27 km (D01), to 9 km (D02 and D04), and finally to 3 km (D03). The last step to 3 km was only performed over the highly mountainous Ecuadorian-Peruvian region. Over the Brazilian region, the final horizontal resolution is 9km.

The selection of tested parameterization schemes was based on literature review and recommendations from local partners. In sum, simulations were conducted for a WRF ensemble of 32 combinations of 4 CU (G3D, Tiedtke, BMJ, KF), 2 MP (WSM3, WSM6), 2 PBL (YSU, ACM2) and 2 RA (RRTMG, RRTM) schemes. To test 32 combinations of physics parameterization for 2.5 years (Jan 2006 – Jun 2008) with nested dynamical downscaling, high performance computing resources are inevitable.

2.5 years for each of the 32 parameterization combinations result in 960 simulation months and required 18.5 Mio core hours including several repeated runs due to model instabilities and incomplete output writing. Using 576 cores (12 nodes) at SuperMUC-NG, the average required computing time for the D01-D04 configuration per simulated time step for all runs was 0.405 s. This required computing time differs especially among different parameterization groups of runs using the respective schemes of microphysics, planetary boundary layer physics and radiation physics (Fig. 2). Maximum resources were required by runs using WSM6, ACM2 and RRTMG.

The already reduced (only surface variables) 6-hourly WRF output for all four domains required in total 25.9 TB of storage capacity (Table 1). Compression of the files allowed a reduction to 5.2 TB. WRF daily diagnostics output resulted in 1.1 TB. Monthly restart files for continuous model runs over the 2.5 years made up in total 8.4 TB. Including more output variables, simulations with the RRTMG scheme doubled the needed storage relative to RRTM runs.

Compared against gridded observation data by Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS) for the main three months (JFM) of the rainy season in 2007 (Fig. 3), the WRF ensemble median (ENS) shows a dominating positive bias. The G3D and KF cumulus groups simulate far too much precipitation, whereas the Tiedtke and BMJ cumulus groups produce less precipitation. Bias patterns of MP (WSM3, WSM6), PBL (YSU, ACM2), and RA (RRTMG, RRTM) physics groups are similar with too wet central regions and too dry northern coasts and southwest areas. The largest impact on modeled precipitation is therewith found for the cumulus and radiation physics

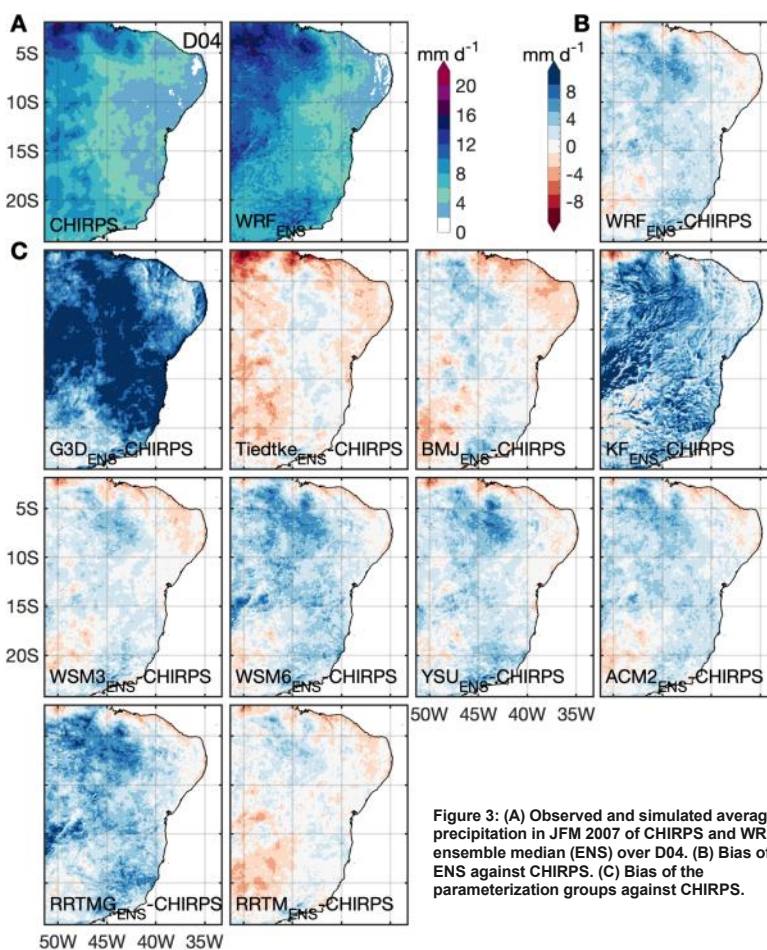


Figure 3: (A) Observed and simulated average precipitation in JFM 2007 of CHIRPS and WRF ensemble median (ENS) over D04. (B) Bias of ENS against CHIRPS. (C) Bias of the parameterization groups against CHIRPS.

groups. Compared to CHIRPS, precipitation biases are lowest for runs with model settings using Tiedtke or BMJ cumulus physics, WSM3 microphysics and RRTM radiation physics.

Ongoing Research / Outlook

The combination of best performing parameterization groups does not have to result in the overall best performing model setup. Single combinations may generate interactions between the physics schemes, leading to unexpected performances. The final optimized WRF setup for each domain is still under investigation and will also depend on the required resources of computing time and storage.

An originally planned evaluation period over 5 years could not be achieved due to the combination of a limited computing time per job of 2 days, during which not all runs were able to simulate 2 months (monthly restarts), and the long job queues at SuperMUC-NG. The switch from SuperMUC Phase 2 to SuperMUC-NG reduced computing times per simulated time step by 15-20%.

References and Links

- [1] Food and Agriculture Organization of the United Nations (2015): <http://www.fao.org/3/a-i4560e.pdf>
- [2] SaWaM Project: <http://grow-sawam.org>
- [3] Martins, E. et al. (2018): Bull. Amer. Meteorol. Soc. 99, S65–S69, DOI: 10.1175/BAMS-D-17-0102.1.
- [4] Dominguez-Castro, F. et al. (2018): Int. J. Clim. 38, 2006–2014, DOI: 10.1002/joc.5312.
- [5] Klein, C. et al. (2015): Clim. Dyn., 45, 2733–2755, DOI 10.1007/s00382-015-2505-5.

High resolution combined global gravity field

modelling

RESEARCH INSTITUTION

Chair of Astronomical and Physical Geodesy, Technical University of Munich

PRINCIPAL INVESTIGATOR

Thomas Gruber

RESEARCHER

Philipp Zingerle

PROJECT PARTNERS

—

SuperMUC Project ID: 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 spheroidal harmonic series, the parameters to be estimated in gravity field modelling are spheroidal 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, this procedure shall result in a set of spheroidal harmonic coefficients representing the global Earth gravity field up to highest possible resolution.

Results and Methods

To achieve an optimal combination of the different data sources, the method of least squares adjustment (LSA) is applied. While LSA is probably the numerically least demanding optimization method, the complexity of solving such (general) LSA problems nevertheless grows with the power of three w.r.t. the number of parameters. As the number of parameters itself increases quadratically with the spatial density, the final (numerical) complexity increases with the power of six with regard to the target resolution of the gravity field. The applicability of the general LSA approach has therefore been effectively limited to a spatial resolution of 15' (about 30km) on SuperMUC Phase 2 (creating/solving/inverting matrices of a size of about 2TB).

As the goal is to further increase the resolution beyond 15', some limitations (such as equally weighted observations) have to be introduced to achieve sparsity in the normal equation system of the LSA. Exploiting this sparsity allows one to raise the spatial resolution up to 1' (about 2km), reaching the practical limits of the spheroidal harmonic parametrization. Combining the sparse approach for the short wavelengths with the general LSA for the longer wavelengths leads to the recently published global gravity field model XGM2019e (see [1], and Fig. 1).

To solve the stated problems the GST (Geodetic Sparse Toolbox) software suite has been implemented with the aim to provide a maximum scalability on distributed memory systems as SuperMUC(-NG). Thus, the software has been written in Fortran (F08) using the MPI-2 standard as primary communication interface and for parallelized I/O-operations. To




Figure 1: Gravity anomalies derived from XGM2019 up to d/o 719. Gravity anomalies reflect the deviation of the gravity acceleration from that of a homogeneous reference ellipsoid). $1[\text{mGal}] = 10^{-8} [\text{m/s}^2]$.




Figure 2: DTU18 MSS on a 1' grid in the area of the Solomon Sea (left). Gravity anomalies derived from DTU18 MSS (and DTU17 MDT) using a direct global spheroidal-harmonic analysis approach (right, same region).

maximize node level performance all calculation intense operations are outsourced to appropriate numerical libraries (such as PBLAS or SCALAPACK) from the Intel MKL collection.

Ongoing Research / Outlook

While on SuperMUC Phase 2 the maximum resolution of the general LSA has been limited to about 15', with SuperMUC-NG it is now feasible to increase this resolution to 5'. Solving such a LSA requires approx. 200TB+ of system memory and more than 7M core-hours, exhausting the capabilities of SuperMUC-NG. The computation of such a system is the main objective of the recently approved follow-on project of pn98be.

The ability to solve large sparse systems that has been developed during the processing of XGM2019e also opens doors for solving other problems. As an example, the problem of inverting geoid heights (in terms of a mean sea surfaces) to gravity anomalies is now directly solvable up to the available resolution of the data (see Fig. 2).

References and Links

- [1] Zingerle, P.; Pail, R.; Gruber, T.; Oikonomidou, X., J Geod 94, 2020. DOI: 10.1007/s00190-020-01398-0.

Advanced Simulation of Coupled

Earthquake-Tsunami Events

RESEARCH INSTITUTION

Technical University of Munich and LMU Munich

PRINCIPAL INVESTIGATORS

Michael Bader, Alice-Agnes Gabriel

RESEARCHERS

Ravil Dorozhinskii, Lukas Krenz, Elizabeth Madden, Leonhard Rannabauer, Sebastian Rettenberger, Thomas Ulrich, Carsten Uphoff, Sebastian Wolf, Stephanie Wollherr

PROJECT PARTNERS

Universität Hamburg and ETH Zurich

SuperMUC Project ID: pr45fi

Introduction

In the framework of the ASCETE (Advanced Simulation of Coupled Earthquake and Tsunami Events) project, the computational seismology group of LMU Munich and the high performance computing group of TUM jointly used the SuperMUC HPC infrastructures for running large-scale modeling of earthquake rupture dynamics and tsunami propagation and inundation, to gain insight into earthquake physics and to better understand the fundamental conditions of tsunami generation. The project merges a variety of methods and topics, of which we highlight selected results and impacts in the following sections.

Results and Methods

We link physics-based models of subduction zone geodynamics and seismic cycling, three-dimensional dynamic earthquake rupture, and tsunami generation, propagation, and inundation (Figure 1, [1]). A long term subduction zone geodynamics and seismic cycling model is used to constrain the initial conditions of 3D dynamic earthquake rupture models. These conditions include the lithology, stress field, fault geometry, and fault strength, which are physically consistent with one another due to their development together over many slip events in the subduction scenario. 3D dynamic rupture simulations are then run and the time-dependent seafloor displacements are used to dynamically source a potential tsunami. Such linked models allow evaluating the effects of certain earthquake characteristics on tsunami behavior. These methods facilitate research into the physical relationships between processes operating across the spatial and temporal time scales of long-term deformation, earthquake rupture, and tsunami propagation.

In the model pipeline, three-dimensional dynamic rupture simulations pose an extreme computational challenge and require extensive supercomputing resources. In order to make these computationally feasible, we heavily optimised the software package SeisSol during the course of this project. These optimizations include a new cache aware wave propa-

gation scheme, optimizations of the dynamic rupture kernels using code generation, a novel clustered local-time-stepping scheme for dynamic rupture and the development of asynchronous outputs to overlap I/O and compute time. Our largest model of the 2004 Sumatra-Andaman earthquake ran on the full SuperMUC Phase 2 for 14 h [3]. Our simulation, featuring an unstructured tetrahedral mesh of 221 million elements and sixth-order accuracy in space and time, solves a discretized PDE system of 111 billion degrees of freedom. This allows the earthquake dynamics to be properly resolved as well as the wavefield up to 2 Hz in the surrounding media.

In follow-up studies of the 2004 Sumatra-Andaman earthquake, we focused on the controlling mechanisms of the event kinematics and dynamics, using new dynamic rupture scenarios and tsunami simulations. Our new scenarios suggest that along-depth variation of trench sediments, including off-fault plastic yielding [4], as well as along-arc variations of regional stresses and tectonic convergence rates are the dominant factors controlling the event's dynamics and kinematics. Depending on the intensity of the plastic wedge failure (Figure 2), the earthquake may have produced a narrow band of shallow and large slip. This would have translated into localized high ground displacement, invisible to tele-seismic and near-field geodetic measurements, but able to modulate locally the tsunami.

We applied combined earthquake-tsunami modeling to the MW 7.5 Palu earthquake, which occurred on September 28th, 2018, ruptured a 180 km long section of the Palu-Koro strike-slip fault system, in Sulawesi, Indonesia. The earthquake ruptured predominantly southward and beneath Palu Bay, a narrow inlet with a length of 30 km, until it stopped after a total rupture time of 30–40 s. The Palu earthquake triggered a local but powerful tsunami that devastated the coastal area of the Palu Bay quickly after the earthquake, characterized locally by inundation depths of over 6 m and run-up heights of over 9 m. Devastating tsunamis associated with submarine strike-slip earthquakes are rare. In fact, their associated ground displacements,




Figure 1: Illustration of the coupling framework between a 2d geodynamic software, a 3D dynamic rupture model and a Tsunami code [1].




Figure 2: Off-fault plastic strain accumulated for a weak (a) and a stronger (b) wedge. A weak wedge results in less slip to the trench, but enhance uplift at an intermediate distance of the trench.





Figure 3: Dynamic rupture multi-physics scenario of the 2018 MW 7.5 Palu earthquake and tsunami. Top: Snapshot of the wavefield (absolute particle velocity in m/s) across the fault network and of the fault slip rate after 15 seconds of simulation time. A shear Mach front, caused by the supershear rupture, is clearly visible. Right: Snapshots of the tsunami scenario at 600 s showing the dynamic mesh adaptivity of the model (from [2]).

predominantly horizontal, not vertical, do not favor tsunami genesis.

Using combined earthquake and tsunami models of the 2018 Sulawesi event (Figure 3), we were able to identify the drivers of the deadly tsunami in Palu last year. The key finding is that mean vertical offset beneath Palu Bay of 1.5 m is directly caused by the earthquake, generating a tsunami that matches field observations without any contribution from landsliding. These findings may lead to rethinking tsunami hazard in other regions worldwide, where similar fault systems cross beneath narrow Gulfs, such as the elongated Bodega and Tomales bays in northern California, USA, which host major segments of the San Andreas fault system.

Ongoing Research / Outlook

There remain unresolved questions in the (unidirectional) coupling of physical models used in ASCETE. E.g., how does one filter out fast waves travelling



along the ocean bottom which cannot be represented by the tsunami model? Or how shall sharp gradients in the bathymetry be handled? In order to answer such questions, we plan to implement fully coupled earthquake-tsunami models in which the ocean is modeled as an acoustic layer with a gravitational boundary condition on top of the solid domain of an earthquake.

References and Links

- [1] E. Madden et al., *Geophys. J. Int.*, 224.1, pp. 487–516 (2021). <https://doi.org/10.1093/gji/ggaa484>
- [2] Ulrich, T. et al., *Pure Appl. Geophys.*, 176.10, pp. 4069–4109 (2019).
- [3] Uphoff, C. et al. (2017). "Extreme Scale Multi-physics Simulations of the Tsunamigenic 2004 Sumatra Megathrust Earthquake". In: *Proceedings of the International Supercomputing Conference SC '17*. Denver, Colorado: ACM, 21:1–21:16.
- [4] S. Wollherr, A. Gabriel, C. Uphoff, *Geophys. J. Int.*, 214.3, pp. 1556–1584 (2018).

Towards an Exascale Hyperbolic PDE Engine

RESEARCH INSTITUTION

¹Technical University of Munich

²Ludwig-Maximilians-University of Munich

³Durham University

⁴University of Trento

⁵Institute for Advanced Study, Frankfurt

PRINCIPAL INVESTIGATORS

Michael Bader¹, Alice-Agnes Gabriel², Tobias Weinzierl³, Michael Dumbser⁴, Luciano Rezzolla⁵

RESEARCHERS

Jean-Matthieu Gallard¹, Leonhard Rannabauer¹, Anne Reinarz¹, Philipp Samfaß¹, Kenneth Duru², Dominic E. Charrier³, Benjamin Hazelwood³, Francesco Fambri⁴, Maurizio Tavelli⁴, Luke Bovard⁵, Sven Köppel⁵

PROJECT PARTNERS

RSC Group, BayFor, Leibniz Supercomputing Centre

SuperMUC Project ID: pr48ma

Introduction

Hyperbolic conservation laws model a wide range of phenomena and processes in science and engineering – from various areas of fluid dynamics via seismic wave propagation in earthquake simulation to extreme objects in astrophysics, such as neutron stars or black holes. The ExaHyPE engine [2] has been developed to solve an as big as possible class of hyperbolic systems of partial differential equations (PDEs) using high-order discontinuous Galerkin (DG) discretisation with ADER (Arbitrary high order DERivative) time stepping and a-posteriori sub-cell finite volume limiting on tree-structured Cartesian meshes. ExaHyPE thus relies on a well-defined numerical model and mesh refinement strategy, but strives for utmost flexibility regarding the tackled hyperbolic model.

The ExaHyPE SuperMUC-NG project accompanied a Horizon 2020 project [1] to develop the ExaHyPE engine [2], together with a suite of example models [3] and two large demonstrator applications from earthquake simulation and from relativistic astrophysics (see, e.g., [4,5]).

The ExaHyPE Engine

Users of the ExaHyPE engine need to specify details (number of quantities, discretisation order, etc) of their desired hyperbolic PDE system and model setup via a specification file. From this, the ExaHyPE Toolkit creates application-specific template classes, glue code and core routines, which are tailored to the application and target architecture. Application developers then implement PDE-specific functionality (flux functions, initial and boundary conditions, e.g.) within this generated code frame. ExaHyPE builds on Peano (www.peano-framework.org) as framework for adaptive mesh refinement, which provides shared- and distributed-memory parallelism via Intel Threading Building Blocks and MPI, respectively. Figure 1 illustrates the components of the ExaHyPE engine.

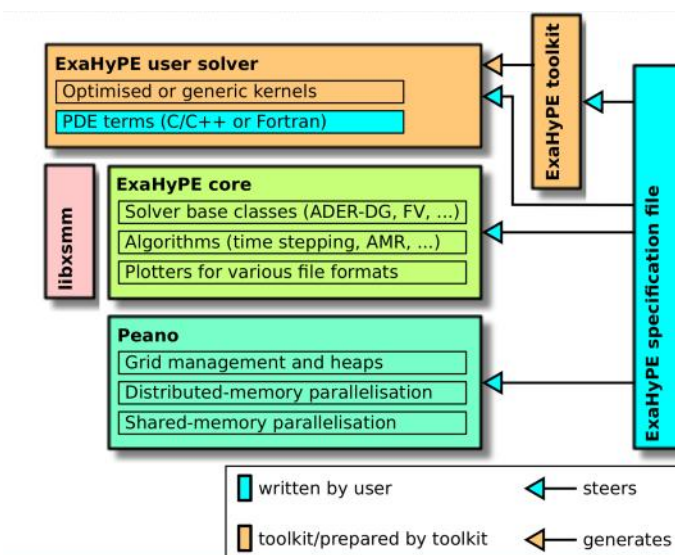


Figure 1: Schematic view of the ExaHyPE software architecture and components.

Earthquake Simulation

One of ExaHyPE's main target applications is the simulation of earthquakes and other scenarios governed by seismic wave propagation. In particular, we study the scattering of seismic surface waves in regions with complex topography properties, which is especially challenging for steep valleys or summits – an example is shown in Figure 2. To capture respective effects with high accuracy it is important to represent the topography as detailed as possible. This is often addressed by geometry- and feature-following meshes, for which the generation is often tedious and requires user intervention.

To avoid such mesh generation issues, we developed two methods that avoid or completely automate this problem. The first is a Diffused Interface Method, which represents complicated geometry via a color function [5], the second is a curvilinear-mesh method




Figure 2: Curvilinear mesh approximating the alpine region around Zugspitze – seismic waves emerging from a kinematic point source.

for which we exploit known properties of a simulated region to generate the mesh in-situ at initialization time. Both methods allow fully automatic geometry approximation, which makes them particularly attractive for Uncertainty Quantification or urgent seismic computing scenarios, where multiple simulation runs with a fast and automatic setup are required.

Relativistic Astrophysics

In astrophysics, the target application was to simulate the collapse of a binary system of two neutron stars. Such a merger will possibly create a black hole and a strong signal of gravitational waves. These scenarios are substantially more complex than the collapse of two black holes, which was the source of the 2016's first experimental observation of gravitational waves. Predicting the exact signal of such events via simulations is substantial to experimentally identify and interpret observed black-hole and neutron-star events.

In the project we developed high-order ADER-DG methods for various PDE models that approximate the problems described above with increasing model complexity. In addition to excellent properties w.r.t. parallel scalability, DG methods promise high order of accuracy for smooth solutions and, in general, solutions that are less dissipative and dispersive than more widely used methods such as finite volumes. However, the development of high-order DG methods for models that contain the full complexity of relativistic astrophysics also pose numerical challenges that have not been solved yet.

To solve respective models in the ExaHyPE engine, the system of PDEs must be recast into hyperbolic, first-order form. Deriving this first-order form of the equation (FO-CCZ4) was non-trivial and one of the early outcomes of the project. Using SuperMUC-NG we have run a sequence of models and scenarios of increasing complexity. We initially achieved the numerically stable evolution of isolated compact astrophysical objects, such as a single black hole with




Figure 3: Numerical simulation of a stable neutron star in anti-Cowling approximation at time $t = 1,000$ using the new hyperbolic GLM curl cleaning approach [4]. Contour surfaces of the conformal factor φ .

the FO-CCZ4 PDE and a single neutron star (TOV solution) with the GR(M)HD PDE, both solved for the first time with the novel ADER-DG scheme implemented in ExaHyPE [4].

Ongoing Research / Outlook

The ExaHyPE engine and the models based on ExaHyPE are further developed in various projects, such as the Horizon 2020 projects ChEESE (centre of excellence for exascale computing related to solid earth, www.cheese-coe.eu) and ENERXICO (energy and supercomputing for Mexico, www.enerxico-project.eu), the project TEAR granted to PI Alice Gabriel as an ERC Starting Grant (www.tear-erc.eu) or the UK-funded project ExaClaw (PI Tobias Weinzierl). For CHEESE, in particular, we develop a pilot demonstrator for a future urgent computing service to assess potential hazards quickly after a detected earthquake, which is also the topic of a follow-up SuperMUC-NG project.

References and Links

- [1] ExaHyPE project: www.exahype.eu
- [2] ExaHyPE software: www.exahype.org
- [3] A. Reinarz, et al., *Comp. Phys. Commun.* 254, 2020.
- [4] M. Dumbser, F. Fambri, E. Gaburro, A. Reinarz, *J. Comp. Phys.* 404, 2020.
- [5] M. Tavelli, M. Dumbser, D. E. Charrier, L. Rannabauer, M. Bader, *J. Comp. Phys.* 386. 158-189. 2019.

So2Sat – 10^{16} Bytes from Social Media to Earth Observation Satellites

RESEARCH INSTITUTION

¹Remote Sensing Technology Institute, German Aerospace Center

²Data Science in Earth Observation, Technical University of Munich

PRINCIPAL INVESTIGATOR

Xiaoxiang Zhu¹

RESEARCHERS

Yilei Shi³, Yuanyuan Wang², Marie Lachaise¹, Cong Luo²

PROJECT PARTNERS

³Chair of Remote Sensing Technology, Technical University of Munich

SuperMUC Project ID: pr53ya, pr45ne

Introduction

The rapid urbanization poses fundamental challenges to our societies across the globe. New phenomena of urbanization, such as megaregions and informal settlements, have raced too far ahead of our current understanding of urbanization, which is mostly based on the United Nation's population figure. Therefore, the scientific question of the project 4D City (pr45ne) and So2Sat (pr53ya) [1] is: how does the global urban geographic figures, including geometry, thematic, population density, evolve over time, and in what detail can we observe and measure them?

Our scientific objectives are: for the first time systematically fuse the remote sensing data, and the massive data available from GIS and social media, to map 3D urban infrastructures and their evolution over time, i.e. 4D, in high resolution and on a global scale.

The outcome will create a first and unique global and consistent 3D/4D spatial data set on the urban morphology of settlements, and a multidisciplinary application derivate assessing population density. This is seen as a giant leap for urban geography research as well as for formation of opinions for stakeholders based on resilient data.

Results and Methods

We use a combination of traditional signal processing methods and deep learning to reconstruction our 3D/4D city models. We have developed the world's first algorithm to reconstruct urban 3D models from a very small TanDEM-X radar image stack. This is a task that usually requires at least 20 images. We employed modern signal processing techniques, including non-local means filtering, as well as compressive sensing (CS) [2], to accomplish this challenging task by using just 5 TanDEM-X interferograms. Figure 1 shows the point cloud of the Munich central station reconstructed by our algorithm. Based on the result in Figure 1, we employed the 2D building footprint from OpenStreetMap (OSM) to reconstruct a level-of-detail 1 (LOD1) building model of Munich. The result is shown in Figure 2. Although most European cities are

well modeled, a LOD1 urban model on a global scale will be a great leap from the state of the art. Therefore, this is a demonstration of one of the outcomes of the So2Sat project, i.e. a global LOD1 urban model. So far, we have collected most of the TanDEM-X images to cover the global urban area. However, the OSM footprint is not globally available. In fact, it is estimated that only 10% of the buildings in the world are covered by OSM. Therefore, we need to retrieve the building footprint from high resolution optical images. This requires the development of a new algorithm. Automatic building extraction from optical imagery remains a challenge due to, for example, the complexity of building shapes. Often, the footprint extraction from optical images is tackled by deep learning. The existing algorithms often lose resolution due to the convolutional operation in the neural network. To overcome these issues, we have proposed a novel framework, "gated graph convolutional networks (GGCN) with deep structured feature embeddings (DSFE)" [3,4]. Without going into the details of the proposed algorithm, we show the predicted building footprint of Munich using our algorithm in Figure 3, which outperforms the state of the art. This algorithm and the abovementioned 3D reconstruction algorithm form the backbone of our global 3D/4D urban modeling pipeline.

Apart from the 3D reconstruction, we also apply deep learning-based settlement type classification on the TanDEM-X data. Currently, we employed the so called local climate zones (LCZs) classification scheme. This scheme classifies the global urban area according to their morphology, i.e. the building density, height, and imperviousness. The LCZ map of Munich classified from TanDEM-X is shown in Figure 4. The colors red, brown, and gray represent different density of the buildings, with red being the densest. We can clearly see that the building density decreases from the center to the outskirts of Munich. We will make this type of map available for every single city in the world.

One challenge is that typical deep learning processing is based on GPUs. However, SuperMUC-NG is fully based on CPUs. Therefore, we have also implemented our deep learning inferencing codes for CPU, and optimized according to the architecture of




Figure 1: Point cloud of the central station of Munich, reconstructed from five TanDEM-X image only, a task usually requires more 20. The color represents the height.




Figure 2: LOD1 city model reconstructed from the point cloud shown in Figure 1. Color from blue to white indicates height from low to high.



Figure 3: The building footprint of Munich, Germany, predicted from PlanetScope low cost satellite images.

SuperMUC-NG thin nodes. We noticed that employing OpenMP-based multi-processing for our problem is more than 10 times slower than on a NVIDIA V100 GPU. Therefore, a new parallelization strategy was developed in our project. We partitioned each image and manually fixed the affinity of the processes of each partition and individual CPUs. By doing so, we were able to improve the performance by a factor of 5, reaching nearly linear scaling.

In summary, about 20 million core hours were consumed in the two projects. The preprocessing of the images, and the code of our 3D reconstruction algorithm were written in C++ and MPI, which scale well up to 1,000–2,000 cores. In the processing step, we typically use 512, 1,024 or 2,048 cores (depending on




Figure 4: Local climate zones classification of Munich using TanDEM-X image.

the available number of nodes). The deep learning codes are all implemented in Python. We use SuperMUC-NG only for inferring the trained model. So far, we have processed over 20,000 TanDEM-X images, occupying over 1 Petabytes of storage in WORK.

Ongoing Research / Outlook

SuperMUC-NG is vital for the success of our projects, because the sheer data volume and the processing hours required by our projects can only be accommodated in SuperMUC-NG. Our projects currently take more than 1 Petabytes of storage in SuperMUC-NG, for which we are grateful. This biggest challenge in our project is also its big data nature. Apart from finishing processing the global 3D city models and settlement type maps, we will also focus on big Earth observation data management in the future.

References and Links

- [1] <https://www.so2sat.eu>
- [2] Y. Shi, X. X. Zhu, W. Yin, and R. Bamler, IEEE Trans. Geosci. Remote Sens., 56(10), 6148–6158, 2018.
- [3] Y. Shi, Q. Li, and X. X. Zhu, IEEE Geosci. Remote Sens. Lett., 16(4), 603–607, 2019.
- [4] Y. Shi, Q. Li, and X. X. Zhu, ISPRS J. Photogramm. Remote Sens., 159, 184–197, 2020.

Supercomputing of Extreme Natural Events

RESEARCH INSTITUTION

¹Ludwigs-Maximilians-Universität München

PRINCIPAL INVESTIGATOR

Dieter Kranzlmüller¹

RESEARCHERS

Antonio Parodi², Elisabetta Fiori², Martina Lagasio², Stephan Hachinger³

PROJECT PARTNERS

²CIMA Research Foundation, Savona, Italy

³Leibniz Supercomputing Centre

SuperMUC Project ID: pr62ve

Introduction

In the era of global warming, mountain and cost areas are particularly affected by extreme, weather-driven natural events. Flash floods and storms are notoriously difficult to simulate and predict in these areas with complex orography. However, reliable predictions are absolutely necessary for mitigating disaster risks, e.g. by activating civil protection mechanisms and warning the population.

SCENE (Super-Computing of Extreme Natural Events) thus aims at testing and improving the forecast skills of state-of-the-art Numerical Weather Prediction (NWP) models for such areas. As a setting, it focuses on the Mediterranean region. Scientifically, SCENE is largely driven by CIMA Foundation (Savona, I). CIMA has a broad background in hydro-meteorological disaster prediction and prevention, in collaboration with entities of the Italian administration and Civil Protection. Thus, the results, obtained in synergy with the supercomputing expertise from LMU/LRZ, are immediately useful for the general public. Data sets and results from SCENE are already in use by the Ligurian Environmental Agency (ARPAL).

The work of this project focuses on two geographical areas and respective extreme events:

Area (A) is Liguria (in northwestern Italy), where three extreme hydro-meteorological events took place in



Figure 1: Streets of Genoa turned into rivers - Flash Flood October 2014 (frame from "ALLUVIONE A GENOVA - LE STRADE DIVENTANO FIUMI", Paolo Provenzale / THE STORM, Youtube – License: CC-BY 3.0).

2011 and 2014. The first event occurred on 25 Oct. 2011 with a flash-flood phenomenon in Cinque Terre, causing the death of 13 people and several millions of euros of damages. Ten days later, Genoa's city centre was affected, with a large amount of damages and the death of six people. After three years (9 Oct. 2014), a third flash flood struck the very same part of Genoa (Figure 1). Our aim is to accurately understand and predict such events with km-scale cloud resolving NWP models, based on a well-established community code (Weather Research and Forecasting – WRF [1]). Very precise NWP in a orographically extremely complex environment is however – as of now – impossible with "best of breed" simulation techniques alone. Therefore, adequate data from weather stations and satellites are conventionally "assimilated" in the models, compensating for inaccuracies and keeping the simulation close to the real situation (as far as it is known). For Liguria, we set up WRF models assimilating Radar Reflectivity CAPPI and Weather-Station data. Models with different assimilation methods are set up and compared to reality and to one another.

Area (B) is situated more southwards, where we focus on more recent events in Livorno (Tuscany, Italy) 9-10 Sept. 2017, and Silvi Marina (Abruzzo, Italy) 14-15 Nov. 2017. For these, high-quality Sentinel satellite data (from the Copernicus programme of the European Space Agency, ESA) are available. Thus, we compare models assimilating different Sentinel products (e.g. land/sea surface temperature, soil moisture, wind over ocean).

Results and Methods

Area (A) – Liguria

In this part of the project, various simulations of the events of 9 Oct. 2014 are compared to one another and to reality. The simulations assimilate different sets of external measurements with different methods.

Our basic WRF model is based on previous set-ups for simulating V-shaped back-building mesoscale convective systems in Liguria [2]. Two nested domains with resolutions of 5 and 1 km have been used (179x200 and 475x475 grid points, respectively). The

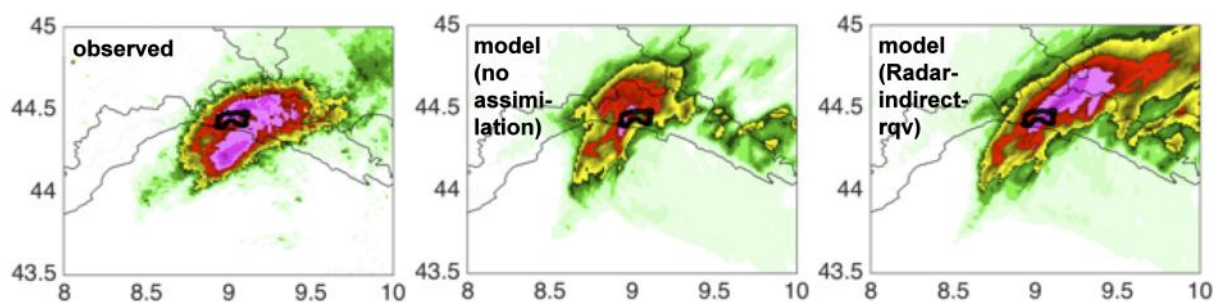


Figure 2: Accurate prediction of the Genoa 2014 flash floods (cf. Fig. 1) crucially depends on data assimilation. The panels compare observed rainfall (Settepani-radar quantitative estimate, left) to WRF model predictions (mid/right). Precipitation is colour coded (white-green-yellow-red: 0-200mm; light-dark pink: 200-500mm); numbers on the axes are latitude/longitude in °(N/E). A model without data assimilation predicts only moderate rainfall (mid panel). With optimised assimilation (method Radar-indirect-rqv, right panel), we predict shape and intensity of the heaviest rainfalls much better (with weaker-rainfall areas smeared out as a small downside), and actually also total precipitation.

number of vertical levels has been set to 50, with higher density in the lower km of the atmosphere. All our grid spacings allow for the explicit treatment of many convective processes.

In the data analysis, we compare in particular the quantitative precipitation forecast (QPF) of models with the Quantitative Precipitation Estimation (QPE) yielded from radar (some examples: Figure 2). An objective assessment of the results of NWP models is not trivial: modelling errors in the flow not only result in weaker or stronger features (e.g. rainfall regions). They also cause features to be spatially shifted and/or distorted. The Method for Object-Based Evaluation (MODE, [3]) aims at quantifying practically relevant characteristics and inaccuracies. On this basis, models can be selected which are, e.g., forecasting rain features such that floods are correctly predicted in the worst-affected zones (Figure 2, Radar-indirect-rqv model). Our best models show a significantly improved rain/flood forecasting capability [4] compared to our and others' previous attempts.

Area (B) – Central Italy

In this ongoing part of our project, satellite data of the prestigious Copernicus programme are assimilated in order to improve rainfall predictions. Also here, it is unclear a priori which data products (e.g. Land Surface Temperature, Figure 3) from the deluge of available Sentinel data improve extreme weather forecasts the most. An assimilation of Zenith Total Delays of navigation satellites seems to improve the predicted rainfall phenomena as well. At the moment, we are applying the MODE method to different simulations of the Livorno events (9-10 Sept. 2017) to delve further into all this.

Resource Usage

The ongoing project has utilized about one million core hours for first case studies in each area, and more than 10 million for modelling the affected regions over a longer time. These efforts allow for statistically assessing the accuracy of the simulations. All simulations in the ensemble use moderate numbers of CPU cores (e.g., ~2,500) and take some hours. The datasets produced in 5-10,000 core-hours with such a simulation have typical sizes of 100GB. Several 100TBs we have thus produced are archived on the SuperMUC tape archive.

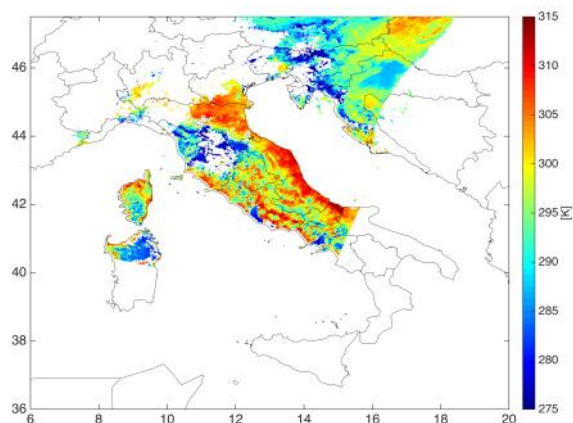


Figure 3: Sentinel 3 Land Surface Temperature data from 9 Sep. 2017 (10UTC), as assimilated for our area B.

Ongoing Research / Outlook

The project will be concluded in 2020 by further evaluation of the results, and updates/publications to our optimum simulation strategies for extreme weather events. As a next project, we plan to assess risk-prediction capabilities using data over the past 30 years. We will produce large-scale WRF simulations and compare, e.g., hydrological outcomes with observations. This unique time series of models can also be used to better understand and predict extreme natural phenomena as they develop with climate change. In the scope of the LEXIS project (www.lexis-project.eu, H2020 GA 825532; lead IT4I/CZ, with CIMA & LRZ participation), our simulations are coupled with further application models (e.g., agricultural yield and forest-fire prediction).

Looking back to SuperMUC as an ideal environment for WRF, we hope for a similar success with SuperMUC-NG.

References and Links

- [1] Skamarock, W. C., et al., 2008, NCAR Tech. Note NCAR/TN-475+STR (doi:10.5065/D68S4MVH).
- [2] Lagasio, M., Parodi, A., Procopio, R., Rachidi, F., and Fiori, E., 2017, *J. Geophys. Res.: Atmos.*, 122(8), 4238.
- [3] Davis, A. C., Brown, B., and Bullock, R., 2006, *Mon. Wea. Rev.*, 134, 1785.
- [4] Lagasio, M., Silvestro, F., Campo, L., and Parodi, A., 2019, *J. Hydrometeorology*, 20(7), 1307.

3-D seismic wave propagation and earthquake rupture

RESEARCH INSTITUTION

¹LMU Munich

²Technical University of Munich

PRINCIPAL INVESTIGATORS

Alice-Agnes Gabriel¹, Heiner Igel²

RESEARCHERS

Lauren Abrahams³, Emanuele Casarotti⁶, Michael Dumbser², Eric Dunham³, Kenneth Duru⁴, Celine Hadziioannou⁵, Lukas Krenz⁵, Bo Li⁶, Duo Li⁶, Elizabeth Madden⁷, Taufiqurrahman⁸, Elisa Tinti⁸, Thomas Ulrich⁹, Mitchell Williams⁹, Sara A. Wirp⁹, Shihao Yuan⁹, Djamel Ziane⁹

PROJECT PARTNERS

³Australian National University, ⁴Stanford University, ⁵University of Trento, ⁶Istituto Nazionale Geofisica e Vulcanologia

⁷King Abdullah University of Science and Technology, ⁸University of Brasília, ⁹University of Hamburg

SuperMUC Project ID: pr63qo

Introduction

The Computational Seismology group at LMU Munich uses SuperMUC HPC infrastructure in a variety of international research projects in computational wave propagation and earthquake rupture covering 3D forward and inverse problems across spatial and temporal scales. State-of-the-art modeling software based on a variety of numerical methods including high-order accurate Finite Difference (FD), Spectral Element (SE) and Discontinuous Galerkin (DG) methods were used to advance our understanding of earthquake physics and pave new ways for observational seismology. The project's research at SuperMUC-NG was awarded with the 2020 PRACE Ada Lovelace Award for HPC [5], a 2019 ERC Starting Grant Award [6], the SC'18 Selected Paper for Student Cluster Competition Reproducibility Challenge and top altmetric scores of publications in Journal of Geophysical Research: Solid Earth and Geophysical Journal International.

Results and Methods

Using SeisSol, a highly accurate and efficient open-source software for earthquake simulations (www.seissol.org), we develop a dynamic rupture model of the 2016 Mw 7.8 Kaikōura (New Zealand) earthquake, considered the most complex earthquake rupture observed to date. The unprecedented geometric complexity of this earthquake challenges standard data-driven approaches. In [2], we present a model which allows unravelling some of the event's riddles in a physics-based manner and provides insight on the mechanical viability of competing hypotheses proposed to explain them. In particular, the model allows constraining puzzling features inferred from high-quality observations including a large gap separating surface rupture traces, the possibility of significant slip on the subduction interface, the non-rupture of the Hope fault, and slow apparent rupture speed.

High frequency seismic shaking is particularly hazardous and can be generated from earthquake source complexity, such as rupture front acceleration in areas of great stress and geometrical variations. We investigate such effects using ensemble simulations of dynamic rupture on fractal rough faults in random heterogeneous media with WaveQLab3D, a high-order-accurate finite difference solver. By studying the simulated rupture on two rough faults embedded in heterogeneous media with increasing level of material heterogeneity, we show that enhanced material heterogeneity leads to higher coda amplitude across the domain (see Fig. 2), lowered average peak ground motions, but also increased on-fault peak ground acceleration (PGA) and off-fault peak ground velocity (PGV).

Seismic noise is used to image Earth's interior and monitor changes in the subsurface, sometimes associated with natural hazards. A solid understanding of the composition of the ambient wave field is required to interpret inferences from seismic noise data. The ambient seismic noise wavefield in the secondary microseism is mainly composed of Rayleigh and Love waves. While the presence and amount of Rayleigh waves is well understood, the generation of Love waves is still under debate. We investigate multiple scattering of surface waves as a possible mechanism that could be responsible for the generation of Love waves in the secondary microseism. We find that the fluctuation strength and correlation length have a major influence on the Love to Rayleigh ratio [4].

We develop new nonlinear mathematical models for the description of complex topography on uniform and adaptive Cartesian meshes. A novel diffuse interface formulation of the linear elastic wave equations describes geometries merely by a scalar volume fraction function [1]. With this approach, it is possible to model shapes of arbitrary complexity and completely without the need of generating anybody-




Figure 1: Dynamic rupture multi-physics scenario of the 2016 Mw 7.8 Kaikōura earthquake. Cut view of the mesh and snapshot of the wavefield (absolute particle velocity in m/s) across the fault network after 44 seconds simulation time.




Figure 2: 3D dynamic rupture simulations on a fractally rough fault in random heterogeneous media. Rupture front contours of simulated earthquake rupture on the same rough fault with varying amount of material heterogeneity (heterogeneity standard deviations varying from 0–10% by steps of 2.5%).

fitted computational mesh. The diffuse interface approach has been used to model the diffracting effect of topography on seismic waves (see Fig. 3).

New nonlinear, thermodynamically consistent mathematical models for large strain elasto-plasticity with material failure have also been developed by the same research group. The Godunov-Peshkov-Romenski (GPR) model of nonlinear hyper-elasticity has been extended to material damage and combined with the previously developed diffuse interface approach for complex geometries. Fig. 4 illustrates an application of the developed mathematical models: the modeling of cracks in a pre-damaged rock-like disc.

Ongoing Research / Outlook

Earthquakes, secondary hazards such as tsunamis [2], and seismic wave propagation are multi-physics and multi-scale processes. In this project we developed and applied novel numerical and computational approaches advancing our understanding of earthquake physics and complex seismic wave propagation. In the future, we will be using supercomputing empowered physics-based numerical forward and inverse modeling to account for the non-linear coupling of source, path and site effects in complementing data-driven natural hazard assessment and sharpen our images of Earth's interior.




Figure 3: Wave propagation in a complex 3D geometry calculated using ExaHyPE and a diffuse interface approach.




Figure 4: Crack generation in a pre-damaged rock-like disc. Stress contour colors and rupture field compared against experimental observation.

References and Links

- [1] M. Tavelli, M. Dumbser, D. E. Charrier, L. Rannabauer, T. Weinzierl, M. Bader, *J. Comput. Phys.*, 386, pp. 158–189 (2019). DOI: 10.1016/j.jcp.2019.02.004.
- [2] Th. Ulrich, A. Gabriel, J.-P. Ampuero, Wenbin, *Nat. Commun.*, 10, 1, p. 1213 (2019). DOI: 10.1038/s41467-019-09125-w.
- [3] Th. Ulrich et al., *Pure Appl. Geophys.*, 176, 10, pp. 4069–4109 (2019).
- [4] D. Ziame, C. Hadziioannou, *Geophysical Journal International* 217.2, pp. 1108–1122 (2019). DOI: 10.1093/gji/ggz056.
- [5] <https://prace-ri.eu/dr-alice-agnes-gabriel-wins-2020-prace-adelovelace-award-for-hpc/>
- [6] www.tear-erc.eu

Variable resolution meshes in climate research with the Model for Prediction Across Scales

RESEARCH INSTITUTION

Karlsruhe Institute of Technology (KIT)
Institute of Meteorology and Climate Research (IMK-IFU)

PRINCIPAL INVESTIGATOR

Gerhard Smiatek

RESEARCHERS

—

PROJECT PARTNERS

—

SuperMUC Project ID: pr94mi

Introduction

Variable resolution meshes are employed in simulation of the African Monsoon with the Model for Prediction Across Scales (MPAS). Computational and technical issues of the CPU demanding climate simulations are one aspect of the project. Another is the practical application of the obtained results. Here the focus was moved to an important environmental question of great concern. This is the potential impact of the planned pan-African Great Green Wall (GGW) for the Sahara and Sahel on the regional and continental climate. The GGW initiative aims at greening of the area receiving 100 mm – 400 mm precipitation per year through afforestation and natural regeneration. The planned wall is at least 15 km wide (Fig. 1). Several countries, however, plan reforestations of much wider areas. In total, 30 woody and 22 herbaceous plants are proposed to cultivate in GGW resilient to arid and semi-arid climatic conditions [1]. These species have been selected taking into concern socio-economic values to the population and important ecological functions.

Results and Methods

Multiple runs of the Model for Prediction Across Scales (MPAS) version 7.0 [2,3] with variable meshes ranging from 60 km, 60–10 km to 60–3 km were performed. The 60 km mesh allows for a computationally effective system setup and test. The substantially more CPU-demanding variable meshes with 10 km, respectively 3 km resolution are indented to give a detailed view over the entire GGW area and selected regions. The highest 3 km resolution applies to a comparatively small area of few degrees of longitude and the transition to the 60 km global mesh occurs here much faster than in the 60–10 km mesh.

In all simulations, two boundary condition are applied. The first one is the sea surface temperature (SST) of the years 2014–2016 as in the ERA-Interim data. The second one is the land use data. Here, one experiment is run with MODIS derived data as in the MPAS standard configuration and another with changed land

use along the planned GGW. From the difference between the two experiments potential influence of the GGW on the regional climate can be derived.

Main change to the applied code (preprocessing, model, postprocessing) was a routine which allows to change the land cover applied in the MPAS model run within the area of the GAW. Here, parts of the area covered by grassland were changed to woody savannas (Figure 2).

In order to account for the internal model variability an ensemble of several simulations is required. Thus, the simulations were initialized with ERA-Interim data on consecutive five days starting at December 10th and run over the entire following year. Further investigations will reveal if the size of the ensemble members is sufficient.

Obtained results reveals differences in the two mesh resolutions in the 2m temperature illustrating much more details in the highest resolution area and the faster transition to the coarser resolution. Similar applies to precipitation. This results from the ability of MPAS to switch from parametrized convection in coarse resolution mesh to convection permitting mode in meshes below 4km resolution.

Figure 3 exemplarily shows results obtained with the MPAS runs at 60 km resolution covering the entire globe for temperature and precipitation in the area of the planned GGW. The simulations apply 55 vertical levels and 4 soil levels and the SST of the period is from 10.12.2014 to 01.01.2016.

They reveal a clear sensitivity of the applied model system to changes in the land use and their impact on key climatological variables. Changes in mean summer temperature due to the GGW may reach 1 °C and changes in the mean summer precipitation may reach more than 100 mm which is in some areas roughly 25% of the total precipitation. While the Sahel becomes wetter with some drier areas in Mali and Burkina Faso, the eastern part close to the Red Sea is experiencing are precipitation reduction.

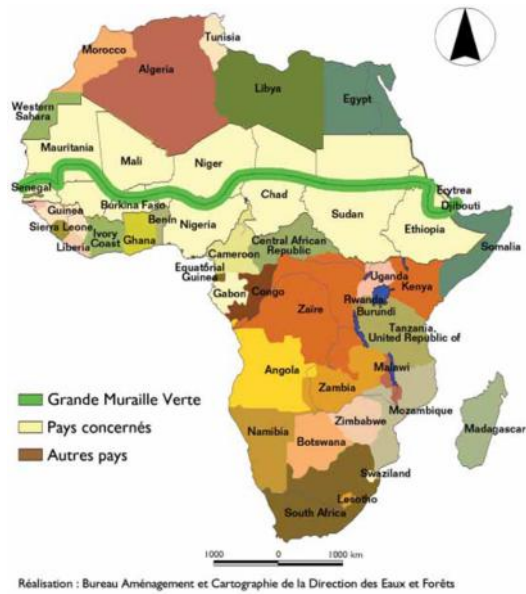


Figure 1: Great African Green Wall.

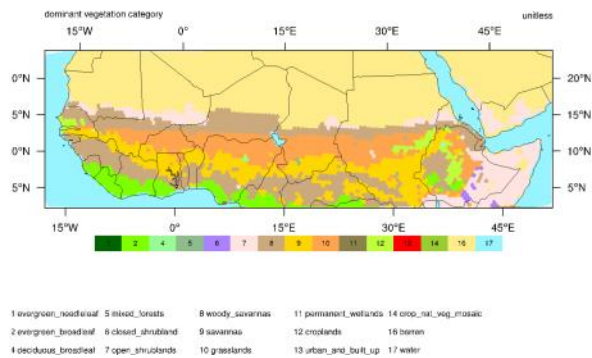


Figure 2: Changed land use applied in the MPAS runs.

Ongoing Research / Outlook

The MPAS model simulates a possible weather that fits to the prescribed sea surface temperature. The number of repetitions runs with different initialization times is at current stage 5. This may be not sufficient to derive all possible weather conditions as small changes in the initial conditions lead after some time to very different weather conditions. Therefore, the number of repetition runs will be increased to 25. It will allow to evaluate the performance of the applied model system in comparison to the observed temperature and precipitation and addition consider important extreme weather events.

Extended model runs and investigations will also address the quantification of the GGW land use change impact, especially in context of the internal model variability. Further work includes the adaptation of the land use data to higher mesh resolutions and evaluation of model runs employing the convection permitting mode.

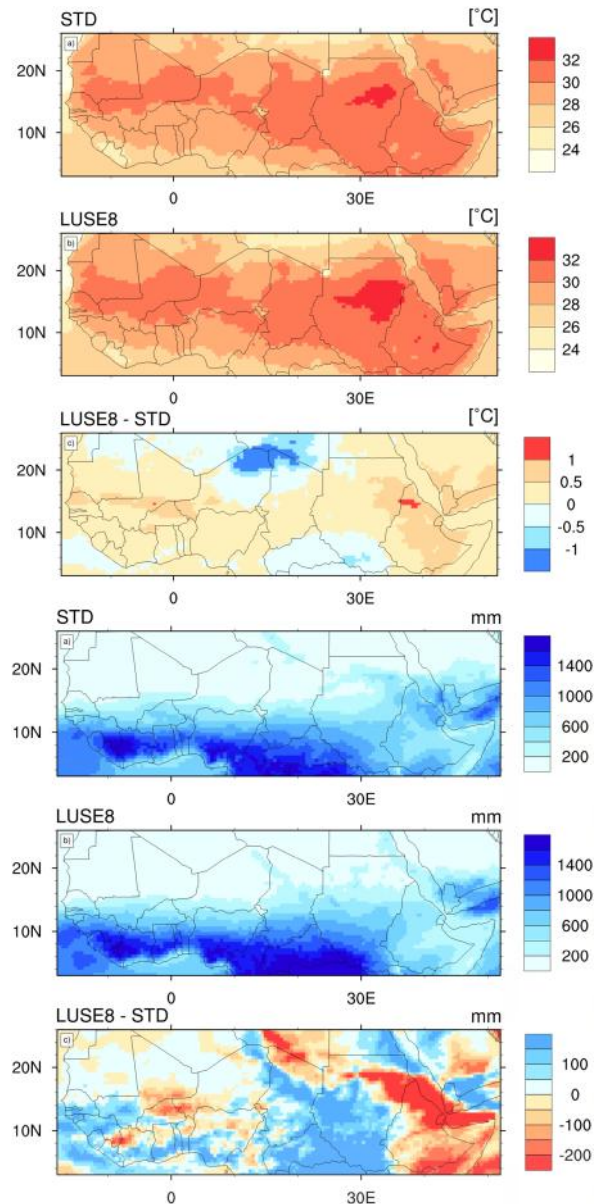


Figure 3: Simulated summer (JJA) mean temperature and precipitation. STD – Standard MPAS simulation. LUSE8 simulation with changed land use and LUSE8-STD difference between der obtained mean values.

References and Links

- [1] Ba, A. T., 2010. Grande muraille verte: choix des especes vegetales. In: Dia, a., Duponnois, R. (Eds.), Le projet majeur africain de la Grande Muraille Verte. IRD, Institut de Recherche pour le D'veloppement, pp. 31–38.
- [2] Heinzeller, D., Duda, M. G., Kunstmann, H., Geosci. Model Dev., 9, 77–110 (2016).
- [3] Skamarock, W. C., Klemp, J. B., Duda, M. G., Fowler, L. D., Park, S., Ringler, D. T., Mon. Wea. Rev. 140, 3090 – 3105 (2012).

Atmospheric Chemistry and Climate

RESEARCH INSTITUTION

¹DLR – Institut für Physik der Atmosphäre, Oberpfaffenhofen, Germany

PRINCIPAL INVESTIGATOR

Mattia Righi¹

RESEARCHERS

Sabine Brinkop¹, Simone Dietmüller¹, Roland Eichinger^{1,2}, Christine Frömming¹, Hella Garny^{1,2}, Phoebe Graf¹, Volker Grewe¹, Johannes Hendricks¹, Patrick Jöckel¹, Bastian Kern¹, Sigrun Matthes¹, Mariano Mertens¹, Andreas Pfeiffer¹, Vanessa Rieger¹, Robert Sausen¹, Franziska Winterstein¹

PROJECT PARTNER

²Meteorologisches Institut, Ludwig-Maximilians-Universität München

SuperMUC Project ID: pr94ri

Introduction

The objective of this project is to improve our understanding of the Earth's atmospheric composition and climate by means of numerical modeling. The Earth system is driven by several dynamical, chemical and physical processes, which determine its composition and evolution and affect the global climate. The numerical representation of such processes constitutes an Earth System Model (ESM). ESMs require considerable computational resources, in order to realistically simulate the different components of the system and their interactions. Moreover, substantial storage capacities are essential to store and analyze the terabyte-sized output produced by the simulations. In addition to the scientific interest, the growing societal concern for topics like the global climate change, the ozone hole and the increasing air pollution in metropolitan areas is further motivating the need for a deep understanding of these processes. The focus of this project is therefore twofold, (i) we aim at a better understanding and representation of the Earth atmosphere in our numerical model and (ii) we apply our model to quantify the impacts of human activities on the atmospheric composition and climate.

We use a highly flexible numerical system, the ECHAM/MESSy Atmospheric Chemistry (EMAC) model (<https://www.messy-interface.org>), which makes it possible to use the same code for tackling numerous scientific topics. The model can be easily configured to set the horizontal and vertical resolutions, the level of interaction between dynamics, chemistry, physics and radiation, the complexity of the chemical kinetics, as well as the parameters that describe the physical processes within the system. The use and application of the same code also facilitates development and documentation, as well as synergies among the users.

Results and Methods

Evaluation of the Lagrangian Tracer Transport Model ATTILA

The Lagrangian (LG) model ATTILA (Atmospheric Tracer Transport In a Lagrangian model) was extended with a LG convection scheme and a diabatic calculation of the vertical velocity. To evaluate this new LG model we performed two global hind-cast simulations with EMAC (grid-point model, GP) over the period 1950–2013 including the LG model ATTILA and the newly developed parameterizations. The LG scheme was evaluated with respect to the transport characteristics in the stratosphere using an age-of-air tracer. The transport in the stratosphere was clearly improved when using a diabatic vertical velocity for tracer transport. For the troposphere we used the short-lived tracer Radon (²²²Rn) to evaluate the vertical transport of the LG scheme. Figure 1 shows an inter-comparison of the mean ²²²Rn distribution in the surface layer at different stations world-wide as simulated with the EMAC model (GP) and from the two LG simulations with either the kinematic or the diabatic vertical velocity, against available observations. Depending on the geographical location, the model results are more or less in agreement with the observations. These results and the description of the newly developed and modified physical routines, as well as the new diagnostic and infrastructure sub-models are published in [1].

Impact of transport emissions on air quality in Europe

The land transport sector is an important emission source of carbon dioxide, nitrogen oxides, carbon monoxide and volatile organic compounds. The latter three compounds are precursors for the formation of tropospheric ozone. An increase of tropospheric ozone reduces air quality and leads to a positive radiative forcing, hence contributing to warming the climate. To quantify the impacts of land transport emissions on air




Figure 1: Monthly mean ^{222}Rn (mBq m^{-3} STP/surface concentrations): GP and LG model results against observations at different sites from [2]. The thick dashed lines include a range within a factor of 2 of the observations. Reproduced from [1].

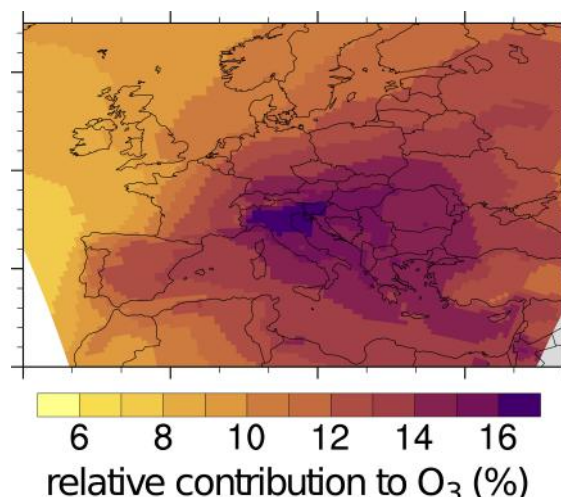


Figure 2: Relative contribution (in %) of land transport emissions to ground-level ozone in the summer months averaged over the 2007–2010 period. Reproduced from [4].

quality in Europe, we applied the global-regional model system MECO(n), which combines the global model EMAC with the regional model COSMO/MESSy.

For this project we used a MECO(2) set-up, i.e. MECO(n) with 2 refinements, over Europe (with 50 km resolution) and over Germany (with 12 km resolution). Using this configuration, we performed two three-year simulations with two different anthropogenic emission inventories and five one-year simulations with systematic variations of anthropogenic and natural emission strengths within the range of uncertainties.

The results of the model simulations set the basis for two publications. The first investigated the impact of the model resolution on the results of ozone source apportionment studies [3]. This work showed that the resolution of the global model is sufficient to diagnose contributions of land transport emissions to ozone on the continental scale. On the regional scale, however, the coarse resolution is not sufficient and therefore regional models are needed for detailed regional assessments. The second study investigates the contribution of land transport emissions to ozone and ozone precursors in Europe [4]. The results showed that land transport emissions contribute up to 70% to ground level reactive nitrogen levels in Europe. In addition, contributions of land transport emissions to ozone of up to 16% were diagnosed. These levels were found in the Po Valley (see Fig. 2), where large emissions and favorable conditions for ozone production coexist.

Ongoing Research / Outlook

Due to its domain cloning approach, the Lagrangian model ATTILA, in particular the convection scheme, is highly demanding with respect to communication between the parallel tasks (MPI). Thus, the scaling on modern HPC architectures is very limited. Since we are in a transition phase replacing the spectral model ECHAM5 by the new gridpoint model ICON, which

requires the adaption of large parts of the code, a new Lagrangian model (with a hand-shaking approach) is currently under development. This will enable us to perform Lagrangian simulations on various spatial scales (LES to global).

The models involving atmospheric chemistry, EMAC and MECO(n), are very demanding with respect to I/O bandwidth and storage requirements, since a large number of prognostic variables need to be output. Limited storage and the currently applied serial I/O (netCDF) limit our large-scale simulations. Therefore we are working on parallel and asynchronous I/O as well as on methods for data compression and/or reduction. With a better scaling behavior and improved I/O, we can further our research and reduce uncertainties, e.g. by applying finer resolved grids, including more detailed process descriptions, and by performing ensemble simulations.

Further improvements to the MECO(n) models are required for detailed evaluation with aircraft observations. In the current version, the results in the finer resolved regional model are relaxed to the results of the global model up to 11 km, but some aircraft observations take place at higher altitudes. A new setup based on [5] has been developed which extends the upper boundary to 28 km. A long-term simulation (30 years) with this new configuration is currently being evaluated and further simulations are planned to include atmospheric chemistry and to cover the periods of measurement campaigns.

References and Links

- [1] S. Brinkop, P. Jöckel, *Geosci. Model Dev.*, 12, 1991–2008 (2019). DOI: 10.5194/gmd-12-1991-2019.
- [2] K. Zhang, H. Wan, M. Zhang, B. Wang, *Atmos. Chem. Phys.*, 8, 2811–2832 (2008). DOI: 10.5194/acp-8-2811-2008.
- [3] M. Mertens, A. Kerkweg, V. Grewe, P. Jöckel, R. Sausen, *Geosci. Model Dev.*, 13, 363–383 (2020). DOI: 10.5194/gmd-13-363-2020.
- [4] M. Mertens, A. Kerkweg, V. Grewe, P. Jöckel, R. Sausen, *Atmos. Chem. Phys.*, 20, 7843–7873 (2020). DOI: 10.5194/acp-20-7843-2020.
- [5] J. Eckstein, S. Schmitz, R. Ruhnke, *Geosci. Model Dev.*, 8, 1839–1855 (2015). DOI: 10.5194/gmd-8-1839-2015.

From New Materials to Planetary Cores

RESEARCH INSTITUTION

¹Carnegie Institution for Science

²Department of Earth and Environmental Sciences Ludwig Maximilians Universität München

PRINCIPAL INVESTIGATOR

Ronald E. Cohen^{1,2}

RESEARCHERS

—

PROJECT PARTNER

—

SuperMUC Project ID: pr92ma (Gauss Large Scale project)

Introduction

First-principles methods use fundamental quantum physics to compute properties of materials without fitting experimental data. In this way conditions of very high pressures and temperatures can be simulated, which are very difficult or impossible to achieve in the laboratory. We can also simulate materials that do not exist yet, and thus design new materials. The underlying theory is known as Density Functional Theory (DFT), which is exact in principle, but in practice must be used with approximations. These computations often require supercomputers, because extreme precision is needed—the chemical energy is in the 8th–12th significant figure of the total energy, so numerically accurate methods are required. In many materials the straightforward application of DFT fails, because electrons are correlated in a non-mean-field way. Then we need to use special methods which are even more computationally intensive. To simulate high temperatures we use first-principles molecular dynamics (FPMD), in which we do many thousands of first-principles calculations, stepping forward in time using Newton's law of motion $F = m \cdot a$, with the force F computed from quantum mechanics, m the mass of each atom, and a its acceleration. Examples of our work on SuperMUC-NG follow, for Earth's core, filled ice on Titan, carbonated melts in Earth's mantle, Polar metal-locenes, and ferroelectric clathrates.

Iron in Earth's Core

Earth's magnetic field arises from turbulent motion of liquid iron alloy in Earth's outer core. It used to be thought that thermal convection driven by Earth's original heat provided the energy for the dynamo, but this was brought into question by first-principles computations that showed too high thermal conductivity in iron to allow for thermal convection [1,2].

We performed extensive simulations of iron (solid and liquid) at Earth core conditions using a variety of methods, density functional perturbation theory with ABINIT for solids, molecular dynamics with QUANTUM ESPRESSO for solids and fluids, and included electron-correlations with dynamical mean field theory (DMFT) with continuous time quantum Monte Carlo using Kristijan Haule's code, in conjunction with the all




Figure 1: Electric resistivity for iron at core density from measurements, geophysical models, and computations (ours are labeled KKR, DFPT, and the red curves. From [1].




Figure 2: Thermal conductivity for iron at core conditions. Our results in Red are higher than experiment, but lower than previous theory. From [4].

electron WIEN2K LAPW code, or SPKKR on MD snapshots. We computed the electrical resistivity and thermal conductivity at Earth core conditions, and at lower pressures and temperatures to compare with experiments. We obtained consistent results with all of these methods. Results are between those measured in the lab for thermal conductivity and electrical conductivity. We find that Earth's magnetic field throughout Earth's history can be explained via a geodynamo driven primarily by thermal convection [3,4].




Figure 3: Electrical resistivity measured at about 100 GPa compared with theory. Our latest FPMD results obtained on SuperMUC-NG are shown as the green star, in perfect agreement with the experiments.




Figure 4: Electrical and thermal conductivity computed for Niobium. Note that the experiments are at atmospheric pressure and the theory is at constant volume, with pressures shown (top).

We have shown that, including electron-electron scattering, it becomes possible for the traditional theory to work (Fig. 1-2) [3,4]. We have continued to work on this problem, to include liquid iron, to study fully converged system sizes, and to include iron alloys. New experiments on electrical resistivity of iron are in excellent agreement with our new FPMD results (Fig. 3) [5].

We are often asked how our method works on other materials than iron, to show that we do not have accidental agreement. We have studied now the element Nb on SuperMUC-NG (Fig. 4) and see excellent agreement for electrical and thermal conductivity at zero pressure.




Figure 5: Layered model for Titan: crust, subterranean ocean, high-pressure water ice layer, mixed ice-rock layer, and a rocky core. A layer enriched in filled ice MH-III may reside above the core (from [6]).




Figure 6: Animation of first-principles molecular dynamics simulation for carbonated silicate melt. The blue atom is a carbon. With Razvan Caracas.

Filled ices in water planets

Filled ices mediate the methane and oxygen contents of the atmospheres of water planets [6]. We have simulated methane filled ice, and computed its equation of state. We found that filled ice will sink in the melt, and form a shell around the core of the planet (Fig. 5).

Carbonated melts in Earth's Mantle

Little is known about the behavior of carbon in melts, because it exsolves as pressure is released and is not found in observed magmas or volcanic rocks. Experiments are difficult because carbonated fluids are very reactive, and the usual method of studying high pressures uses the diamond anvil cell, which contains already C. We have simulated silicate melts containing carbon and have computed, for example, the diffusivities of different elements in such melts (Fig. 6,7) [7].

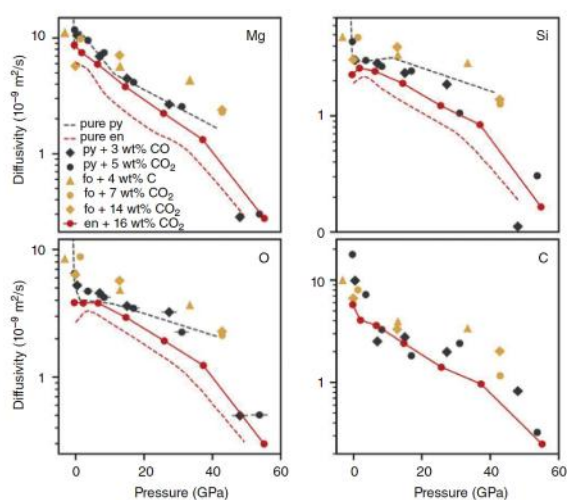


Figure 7: Elemental diffusivities of magnesium, silicon, oxygen, and carbon at 3,000 K. From Ref. [7].



Figure 8: Cover of the journal *molecules*, showing our proposed polar metalloenes.




Figure 9: Polar azaferrocene with N replacing a C.




Figure 10: Predicted ferroelectric clathrate ScB_3C_5 .

Polar Metalloenes

Metalloenes are interesting and useful organometallic sandwich molecules, with a transition metal ion between two cyclopentadienyl (Cp , or C_5H_5^-) rings. Normally the molecules are symmetric, with the two rings the same, but we proposed to make asymmetric molecules, and crystallize them to form a new class of polar materials (Fig. 8,9) [8].

Ferroelectric Clathrates

Recently we developed, computationally and in the laboratory, a new family of clathrates [9]. Most of this family is metallic and centrosymmetric. We considered how to make a polar, possibly ferroelectric, insulating sp^3 bonded clathrate and succeeded (Fig. 10) [10].

In summary, the SuperMUC-NG has been used to study a wide variety of problems in materials, ranging from planetary cores to designing new materials. This work was supported by the ERC Advanced Grant ToMCaT, The US NSF grant EAR-1901813, U. S. Office of Naval Research Grants No. N00014-12-1-1038, N00014-14-1-0516, and N00014-17-1-2768, and the Carnegie Institution of Washington.

References and Links

- [1] Pozzo, M., Davies, C., Gubbins, D. & Alfe, D., *Nature* 485, 355-U399, DOI: 10.1038/Nature11031 (2012).
- [2] de Koker, N., Steinle-Neumann, G. & Vlcek, V. *PNAS* 109, 4070-4073, DOI:10.1073/pnas.1111841109 (2012).
- [3] Zhang, P., Cohen, R. E. & Haule, K., *Nature* 517, 605-607, DOI:10.1038/nature14090 (2015).
- [4] Xu, J., et al., *Phys Rev Lett* 121, 096601, DOI:10.1103/PhysRevLett.121.096601 (2018).
- [5] Zhang, Y et al., *Phys Rev Lett* 125, 078501, DOI:10.1103/PhysRevLett.125.078501 (2020).
- [6] Levi, A. & Cohen, R. E., *The Astrophysical Journal* 882, DOI:10.3847/1538-4357/ab2f76 (2019).
- [7] Solomatova, N., Caracas, R. & Cohen, R. E. in *Carbon in Earth's Interior Geophysical Monograph Series* (eds Craig E. Manning, Jung-Fu Lin, & Wendy L. Mao) 179-194 (2020).
- [8] Zhang, H., Yavorsky, B. Y. & Cohen, R. E., *molecules* 24, doi:10.3390/molecules24030486 (2019).
- [9] Zhu, L., et al., *Sci Adv* 6, eaay8361, DOI:10.1126/sciadv.aay8361 (2020).
- [10] Zhu, L., Strobel, T. A. & Cohen, R. E. Prediction of an extended ferroelectric clathrate. *Physical Review Letters* *in press* (2020).

Deciphering Large-scale Geologic Events Through Retrodictions of Past Mantle Flow

RESEARCH INSTITUTION

¹Geophysics Section, Department for Earth and Environmental Sciences, Ludwig-Maximilians-Universität München

PRINCIPAL INVESTIGATORS

Hans-Peter Bunge¹, Bernhard Schubert¹

RESEARCHERS

Lorenzo Colli², Siavash Ghelichkhan^{1,3}

PROJECT PARTNERS

²University of Houston, USA

³Australian National University, Canberra, Australia

SuperMUC Project ID: pr48ca

Introduction

Earth's mantle, although solid on short time scales, can flow like a highly viscous fluid over the course of geologic eras. Heat coming from the underlying core and internal heat production due to radioactive decay turn Earth's mantle into a convecting heat engine. 'Mantle convection' provides the primary force in driving the tectonic plates and dictates the long-term evolution of the Earth: it controls the distribution of continents and oceans and their topographic elevation; it determines the formation of mountain ranges, shallow seas and land bridges between continents; and it is the cause of Earth's seismicity and volcanic activity. As such, it has a broad impact on many aspects of the Earth system, ranging from its oceanic and atmospheric circulation to its climate, from its hydrosphere to the erosion and deposition of sediments, from the location and abundance of natural resources to the evolution of life.

There has been much progress over the past decades in modelling mantle convection processes by means of fluid dynamical simulations. These simulations involve numerical solutions for the principal conservation equations that govern mantle flow. Due to the non-linear and complex nature of these equations it is necessary to employ high performance computing techniques and infrastructure to solve these equations.

Results and Methods

A long-standing challenge in fluid dynamical mantle simulations is integration of the ever-growing disparate datasets in various disciplines of Solid Earth Sciences. For this reason, geodynamicists have turned to large scale optimization methods where models are optimized with regards to various observational datasets. An example is reconstructions of past mantle flow that involve the solution of a geodynamic inverse problem through the adjoint method. This inverse problem aims at finding the (unknown) state of the mantle in the past that naturally evolves into its (known) present-day state by iteratively minimizing the difference between the observed present-day mantle

structure and the prediction of a geodynamic model. The adjoint method is a powerful technique to compute sensitivities (Fréchet derivatives) with respect to model parameters, allowing one to solve inverse problems where analytical solutions are not available or the cost to determine many times the associated forward problem is prohibitive. In Geodynamics it has been applied to the restoration problem of mantle convection so that poorly known mantle flow parameters can be tested against observations gleaned from the geological record. By enabling us to construct time dependent Earth models the adjoint method has the potential to link observations from seismology, geology, mineral physics and palaeo-magnetism in a dynamically consistent way, greatly enhancing our understanding of the solid Earth system.

The iterative nature of the approach, however, imposes a computational challenge, but provides a promising tool to improve our understanding of Earth's history. The geodynamic group at LMU is a world leader in generation of such models, and access to SuperMUC Phase 2 and SuperMUC-NG enabled us to further explore the methodology and to address fundamental questions on the evolution of Earth's mantle.

In this project, we built upon our pioneering first mantle flow retrodictions [2], and explored the parameter space for such reconstruction models by generating a suite of geodynamically plausible, compressible, high resolution Earth models that go back to 50 million years ago. For this project we used the parallel finite element code TERRA using a sufficiently high resolution, which is obtained by dividing the volume of the mantle into ~670 million finite elements, for a maximum grid spacing of ~11 km at Earth's surface. An adjoint iteration for this model over 50 million years requires between 75 and 150 thousand Core-hours, equivalent to 36 to 72 hours of computation using 2,048 Cores. The initial condition is optimally recovered after 5 to 10 iterations, leading to a total of about one million Core-hours per retrodiction. Synthetic experiments demonstrate for the ideal case of no model error and no data error that the adjoint method restores mantle flow over timescales on the




Figure 1: 3-D views, centred on the African hemisphere, showing temperature anomalies, in Kelvin, of the initial condition and final condition of the reference 'true' history of mantle convection, together with the reconstructed initial and final conditions for a case where only the long-wavelength part of the final state is assumed to be known. The rightmost column shows the absolute difference between the true and the reconstructed state of the mantle [3].

order of a transit time (~ 100 Myr). But in reality, unavoidable limitations enter the inverse problem in the form of poorly known model parameters and uncertain state estimations, which may result in systematic errors of the reconstructed flow history. In a first part of the project, we used the high-resolution 3-D spherical mantle circulation models to perform a systematic study of synthetic adjoint inversions, where we insert on purpose a mismatch between the model used to generate synthetic data and the model used for carrying out the inversion [3]. By considering a mismatch in rheology, final state estimate and history of surface velocities we found that mismatched model parameters do not inhibit misfit reduction: the adjoint method still produces a flow history that fits the estimated final state. However, the recovered initial state can be a poor approximation of the true initial state, where reconstructed and true flow histories diverge exponentially back in time and where for the more divergent cases, the reconstructed initial state includes physically implausible structures, especially in and near the thermal boundary layers. Consequently, a complete reduction of the cost function may not be desirable when the goal is a best fit to the initial condition. When the estimated final state is a noisy low-pass version of the true final state, choosing an appropriate misfit function can reduce the generation

of artefacts in the initial state (see Fig. 1). While none of the model mismatches considered in this study, taken singularly, results in a complete failure of the recovered flow history, additional work is needed to assess their combined effects.

In a second part of the project, we have retrodicted past mantle flow for a suite of scenarios, combining five different present-day seismic tomography models with four viscosity profiles for the geodynamic model. We then investigated the related implications in terms of surface observables in these models. We found that the restoration of mantle flow is sensitive to the assumptions about the present-day mantle state and its viscosity, but in many cases consistent with geologic findings about our planet's history. Specifically, we have used our retrodictions to model one of the most outstanding geologic events of the entire Cenozoic era (the past 65 million years), namely the rise of the eastern half of the African continent from near sea-level to well over 1.5 km of elevation. It is commonly presumed by geologists that this continental scale uplift occurred in response to the rise deep within Earth's mantle of a huge thermal plume. Our retrodiction show—for the very first time—that this mechanism is feasible and that it can occur at the required time scales in geodynamic Earth models

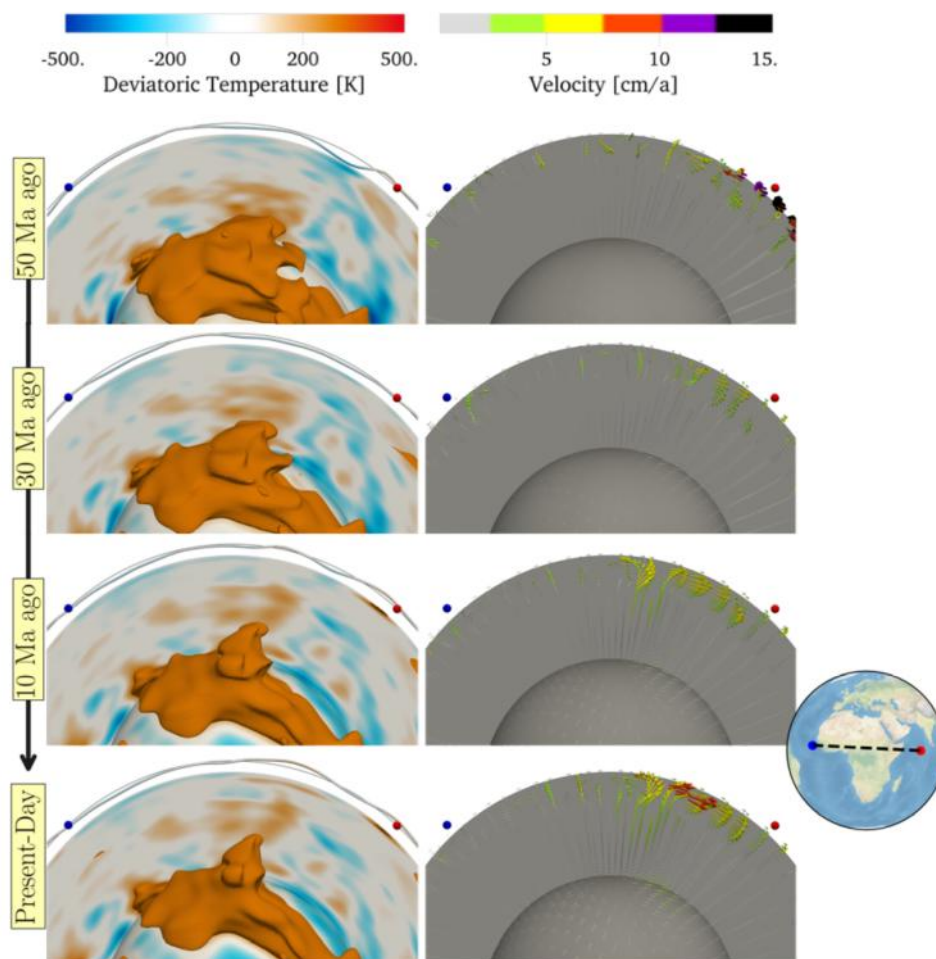


Figure 2: Reconstructed evolution of sub-lithospheric mantle temperature, shown as an iso-surface, and convectively induced surface topography (left), together with the computed mantle flow velocity field shown as arrows (right). The visualization shows the restoration of a plume system acting below the East African Rift System from an iterative flow optimization. The cross-section through the computational domain is taken along the great circle passing through the red and blue dots (see inset map) across the African continent. Note the rise of the large thermal plume anomaly (red iso-surface) through the mantle beneath the continent.

having the appropriate parameters (see Fig. 2). We find that the material flow within the Earth's mantle is not only responsible for uplift of eastern Africa. The flow also induces substantial stresses at the interface of the flowing (ductile) mantle and the mechanically strong tectonic plate above. It is likely that these flow induced stresses are involved in the development of the East African Rift System by which parts of eastern Africa are currently breaking away from the main part of the African continent. This outstanding result is currently prepared for publication.

Ongoing Research / Outlook



Geodynamic retrodictions require estimates of its present-day state from seismic tomography, and they would benefit greatly from knowledge of the associated error bars. More robust inferences could be drawn from running a suite of simulations instead of a single model realization. Some recent tomographies provide information on both resolution and model uncertainty, but so far it is unclear how to make practical use of the uncertainties in simulations of mantle convection, given that only a limited set of retrodictions is computationally feasible. We plan to

perform a systematic uncertainty quantification that will provide important insight into the evolution of seismic-data related errors in adjoint retrodictions. Knowledge of uncertainties back in time is crucial to be able in future to draw robust conclusions on the adequateness of model parameters in comparison to observations. Answering the above questions thus is a crucial step in the effort of generating data-driven models of Earth's mantle that provide quantitative information to other geoscience disciplines.

References and Links

- [1] www.geophysik.uni-muenchen.de
- [2] L. Colli, S. Ghelichkhan, H.-P. Bunge and J. Oeser., *Gondwana Research* 53, 252–272 2017. DOI: <https://doi.org/10.1016/j.gr.2017.04.027>
- [3] L. Colli, H.-P. Bunge and J. Oeser, *Geophys. J. Int.*, 221 (1), 617–639, 2020. DOI: <https://doi.org/10.1093/gji/ggaa023>

Elementary Particle Physics



LPC $P_1^Z/P_2^Z = 4/2$ (200

LPC $P_1^Z/P_2^Z = 4/3$ (2005.14572)

Bernstein (2003.06063)

Hermite (2003.06063)

$h_1(P^+ = 2.27\text{GeV})$

$h_1(P^+ = 1.25\text{GeV})$

$h_1(P^+ = 2.27\text{GeV})$
 $h_1(P^+ = 1.74\text{GeV})$

$h_1(P^+ = 1.74\text{GeV})$
 $h_1(P^+ = 1.25\text{GeV})$

The QCD phase diagram and equation of state

RESEARCH INSTITUTION

¹University of Wuppertal

PRINCIPAL INVESTIGATOR

Zoltan Fodor^{1,2,3}

RESEARCHERS

Szabolcs Borsanyi¹, P. Parotto¹, C. Ratti⁴

PROJECT PARTNERS

²Jülich Supercomputing Centre, Forschungszentrum Jülich

³Eötvös University, Budapest

⁴University of Houston, Texas, USA

SuperMUC Project ID: pn34mu (Gauss Large Scale project)

Introduction

Only microseconds after the Big Bang, the whole Universe consisted of a tiny droplet of the hottest and most dense *liquid* known to humankind: the Quark Gluon Plasma (QGP). Under normal conditions, such as those in our everyday life, the quarks and gluons - the most fundamental building blocks of matter - are confined into protons and neutrons. When matter is extremely hot, or extremely compressed, these nucleons melt and quarks and gluons form the Quark Gluon Plasma, in which they are free to move. When the Universe expanded and cooled down, eventually the quarks and gluons coalesced into nucleons, that in turn formed the atoms, then molecules and slowly formed matter as we see it today.

Thanks to extraordinary experimental efforts, the QGP is now routinely produced in particle accelerator facilities such as the Relativistic Heavy Ion Collider - in the US, and the Large Hadron Collider at CERN - in Switzerland. Heavy nuclei are accelerated to almost the speed of light and collided; the enormous amount of energy produced is sufficient to create the QGP in the laboratory, which then undergoes the same transition as the early Universe itself.

Intense theoretical as well as experimental investigations are in place in order to study the properties of this new form of matter, and in particular its transition to ordinary matter. A delicate question regards the study of the nature of this transition. In the early Universe there was a virtually perfect symmetry between matter and anti-matter. This is not the case in today's universe, especially so in ultra-dense astronomical objects such as neutron stars, which are also the subject of intense investigations. A key question is today how does this transition change when conditions more similar to these high-density objects are considered - i.e. when the density of quarks is much larger than that of the early Universe.

The theory of strong interactions, which governs the behavior of quarks and gluons - and thus of protons




Figure 1: A sketched view of the phase diagram of strongly interacting matter, in the plane of temperature and (quark) chemical potential.

and neutrons - is called Quantum Chromodynamics (QCD). The nature of strongly interacting matter under different conditions is commonly depicted in the so-called QCD phase diagram, which is sketched in Fig. 1. The different phases of matter are shown in the plane of temperature and chemical potential. The chemical potential is a thermodynamic coordinate like the temperature, which can be best understood as the energy required by the system to change its chemical composition. It is tightly connected with the density of quarks: when the former is zero, the latter is zero as well. This two-dimensional depiction is completely analogous to what is done in the phase diagram of water: we know that for certain values of the thermodynamic coordinates water is liquid, gaseous (vapor) or solid (ice).

Most of the features one sees in Fig. 1 are in reality conjectured, based on model-dependent calculations and not on first-principle determinations. The major first-principle tool of investigation is represented by the sophisticated numerical simulations - like the ones we perform - of lattice QCD. Most notably, they have led to the outstanding discovery that the transition at zero quark density is extremely smooth, very much unlike e.g., the boiling of water [1]. On the other hand, this transition is expected to become much abrupt at larger




Figure 2: Extrapolation strategy using imaginary values of the quark chemical potentials. We are interested in positive μ^2 values (right hand side), but this parameter range is not available from simulations. If μ^2 is negative (or if μ is imaginary), simulations are possible (left hand side). We seek an analytical continuation from the left side to the right.

chemical potentials, which would imply the presence in the phase diagram of a so-called *critical point*. The search for such point is possibly the main endeavor in today's high-energy nuclear physics research.

Direct simulations at non-zero chemical potentials are hindered by the infamous, so-called “sign problem”, which in essence is a problem of calculating averages from quantities that fluctuate violently. However, tools have been developed to reconstruct the physics at non-zero chemical potential from available calculations, performed at zero (or imaginary) chemical potential. (Fig 2.) The impact of these results is mainly twofold. Firstly, knowledge of the phase diagram of QCD is inherently a fundamental achievement, which deepens our understanding of matter and the Universe, with applications that range from the smallest (nuclear physics) to the largest scales (astrophysics and cosmology). Secondly, they are crucial input for many other simulations, such as the evolution of the matter created in nuclear collisions, the early Universe or dense objects as e.g. neutron stars.

Results and Methods

In 2006, we calculated the transition temperature at zero net baryon density, as it was realized in the Early Universe. In that work we have seen that although there is no real transition, one can find a transition range around 150 MeV temperature [2].

In our latest work [3] we determined the so-called “transition line” of QCD. In short, we calculated with very high precision the temperature at which the QCD transition takes place, for quite a broad range of small but non-zero chemical potential. In order to do so, we have exploited simulations at zero and imaginary chemical potential, which do not suffer from the “sign problem”. Studying observables which are sensitive to the phase transition, we have determined the temperature at which the transition occurs for each chemical potential we simulated. We then analytically continued the dependence of this transition temperature on the chemical potential, extending it to the relevant regime of real chemical potential. The extrapolated transition




Figure 3: The transition line of QCD at non-zero chemical potential. The different colored bands represent the uncertainty in different approximations of the transition temperature.

line is shown in Fig. 3 with a green band which represents the uncertainty of the result. Other comparable results, although obtained with other, non first-principle methods, are shown along with it. It is worth mentioning that this determination is the most precise both at zero and non-zero chemical potential available in the literature.

Ongoing Research / Outlook

The main focus of our further research is the study of the so-called equation of state (EoS) of QCD. This is simply the relation between quantities such as the pressure and density, and how they change in different regions of the phase diagram. Because of the “sign problem”, results at non-zero density are difficult to obtain, but nonetheless extremely valuable, especially in a new era of astrophysical observations inaugurated by the very first detection of gravitational waves.

References and Links

- [1] Y. Aoki et al, Nature 443 (2006) 675-678.
- [2] Y. Aoki et al, Phys. Lett. B 643 (2006) 46-54.
- [3] S.Borsányi et al., Phys. Rev. Lett. 125 (2020) no.5, 052001, DOI:10.1103/PhysRevLett.125.052001

Hadron structure observables on a fine lattice

at the physical point

RESEARCH INSTITUTION

Institute for Theoretical Physics, University of Regensburg

PRINCIPAL INVESTIGATOR

Sara Collins

RESEARCHERS

Gunnar Bali, Marius Löffler, Andreas Schäfer, Jakob Simeth, Wolfgang Söldner, Simon Weishäupl, Thomas Wurm

PROJECT PARTNERS

—

SuperMUC Project ID: pn34xo

Introduction

Protons and neutrons, known collectively as nucleons, are the building blocks of visible matter. Since their discovery in alpha particle scattering experiments in 1919 by Rutherford and in 1932 by Chadwick, for the case of the proton and neutron, respectively, their basic properties have been mapped out very accurately. For example, their masses and spins are known to a precision of at least one part in 10^{-8} . However, precise information on their internal structure has not yet been provided. Thanks to deep inelastic scattering (DIS) experiments at SLAC in the USA in the late 1960s and others we know that nucleons are comprised of fundamental constituents, quarks and gluons. A very simple picture of the nucleons is one where they are made up of three light (u and d) quarks, uud (the proton) and udd (the neutron). However, comparing the masses of the quarks (of the order of a few mega-electron volts (MeV)) with that of the nucleons (of the order of a thousand MeV) indicates that the internal structure of nucleons is far more complex and, for example, the gluon degrees of freedom play an important role. Similarly, early deep inelastic muon-nucleon scattering experiments showed that the spin of the u and d quarks may contribute as little as 20% to the spin of the nucleon, starting the so-called spin “crisis” or “puzzle”.

Extracting precise information on nucleon structure, how the quark and gluon constituents account for the properties of nucleons is difficult due to the nature of the strong interaction between them. At large distances (or equivalently low energies), the strong force increases such that the quarks and gluons are always confined to “hadronic” bound states. At present, the only first principles approach to calculating hadronic properties, i.e. using the theory of the strong interaction, Quantum Chromodynamics (QCD), without any additional assumptions, is via large scale numerical simulations (lattice QCD). Predictions of nucleon properties are a necessary input for interpreting the many experiments which use nucleons (or nuclei) as probes or targets. These experiments are searching for evidence of the physics that lies beyond our current understanding of particle physics (that is en-




Figure 1: The masses of the (top to bottom) Ξ , Λ , Σ and N baryons as a function of the pion mass squared. All hadron masses have been converted into dimensionless units using the lattice scale t_0 . The experimental masses are shown as filled circles [2] while the shaded bands indicate a simultaneous fit to all the data. The green data points are obtained from simulations where the strange quark mass is kept approximately fixed to its physical value, the blue data points are from simulations where the flavour average of the light and strange quark masses is held fixed and the red data points indicate that the strange and light quark masses are equal.

capsulated in the Standard Model). The Standard Model has a number of limitations, most notably that it only involves three of the four forces of nature and does not explain the origin of dark matter and energy. Astrophysical evidence suggests that dark matter (which does not interact electromagnetically and so is difficult to detect) comprises 85% of the total mass of the Universe. This project is concerned with computing nucleon properties relevant for experiments aiming to directly detect dark matter particles, for those experiments investigating the poorly understood neutrino sector of the Standard Model and for precision experiments searching for signals of new interactions in beta decay.

Results and Methods

In order to calculate the properties of nucleons, QCD is formulated on a space-time Euclidean grid, with a finite lattice spacing (a) and volume (V). Representative (gauge) configurations of the gluon and quark fields are generated by hybrid Monte Carlo simulation and the quantities of interest can be extracted (after a statistical average) from “correlation” functions which




Figure 2: The axial form factor of the nucleon as a function of the momentum transfer squared. The curved bands are our results for the form factor after extrapolation to the physical light quark mass point and to the continuum limit at infinite volume. Two different fit forms were used to parameterise the lattice data (differentiated by the colours blue and green in the figure), see Ref. [4] for more details. The straight lines indicate the slope of the form factor at $Q^2 = 0$, which is related to the axial radius of the nucleon (analogous to the charge radius). The experimental value for G_A in the forward limit, is shown as a black circle [2].

are computed on top of the gauge configurations. The simulations must be repeated for several values of the lattice spacing to enable the continuum limit (a to zero) to be performed. Similarly, finite (spatial) volume effects must be explored. A large part of the computational cost is due to inverting large matrices of size $12V \times 12V$ (where V is of the order of $64^3 \times 192$) that are related to the propagation of quarks across the lattice. The condition number of these matrices increases inversely proportional to the mass of the light (u and d) quarks and for this reason unphysically large light quark masses are often employed in the simulations. This necessitates an additional extrapolation to the physical light quark mass value. Note that the light quark mass is proportional to the pion mass squared and this is used when displaying the light quark mass dependence of physical quantities.

This project is part of a larger analysis involving the computation of correlation functions relevant for determining nucleon structure observables. A unique feature of our calculation is that we use configurations generated with open boundary conditions which enable lattice spacings as low as 0.04 fm to be realized while maintaining ergodicity in the simulation [1]. We vary the light quark mass such that the physical point is approached along two trajectories with an additional trajectory where the light and strange quark masses are equal. This project completes the analysis by computing correlation functions on a physical quark mass ensemble with a fine lattice spacing of 0.06 fm. We are still in the process of analysing much of the data, however, first results are presented in Figs. 1 to 3. The masses of the Ξ , Λ and Σ baryons have been computed in addition to that of the nucleon. We are able to precisely fix the lattice spacing using the Ξ mass, with the other masses at the physical point being predictions. Note that the data points generated as part of this project (directly at the physical point) are consistent with the experimental results (filled circles in Fig. 1). The nucleon sigma term, given by the slope of the nucleon mass with respect to the light quark mass will be extracted from this analysis. The sigma term is required to compute




Figure 3: The induced pseudoscalar form factor of the nucleon from the same analysis as in Fig. 2 and Ref. [4]. Note that the form factor has been multiplied by the mass of the muon and divided by twice the nucleon mass. This combination is relevant for comparing with the experimental results (black circle) for the form factor obtained in muon capture, see Ref. [4].

dark matter-nucleon scattering cross-sections in direct dark matter detection experiments.

Figures 2 and 3 display the axial and induced pseudoscalar form factors of the nucleon after extrapolation to the continuum limit and to physical light quark mass. The axial form factor is a crucial input to Monte-Carlo generators which determine the energy distributions of neutrinos in long baseline neutrino oscillation experiments. The two form factors are linked via symmetry relations and our analysis is one of the first to show the lattice results are consistent with these relations and also agreement with the experimental results for the induced form factor at one particular point where this exists.

The data were generated on SuperMUC-NG using batch jobs involving 432 nodes with 48 tasks per node. While we compute many measurements per configuration the data is stored in HDF5 format and only around 1,000 files are generated, which were stored on the project WORK directory. As mentioned above, the main computational expense is the inversion of the large sparse Dirac matrices for which we used the locally deflated domain decomposition solver of openQCD [3].

Ongoing Research / Outlook

The data generated during the project are still being analysed. Final results, incorporating also earlier simulations, will include the contribution of the light and strange quark spin to the spin of the nucleon (needed to understand the spin puzzle) and the scalar and tensor charges which are used to provide bounds on the coupling strength of beyond the Standard Model interactions contributing to beta decay.

References

- [1] M. Bruno et al., JHEP 02 (2015) 043.
- [2] Particle Data Group collaboration, M. Tanabashi et al., Phys. Rev. D98 (2018) 030001.
- [3] <http://luscher.web.cern.ch/luscher/openQCD/>
- [4] RQCD collaboration, G. Bali et al., JHEP 05 (2020) 126.

Isospin breaking effects in QCD

RESEARCH INSTITUTION

¹Department of Physics, University of Wuppertal, Germany

²Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany

PRINCIPAL INVESTIGATOR

Kalman K. Szabo^{1,2}

RESEARCHERS

Sz. Borsanyi¹, Z. Fodor^{1,2,3,4,5}, J. N. Guenther⁶, C. Hoelbling¹, S. D. Katz³, L. Lellouch⁷, T. Lippert^{1,2}, K. Miura^{7,8,9}, L. Parato⁷, F. Stokes², B. C. Toth¹, Cs. Torok², L. Varnhorst¹

PROJECT PARTNERS

³Institute for Theoretical Physics, Eötvös University, Budapest, Hungary, ⁴University of California, San Diego, USA

⁵Pennsylvania State University, Department of Physics, USA, ⁶Department of Physics, University of Regensburg,

⁷Aix Marseille Univ, Université de Toulon, CNRS, CPT, Marseille, France

⁸Helmholtz Institute Mainz, Germany, ⁹Kobayashi-Maskawa Institute, Nagoya University, Japan

SuperMUC Project ID: pn56bu (Gauss Large Scale project)

Introduction

The standard model of particle physics describes the vast majority of experiments and observations involving basic constituents of visible matter. Any deviation from its predictions would be a sign of entirely new, fundamental physics.

A particularly important, long-standing discrepancy concerns the magnetic moment of an elementary particle called the muon. The current measurement and theoretical calculations of this property have similar precision, but disagree significantly.

To transform this disagreement into an actual discovery of new physics, an ongoing experiment at Fermilab (Batavia, USA), and one planned at J-PARC (Tokai, Japan), are aiming to reduce the measured uncertainty by a quite large factor of four.

On the theory side, the largest part of the error comes from the leading-order, hadronic vacuum polarization (LO-HVP) contribution. This term accounts for important quark, gluon and photon quantum fluctuations in the vacuum, which are described by quantum chromodynamics (QCD) and quantum electrodynamics (QED). Up until now, the most precise computations of this contribution have been performed using dispersion relation based phenomenological analyses, with input from electron-positron annihilation experiments.

Here we present a completely independent, ab-initio computation of the LO-HVP term, in which the equations of QCD and QED are solved using supercomputers and state-of-the-art numerical techniques. We reach a precision similar to that of the phenomenological approach for the first time. Surprisingly, our result leads to a standard model prediction for the muon's magnetic moment that is in agreement with the current experimental

measurement, suggesting that new physics is not needed to explain this measurement at present levels of precision.

The methods developed here will be useful to continue improving the accuracy of the standard model prediction, as will be required to pursue the search for new physics in ongoing and future experiments designed to measure the magnetic moment of the muon.

Isospin breaking effects

The ab initio calculation means in the above context lattice quantum field theory. A space-time grid is introduced and at every point of it the time evolution of various quantum operators are determined (to be more specific a path integral formalism is applied to that end). In some sense it reminds us to meteorology. Usually, people also introduce a three-dimensional grid, temperatures, pressures and wind velocities are measured and using the underlying equations the time evolution is determined. In both cases it is a heroic effort.

Let us come back to the magnetic moment of the muon. Obviously, the most important goal is to reach an accuracy, which is compatible with the expected experimental errors. Only reaching this accuracy guarantees that the experimental findings of several hundred million dollars are fully utilized and only with this accuracy can we decide if and what sort of new physics is there. When we speak about precision a sub-percent error is needed. Electromagnetic and strong isospin breaking effects are on the percent level. Thus any reasonable result needs the inclusion of these effects. This is a very hard task. The electromagnetic interaction is weak and long-ranged, whereas the strong interaction is strong and short-ranged. Keeping both of them in a system is more than just challenging.

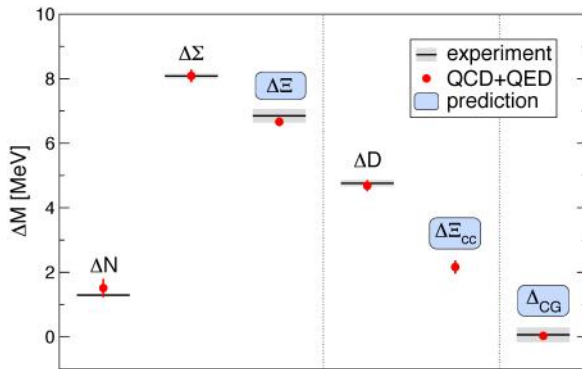


Figure 1: Isospin splittings for various baryons.

Proton-neutron mass difference

In 2015, we published a paper in Science [1], in which we included both theories in order to determine the mass splittings between the proton and the neutron (and for many similar hadron pairs). The building blocks of these nucleons are up and down quarks. Most of the lattice QCD calculations are using the same, very small mass for both of them, which is a very good approximation in most cases (no strong isospin breaking scenario). In this isospin symmetric case without QED the protons and the neutrons have the same mass, they are mass degenerate. However, if one introduces electromagnetism (QED) the proton becomes slightly heavier than the neutron. The world, as we know it today, would not exist. No stars, no people, not even atoms. This is in full contradiction with observations. One has to introduce strong isospin breaking, too. These two extensions of the isospin symmetric case are contributing similar amounts, but there is a huge cancellation between them. In our 2015 paper, we have shown how to include both effects. This also had a deep theoretical outcome. It was shown how to use renormalization theory simultaneously for QCD and QED. The calculation was fundamental and expensive. The main result, the isospin splittings for various baryons, is shown in Figure 1.

This figure is shown in many textbooks and year by year belongs to the 0.3% best cited lattice papers ever written. Nobel Prize Laureate Frank Wilczek writes about this work “This progress encourages us to predict a future in which nuclear physics reaches the level of precision and versatility that atomic physics has already achieved, with vast implications for astrophysics, and conceivably for technology. We can look forward to much more accurate modelling of supernovae and neutron stars than has so far been possible, and entertain dreams of refined nuclear chemistry, enabling, for example, dense energy storage and ultrahigh-energy lasers.”

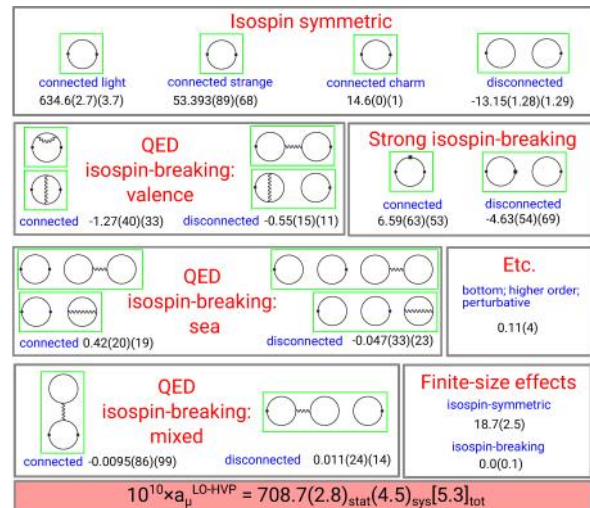


Figure 2: Isospin breaking.

The muon’s magnetic moment

Indeed, a similar level of precision is needed to decide if there is a new physics signal in the magnetic moment of the muon. After we have introduced full dynamical QED into our work, we were using it in a step-by-step procedure in our 2020 paper on the muon [2]. Electromagnetism and isospin breaking is introduced and controlled in a Taylor-series manner. This approach made it possible to shuffle around with the CPU demands in a way that the CPU was invested in those parts of the calculation which minimized the overall error. Surprisingly, our results are in agreement with the experimental value for the magnetic moment of the muon and the long-standing discrepancy seems to have disappeared. Clearly, this work should be cross-checked and repeated by other groups.

After many years of hard work and with the inclusion of strong isospin breaking and electromagnetism, we reached a sub-percent accuracy, comparable to the experimental errors on the muon’s magnetic moment. We are waiting for the new experimental results from Fermilab (USA), which should come out in a few months, to see if there is new physics beyond the standard model.

References and Links

- [1] Borsanyi et al, Science 347, 14521455 (2015).
- [2] Borsanyi et al, arxiv:2002.12347 (2020).

Charm sea effects on heavy flavor mesons

RESEARCH INSTITUTION

¹Bergische Universität Wuppertal

PRINCIPAL INVESTIGATOR

Francesco Knechtli¹

RESEARCHERS

Salvatore Cali², Jochen Heitger³, Tomasz Korzec¹

PROJECT PARTNERS

²Jagiellonian University, Cracow, Poland

³Westfälische Wilhelms-Universität Münster

SuperMUC Project ID: pn56fo

Introduction

Quantum Chromodynamics (QCD) is the theory proposed in the early 1970's to explain the properties of the strong interactions. The latter hold quarks together to form composite particles called hadrons. Quarks are elementary particles of spin 1/2 and represent, along with leptons, the smallest building blocks of matter, according to the Standard Model of Particle Physics. The strong interactions are mediated by spin 1 particles, the so-called gluons. One of the main features of QCD is confinement. Quarks and gluons are never observed in isolation, but only within hadrons, that in turn are usually classified into mesons (composite particles made of a quark and an antiquark) and baryons (composite particles made of three quarks). The newly discovered states with a charmonium component could hint at the existence of exotic states such as tetra- or penta-quarks, but this is still under debate.

Quark confinement and hadron properties cannot be understood using perturbative methods and nowadays lattice QCD represents one of the most suitable tools to investigate QCD properties starting from first principles. In this approach QCD is discretized on a Euclidean four-dimensional space-time and the quantities of interest, like masses and decay constants of hadrons, can be computed numerically via Monte Carlo methods. This kind of study requires a huge computational effort, especially when considering a theory with dynamical quarks, and the use of supercomputers is necessary if we want to achieve results that can be compared with experiments.

QCD encompasses six flavors of quarks (up, down, strange, charm, bottom, top). However, since quark masses cover a large range of values that can differ by several orders of magnitude, lattice QCD simulations often include the effects of only two, three or at most four flavors in the sea. In this project we estimate the effects of including a sea charm quark. Its inclusion means a significant effort in tuning of simulations and we want to know at which level of precision it matters. In order to do that we consider

QCD with just a single species of quarks, the charm quarks, and we compare the results obtained with this simplified model to a theory without dynamical quarks (often called quenched QCD). This gives the possibility to use moderately large lattice volumes and perform reliable extrapolations to zero lattice spacing.

We study charmonium states, which are composite particles made of a charm quark and a charm antiquark. The charmonium system, frequently characterized as the “hydrogen atom” of meson spectroscopy owing to the fact that it is non-relativistic enough to be reasonably well described by certain potential models, is the perfect testing ground for a comparison of theory with experiment. Over the last years, there has been a renewed interest in spectral calculations with charmonia because of the experimental discovery of many unexpected states, e.g. the so-called X, Y, Z states, which highlight the need for a more complete theoretical understanding. However, to accurately understand the charmonium spectrum, one must also investigate properties other than masses, such as decay rates. In the charmonium system, the lowest-lying states lie below the $D\bar{D}$ threshold, resulting in relatively narrow widths due to the absence of OZI allowed strong decays. This means that radiative transitions, i.e. transitions from an initial state to a final state via the emission of a photon, can have significant experimentally accessible branching ratios. Therefore, a lattice calculation of the decay constants addressed here provides valuable theoretical insight for experiment at the fully non-perturbative level such that we consider them as natural and representative observables to quantify charm sea quark effects in hadron physics beyond the mass spectrum.

Results and methods

With $N_f = 2$ charm quarks, the calculation has been performed on a total of five ensembles that differ in the lattice spacing, while in the quenched ($N_f = 0$) case four ensembles were used [1]. The measurements of masses and decay constants over a wide range of

lattice spacings, including very fine ones, enables a safe continuum extrapolation for both theories, which is shown in Fig. 1 for the decay constants of the η_c and J/ψ mesons [2]. In the $N_f = 2$ theory we consider only ensembles with lattice spacings $a \leq 0.05$ fm, which is typically the smallest value which can be achieved in simulations including the light quarks. Our simplified setup allows to simulate lattice spacings smaller by a factor two. They are essential to control the continuum extrapolations which show the expected linear behavior in the lattice spacing squared (a^2). The comparison of their values in the continuum limit allows to compute the impact of charm sea quarks on these observables. For the decay constant f_{η_c} the relative difference between $N_f = 2$ and $N_f = 0$ is 0.48(34)%, for the decay constant $f_{J/\psi}$ it is 0.12(77)%. These results mean that the decoupling of the charm quark at low energies works very well at an energy scale as large as about 500 MeV. While the decay constant f_{η_c} has not been measured experimentally, the experimental value for the J/ψ meson is $f_{J/\psi} \cong 407(4)$ MeV (obtained from the partial decay width of J/ψ into an electron-positron pair [3]). The discrepancy with our value is due to several effects (light sea quarks, charm annihilation, electromagnetism, number of charm quarks, meson mass).

We also study the effects of a dynamical charm quark on the hyperfine splitting of a B meson made of a bottom quark (anti-quark) and a charm anti-quark (quark) [2]. In particular we focus on its pseudo-scalar (B) and vector channel (B^*). The properties of the charmed B meson system are of special interest in spectroscopy because they are the only heavy mesons consisting of heavy quarks with different flavors. Because they carry flavor, they cannot annihilate into gluons and so are more stable with widths less than 100 keV. At the LHC, with its higher luminosity, the spectroscopy and decay of B^* mesons can now be experimentally measured such that on the theory side complementary precision studies of these meson states by means of lattice QCD become increasingly important. Our simplified setup provides a first estimate of the charm sea effects on the ratio of the B^* to B meson mass. For a “bottom” quark with a mass twice of the physical charm quark we find a relative difference for this mass ratio of 0.092(50)% between the $N_f = 2$ and $N_f = 0$ theories, which is comparable to what we found for the case of charmonium [1].

The computation of the meson masses and decay constants was performed with our measurement program [4], which invokes the openQCD [5] solvers for the inversions of the Dirac operator. The latter are state-of-the art Krylov space solvers that support even-odd preconditioning, an algebraic multi-grid method, mixed precision calculations and various further improvements. Optimized versions of the solvers for SuperMUC-NG (AVX instructions) exist. Our package has been developed for a CPU environment and includes an MPI parallelization. For the inversion of the Dirac operator of the heavy quarks we use a modified version of the SAP GCR solver, which includes the distance preconditioning method.

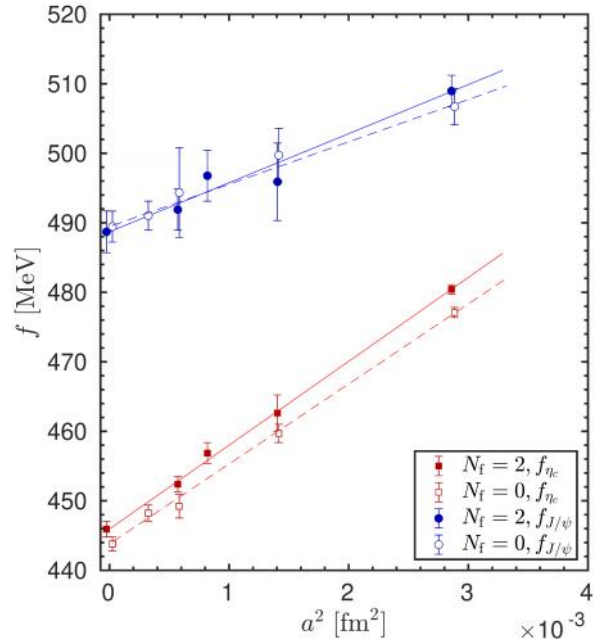


Figure 1: Continuum extrapolations at $a = 0$ of the meson decay constants f_{η_c} and $f_{J/\psi}$ on our $N_f = 2$ (full markers) and $N_f = 0$ (empty markers) ensembles. The continuum limit values are slightly displaced horizontally.

Ongoing Research / Outlook

In this project, we have neglected the effects of charm annihilation in charmonium. They will be computed in a follow-up project, where also more channels (J^{PC} quantum numbers) and mixing with glueballs will be investigated.

References and Links

- [1] S. Cali, F. Knechtli, T. Korzec, Eur. Phys. J. C, 79 (7), 607, 2019.
- [2] S. Cali, K. Eckert, J. Heitger, F. Knechtli, T. Korzec, in preparation, 2021.
- [3] D. Hatton, C.T.H. Davies, B. Galloway, J. Koponen, G.P. Lepage, A.T. Lytle, Phys. Rev. D 102, 054511, 2020.
- [4] <https://github.com/to-ko/mesons>
- [5] <http://luscher.web.cern.ch/luscher/openQCD/>

TMDs and parallel transport in QCD

RESEARCH INSTITUTION

¹Regensburg University

PRINCIPAL INVESTIGATOR

Andreas Schäfer¹

RESEARCHERS

Michael Engelhardt², Piotr Korcyl³, Maximilian Schlemmer¹, Thomas Wurm¹, Alexey Vladimirov¹, Christian Zimmermann¹

PROJECT PARTNERS

²New Mexico State University, Las Cruces, USA

³Jagiellonian University, Krakow, Poland

SuperMUC Project ID: pn56yo (PRACE project)

Introduction

January of this year, the US Department of Energy announced plans to build a new Electron Ion Collider (EIC) at Brookhaven National Laboratory. One of the primary goals of this accelerator will be the exploration of Transverse Momentum Dependent parton distributions (TMDs) in the proton. These functions have many fascinating properties. For example they parameterize angular asymmetries for mesons produced in the scattering of electrons and transversely polarized protons, as first observed by the HERMES collaboration in 2005 [1]. These asymmetries are naively counter intuitive as they seem to violate fundamental time reversal symmetry which is a property of Quantum Chromodynamics (QCD), the theory of quarks and gluons. This asymmetry can be related to another counter-intuitive but also well-understood quantum mechanical effect, called the Aharonov-Bohm effect, see Figure 1: The interference pattern of a two-slit experiment with electrons moves, if a magnetic field is placed between the slits, even if it is shielded such that no electron can possibly penetrate into it. While there are differences between these two phenomena the fundamental origin of them is the same and applies to all gauge theories: The gauge degrees of freedom lead to non-trivial parallel transport which introduces phase factors and interference effects which result in a surprising behavior of reaction probabilities. Gauge theories have a geometric interpretation which links these phases to so-called gauge links. The non-trivial parallel transport of gauge links is generated by the field-strength tensor of the gauge theory in a similar way as non-trivial parallel transport in General Relativity is generated by the energy momentum tensor. General Relativity is substantially more complex than Special Relativity and the physics of TMDs is substantially more complex than that of PDFs (parton distribution functions) or purely collinear physics, which depends only on longitudinal coordinates and momenta. In fact, for a purely collinear process all gauge links can be set to one and thus simply disappear. In an analogous manner any process which probes only a straight line cannot detect the warping of space-time in General Relativity. TMDs depend on both, longitudinal and transverse momenta and are, therefore, sensitive to all such effects in QCD.




Figure 1: Sketch of the Aharonov-Bohm effect. The interference pattern of electrons traveling through two slits is shifted by the magnetic field B , even if it is completely shielded. This is an example of non-trivial parallel transport.

This makes them both very complicated and very interesting objects.

Of the many fascinating new properties of TMD physics one of the most characteristic ones is the appearance of “soft factors”. Gauge links can be interpreted as a coherent superposition of arbitrarily many soft gluons and the interactions caused by these gluons cannot be described by standard rules of quantum field theories, which necessitated the introduction of new rules and a new dependence on a parameter called rapidity. To compare measurements at different rapidities requires knowledge of the so called “rapidity anomalous dimension”, which is a function of the transverse displacement, b , of the parallel transport. For values of b which are not very small this crucial function could until very recently only be guessed but not calculated. This has changed this year, see next section.

Before discussing these results a general remark on the status of Lattice QCD is in order. Because both the validity of QCD as such and the soundness of Lattice QCD are established beyond reasonable doubt, Lattice QCD has acquired a similar status as experiment and consequently the control of all sources of error, especially systematic errors has also become equally important. Just as the determination of systematic uncertainties is usually the most demanding task of a high energy experiment so it is for Lattice QCD. The calculations we performed are especially well suited to control these uncertainties.

Results and Methods

Because TMDs have properties which differ significantly from those of other lattice observables and because so many of their properties are still unknown, the emphasis of research for the different groups working on TMDs lies on the development of methods to extract the quantities of interest from a lattice simulation. In our project we used a new method to analyze the TMD correlators we have calculated to extract the rapidity anomalous dimension, which is also called Collins-Soper kernel. This method was only published recently [2]. Among other advantages this method allows to extract the same information from, in principle, 16 independent correlations. While for many of them signal over background is probably too small for practical use this should provide a valuable tool to estimate systematic uncertainties. So far, we have analyzed only three of these correlations and the results agree within errors. Also, we have analyzed different lattice spacings which allows us to study carefully the continuum limit which is usually the least controlled source of systematic error. This analysis is still ongoing. Figure 2 shows only one sample plot. This is a very busy plot because it contains results from five different approaches for the rapidity anomalous dimension as a function of the transverse distance b . The main message of this figure is that all of these methods give compatible results. The blue-shaded band is the result of a fit to experimental data. Because this data does not really constrain its form for b larger than about 0.4 fm the shape of the band is determined primarily by the chosen parameterization. The result of a perturbative QCD calculation in next-to-next-to-next-to leading order (NNNLO) plus resummation is also unreliable beyond 0.4 fm. The MIT results (labeled Bernstein and Hermite) is from a quenched simulation and has hard to estimate systematic uncertainties in addition to the plotted ones.

The US-German-Chinese LPC collaboration, to which we also belong, uses a completely new approach to extract the rapidity anomalous dimension from CLS configurations but was not part of this project. LPC has attributed a generous systematic error to these results in view of several conceptual issues which are still debated.

The pink, purple and brown dots are results of the present project. Ours were full QCD simulations (in contrast to quenched) and we also included so called higher twist effects in our analysis for which we, therefore, also obtained quantitative estimates. The large systematic errors are primarily driven by these higher twist effects. All lattice results came out only 2020. In sum they demonstrate that this crucial fundamental function of modern hadron physics should be known with high precision in very few years.

This PRACE project was granted 44 M Core-hours, starting on SuperMUC Phase 2 and finishing on SuperMUC-NG (about 33 M Core-hours). In addition




Figure 2: Results of five different approaches for the rapidity anomalous dimension. For explanations see text.

we profited from friendly user running in the starting phase of SuperMUC-NG. Typically we used 24 nodes with 1,152 cores. The software we use is largely public domain within the lattice QCD community. It goes under the name of Chroma and is written primarily in C++. It is based on the QDP++ library. The main computational effort in lattice simulations goes into the inversion of huge sparse matrices for which one can choose between a large variety of highly optimized code using different algorithms (we used e.g. the OpenQCD multigrid solver [4]). We did not optimize these solvers but used them as black box. In this project we also used also code we have co-developed as members of the TMD-Collaboration in the US several years ago for the specific purpose of analyzing TMDs, see [3]. The configurations analyzed were generated by the CLS collaboration (partly by us) in a large and long-term effort. These data are stored in Regensburg, Mainz and Berlin. Their generation was not part of this project. For data handling we used primarily HDF5.

Ongoing Research / Outlook

As illustrated by Figure 2 the analysis of TMDs is an internationally highly competitive research topic such that having access to competitive computer resources is crucial. Presently, it is not yet decisive because the optimal approach to extract the Collins-Soper function is still disputed but it will be once this discussion has converged, which in view of the rapid progress in this field cannot take long. Therefore, we expect to submit soon a new TMD proposal.

References and Links

- [1] A. Airapetian et al. (HERMES Collaboration), Physical Review Letters 94 (2005) 012002
- [2] A.A. Vladimirov and A. Schäfer, Physical Review D101 (2020) 074517
- [3] B. Yoon et al. (TMD Collaboration), Physical Review D96 (2017) 094508
- [4] <https://luscher.web.cern.ch/luscher/openQCD/>

How large is the color magnetic field in a proton?

RESEARCH INSTITUTION
Regensburg University

PRINCIPAL INVESTIGATOR
Andreas Schäfer

RESEARCHERS
Gunnar Bali, Jacques Bloch, Simon Bürger, Sara Collins, Meinulf Göckeler, Marius Löffler, Thomas Wurm

PROJECT PARTNERS
—

SuperMUC Project ID: pn69ma

5

Introduction

With the advent of high luminosity spin polarized particle accelerators exploring the spin structure of hadrons, i.e. bound states from quarks and gluons, became a major research topic. In the beginning experimental results contradicted naive expectations to such an extent that this physics was often summarized under the title “spin crisis”. To be precise, this term summarized the fact that the distribution of the total proton or neutron (collectively called nucleon) spin of $\frac{1}{2}$ on quark and gluon spin as well as quark and gluon orbital angular momentum was hotly disputed. Over the years, however, to a large extent thanks to lattice calculations, a quite detailed understanding was developed, such that today more subtle questions have moved into the focus of attention. Quantum Chromodynamics (QCD), the theory of quarks and gluons, has strong similarities with Quantum Electrodynamics (QED) the quantum theory of the electromagnetic interaction but there exist also fundamental differences. As in QED spin couples to the (color) magnetic fields and (color) charged quarks experience a (color) Lorentz force in such a (color) magnetic field. In contrast to QED, however, all quarks and gluons in a nucleon are “confined”, i.e. they cannot be extracted in isolation from a nucleon making access to the color analogs of well-known electromagnetic effects far more difficult. One can only observe certain color neutral correlations of quarks and gluon fields which are classified by abbreviations like d_2 the correlation we are concerned with in this project. This correlation is judged so important that it motivated several large scale experiments and many theory investigations. Figure 1 shows an up to date compilation, with various model predictions to the right and various experimental data as well as our old lattice results [2] to the left. Solid symbols denote results for the proton and open ones those for the neutron. The most important fact presented in this figure is that d_2 is small. (Naively one could have expected numbers anywhere in the range from -0.02 to $+0.02$, as is illustrated by the scattering of model predictions.) Because d_2 can be related to the color Lorentz force on quarks in a color magnetic field, which is in turn oriented along the spin vector direction of the studied

hadron, this implies that this force is weaker than one could have expected. Another remarkable feature of this plot is that our old lattice results [3] (plotted in green) are displayed together with experimental data, illustrating the status lattice QCD calculations have gained meanwhile in hadron physics. In contrast model predictions of various sorts are depicted by the symbols on the right. They scatter widely over the phenomenologically plausible region. While this kind of plotting is kind of flattering for us, it also made us feel a bit uneasy, because at the time (2005), lattice QCD techniques were far less developed, and the error analysis leading to the displayed error bars does not fulfill our present standards. In particular the most problematic uncertainty, caused by the required extrapolation to zero lattice constant, i.e., the physical space-time continuum, could not be determined at the time. (In lattice QCD continuous space-time is substituted by a four-dimensional hypercubic lattice of constant lattice spacing.) This we have now improved.

Results and Methods

The problem of controlling the continuum limit is not only one of required computer resources (simulations with smaller lattice constants require far more lattice points. Actually, their number scales like the fourth power of the inverse lattice spacing) but also a fundamental theoretical one. QCD has different topological sectors and different ones have to be sampled to get correct results. For lattice constants below 0.05 fm this becomes unfeasible, a problem known as “topological freezing”. On the other hand, for lattice spacings above 0.1 fm discretization errors become large such that the typical lever arm for a continuum extrapolation is only about a factor of two. (The radius of a proton is roughly 0.7 fm, so one needs fine lattices to resolve its structure.) Within the very large CLS collaboration of collaborations we have implemented an algorithmic formulation which circumvents this problem. Figure 2 and Figure 3 show some of the preliminary results we obtained. Such calculations proceed in two steps. First one generates large ensembles of field configurations (typically each ensemble contains several thousand configurations) which contain information on the structure of all hadrons (e.g. bound states of quarks and




Figure 1: Compilation of the experimental and theoretical results on d_2 taken from a recent publication of the SANE collaboration [1], downloaded from INSPIRE-HP [2]. The red and blue bands at the bottom of the figure give estimates for systematic uncertainties of the experimental results.

gluons). From these one then extracts information on specific aspects of specific hadrons, like d_2 of the proton. The main results of this project are updates of the green lattice points in Figure 1. Looking at our results in Figure 2 and 3 the most important observation is that the lattice constant dependence is significant for the proton, making this a showcase example for how important control of the continuum limit is and thus adding to the justification of the whole CLS effort. Although our results contain besides all other systematic uncertainties also the uncertainty of the continuum extrapolation the resulting error bars are significantly smaller than for the old simulations. In particular the result for the proton is now distinctly non-zero. (Its size actually fits now better to phenomenological expectations.) Our results provide the most precise calculation of the color Lorentz force, e.g., the force exerted by the color magnetic field on a quark moving transversely to the color magnetic field direction, which in turn is oriented along the hadron spin, as of today.

Ongoing Research / Outlook

While our result for the proton is not incompatible with the experimental data there is a certain degree of tension which adds to the motivation of future experiments and lattice studies. While along the lines of standard Operator Product Expansion techniques, which we used in this work, it will be very demanding to substantially improve on the precision reached, let us note that A. Schäfer is also part of a large scale US-Chinese-German collaboration called LPC which explores since a few years a completely different approach in Lattice QCD, going under the name of quasi- or pseudo-distributions which was shown in [4] to be most promising to study d_2 and related quantities on the lattice. It would be most interesting to perform such simulations on the same configurations used in the present study and to compare the results. This project, providing simultaneously much more novel information on nucleon structure, would probably require a comparable amount of computer resources (without serious benchmarking this is impossible to say more precisely). SuperMUC-NG is an ideal HPC platform to serve our purpose.




Figure 2: Preliminary results for d_2 of the proton, continuum extrapolated to the value at the left. Several points have been displaced a bit horizontally to avoid overlapping error bars. The blue band is the results of a global fit to all data, including all dependencies on simulation parameters.




Figure 3: The same for d_2 of the neutron.

The software we use is largely public domain within the lattice QCD community. It goes under the name of Chroma and is written primarily in C++. It is based on the QDP++ library [5]. The main computational effort in lattice simulations goes into the inversion of huge sparse matrices for which one can choose between a large variety of highly optimized code using different algorithms (e.g. multigrid techniques) and usually programmed in machine language. We did not optimize these solvers but used them as black box.

The configurations analyzed were generated by the CLS collaboration (partly by us) in a large and long-term effort. These data are stored in Regensburg, Mainz and Berlin. Their generation was not part of this project. For data handling we used primarily HDF5. The project was funded by LRZ with 33 M core-hours and profited in addition from the friendly user phase of SuperMUC-NG. This project did not face any severe specific problems. The underlying theory and algorithms were well understood from the very beginning. It just required the processing of hundreds of Terabytes of data.

References and Links

- [1] W. Armstrong et al. (SANE Collaboration) Phys. Rev. Lett. 122 (2019) 022002.
- [2] <https://inspirehep.net>
- [3] M. Göckeler et al. Phys. Rev. D 72 (2005) 054507.
- [4] S. Bhattacharya et al. ArXiv:2004.04130.
- [5] <https://github.com/usqcd-software/qdpxx>

A new twist on plasma-based accelerators

RESEARCH INSTITUTION

¹Ludwig-Maximilians-Universität München

²Max-Planck-Institut für Quantenoptik, Garching

PRINCIPAL INVESTIGATOR

Andreas Döpp^{1,2}

RESEARCHERS

Jorge Vieira³, Stefan Karsch^{1,2}

PROJECT PARTNERS

³GoLP/Instituto de Plasmas e Fusão Nuclear, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal

SuperMUC Project ID: pn69ri

Introduction

Over the last century, high-energy particle colliders such as the LHC have played a crucial role in improving our understanding of the universe. But with accelerators aiming for ever higher collision energies, the size and cost of these facilities has drastically increased. To reach the tera-electron-volt frontier, accelerators using radio-frequency technology require tens of kilometers of acceleration length are required and such accelerators will eventually become too expensive to be built. Accordingly, a number of alternative accelerator concepts have been explored over the last decades. One of the most promising is wakefield acceleration in plasmas.

The technology relies on an intense drive beam, either a laser pulse or a particle bunch, ploughing through a plasma medium and pushing the electrons aside by their ponderomotive force or Coulomb repulsion, respectively. This sets up charge separation fields that pull electrons back and cause them to oscillate around their equilibrium position. As the driver travels through the medium with a velocity close to the speed of light, the field structure ("wakefield") follows it at the same speed. Electrons injected into this moving wakefield can be accelerated as a so-called witness bunch. The accelerating field structure has a typical length scale given by the plasma wavelength, and both accelerating and focusing gradients are several orders of magnitude larger than in radio-frequency (RF) accelerators, leading to very dense and ultra short accelerated particle bunches (some tens of femtoseconds duration). Depending on the driver type, the process is either called laser wakefield acceleration (LWFA) or particle wakefield acceleration (PWFA). While the former can be studied at many high-intensity laser laboratories at drive pulse powers starting at the terawatt level, the latter typically requires drive bunches that can only be provided by a national-laboratory scale RF accelerator.

In this project „Hybrid wakefield concepts for electron acceleration and radiation generation“ [1] we are investigating new concepts for wakefield acceleration in plasmas. In particular, we focus on improving our understanding of laser-driven wakefields as potential

X-ray source for different laser parameters and show how a laser-driven accelerator can provide the seed for a beam-driven wakefield accelerators, thus enabling PWFA studies on a laboratory scale.

Results and Methods

The first prerequisite for a plasma wakefield accelerator is the presence of a plasma wave itself. We have measured the generation of such wakefields in experiments using phase contrast imaging of the density modulations of plasma waves. This technique has been previously employed to measure laser-driven plasma waves and we have for the first time used it to observe electron-driven plasma waves [2]. To make this distinction we use a thin foil acting as plasma mirror to reflect the laser and assure that only the electron beam can propagate into a second gas stage. A first observation was that the phase contrast signal in this second stage is much more pronounced in case of a pre-ionized plasma. We simulated both the initially neutral and the pre-ionized case and found that beam self-focusing is much weaker in the case without pre-ionization. This effect can be understood when considering the field structure of a relativistic electron bunch. Due to the vanishing fields on axis, the head of the bunch does not ‘see’ a plasma but only gas and will therefore not self-focus. We also noticed an unexpected difference of the overall structure of the plasma on a picosecond timescale in the laser and beam-driven cases. We have used SuperMUC to perform particle-in-cell simulations of this system. The feature could initially not be reproduced in simulations until we realized that ion motion (which is usually neglected in laser wakefield acceleration) can be significant on picosecond timescales. Including a separate ion species immediately showed the appearance of a cone-like feature (see Fig.1), which is caused by the ponderomotive force of the plasma wave on the plasma ions. This leads to the formation of a density peak on axis and an expanding ion ring off axis. Remaining quasi-neutral, the plasma electrons follow the ion motion, thus yielding the observed diffraction pattern of the probe beam.

Using SuperMUC we have also performed detailed three-dimensional numerical simulations that demon-

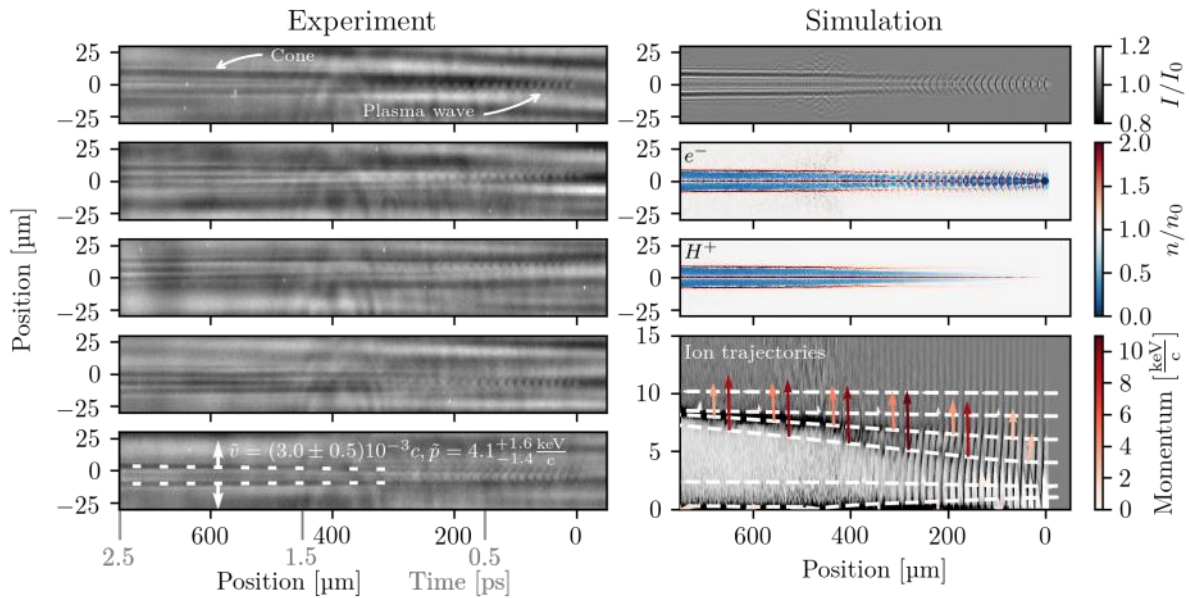


Figure 1: Shadowgram of picosecond ion motion due to the beam-driven plasma wave. Left: Experimental data. Right (top to bottom): Simulations of the shadowgrams, electron and ion distributions. Trajectories of ions are shown in the bottom right plot.

strate a novel configuration for a laser-wakefield betatron source that increases the energy of the X-ray emission and also provides increased flexibility in the tuning of the X-ray photon energy. This is made by combining two Laguerre-Gaussian pulses with non-zero net orbital angular momentum. We noticed that, because the phase velocity of a focused laser pulse depends on its orbital angular momentum (OAM), the laser intensity pattern associated with a superposition between multiple higher order modes, with different amounts of OAM, will exhibit a rotation about the propagation axis. Using these simulations, we were also able to understand this rotation from a perspective of the motion of light rays. Each photon in an OAM laser pulse will have an azimuthal component on its velocity. This rotation cannot be clearly seen if only a single mode is present, because each pure OAM mode is cylindrically symmetric. Adding several modes together will break the azimuthal symmetry of the intensity pattern of the resulting laser pulse. This way, the rotational motion of the photons can then become apparent. Naturally, this rotation only occurs if the combined OAM is finite. The simulations revealed that

this rotation on the laser intensity also induces a corresponding a rotation of the driven plasma waves, thus inducing a helical motion of the relativistically accelerated electrons. These results may pave the way to new electron beam configurations to drive plasma wakefields.

Ongoing Research / Outlook

Our research paves the way towards new compact radiation sources. We have recently observed first electron injection and acceleration in a laser-seeded PWFA [4]. Future experiments may lead to the production of high-brightness electron beams or beams with tailored OAM for radiation generation.

References and Links

- [1] <https://www.lrz.de/projekte/hlrp-projects/000000000F43AF8.html>.
- [2] M. F. Gilljohann, et al. Phys. Rev. X 9, 011046 (2019).
- [3] J. L. Martins, et al., Sci. Rep. 9:9840 (2019).
- [4] T. Kurz et al., arXiv:1909.06676 (2019).




Figure 2: 3D simulation results showing the trajectories of a sample of relativistic electrons from the rotating wakefield (left panel) and non-rotating (right panel) wakefields scenarios, coloured according to the energy. Electrons move towards higher values in the x1 axis.

Study of anomalies in B decays

as a probe of New Physics

RESEARCH INSTITUTION

¹University of Münster

PRINCIPAL INVESTIGATOR

Jochen Heitger¹

RESEARCHERS

Benoît Blossier², Savvas Zafeiropoulos³

PROJECT PARTNERS

²Laboratoire de Physique des 2 Infinis Irène Joliot-Curie, Orsay

³Centre de Physique Théorique, Marseille

SuperMUC Project ID: pn72gi

Introduction

The starting of Large Hadron Collider (LHC) in 2008 has been a key moment in the history of high energy physics. Even if the Standard Model of particle physics (SM) revealed very robust versus numerous experimental tests performed the past 50 years, some questions are still open, especially in the physics of quarks: such as the origin of the strong hierarchy observed among quark masses, the dynamics at work in the mixing pattern among quark flavors (u, d, c, s, t, b), the excess of the baryon-antibaryon asymmetry observed in the Universe with respect to what is expected by theoretical predictions of the SM. Thus, one has in mind to test the full set of New Physics models (NP), whose goal is to describe interactions at higher energy than the electroweak scale, 100 GeV, versus experimental measurements that are available at the moment or in the next future. To fully exploit experimental data in flavor physics, the goal is to detect deviations from the SM and then constrain efficiently the NP scenarios. In order to do so, theorists have to reduce as much as possible the uncertainties coming from the confinement of quarks in hadrons (mesons and baryons like the proton, the neutron, the pion, etc). Let's remind that 99% of the mass of the proton comes from the strong interaction mediated by the exchange of gluons, whilst only 1% of the observed amount of mass finds its origin in the presence of the Higgs boson among fundamental particles. The estimate of the confinement effects is really difficult in theory and the only robust way to do so is via lattice QCD simulations [1]. Under that general respect the purpose of our project is to provide hadronic inputs to study how various beyond the SM (BSM) scenarios [2, 3] are constrained from $b \rightarrow c$ quark decays.

Results and Methods

Semileptonic decays of $B_s \rightarrow D_s^{(*)}$ are particularly interesting in phenomenology. Mediated in the SM by a W weak boson, the $b \rightarrow c$ current is left-handed with the Lorentz structure (V-A). Scenarios beyond the SM with non-vanishing right-handed couplings have been

proposed. Such currents are helicity suppressed, i.e. they are reinforced in the case of a lepton emission with respect to the muon or the electron, typically 10 and 2,000 times lighter than the τ , on the leptonic side of the process. Thus, it is appealing to include it in constraints of that corner of NP scenarios. The aim of our proposal is to extract form factors associated to the hadronic transitions $B_s \rightarrow D_s^{(*)} l \nu_l$ mediated by currents allowed in SM and in NP scenarios. The most elegant approach to deal with the quark confinement in hadrons is lattice QCD. One regularizes the field theory of the strong interaction, quantum chromodynamics (QCD), by introducing a hard cut-off, the lattice spacing. A Monte-Carlo simulation gives a set of background gluon fields with a probability distribution derived from the QCD Lagrangian (thus, by construction, it incorporates quark loop effects of the sea) and correlation functions are computed from an average over that set of fields: it precisely corresponds to the path integral formalism of quantum field theory. The only parameters of the simulation are the bare strong coupling constant and the bare quark masses. The final result is obtained after an extrapolation to the continuum limit (the lattice spacing is made smaller and smaller). The numerical work done on SuperMUC-NG consists in inverting the Dirac operator with a background field corresponding to a set of gauge fields among the sets of different gauge ensembles that we analyze. The inverse of the Dirac operator is a quark propagator. The source can be itself a quark propagator, for instance to compute 3-pt correlation functions. Those propagators have color, spinor and space-time indices. A 2-pt correlation function is given by the trace of the product of two quark propagators while a 3-pt correlation function is given by the trace of the product of 3 quark propagators. The Dirac matrix is sparse and ill-conditioned. We take advantage of the former property by using domain decomposition methods that are particularly suitable on massively parallel machines. The latter property needs sophisticated algorithms. We use a solver based on even-odd preconditioning, Schwarz alternating procedure (right preconditioning), deflation subspace projection (left preconditioning) and the Generalized Conjugate Residual algorithm to solve in mixed




Figure 1: Effective mass of 2 heavy-strange mesons made of a strange quark and either a charm quark or a “light” bottom quark.




Figure 2: Form factor of the D_s elastic transition, in the heavy quark limit, in function of the relative velocity.

precision. The simulation was a subset of 30 configurations of size $64^3 \times 128$ points, each of them made of four 3×3 complex matrices [4]. For each configuration, 792 Dirac equations were solved in 20 separate jobs. Each job was using 2,048 cores. Per configuration, the output consisted in 168 ASCII files containing 2-pt and 3-pt correlation functions. The total size of output stored on \$SCRATCH is 10.6 GB. The simulation stored on \$WORK had a size of 560 GB. We have used 16 M core-hours up to now.

Ongoing Research / Outlook

SuperMUC-NG has helped us to get values of form factors at one point of the parameter space (lattice spacing, pion mass) among the 8 we will consider in the full analysis. The simulation performed here has a very large volume compared to others used in our global scientific project, needing such a large HPC facilities like SuperMUC-NG. Petascale supercomputers such as SuperMUC-NG have clearly a transformational impact on our field. We plot in Figures 1 and 2 the very first quantities extracted from the simulation: 2 heavy-strange meson masses obtained from 2-pt correlation functions and the elastic form factor in function of the relative velocity at the current. The statistics analyzed so far is quite low because our jobs are very expensive. To assess systematic effects a lot of correlation functions at different quark masses, momenta and source kinds have to be computed. The complete project takes more than two years, computations are shared among several HPC systems in France and Germany. That is why, at this stage, it is too early to get the answer to the scientific question we asked about the constrains of NP scenarios coming from $b \rightarrow c$ quark decays.

References and Links

- [1] For an introduction to the subject, see for instance the textbook *Quantum Chromodynamics on the Lattice : an introductory presentation* by C. Gattringer and C. Lang or *Lattice Quantum Chromodynamics* by F. Knechtli, M. Günther and M. Peardon.
- [2] S. Fajfer, J. F. Kamenik and I. Nisandzic, Phys. Rev. D 85, 094025 (2012).
- [3] L. Di Luzio, A. Greljo and M. Nardecchia, Phys. Rev. D 96, no. 11, 115011 (2017).
- [4] <https://wiki-zeuthen.desy.de/CLS/>

The QCD phase diagram in the quark mass plane

RESEARCH INSTITUTION

¹University of Wuppertal

PRINCIPAL INVESTIGATOR

Zoltan Fodor^{1,2,3}

RESEARCHERS

Szabolcs Borsanyi¹, C. Hölbling¹, B. Tóth¹, L. Varnhorst

PROJECT PARTNERS

²Jülich Supercomputing Centre

³Eötvös University, Budapest

SuperMUC Project ID: pn72tu (Gauss Large Scale project)

Introduction

Quarks are the constituents of the massive basic building blocks of visible matter. These building blocks are the hadrons, more precisely protons and neutrons, which are about 2,000 times heavier than electrons. The most important of these quarks are the up and the down quarks, whereas the strange quark also plays some role in these hadrons. A proton contains two up quarks and one down quark, whereas the neutron consists of one up quark and two down quarks (so-called valence quarks). In normal conditions, as on earth or even in the sun, the quarks are confined in hadrons, they can not get out. The confining force is called the strong force, the underlying theory is the theory of the strong interactions, or in other words, Quantum Chromo Dynamics, or in short QCD.

It is important to note that quark masses are constants of nature. They are what they are, and experiments can not change them. In nature, the up and down quark masses are just a few per-mill of the proton's mass, whereas the strange quark mass is about one tenth of the proton's mass.

In the early universe, at very high temperatures, the quark ingredients of protons and neutrons were freed. We call this state the deconfined phase. They were not confined in protons or neutrons, in some way these hadrons were boiled and the ingredients got liberated. As the universe expanded and cooled down, the quarks got confined in hadrons. We call this state the confined phase. It is an interesting question to tell the order of the transition between these two—confined and deconfined—phases. Is it a first order phase transition, as in the case of boiling water? Or is it a second order phase transition? Or it is not a phase transition at all, just an analytic transition, as in the case of melting butter? It was believed for a long time that the transition between free quarks and confined quarks is a first order phase transition. It is easy to understand why physicists conjectured this sort of transition type. When we boil water, the attractive forces between the water molecules are not strong enough to keep them together, and the water molecules get liberated. Similarly, it was believed that the qualitative difference between the confined and

liberated quarks is so huge that it must be accompanied by a first order phase transition. What would that mean for the early and even for the present universe? A first order phase transition is always a place for dramatic changes. Bubbles or droplets appear, they grow and finally fill out the whole space. This happens for the water vapor phase transition, and this would be the case for a first order QCD phase transition, too. Bubbles would appear, they would collide, and the low temperature confining phase would fill the whole universe. During the course of the phase transition, various remnants would appear, such as quark nuggets, or primordial black holes, just to mention two of them. These remnants could have cosmological consequences, or they could be even detectable by experiments. Despite the efforts of three decades, no such remnants were detected. Are our experiments not sensitive enough (this is the more plausible scenario)? Or even worse: is the underlying picture false (this seemed to be a less plausible scenario)?

Lattice methods

It was a heroic effort to determine the nature of the phase transition for physical quark masses, which our group carried out in 2006 and published in Nature [1] (this publication turned out to be the most cited lattice paper of the present century). This finding has fundamental consequences for the early universe, and for the possible remnants we might detect even today. Note, however, that the result not only has relevance for the early universe (Big Bang), but also for heavy ion collisions (Little Bang), which are carried out at the RHIC (Brookhaven, USA) and LHC (Geneva, Switzerland) accelerators.

Since it is such a fundamental result, it is important to repeat it with present day techniques and computational resources, which are far better than they were two decades ago. Most importantly, we should work with a discretization scheme which does not take roots of the determinants naturally appearing in the calculations, but uses the exact chiral symmetry of the theory (chiral symmetry appears when quark masses are zero). Unfortunately, these sort of discretization schemes are at least two orders of

magnitude more CPU demanding than the discretization scheme of the old calculation from 2006, which used the so-called staggered formalism (which also possesses some of the original chiral symmetry of the underlying QCD theory).

One important remark is in order. For many quantities chiral symmetry is less important. For instance, in 2008, we determined the mass of the proton and about a dozen other hadrons [2] (this work became one of the breakthroughs of the year of 2008 in all sciences). On the one hand, the masses of these particles do not rely on chiral symmetry and the calculation did not have any difficulty related to chirality. On the other hand, the quark masses are so small (as we mentioned, they are only in the per-mill range of the proton's mass) that the behavior of the QCD transition is close to the massless case. In addition, during the transition, the rapid change in the order parameters are all related to chiral quantities (for completeness: chiral condensate and chiral susceptibility). Therefore, one expects that an exactly chiral discretization scheme would lead to a cleaner description than other discretization schemes. Having said that, in principle the technique of the old results from 2006 should be correct, since all discretization schemes should give the same results as we approach zero with the lattice spacing. The only question is how fast one can reach the so-called scaling regime, where the deviation from the continuum result is fully controlled and it behaves as a power of the lattice spacing.

Results and further plans

Since the nature of the transition is an extremely crucial, fundamental information and all of the quantitative features (transition temperature, equation of state etc.) depend on it, it is very important to pin down the transition's nature for various points on a hypothetical light quark versus strange quark mass plane. This will embed the finding for the physical point and gives it further support. Here it is of particular importance to locate two characteristic points. The small mass and the large mass degenerate cases. For zero mass we expect a first order phase transition, which becomes second order if we increase the masses, after which it is an analytic crossover, probably all the way to the physical point. For infinitely large masses, the transition is known to be first order. Again reducing the masses all the way down to the physical point, with analytic crossover one should find a mass value with a second order phase transition. These cases are illustrated in Figure 1. Our goals are threefold. a) Determine the nature of the transition at the physical point with chiral fermions and with very fine lattices using so-called staggered formalism; b) locate the small mass, second order phase transition point (this task needs probably chiral fermions, too, but the first steps can be done with staggered fermions); c) locate the large mass second order phase transition point (in this case chiral symmetry is not important—staggered fermions can be used).

The diagram in Figure 1 is a modern version of what was first sketched by the Columbia group in Ref. [3].




Figure 1: Location of three exemplary potential decen-tralized retention basins. The basins are situated on the brook and are dammed up by notional dams.

Nature with its given quark masses is represented by a single point in this Figure. This physical point is in the region where the confined and plasma phases are not separated by a real transition. Those quark masses where such a transition is to be expected to be second order are marked red in the diagram. Our project aims to calculate these lines, since the position of these are, at present, unknown in the continuum limit.

Thus, we make simulations on LRZ's SUPERMUC-NG system on a wide range of lattice resolutions and sizes. We determine the order of the transition by monitoring the fluctuations of the order parameter in a finite volume scaling analysis. Our first round of simulations focused on the top right corner of the phase diagram in Figure 1, where we identified a boundary line between the first order and crossover behavior.

Through the study of the phase structure of the strongly interacting matter in this enlarged parameter space, we will also learn something about QCD with Nature's choice of the quark masses. We'll be able to make connections to the phase structure of the strongly interacting matter in actual collider experiments. For this, we'll quantify the relation of the critical quark masses (red lines in Figure 1), and the quark densities in high energy accelerator experiments.

References and Links

- [1] Y. Aoki et al, Nature 443 (2006) 675-678.
- [2] S. Dürr et al, Science 322 (2008) 1224-1227.
- [3] F. R. Brown et al, Phys.Rev.Lett. 65 (1990) 2491-2494.

Simulating bound states of gluons and gluinos

RESEARCH INSTITUTION

¹Institute for Theoretical Physics, University of Jena

PRINCIPAL INVESTIGATOR

Georg Bergner¹

RESEARCHERS

Gernot Münster², Istvan Montvay³, Philipp Scior², Stefano Piemonte⁴, Sajid Ali², Juan Camilo Lopez¹

PROJECT PARTNERS

²University of Münster

³DESY Hamburg

⁴University of Regensburg

SuperMUC Project ID: pr27ja (Gauss Large Scale project), pr27re

Introduction

Quantum chromodynamics (QCD) is the theory that describes how the elementary quarks and gluons are bound by the strong force to form nuclear matter. The exchange particles of the strong force are gluon fields.

One particular feature of the strong interactions is a phenomenon called confinement: only bound states of the elementary fields can be observed at low energies due to the strength of the binding. This phenomenon occurs already in the pure gluonic part of QCD, called $SU(3)$ Yang-Mills theory (YM), since it includes a self interaction of the gluon fields. The low energy effective theory of $SU(3)$ YM is described by glueball states that are bound states of gluons.

There is so far no analytical understanding of the mechanism behind confinement. Therefore only large scale numerical simulations provide insights into the low energy properties of the theory.

We investigate a supersymmetric version of $SU(3)$ Yang-Mills theory. Supersymmetry relates fermionic and bosonic particles. Thus the theory requires fermionic partners of the gluon fields, called gluinos. These need to have similar properties like the gluons and are therefore described by Majorana fermions in the adjoint representation of the gauge group. The theory can thus be seen as a version of QCD with a different fermion content: the number of fermions is reduced and instead of the quark fields that transform in the fundamental representation, fermions in the adjoint representation are considered.

The investigations of supersymmetric theories have several motivations. The most prominent one is the relation to extensions of the Standard Model of particle physics. The Standard Model describes all known particles and their interactions, except gravity. It is a well-established theory, but it has several limitations. For example, a large amount of matter in the universe is currently not described by the Standard Model. Astronomical observations provide good indications for the existence of this additional dark matter. Furthermore, there are conceptual drawbacks like the Hierarchy problem. The search for an extension of the Standard Model towards a more complete description of nature has been going on for several decades and supersymmetric theories are one of the most prominent candidates for such extensions. Figure 1 shows the basic idea of supersymmetric extensions. The bosonic fields are complemented by fermionic partners and fermion fields get bosonic counterparts. Since the partner particles have so far not been observed in experiments, there must be a non-perturbative breaking mechanism for supersymmetry at low energies. This is one motivation for studying supersymmetric theories with non-perturbative methods such as numerical lattice simulations.

The second recent motivation for our investigations of $SU(3)$ supersymmetric Yang-Mills theory (SYM) are theoretical considerations of supersymmetric gauge theories. These theories allow, due to the extended symmetry, for analytic approaches to describe the confined phase and a possible general understanding of the confinement mechanism.

One very interesting question that can be answered only by non-perturbative methods is how the super-

symmetry of the fundamental theory is reflected in the bound states at low energies. It has been conjectured that the bound states form certain chiral supersymmetry multiplets composed out of a fermion together with a scalar and pseudoscalar bosonic parti-

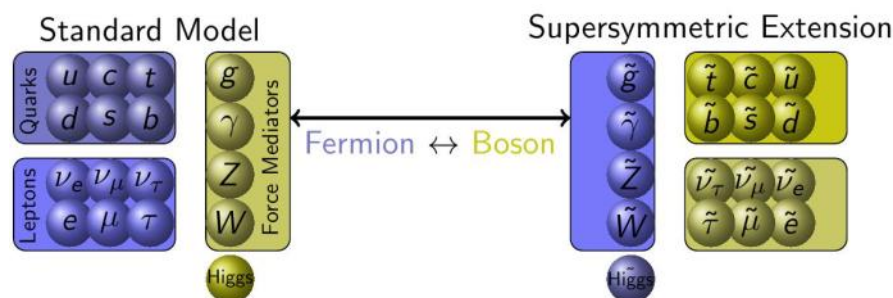


Figure 1: Simplified picture of the supersymmetric extension of the standard model: the bosonic particles get additional fermionic partners and vice versa. In the standard model the interactions between the quarks and leptons are mediated by the gauge bosons: the strong interactions by the gluons, the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




Figure 2: The measurement of lightest and first excited states in $SU(2)$ SYM. The expected multiplet structure suggests scalar (0^{++}), pseudoscalar ($a-\eta'$, gb^-), and fermionic (gg) particles with the same mass. The ground states are marked by (0). We found that the (gb^-) has better overlap with the first excited state. The results are presented in terms of the gradient flow scale w_0^3 . Further details are presented in [3].

cle [1,2]. These particles are supposed to be glueballs, gluino-balls, and gluino-gluon-balls. Gluino-balls are meson-like particles and the gluino-gluon-balls are exotic particles that have no similar counterpart in QCD. The formation of these multiplets can only be investigated based on numerical lattice simulations.

Results and Methods

Numerical lattice simulations

The method of numerical lattice simulations is based on a discretized version of the continuum theory. The discretization breaks supersymmetry since it is tied to the symmetries of space-time. In a classical theory, the recovery of the symmetry is straightforward when going to finer and finer discretizations corresponding to the continuum limit. In a quantum field theory, this is in general not guaranteed and requires detailed theoretical considerations of operator mixing. In SYM this can be solved by a specific tuning of the gluino mass. This implies that two limiting procedures are required: the extrapolation of the gluino mass to the so-called chiral limit and the continuum limit.

Besides these theoretical complications, there are also technical challenges in the simulations of SYM. The bound states provide quite noisy signals and we had to develop our own tools and simulation package for this theory. Despite these challenges, we were able to obtain the first final results extrapolated to the chiral continuum limit of the bound state particle spectrum in $SU(3)$ SYM.

Determinations of bound states

The bound states of the theory are defined by their quantum numbers, their transformation properties under the basic symmetries of the theory. How they are composed from the elementary fields and what kind of effective size they have is determined by the non-perturbative dynamics. In order to get a good signal for the masses, the lattice operators are required to have a reasonable overlap with the considered state of the theory. This can be achieved by a numerical optimization of the operators. We have tested our methods with an investigation of the bound states in $SU(2)$ SYM. As shown in Figure 2, we have been able to extract even data for the first excited state in each channel.




Figure 3: The final results of our simulations of $SU(3)$ SYM extrapolated to the chiral continuum limit. In this case a quadratic form of the corrections by the finite lattice spacing (a) is assumed. The data show the formation of a multiplet of gluino-gluon (gg), scalar (0^{++}), and pseudoscalar ($a-\eta'$) particle. Further details can be found in [4].

After we have established our methods for the bound state determination, we were able to approach our final goal, the determination of the bound states in $SU(3)$ SYM. The simulations are more expensive than for $SU(2)$ SYM or QCD. Therefore we have optimized our simulation strategy using a one-loop improved fermion operator. This allowed reliable simulations already with rather small lattice sizes. The final results extrapolated to the combined chiral and continuum limit are shown in Figure 3.

Our results have thus confirmed for the first time the multiplet formation in the low energy effective theory of $SU(3)$ SYM. This important observation also shows that the lattice simulation of the theory and reliable extrapolations towards the supersymmetric continuum limit are possible.

Ongoing Research / Outlook

Based on these findings, we are now able to investigate further properties of the theory. In particular interesting are the phase transitions and the formation of a gluino condensate at low temperatures, since there have been several theoretical studies about this subject.

Furthermore, we are currently starting to explore the simulations of supersymmetric QCD, which requires the addition of quark fields and their partners, the bosonic squarks. In addition, we are investigating better implementations of the symmetries of the theory on the lattice using overlap fermions [5].

References and Links

- [1] G. Veneziano and S. Yankielowicz, Phys. Lett. B 113 (1982) 231.
- [2] G. R. Farrar, G. Gabadadze and M. Schwetz, Phys. Rev. D 58 (1998) 015009;
- [3] S. Ali, G. Bergner, H. Gerber, S. Kuberski, I. Montvay, G. Münster, S. Piemonte, and P. Scior, JHEP 1904, 150 (2019); arXiv:1901.02416 [hep-lat].
- [3] S. Ali, G. Bergner, H. Gerber, I. Montvay, G. Münster, S. Piemonte, and P. Scior, Phys. Rev. Lett. 122, no. 22, 221601 (2019); arXiv:1902.11127 [hep-lat].
- [5] S. Piemonte, G. Bergner and C. López, Phys. Rev. D 102, no.1, 014503 (2020); arXiv:2005.02236 [hep-lat].

A strong side of weak decays: How beauty and charm resonate in pions and kaons

RESEARCH INSTITUTION

¹Helmholtz-Institut für Strahlen- und Kernphysik, Rheinische Friedrichs-Wilhelms-Universität Bonn

PRINCIPAL INVESTIGATOR

Marcus Petschlies¹

RESEARCHERS

Constantia Alexandrou², Stefan Krieg³, Luka Leskovec⁴, Thomas Lippert³, Stefan Meinel⁵, John Negele⁶, Srijit Paul⁸, Andrew Pochinsky⁶, Gumaro Rendon⁷, Giorgio Silvi³

PROJECT PARTNERS

²Cyprus Institute & University of Cyprus, Cyprus, ³Forschungszentrum Jülich & Bergische Universität Wuppertal

⁴Thomas Jefferson National Accelerator Facility & Old Dominion University, USA, ⁵University of Arizona, USA

⁶Center for Theoretical Physics, Massachusetts Institute of Technology, USA, ⁷Brookhaven National Laboratory, USA

⁸Johannes Gutenberg-Universität Mainz

SuperMUC Project ID: pr27yo (Gauss Large Scale project)

Introduction

The transition amplitudes for B-meson (and D-meson) decays $B \rightarrow \pi\pi$ and $B \rightarrow K\pi$ determined in pure QCD are an essential ingredient to the investigation of semileptonic B decays in the effective electroweak sector of the Standard Model. These flavor changing processes, involving $b \rightarrow u \ell \bar{\nu}$ and $b \rightarrow s \ell^+ \ell^-$, are gold-plated candidates at the high-precision frontier in the search for physics beyond the Standard Model. The rigorous calculation of such matrix elements must take into account the interaction due to elastic scattering of the final $\pi\pi$ and $K\pi$ 2-hadron state. In particular, enhancements for decays via an intermediate resonant excitation in the elastic regime, such as $B \rightarrow (\rho \rightarrow \pi\pi)$ and $B \rightarrow (K^* \rightarrow K\pi)$, require a fully non-perturbative treatment. These are examples of so-called $1 \rightarrow 2$ transitions and their conversion from lattice to infinite volume amplitudes requires the technique originally introduced in [1], and recently generalized in [2].

Results and Method

The major steps in our lattice QCD calculation are spectroscopy and mapping to scattering amplitudes (1) for $\pi\pi$ isospin $I = 1$ with focus on P-wave (ρ resonance) channel and (2) for $K\pi$ $I = \frac{1}{2}$ with focus on S-wave (κ resonance) and P-wave (K^* resonance). Step (3) is the production of 3-point functions for B, D-meson and $\pi\pi$ and $K\pi$ initial and final states and mapping of extracted matrix elements to infinite-volume amplitudes, based on items (1) and (2). The required Wick contractions to construct 2-point correlation functions and matrices thereof for variational analysis are calculated based on SuperMUC-NG partitions with 1,536 (C13) and 2,592 (D6) cores and approximately 25,000 core-hours per job, and using the C/C++ QLUA software suite [3] and QOPQDP-multigrid inverter. With SuperMUC Phase 2 and

SuperMUC-NG based data, we analyzed $K\pi$ elastic scattering in $(I)J^P = (\frac{1}{2})1^-$ and $(I)J^P = (\frac{1}{2})0^+$ channels on ensembles C13 and D6 with parameters given in Table 1.

Label	L (fm)	a (fm)	m_π (MeV)	m_K (MeV)
C13	3.6	0.114	≈ 320	≈ 530
D6	4.2	0.088	≈ 180	≈ 515

Table 1: Parameters for C13 and D6 simulations.

We fit energy levels predicted by the Lüscher quantization condition to our observed spectrum. From combined analysis for S- and P-wave with various K-matrix parametrizations follow the S- and P-wave phase shifts in Figure 1, see also [4].

To characterize the κ resonance, we extract the complex pole locations from several K-matrix parametrizations. We show current determinations of the κ pole as given by PDG [5] in Figure 2, together with our pole values for Bugg's and conformal parametrization (violet symbols). We can check the onset of tracking the pole location with pion mass: from 320 MeV to 180 MeV. Our results—though still at larger than physical pion mass—become consistent with the PDG average (black clock symbol).

Ongoing Research / Outlook

With $\pi\pi$ and $K\pi$ scattering amplitude information in hand, we simulate the beauty and charm meson 3-point functions for the transitions $B, D \rightarrow \pi\pi, K\pi$ with analysis parallel to the $\pi\gamma \rightarrow \pi\pi$ case in [6]. With the resources of SuperMUC-NG, we can continue to probe the effect of coupled decay channels ($\bar{K}K$ for $\pi\pi$, $K\eta$ for $K\pi$) and investigate systematic uncertainties (non-zero lattice spacing, finite volume, heavier-than-physical up and down quark) to extrapolate to the physical world.




Figure 1: From combined analysis for S - and P -wave with various K -matrix parametrizations follow the S - and P -wave phase shifts. See also [4].




Figure 2: Current determinations of the κ pole as given by PDG [5], together with our pole values for Bugg's and conformal parametrization (violet symbols). We can check the onset of tracking the pole location with pion mass: from 320 MeV to 180 MeV. Our results—though still at larger than physical pion mass—become consistent with the PDG average (black clock symbol).

References and Links

- [1] L. Lellouch and M. Luscher, Commun. Math. Phys. 219 (2001) 31.
- [2] R. A. Briceño, M. T. Hansen and A. Walker-Loud, Phys. Rev. D 91 (2015) no.3, 034501.
- [3] USQCD software Qlua package, [https:// usqcd.lns.mit.edu/w/index.php/QLUA](https://usqcd.lns.mit.edu/w/index.php/QLUA).
- [4] G. Rendon, L. Leskovec, S. Meinel, J. Negele, S. Paul, M. Petschlies, A. Pochinsky, G. Silvi and S. Syritsyn, [arXiv:2006.14035 [hep-lat]].
- [5] M. Tanabashi et al. [Particle Data Group], Phys. Rev. D 98, no.3, 030001 (2018).
- [6] C. Alexandrou, L. Leskovec, S. Meinel, J. Negele, S. Paul, M. Petschlies, A. Pochinsky, G. Rendon and S. Syritsyn, Phys. Rev. D 98 (2018) no.7, 074502.

$N=1$ Supersymmetric gauge theories on the lattice

RESEARCH INSTITUTION

¹Theoretisch-Physikalisches Institut, Friedrich-Schiller-Universität Jena

PRINCIPAL INVESTIGATOR

Andre Sternbeck

RESEARCHERS

Marc Steinhauser, Andre Sternbeck, Björn Wellegehausen, Andreas Wipf

PROJECT PARTNERS

—

SuperMUC Project ID: pr48ji

5

Introduction

The standard model (SM) of particle physics is experimentally well established. It very successfully describes all processes mediated by the electromagnetic, weak and strong forces. Nonetheless, several open questions remain unanswered and the SM is thus seen as an effective theory of a more fundamental theory. For example, the SM does not contain gravity and cannot explain the large amount of dark matter seen in our universe. Furthermore, the Higgs boson is unreasonably light; a boson mass of the order of the Planck mass would be expected.

Supersymmetric extensions are an interesting step beyond the standard model. They provide an unification of gauge couplings at the GUT (Grand Unified Theory) scale and can explain why the typical energy scales of particle physics are much smaller than the Planck scale. Supersymmetric models also naturally provide a dark-matter candidate.

A straightforward extension of the SM is the minimal supersymmetric standard model (MSSM). This project addresses the strongly coupled subsector of this model, also known as $N=1$ Super-Yang-Mills (SYM) theory. It describes interactions of gluons and their fermionic superpartners, the massless gluinos, and allows for a numerical treatment using lattice Monte Carlo simulations, similar as for lattice regulations of the strong interaction.

This non-perturbative treatment, however, is non-trivial. Lattice formulations explicitly break supersymmetry and consequently the mass degeneration within a super-multiplet at any finite lattice spacing. We start with a lattice formulation introduced by Curci and Veneziano [1] which is based on Wilson-type Majorana fermions. The Wilson term explicitly breaks supersymmetry and chiral symmetry at finite lattice spacing, but this breaking leads to a counter-term which is proportional to the gluino mass term. By adding an explicit gluino mass term this can be compensated, such that the renormalized gluino becomes massless in the continuum limit. Since the gluino mass term is the only relevant operator, both supersymmetry and chiral symmetry will be restored in the continuum limit.

This approach is straightforward, but it requires a careful treatment of the continuum limit and implies studying a wide range of lattice sizes and spacings. This is numerically demanding and any improvement of the lattice regularization, as well as of the numerical methods, are thus highly welcome.

We introduce two novel concepts [2]: Firstly, we modify the original lattice formulation by adding a mass-like term to the fermionic part of the action, similar to an one-flavor formulation of twisted-mass QCD. By tuning the two mass parameters of the lattice action, a considerably improved mass degeneracy of the chiral partners in the Veneziano-Yankielowicz supermultiplet can be achieved, already at finite lattice spacing. Secondly, we apply an adaptive aggregation-based domain decomposition multigrid (DD α AMG) algorithm [3] for the calculation of two-point correlator functions. This reduces the amount of CPU time for the numerous inversions considerably.

Results and Methods

For our study we perform lattice Monte Carlo simulations of $N=1$ Super-Yang-Mills (SYM) theory and estimate bound-state masses of the different supermultiplet partners. To this end, we calculate two-point functions on gauge ensembles, which are thermalized with respect to the Lüscher-Weisz gauge action plus a fermionic part with the twisted Wilson-Dirac operator. For the gauge coupling three values are chosen and for each coupling a range of gluino mass values for the extrapolation to the critical point (at vanishing gluino mass). In addition the twist mass parameter is varied to study the effect of the additional twist term on the bound-state correlators.

The lattice calculations are numerical demanding, because of the adjoint representation of the gluino, but a lattice size of either $8^3 \times 16$ or $16^3 \times 32$ is possible. Most bound-state correlators have connected and disconnected contributions. In particular the latter require large gauge ensembles to reduce the statistical noise. But also connected contributions of correlator functions with gauge fields, like for the gluino-gluon, are afflicted by large statistical uncertainties.




Figure 1: Mass estimates for the adjoint- η' and the adjoint- f_1 (connected contributions only) for different gluino masses. For a twist angle of 45° degree (middle) both states have equal mass as in the supersymmetric limit.

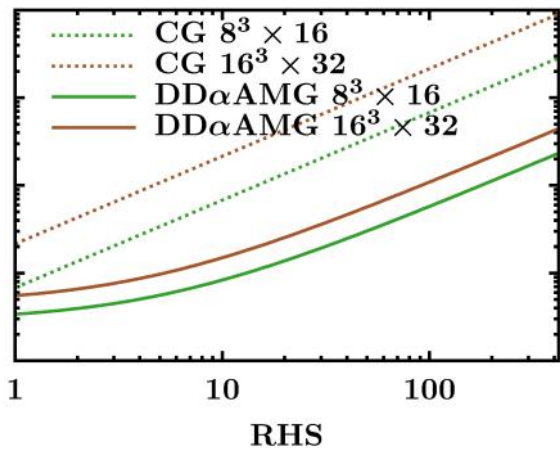


Figure 2: Time to solution vs. the numbers of right hand sides (RHS) of the linear system. The figure compares the timing for a standard CG solver with the DDαAMG solver on lattices of size $8^3 \times 16$ and $16^3 \times 32$.

Performing calculations for different gauge couplings and gluino masses we find that the mass splitting of the chiral partners disappears at a twist angle of 45° degrees (see Figure 1). We can reproduce this for different gauge couplings and gluino masses. Furthermore, the superpartner, the so-called gluino-gluon, has a similar mass at finite lattice spacing. Our findings promise that both chiral and supersymmetry can be improved when using a twisted Wilson-Dirac operator at a special twist and an improved continuum extrapolation may be possible.

We also look at numerical aspects and analyze the performance gain due to the DDαAMG algorithm [3]. This algorithm is based on two ingredients: (1) a Schwarz alternating procedure (SAP) utilizes domain decomposition and deals with the UV-modes. (2) A coarse grid correction acts as an interpolation operator which approximates the small eigenvalues to tackle the IR-modes.

This algorithm has been successfully employed in lattice QCD studies and similar performance gains may be possible for a $N=1$ SYM theory, if the low-lying eigenvalues of the Dirac operator show a similar local coherence as in QCD.

To study the performance we use the DDαAMG library [4] but we had to generalize the hard-coded SU(3) gauge group to SU(Nc) with arbitrary Nc and representation. For the benchmark we choose the following setup: gauge group SU(3) in the adjoint representation, lattice sizes $8^3 \times 16$ and $16^3 \times 32$, two multigrid levels, block size 2^4 , mixed precision and solver combination FGMRES + red-black Schwarz. Figure 2 shows the timings for up to 100 stochastic estimators and 5 point sources and a comparison to the conjugate gradient (CG) algorithm, which typically has been the standard choice in the past. For these scenarios we achieve a speed-up factor of 9 to 20, which is an enormous speed up. Consequently, we use this algorithm now for the calculation of two-point functions of the adjoint- η' and adjoint- f_1 , where many stochastic estimators are required. It has much reduced the amount of required CPU time.

Ongoing Research / Outlook

The results are available as preprint [5] and have been accepted for publication in JHEP. The character of our study remains exploratory, however. With the granted CPU time, we were able to analyze the benefits of using a twisted-mass Dirac operator for lattice studies of $N=1$ SYM theory and to demonstrate the performance gains when using the DDαAMG algorithm. For a reasonable estimate of the associated supermultiplets, however, a follow-up study is required. In particular, larger volumes and ensemble sizes are needed to reduce systematical and statistical errors.

References and Links

[1] Curci and Veneziano, Nucl. Phys. B292 (1987) 555.
 [2] Steinhauser et al., PoS LATTICE2019 (2019) 201 [eprint:1912.09979]; PoS LATTICE2018 (2018) 211 [eprint:1811.01785]; EPJ Web Conf. 175 (2018) 08022 [eprint: 1711.05086].
 [3] Alexandrou et al, Phy.Rev. D94 (2016) 114509.
 [4] Bacchio et al., DDalphaAMG library including twisted mass fermions, github.com/sbacchio/DdalphaAMG
 [5] Steinhauser et al., JHEP 01 (2021) 154 [eprint: 2010.00946].

Simulating Interactions for the ATLAS experiment at the LHC

RESEARCH INSTITUTION

¹Physics Faculty, LMU Munich

PRINCIPAL INVESTIGATOR

Günter Duckeck¹

RESEARCHERS

Rodney Walker¹, ATLAS Collaboration²

PROJECT PARTNERS

²CERN, Geneva

³MPP 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-NG. From 2015 to 2018 LHC was running in its 2nd phase (LHC Run-2) at a centre-of-mass energy of 13 TeV. The LHC operation was very stable and had an excellent efficiency close to or even beyond the original design intensity. The recorded data volume for Run-2 exceeds the Run-1 (2010-2012) volume by a factor 6.

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. The ATLAS collaboration has already published more than 900 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.

Moreover, extensive studies for detector upgrades in preparation for the future Run-3 (2022-24) and the high-luminosity Run (planned to start in 2027) are on-




Figure 1: Example of a simulated event in an upgrade study, comprising 40 signal muons overlaid with minimum bias events consistent with an average pile-up of 140 collisions per bunch crossing.

going. Figure 1 shows an example of a simulated event under extreme pile-up conditions where in a single LHC bunch-crossing about 140 simultaneous collisions take place. To record and disentangle such a high number of collisions is a complex problem and requires careful studies and preparation.

Such simulations are computationally expensive. The CPU time per event depends strongly on the complexity of the event and ranges from few 100 seconds in simple cases up to ~hours in case of complex, high-pileup settings. Both, for detailed analysis of the data and upgrade studies large simulated samples are required matching the about 10 billion recorded events.

The main focus of the future LHC program are detailed searches for 'New Physics' processes, i.e. signatures from particles or interactions predicted in theoretical models beyond the Standard Model of particle physics. Such searches require additional samples of simulated events for these processes, typically for multiple settings of parameters specific for these models. Also the absence of new signal events provide important information, they can be used to set stringent constraints on these theories. One recent example of an exclusion plot for particles predicted in Super-Symmetric Models is shown in Figure 2 [5].

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 sufficiently large simulated samples. Therefore using CPU resources at HPC systems such as SuperMUC-NG 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, automatized submis-




Figure 2: The plot above shows the observed and expected exclusion contours at 95 % confidence level from different searches for top squarks. The search for top squarks using decays to one lepton is shown in orange [5].

sion 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-NG login nodes. Submission into the batch system and subsequent monitoring proceeds using commands run via ssh. The workloads we are running on HPC systems are a well-defined subset of ATLAS central production workflows, in particular detector simulations based on Geant4 [3]. Geant4 is a toolkit for the simulation of the passage of particles through matter. It is CPU-limited and dominated by integer arithmetic operations. Simulating the passage of particles through matter is serial by nature. However, one can naturally parallelize the processing on the level of individual proton-proton collisions. The simulation process requires large memory, in particular for the detailed mapping of the complex detector structure which exceeds the memory capacity per core for independent processes. ATLAS developed a means to usefully use multiple CPU cores: after an initialization step, the process forks to N sub-processes using 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 large many-core nodes as on SuperMUC-NG.

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. If a job is killed by pre-emption, then the events already produced are merged and stored by the ARC CE, thus only the events in-flight are lost. No memory dump of the processes is performed. On restart, the simulation just continues with the next event. This pre-emptable workload is now used exclusively.

This opportunistic use works well and is an important contribution to ATLAS simulations.

An important improvement for ATLAS operation is the option to deploy the software environment in a large container image at SuperMUC-NG .

Another new feature is the use of GlobusOnline [4] to stage in and out the data sets needed and produced by the processing. Further improvements could be made on the submission of multi-node jobs and the backfill submission.

Ongoing Research / Outlook

The LHC Run-3 is planned to start in 2022, and we expect an increase of the data volume by about a factor 3 compared to Run-2. During the current 3-year break of LHC operation, the analysis of the recorded Run-2 data is actively ongoing. At the same time, improvements and upgrades are done for many components of the ATLAS detector. There is a continuous demand for event simulation, both to provide sufficient samples for ongoing Run-2 analyses, and new simulation samples, adapted for the detector upgrades in Run-3. We would be grateful if SuperMUC-NG would continue to contribute to this effort.

The option to use containers at SuperMUC-NG facilitates enormously the deployment of the software and configuration and opens up new prospects to run additional workflows such as reconstruction and complex analyses at SuperMUC-NG and other high performance computing centres, which are actively developed and tested.

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/document.
- [3] Geant4, Nuclear Instruments and Methods in Physics Research A 506 (2003) 250-303.
- [4] GlobusOnline, <https://www.globusonline.org/>
- [5] ATLAS collaboration, Searching Natural Supersymmetry.

Decoding the spin of the proton

from quarks and gluons

RESEARCH INSTITUTION

¹NIC, DESY, Zeuthen

PRINCIPAL INVESTIGATOR

Karl Jansen¹

RESEARCHERS

C. Alexandrou², S. Bacchio², K. Cichy³, M. Constantinou⁴, J. Finkenrath², K. Hadjiyiannakou², G. Koutsou², B. Kostrzewa⁵, M. Petschlies⁵, A. Scapellato², F. Steffens⁵, C. Urbach⁵

PROJECT PARTNERS

²Cyprus Institute, Cyprus, ³Faculty of Physics, Adam Mickiewicz University, Poland

⁴Physics Department, Temple University, Philadelphia, USA

⁵Helmholtz-Institut für Strahlen- und Kernphysik and Bethe Center for Theoretical Physics, Universität Bonn

SuperMUC Project ID: pr74yo (Gauss Large Scale project)

Introduction: the spin of nucleons

The proton and the neutron form the nucleus of all atoms we have discovered so far. Our picture of the inner most structure of these two *nucleons* is that they are bound states of partons, the quarks and the gluons, see Figure 1. The proton takes a spin of $\frac{1}{2}$, which must be composed from the individual spins of its constituents. Quark models predicted that almost all of the proton's spin is actually carried by the quarks. It came therefore as a very big surprise that an experiment, the European Muon Collaboration (EMC), found that only half of the spin of the proton originates from the quarks [2]. This finding was named the *spin crisis*, and attempts to explain this puzzle led to an enormous amount of work.

The great difficulty to compute the spin of the proton directly from our theory of the strong interaction between quarks and gluons, quantum chromodynamics, is that this is a fully non-perturbative question. A principle way out is to formulate QCD on an Euclidean space time grid, so-called lattice QCD, which allows to




Figure 1: The inner structure of the proton, the spiral lines denote gluons which bind the quarks, the blue and green balls, together in the proton. Taken from the DESY figure database.

perform numerical simulations of the theory on supercomputers. However, the computational demand of such simulations turned out to be so large that only with the advent of recent supercomputer architectures and major developments of the employed algorithms realistic simulations became possible. In this project [3], we were indeed able to carry out lattice QCD calculations in fully physical conditions and provided an ab-initio computation of the individual contributions of the quarks as well as the gluons. In particular, we find that the gluons indeed carry a surprisingly large amount of the proton spin of about 40%. This result is consistent with the above mentioned experimental result which found that a large fraction of the spin must originate from the gluons. Our work is therefore a major step forward to solve the very long-standing puzzle of the proton spin.

Results and Methods

The simulations of lattice QCD are based on Markov chain Monte Carlo methods. In particular, we have used a Hybrid Monte Carlo algorithm which combines a Metropolis accept/reject step with a molecular dynamics time evolution of the gluon fields. This algorithm is used to generate gluon field configurations which are stored and then used to compute physical observables. Typical numbers of so generated gluon field configurations are 2,000–5,000. The lattice size we employ are $64^3 \times 128$ and $80^3 \times 160$ lattice points, and since the internal degrees of freedom of the quark and gluon field is 12, this amounts to a size of a single gluon field configuration of 19 and 44 GBytes, respectively.

A basic ingredient to compute a physical observable are so-called quark propagators from which e.g. the proton mass or, as in this project, the individual contributions of quarks and gluons to the proton spin can be calculated. The quark propagators are computed by solving a very large linear system of equations with a coefficient matrix of 40 million time 40 million for our smaller lattice size. On first sight, this sounds imprac-




Figure 2: The dependence of the lattice extent L on the cost for generating a new gauge configuration is shown. Due to our algorithm improvements we are now able to simulate physical light, strange and charm quark masses.

tical, but what helps is that this coefficient matrix is sparse with only the diagonal and a few sub-diagonals being filled. This allows to hard code the required matrix time vector multiplications as required in linear solver algorithms, e.g. conjugate gradient. In fact, in the early days of lattice QCD, the conjugate gradient algorithm has been the method of choice. However, since then substantial improvements could be achieved and nowadays algorithms, partly developed by us [4], based on multi grid methods are used which led to orders of magnitude reduced computational cost. It is only through these most significant algorithmic improvements in combination with new, ever improved supercomputer architectures, which led to the success of this project.

The HMC simulations performed for the larger volume were performed on 125 nodes using a hybrid parallelization employing 16 MPI tasks with 3 OpenMP threads per node. By using the above mentioned highly optimized numerical algorithms like algebraic multigrid methods and 4th order force gradient integrators for the molecular dynamics of the HMC algorithm, the time per new gauge configuration could be reduced to 3 hours and 15 mins, resulting in 19,500 core-hours per trajectory. This requires a total of 60 M core-hours for one physical point ensemble, producing 3,000 gauge configurations of 132 TB in two parallel streams written during 200 days on the WORK partition of SuperMUC-NG. For the scaling of our algorithm, see Figure 2.

In a dedicated effort, which is embedded in the research programme of the Extended Twisted Mass Collaboration (ETMC), we were able to disentangle the various contributions to the proton spin. Since nature has decided—for an unknown reason—that besides the light up and down quarks also heavier strange and charm quarks exist, we included also the contributions of these quarks for the proton spin in order to obtain a comprehensive picture. The results can be seen in Figure 3, where we show the absolute size of the individual contributions as well as the percentage the different quarks and gluons contribute. Our finding that the effect of the gluons is rather large and that all contributions add up to the expected value of $\frac{1}{2}$ provides a nice confirmation of the experimental results mentioned above. The proton spin puzzle is going to be put together and we start to see the picture it reveals.




Figure 3: The decomposition of the proton spin J . The dashed horizontal line indicates the observed proton spin value and the percentage is given relative to the total proton spin. We show the average of the quark and anti-quark contributions of the up (u), the down (d) the strange (s), and the charm (c) quarks as well as their sum and the gluon (g). All contributions add nicely up to the spin of $\frac{1}{2}$ of the proton, which is a major result towards the resolution of a long-standing puzzle. Figure taken from ref.[3].

Besides the spin, the proton also has an average momentum with a value of 1. In the same setup as was used for the proton spin we could also disentangle the various contributions of the quarks and the gluons to the proton average momentum, and we could nicely see that the different contributions add up to the value of 1, as observed in nature.

Ongoing Research / Outlook




In this project [1] we could address a long-standing puzzle found already by the EMC collaboration [2], namely, why is the contribution of quarks of the proton spin so small, while quark models predict a much larger value? Addressing this problem through a dedicated effort and using large scale simulations we could demonstrate that it is the gluon contribution which brings the total spin of all constituents to the value observed in nature to $\frac{1}{2}$.

Although our work [3] is clearly a major step to completely solve the long-standing puzzle of the proton spin, for claiming victory still other steps have to be taken. The most relevant of these is to take the continuum limit. In lattice QCD, we use a finite 4-dimensional grid with a non-zero lattice spacing, denoted by the letter a , to perform our non-perturbative simulations. However, what is really needed are results in the continuum limit, meaning that the lattice spacing is sent to zero, while keeping the physical size of the system fixed. This means that we have to perform further simulations at smaller values of the lattice spacing on correspondingly larger lattices. Hence, we plan simulations on lattices of size $96^3 \times 192$ or even larger. For this, even more powerful computational resources than the presently available are needed. Nevertheless, we consider our work a most important and major step forwards towards solving the proton spin puzzle.

References and Links

- [1] <https://www.lrz.de/projekte/hlr/projects/000000000F4397B.html>
- [2] J. Ashman et al. (European Muon), Phys. Lett. B206, 364 (1988), 340(1987).
- [3] C. Alexandrou et al., Phys.Rev.D 101 (2020) 9, 094513.
- [4] C. Alexandrou et al., Phys.Rev.D 94 (2016) 11, 114509.

Life Sciences



Structure and Dynamics on the Photosynthetic

Complex I

RESEARCH INSTITUTION

Department of Chemistry, Technical University of Munich

PRINCIPAL INVESTIGATOR

Ville R. I. Kaila

RESEARCHERS

Patricia Saura, Ana P. Gamiz-Hernandez, Max E. Mühlbauer

PROJECT PARTNERS

—

SuperMUC Project ID: pn34he

6

Introduction

Energy conversion in nature is catalyzed by membrane-bound proteins that transform light or chemical energy into an electrochemical gradient stored across the membrane [1]. These processes take place in the photosynthetic and respiratory chains, but in cyanobacteria, they coexist within the same cellular compartment. To this end, cyanobacteria express the so-called photosynthetic complex I or NDH-1 complex, a membrane-bound protein machinery that has evolved special adaptation to participate in both energy-conversion processes. This 0.5 MDa redox-driven proton pump catalyzes the electron transfer from reduced ferredoxin (Fd) that originates from the photosynthetic reactions in photosystem I, to plastoquinone (PQ) and employs the released free energy to pump protons across the thylakoid membrane [2]. Four different isoforms of NDH-1 in the cyanobacterial cells are known: the NDH-1L/L' complexes employ the redox energy to pump protons and to increase ATP levels in the cell, whereas the NDH-1MS/MS' isoforms are utilized for inorganic carbon concentration from CO₂, that is further transformed into organic molecules. This latter process takes place in the CupA/S subunits of NDH-1MS (Fig. 1b) with carbonic anhydrase (CA) activity. To gain insight into how the CO₂ concentration process takes place in the photosynthetic complex I, we performed both quantum chemical calculations and large scale classical atomistic simulations in combination with cryo-electron microscopy (cryo-EM) and biochemical experiments to elucidate the molecular principles of this new type of enzyme [3]. The work was done in collaboration with Jan Schuller from Max Planck Institute of Biochemistry, and Marc Nowaczyk from Ruhr-University Bochum, Germany.

Results and Methods

To understand the molecular principles of the carbon concentration process catalyzed by the photosynthetic complex I, we resolved the molecular structure of NDH-1MS by cryo-EM, and investigated its function by large scale classical and quantum simulations.

The enzyme shows a global U-shape architecture, with the soluble electron transfer and CO₂ concentrating modules located on the same sides of the membrane domain (Figure 1a). The CupA/S CO₂ concentrating module presents a unique fold comprising mainly α -helices that drastically differs from canonical CAs. The catalytic site contains a Zn²⁺ ion with an exotic first-coordination sphere formed by His130/Arg135, and a OH⁻/H₂O ligand (Figure 1b). Our calculations suggested that CupA/S can catalyze the conversion of CO₂ into HCO₃⁻ with similar energetics as the canonical CAs [3].

The CupA/S module binds on the top of the NdhF3 subunit of the membrane domain. In this subunit, we identified a non-polar cavity that could channel CO₂ into the CupA active site (Figure 1c) and could further increase the catalytic efficiency of the carbon concentration process. Our MD simulations show that CO₂ can diffuse along this cavity (Figure 1c) [3].

The NdhF3 subunit is connected to the adjacent membrane-subunit by buried charged residues. Moreover, a central axis of buried ion-pairs and charged/polar residues extends along the membrane domain from the PQ binding site to the NdhF3 interface, that could transmit the energy signal from the electron transfer process to power the CO₂ concentration process. These findings support how cyanobacteria has evolved highly efficient carbon concentration mechanisms to survive under low CO₂ conditions, by unique modular adaptations of a fundamental protein machinery involved in photosynthesis and respiration.

The MD simulations were built based on our resolved cryo-EM structure. Missing loops were modelled and we obtained a relaxed structure by using molecular dynamics flexible fitting (MDFF).

To probe the dynamics and mechanism of the NDH-1MS complex we performed atomistic MD simulations of the whole cryo-EM structure embedded in a lipid membrane, and solvated the system with water/ions to simulate its biological environment. The MD simulations were performed at T = 310 K and p = 1 bar within

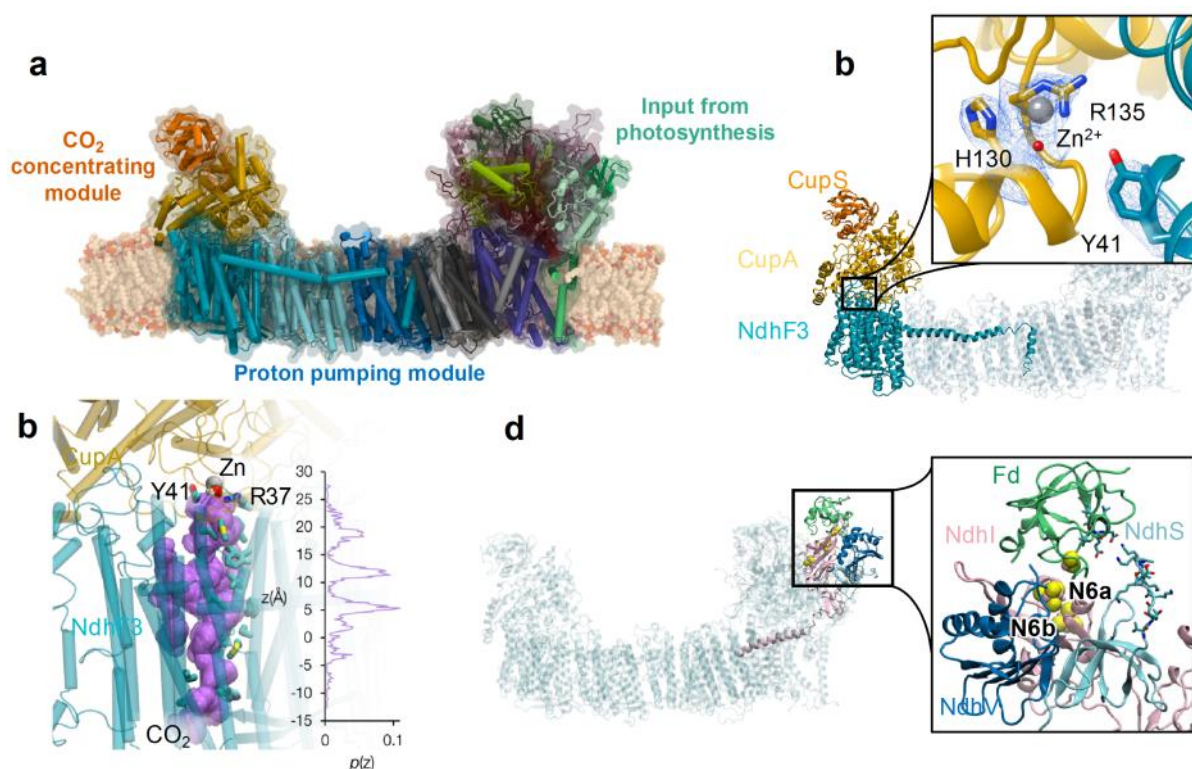


Figure 1: a) Structure of the NDH-1MS complex showing the soluble domains that catalyze electron transfer (right) and carbon concentration (left) processes, and the membrane domain that is involved in proton pumping. b) Structure of the carbon-concentration module, formed by the protein subunits CupA/S that bind to NdhF3. The catalytic site contains Zn^{2+} and presents carbonic-anhydrase activity. c) A non-polar cavity located in the NdhF3 subunit could channel CO_2 into the CupA/S module to increase the catalytic efficiency of the carbon concentration process. d) The electron donor, ferredoxin (Fd) binds at the electron-transfer module, at the interface of NdhS/IV subunits. Figure adapted from reference [3].

an NPT ensemble using an integration time step of 2 fs. Periodic boundary conditions (PBC) were employed, and long-range electrostatic interactions were treated by the particle mesh Ewald (PME) approach. All the MD simulations were performed with the NAMD program package [4]. To study the catalytic mechanism of the CA- module, we employed quantum chemical DFT and hybrid quantum/classical QM/MM calculations. The DFT calculations were performed using the TURBOMOLE program, and the QM/MM calculations were performed using a TURBOMOLE/CHARMM python interface. The MD simulations comprise 580,000 atoms and they are performed using ca. 900 cores per replica. The work resulted in a publication in Nature Communications [3].

Ongoing Research / Outlook

The current SuperMUC project provided us with computational resources to study the molecular basis of the CO_2 concentration mechanism by the photosynthetic complex I that is a fundamental process in biology. The on-going work is focused on understanding how electrons arrive to the photosynthetic complex I. To this end, we resolved the molecular structure of NDH-1MS in complex with its electron donor, ferredoxin (Fd). To probe the dynamics of the electron-transfer complex, we perform large-scale atomistic MD simulations using the same setup as before. Our preliminary data suggest that Fd could bind at the top of the electron-transfer module, in a binding site formed at the interface of the NdhS, NdhV and NdhI subunits

(Figure 1d). The reduced Fd is further stabilized by specific interactions with loops that seem to be sensitive to the redox state of the system. Our future work will also investigate how the electron transfer is coupled to the proton pumping in the membrane domain, a process also relevant for the respiratory complex I.

References and Links

- [1] <https://villekaila.com/>
- [2] Peltier, G., Aro, E.-M., Shikanai, T., Annu. Rev. Plant Biol. 2016, 67, 55–80.
- [3] Schuller, J. M., Saura, P., Thiemann, J., Schuller, S. K.; Gamiz-Hernandez, A. P.; Kurisu, G.; Nowaczyk, M. M., Kaila, V. R. I. Nat. Commun. 2020, 11, 494.
- [4] Phillips JC, Hardy DJ, Maia JDC, et al. J Chem Phys. 2020, 153 (4):044130. doi:10.1063/5.0014475.

Let there be light: the case of the major

Light Harvesting Complex II

RESEARCH INSTITUTION

¹Cyprus University of Technology

PRINCIPAL INVESTIGATOR

Vangelis Daskalakis¹

RESEARCHERS

Sotiris Papadatos¹, Taxiarchis Stergiannakos¹, Eleni Navakoudis^{1,2}

PROJECT PARTNERS

²University of Crete

SuperMUC Project ID: pn34we (Prace project)

Introduction

In our lab, we probe the biophysics of Photosynthesis [1]. Photosynthesis fuels the metabolic pathways of numerous organisms in our biosphere. This is achieved by storing the solar energy into chemical bonds and by contributing to the atmospheric oxygen cycle. However, the diurnal cycle, or the environmental conditions induce fluctuations in the light intensity or quality, and thus they render it unreliable as a source of energy [2]. Higher plants have to cope with such fluctuations to sustain their homeostasis; too little light and photosynthesis cannot occur, too much and the photosynthetic apparatus is at risk of damage, due to oxidative stress [2]. Absorption of light and tunnelling of the associated energy towards the reaction centres of the photosynthetic apparatus are finely-tuned in this context.

The photosynthetic apparatus has a central component: the major Light Harvesting Complex (LHCII) of Photosystem II (PSII). In green plants, LHCII is found as a trimer of polypeptide scaffolds with densely packed networks of pigments within (chlorophylls and carotenoids). LHCII is able (a) to capture the solar energy under normal light conditions, and (b) to quench the excess energy when the light absorption exceeds the capacity of the reaction centres (Fig. 1). LHCII within the thylakoid membrane has the potential to rapidly and reversibly switch between these roles via elusive conformations. At the quenched state, the absorbed energy in excess can be thermally dissipat-

ed under mechanisms that are cumulatively referred to as non-photochemical quenching (NPQ) of Chl fluorescence [2].

Several factors have been associated with NPQ: the increase of the trans-thylakoid membrane proton gradient (ΔpH) by lumen acidification, ion flows in the lumen-stromal areas, changes in the carotenoid content of LHCII from Violaxanthin to Zeaxanthin by the xanthophyll cycle, and finally the photoprotective protein PsbS that binds to LHCII, and can induce LHCII aggregation [4]. Research into these factors has led to the development of more efficient systems in the field of artificial photosynthesis and the increase of crop yields. Our project has focused on the LHCII-PsbS interaction that affects (i) the LHCII aggregation within the thylakoid membranes of higher plants and (ii) the inter-pigment interactions within LHCII induced by LHCII conformational transitions [3,5].




Figure 1: The chain C monomer of the major Light Harvesting Complex II (LHCII) is shown in green-blue cartoons (A). Chlorophylls are shown in green and carotenoids in red licorice. Under normal light, pigments absorb the solar energy and direct it to the reaction center of Photosystem II. Under excess light, LHCII switches to a quenched state where the absorbed solar energy in excess is dissipated as heat (B). Blue regions on the cartoons of A indicate important residues for the switch [3].




Figure 2: A liposome is shown with two LHCII trimers and four PsbS monomers embedded within (A). LHCII trimer aggregation in the presence of PsbS. Trimers are shown in blue, red and orange surfaces, while PsbS is shown in ice-blue surface (B).

Results and Methods

Several LHCII-PsbS models were built at all-atom (AA) resolution [3], and at the coarse-grain (CG) level [5]. These were embedded within a thylakoid membrane patch, or a thylakoid liposome (Fig. 2). We have proposed that in the transition of the LHCII trimer from the light harvesting to the quenched state, a change in




Figure 3: Conformational Changes for the LHCII transition from the light harvesting (green cartoons) to the quenched state (transparent cartoons). Arrows indicate the transitions. Selected LHCII helices are indicated for reference.

the shape of the protein is induced, that is associated with a thinning of the thylakoid membrane, also observed experimentally [5]. The change in the LHCII shape induces an accumulation of galactolipids around LHCII to lower the cost of hydrophobic mismatch. When PsbS is activated, it binds to LHCII proteins, the galactolipids around LHCII are expelled and thus LHCII is aggregating due to solvent (lipid) depletion. The proposed mechanism clearly explains why an enhanced pH gradient (ΔpH) is regulating the photosynthetic process. PsbS is able to control the LHCII trimer diffusion and aggregation within the thylakoid membranes through the manipulation of the induced hydrophobic mismatch at enhanced ΔpH [5]. The LHCII-PsbS models at AA resolution, along with numerous others, have also provided the complete computational framework to drive the LHCII conformation between, previously elusive, different states, associated with light harvesting and quenched modes [3].

Thus, we have probed for the first and at the same time the effects of all the key NPQ factors, on the LHCII trimer conformation. The NPQ-related LHCII main conformational transitions reported (Fig. 3) are in line with the NPQ literature. However, even with the extensive sampling reported herein for the protein scaffold transitions, a hard switch between light harvesting and quenching is still absent in the induced inter-pigment interactions probed by a common approach in computational chemistry [3]. Thus, this study has shifted the attention from the extensive conformational sampling, to the problem of accurately describing inter-pigment interactions within LHCII (i.e. by accurate *ab initio* methods) [3].

For this project, elaborate computational methods were employed to achieve our goals described in detail elsewhere [3,5]. In brief, PTmetaD-WTE, or Parallel Tempering Metadynamics Simulations were employed at the Well-tempered Ensemble (a variant of Replica Exchange) to probe how LHCII trimers (CG) aggregate within the thylakoid membrane of higher plants, or how LHCII conformation is altered upon LHCII-PsbS interaction at AA resolution (Fig. 3). A cumulative sampling of more than 150 μs at AA/ CG levels of LHCII models was made possible by the SuperMUC-NG allocated resources (17,000,000 core-

hours). This has yielded a total of around 210GB production data in around five hundred files for analysis. A typical job was run on 3,072 cores.

Ongoing Research / Outlook

The gained understanding of LHCII-PsbS interaction and LHCII aggregation can further be exploited to guide future experimental or computational studies on plant photoprotection. We thank PRACE for awarding us access to the Petascale resource SUPERMUC-NG. The work was co-funded by the European Regional Development Fund and the Republic of Cyprus through the Research and Innovation Foundation (Project: POST-DOC/0916/0049). Without the SuperMUC-NG resources, the extensive sampling by the PTmetaD-WTE method would not have been possible. The main limitations/obstacles for this study were the convergence of the results, which we overcame by the most extensive sampling of such systems up to date (150 μs). In addition, the important shift in focus for the research in NPQ, that we proposed (from conformational sampling to accurately describing inter-pigment interactions), would be impossible to predict without the allocated amount of resources. SuperMUC-NG could be employed for future studies to describe such inter-pigment interactions by *ab-initio* methods. The setup of such simulations is under way, along with proposals for submission to future PRACE/DECI calls. Unpublished results from this project (mainly related to the LHCII-lipid interactions in liposomes) are still under preparation for future publications.

References and Links

- [1] CEM lab. <https://cemlab.wordpress.com/>.
- [2] Ruban, A. V., Johnson, M. P. & Duffy, C. D. P., *Biophys. Acta (BBA)-Bioenergetics* 1817, 167–181 (2012).
- [3] Daskalakis, V., Papadatos, S. & Stergiannakos, T., *Chem. Commun.* (2020) doi:10.1039/D0CC04486E.
- [4] Ruban, A. V. Light harvesting control in plants. *FEBS Lett.* 592, 3030–3039 (2018).
- [5] Daskalakis, V., Papadatos, S. & Kleinekathöfer, U., *Biochim. Biophys. Acta - Biomembr.* 1861, 183059 (2019).

Simulating Blood Flow in the Virtual Human

RESEARCH INSTITUTION

¹Leibniz Supercomputing Centre

PRINCIPAL INVESTIGATOR

Dieter Kranzlmüller¹

RESEARCHERS

Alexander Patronis², Brian Wylie², Rene Halver², Godehard Sutmann², Thomas Odaker¹, Martin Schulz³, Jon McCullough⁴, Robin Richardson⁴, Peter Coveney^{4,5}, Ioannis Zacharoudiou⁴, Ryan Marshall⁶, Martin Ruefenacht⁶, Anthony Skjellum⁶, Bryn Lloyd⁷, Esra Neufeld⁷

PROJECT PARTNERS

²Jülich Supercomputing Centre, ³Technische Universität München, ⁴University College London, ⁵University of Amsterdam, ⁶University of Tennessee at Chattanooga, ⁷Foundation for Research on Information Technologies in Society (IT²S)

SuperMUC Project ID: pn72qu (Gauss Large Scale project)

Introduction

This Large Scale Project was awarded 35 million core-hours and has been focused on the realisation of studying systemic 3D blood flow at full human-scale. To achieve this we have utilised the lattice Boltzmann fluid flow solver HemeLB [1] to conduct simulations within vascular geometries generated from full human-scale MRI data. A key feature of this work was that two instances of HemeLB were coupled together to allow arterial and venous flow to be simulated simultaneously.

This work demands the extensive computational resources of a machine like SuperMUC-NG to not only conduct simulations within an acceptable timeframe but also to generate and store the necessary input and output data. At the largest scale of a geometry consisting of multiple billions of data locations, files of multiple terabytes in size are generated. Achieving results at this scale requires a code that continues to scale well to large core counts. HemeLB, both in its single and coupled versions, possesses excellent scaling characteristics and these have been reinforced during this project.

The work of this project contributes to the work of CompBioMed [2], a European Commission funded Horizon 2020 Centre of Excellence, in developing HPC methods in the field of biomedicine and, in particular, the development of a virtual human to assist clinicians in optimising the diagnosis and treatment of patients.

Results and Methods

A significant achievement that was made in the first year of this project has been to demonstrate strong scaling of the single version of HemeLB to almost 310,000 cores on SuperMUC-NG (see Fig. 1). This illustrates that our code can be efficiently deployed on this architecture at all scales regularly available (i.e. up to half-machine scale) and can be extended to full-machine scale when block operations and stability of hardware permits. We were also able to demonstrate

that the self-coupled version of HemeLB was able to scale linearly to almost 30,000 cores with a human-scale test geometry (see Fig. 2) – an excellent result for a coupled code in an initial attempt. We expect that this could be further improved with a larger geometry and better load balancing of cores between the coupled geometries.

The focus of our work in the current allocation period was the development and deployment of coupled models of arterial and venous blood flow on 3D human-scale vascular geometries. In particular, these firstly concentrated on generating valid flow geometries from full body data. The second component of work coupled these geometries with sub-scale models for representing the capillary beds between arteries and veins.

In the early stages of the project we were able to generate flow geometries of full systemic arterial and venous networks at a resolution of 60µm (see Fig. 3). This initial processing resulted in geometries consisting of 506,859,052 sites for the arterial domain and 1,558,375,173 sites for the venous domain. This resolution was found to be the approximate limit for the SuperMUC-NG before memory or wall-clock limits were encountered. Recently, in collaboration with LRZ staff, we were able to develop a refinement tool that could sub-scale these domains to higher resolutions without utilising the most resource intensive components of the geometry generation process. This will halve the resolution and increase site counts for both geometries by a factor of eight. Developing 3D geometries of this scale, and particularly of the venous network, remains at the forefront of vascular simulation. In the near future we anticipate being able to make use of improved geometries to further our full-scale human simulations.

Significant work has also gone into developing and improving sub-scale models to represent capillary beds. In our work, these provide the crucial coupling links between the 3D arterial and venous simulations. We have two models we are comparing. The first is based on assigning velocity coupling factors between




Figure 1: Strong scaling behaviour of HemeLB up to full-machine scale on SuperMUC-NG. This data was obtained with a problem size of $\sim 10^{10}$ fluid sites. Markers indicate the ideal, observed and 80% of ideal performance.




Figure 2: Strong scaling behaviour of self-coupled HemeLB. This data was obtained with a total problem size of $\sim 2 \times 10^9$ fluid sites. Markers indicate the ideal, observed and 80% of ideal performance.

linked inlets and outlets based on geometric properties and an assumed pressure drop between them. The second implements a resistor-capacitor representation of capillary beds as the coupling mechanism between simulations, this is based on models commonly used in 1D approaches. We are currently using a system of vessels from the left forearm to examine and compare the behaviour of these models. This domain is of medical interest for those suffering from kidney failure who require haemodialysis to remove waste products from the bloodstream. The insertion of an arteriovenous fistula, often in the left forearm, provides an access point to the bloodstream for dialysis to take place. We plan to use our model to simulate both 'normal' vessel behaviour and when specific arteries and veins are linked to form a fistula.

Simulation of vascular flow did not exhibit the extreme memory requirements observed in geometry generation. Here, typical jobs could be executed on a minimum of 2880 cores with larger and longer jobs taking advantage of HemeLB's scaling characteristics. The calculation speed was dependent on quantity of output information required and properties of the coupling scheme. Full data output for a large geometry could easily demand storage on the order of tens or hundreds of gigabytes depending on output frequency and data requirements.

Ongoing Research / Outlook

The simulation demands of 3D full human blood flow can only be met by HPC infrastructure like SuperMUC-NG. Even so, we implemented more efficient refinement techniques to overcome wall-clock and memory limitations.

We plan to continue the work of this project with the introduction of a 3D heart model (Alya, [3]) to the HemeLB studies and use the models developed to investigate a range of cardiovascular conditions through simulation.

References and Links

- [1] <http://hemelb.org>
- [2] <https://www.compbioimed.eu/>
- [3] M. Vázquez et al., Journal of Computational Science, 14, 15-27, 2016.




Figure 3: Systemic arteries (red) and veins (blue) that we are attempting to simulate at high resolution as part of our project. The smallest vessels captured here have a diameter of approximately 0.4mm. This figure indicates the significantly greater complexity of the venous network.

Mitigating the COVID-19 pandemic with an innovative approach to accelerate drug discovery

RESEARCH INSTITUTION

¹Ludwig Maximilian University, ²University College London

PRINCIPAL INVESTIGATOR

Dieter Kranzlmüller¹, Peter V. Coveney²

RESEARCHERS

Shunzhou Wan², Agastya P. Bhati², Dario Alfe², Shantenu Jha³

PROJECT PARTNERS

²Department of Computer Engineering, Rutgers University, USA

SuperMUC Project ID: pn98ve

Introduction

Our current project [1] aims to deploy INCITE software environment on SuperMUC-NG, and to predict binding free energies for a large set of compounds bound to several COVID-19 protein targets. We thank GSC and LRZ for awarding us this project with 30M core-hours on SuperMUC-NG at such a short notice to help us and the CompBioMed CoE to mitigate the COVID-19 pandemic [2]. Since the outbreak of COVID-19, researchers around the globe have been attempting to develop drug which targets specific viral protein that is vital for its propagation [2]. However, the drug discovery process employed in the pharmaceutical industry typically requires about 10 years and \$2-3 billion for a single new drug which is obviously not useful in emergencies like this pandemic. The current process involves exploring the huge chemical space with in silico methods followed by in vitro and finally in vivo processes on the selected few compounds with filtering at each step. Irrespective of their level of accuracy, in silico methods involved rely heavily on human intelligence (HI) for applying chemical knowledge to filter out or suggest structural features that should be incorporated into proposed ligand molecules to improve their binding interaction with a target protein. This makes the process slow and is a major bottleneck in the drug discovery process. Machine learning (ML) techniques are increasingly being used to overcome this bottleneck. Recent developments in deep learning (DL) allow generation of novel drug-like molecules in silico by extensive sampling of the chemical space of relevance. However, their reliability depends heavily on the training data available; when insufficient reduces their effectiveness. Fortunately, physics-based and ML methods are complementary and, hence, combining these two should provide a very efficient way to predict binding affinities. This is precisely the goal of this project. Here, we are developing and implementing a novel in silico drug design method coupling ML and physics-based molecular dynamics (MD) methods (Fig. 1). Candidate compounds are selected from the output of a DL generative algorithm. The selected compounds are scored using physics-based methods based on the binding free energies calculated and this

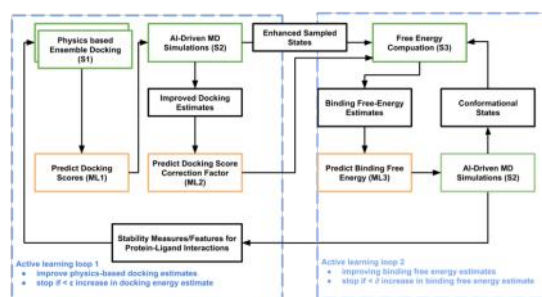


Figure 1: ML meets MD. Coupling the complementary molecular dynamics (MD) and machine learning (ML) methods for drug development in COVID-19.

information is then fed back to the DL algorithm for active learning, thereby refining its predictive capability. This loop proceeds iteratively involving a variety of physics-based scoring methods with increasing level of accuracies at each step ensuring that the DL algorithm gets progressively more accurate in its predictions. Augmenting HI with artificial intelligence (AI) by supplementing chemists' knowledge can substantially reduce the throughput time for exploring this huge chemical space and hence improve the efficacy of exploration of real and virtual chemical libraries. The ongoing COVID-19 crisis has exposed severe limitations in the current pharmaceutical mode of drug discovery and it is imperative to overturn it to urgently develop a drug. Our project is designed to accelerate the required transformation. This project is a theory and compute led initiative to be executed by a consortium consisting of experts of MD and ML methods in collaboration with experimentalists specialising in protein structure determination, chemical synthesis and binding affinity estimation [1]. It is part of the wider CompBioMed collaboration with partners in University of Chicago (UC), Argonne National Laboratory (ANL) and Rutgers university (RU). The team at University College London (UCL) leads the MD front of this project whereas the UC/ANL team performs ML. Our approach will have direct applicability in the pharmaceutical industry for quick identification of potent binders for a given target protein and binding pocket. In addition to the drug design aspect, our method can also be applied to diagnose the impact of protein mutations on the action of different drugs and




Figure 2: The four protein targets studied within this project. From left to right: 3CLPro, PLPro, ADRP and NSP15. The proteins are shown in cartoon representation, and compounds in stick representation.

help in tackling drug resistance which is also anticipated to arise in COVID-19.

Results and Methods

The major use of LRZ resources under project pn98ve was to perform large scale simulations and calculations related to two MD methods for binding affinity estimation, namely, enhanced sampling of molecular dynamics with approximation of continuum solvent (ESMACS) and thermodynamic integration with enhanced sampling (TIES) [3]. Both of them are ensemble simulation based methods which have been developed over past few years by the UCL team to estimate reliable and reproducible binding affinity estimates. ESMACS is an end-point method that has a wider applicability domain and is useful in the hit-to-lead step of the drug discovery process, whereas TIES is an alchemical method that is applicable only for estimating relative binding affinities for a pair of congeneric compounds and hence is applicable during the lead optimisation step. To perform modelling and calculation with optimal efficiency, we have been using the Binding Affinity Calculator (BAC) [4], a highly automated molecular simulation based free energy calculation workflow tool. With the resource allocation in the current project, we have applied ESMACS and TIES to study ~700 compounds binding to four of the main protein targets in COVID-19: 3CLPro, PLPro, ADRP and NSP15 (Fig. 2). These proteins have diverse functions in replication and transcription of the virus, and are important targets for pharmaceutical drug design and discovery. We have performed a large number of ESMACS calculations on the lists of structures generated by our ANL colleagues using their ML algorithm and shortlisted using docking methods. For the four targets we investigated, about 4-19% of the compounds studied show promising binding free energies, that is, more negative than -8.24 kcal/mol (corresponding to a K_D value on the nM scale). Our simulations also show that docking scores are not reliable for predicting the binding potency. We, therefore, need physics-based ESMACS approach for precise and reproducible binding free energy estimation. So far more than 12 million Core-hours, ~41% of the total allocation, has been consumed.

Ongoing Research / Outlook

RU team provides us middleware infrastructure support related to workflow management with their "Radical Cybertools". The implementation of Radical Cybertools on SuperMUC-NG is an ongoing work. We thank LRZ team for the quality of their professional services in our support. In order to train ML algorithm, a large dataset is required. Therefore, under this project, we plan to simulate 10,000 compounds selected with a combination of docking and ML for each of the four targets in the next iteration of our MD-ML loop. For optimal utilisation of computational resources, a coarse-grained ESMACS protocol will be applied for the first few iterations, whereby a reduced ensemble size and a smaller distance for non-bonded interactions will be employed. The MD study will provide an ensemble of protein conformations, much larger than the handful number of structures from X-ray crystallography. More insights will be obtained on binding sites of proteins and about interactions between proteins and compounds bound to them. Such information is expected to significantly improve predictive accuracy of the ML approach. After sufficient number of iterations, we should have identified relevant scaffolds for potent lead compounds. Then we will implement the TIES protocol to calculate relative binding free energies for selected pairs of compounds to perform lead optimisation. TIES is computationally costlier than ESMACS, but at the same time it is also theoretically more sound and accurate method. It will help us identify potential binders with greater confidence that may be passed to experimentalists for binding assay development.

References and Links

- [1] <https://www.compbioed.eu/the-consortium-on-coronavirus>.
- [2] <https://www.gauss-centre.eu/news/research-highlights/article/gcs-centres-support-research-to-mitigate-impact-of-covid-19-pandemic>.
- [3] S. Wan, A. Bhati, S. Zasada P. Coveney. 2020. Rapid, accurate, precise and reproducible ligand-protein binding free energy prediction, *Interface Focus*, in press.
- [4] S. Sadiq, D. Wright, S. Watson, S. Zasada, I. Stoica, P. Coveney. 2008. *J. Chem. Inf. Model.* 48, 1909-19.

Transcription factors pioneer nucleosomes

in motion

RESEARCH INSTITUTION

¹Hubrecht Institute, Utrecht, the Netherlands;

²Max Planck Institute for Molecular Biomedicine, Münster

PRINCIPAL INVESTIGATOR

Vlad Cojocaru^{1,2}

RESEARCHER

Jan Huertas^{1,2}, Caitlin MacCarthy³, Hans Schöler³, Abdenour Soufi⁴

PROJECT PARTNERS

³Max Planck Institute for Molecular Biomedicine, Münster

⁴University of Edinburgh

SuperMUC Project ID: pr27fi (Gauss Large Scale project)

Introduction

Transcription factors are proteins that directly or indirectly bind to DNA in order to transcribe genetic information into RNA. In most cases, accessibility to DNA is a prerequisite for binding of transcription factors. However, in the nucleus, the DNA is packed into chromatin, making it often inaccessible for transcription factor binding. The fundamental unit of chromatin is the nucleosome, which is formed by wrapping 147 DNA base pairs around a core of eight histone proteins. The histones have a structured core and disordered terminal tails. Interestingly, a series of transcription factors are able to bind to closed chromatin states, recognizing their binding sites even in the presence of nucleosomes. These factors, known as “pioneer transcription factors”, can help open chromatin, increase DNA accessibility, and support binding of other transcription factors. It has been reported that many of the transcription factors involved in transitions between different cellular states are pioneer factors. In particular, in a Nobel Prize-awarded discovery, it has been shown that three of those factors, Oct4, Sox2, and Klf4 are required to convert a somatic skin cell into a pluripotent stem cell, a process known as cellular reprogramming. When introduced in skin cells, these transcription factors recognize binding sites in DNA wrapped in nucleosomes¹. They play a major role in the large rearrangements of chromatin that occur during such cell fate transitions. In this project, we studied how the master regulator of stem cell pluripotency Oct4 binds to nucleosome. We aimed to answer 3 major questions: (i) what are the structural features and dynamics of genomic nucleosomes bound by Oct4; (ii) what is the binding mode of Oct4 on different genomic nucleosomes?, (ii) does the binding of Oct4 increase the structural flexibility of the nucleosomes ?

Results and Methods

The 40 million core-hours awarded to this project allowed us to perform extensive sampling of the conformational space of genomic nucleosome in

atomistic molecular dynamics simulations and to test different Oct4 binding modes to different nucleosomes in microsecond long MD simulations. All our simulations amounted to a total of approximately 35 μ s. We used NAMD on 1,024 or 2,048 core per 24 h job. The raw data we generated amounted to roughly 60 TB which we transferred to our local storage space and analyzed locally.

Genomic nucleosomes in motion

Most of the structural studies and MD simulations on nucleosomes to date were performed using nucleosome structures with DNA sequences engineered for strong positioning of the histone core. In this project, we performed a total of 18 μ s atomistic MD simulations of 1 engineered and 2 genomic nucleosomes known to be bound by Oct4. Figure 1 shows the ensemble of structures obtained from a single 1 μ s simulation of the engineered nucleosome. From these simulations, we demonstrated how the interplay between 2 histone tails cooperate to control nucleosome motions. We monitored 2 rare large opening events, one occurring in each genomic nucleosome, and found that nucleosome conformations with different degree of opening display specific conformations and positions of the 2 histone tails (Figure 2). Moreover, nucleosome opening and closing is regulated by specific pattern of interactions between residues in the tails that can act as epigenetic regulatory sites, and the DNA. Because the distribution of open and closed nucleosomes in chromatin fibers impacts their structure and compaction, the mechanism we observed is key to understanding chromatin dynamics [2,3]).

Binding modes of Oct4 to nucleosomes

To study how Oct4 interacts with nucleosomes, we first built models of Oct4 bound to the 2 genomic nucleosomes based on the experimental data available from genomic studies of Oct4 binding during reprogramming of skin cells to stem cells. We probed the stability of these models in relatively short MD simulations. From these, we showed that some noncanonical configurations, but not all tested




Figure 1: Nucleosome motions on the μ s time scale (taken from Huertas et al [3]). The representation was obtained by the superposition of every 5 ns of the 3 μ s simulation ensemble of a nucleosome. The histone core is depicted in gray, DNA in yellow, and the histone tails of H3, H4, H2A and H2B are in green, blue, red, magenta and cyan respectively.




Figure 2: Histone tail mediated nucleosome opening (taken from Huertas et al [3]). Histone core and DNA are in gray cartoons. The linker DNA arm that undergoes large opening is in yellow. The histone tails are in green and red with the non hydrogen atom highlighted and colored by atom name.

configurations of Oct4-nucleosome complexes were stable and not affected by DNA curvature or DNA-histone interactions [2] (2 examples shown in Figure 3). We demonstrated that MD simulations can be used to identify nucleosome-compatible binding modes of transcription factors.

Oct4 binding and nucleosome dynamics

With the models we generated that were stable in the initial short MD simulations, we proceeded and performed extensive μ s long MD simulations. From these, we revealed alternative mechanisms by which Oct4 interprets and induces nucleosome motions. We found that Oct4 captures and stabilizes partly open conformations of nucleosomes using both its DNA binding domains. Only one of the domains, but not always the same, binds to nucleosomes in a sequence specific manner, whereas the other explores the DNA. In certain circumstances, Oct4 induces a metastable large opening of nucleosomes. Our findings explain how histone tails and pioneer transcription factors impact on intra nucleosome dynamics, thus contributing to the understanding of chromatin opening and gene regulation (unpublished data; manuscript in preparation) [4].

Ongoing Research / Outlook

The computer resources awarded on SuperMUC for this project enabled us to extend the length of MD simulations beyond the state of the art and to perform




Figure 3: Binding modes of Oct4 to nucleosomes (taken from a review by Huertas et al. submitted for publication). Histones are in gray, DNA in yellow. The 2 DNA binding domains of Oct4 are in green (the POU_S domain) and in cyan (the POU_{HD} domain).

multiple simulations for each system. This was crucial for 2 reasons: (i) the functionally relevant dynamics of nucleosomes occur on time scales longer than those typically available in MD simulations and therefore it is essential to perform multiple, long simulations of nucleosomes to understand these dynamics, and (ii) the binding modes of Oct4 to nucleosomes were not known when we started the project and therefore we had to probe for many different configurations that fitted the experimental data available. Currently, we are finalizing the analysis of the long Oct4-nucleosome simulations and we are preparing the article presenting these findings [4]. During our project, experimental structures of Oct4 and other transcription factors bound to a nucleosome were reported [5]. These structures confirmed that transcription factors induce structural flexibility in the nucleosomes. We are planning to perform extensive MD simulations similar to the ones presented here for other transcription factors. In addition, we will investigate the application of enhanced sampling simulations in which processes such as nucleosome unwrapping are accelerated to reveal how transcription factors affect nucleosome structure and dynamics.

References and Links

- [1] A. Soufi, M.F. Garcia, K.S. Zaret, et al., Cell, 161,555-568 (2015).
- [2] J. Huertas, C. M. MacCarthy, H. R. Schöler, V. Cojocaru, Biophys. J. 118:9,2280-2296 (2020).
- [3] J. Huertas, C. M. MacCarthy, H. R. Schöler, V. Cojocaru, Biorxiv (2020). DOI: 10.1101/2020.09.04.282921.
- [4] J. Huertas, C. M. MacCarthy, H. R. Schöler, V. Cojocaru, Oct4 interprets and induces nucleosome structural flexibility. In preparation.
- [5] A. K. Michael et al., Science 368(6498):1460-1465 (2020).

Understanding Substrate Recruitment by

Intramembrane Proteases: Progress and Problems

RESEARCH INSTITUTION

Physics of Synthetic Biological Systems, TU Munich

PRINCIPAL INVESTIGATOR

Christina Schamagl

RESEARCHER

Simon Menig

PROJECT PARTNERS

—

SuperMUC Project ID: pr27wa

Introduction

Cells have developed a variety of mechanisms to communicate with their environment and to adapt to changes in the extracellular space. One such mechanism is regulated intramembrane proteolysis which controls the activity of membrane proteins such as signal transduction and cell differentiation or degrades faulty proteins. In contrast to soluble proteases, intramembrane proteases (IMP) provide a hydrophilic catalytically active site within the membrane. This allows the IMP to cleave peptide bonds of their transmembrane (TM) protein substrates within the hydrophobic environment of the lipid bilayer. In order to expose the scissile bond for hydrolysis, the protease must find, recognize, bind, and locally unwind the α -helical TM domain of the substrate [1,2].

In many cases, the IMPs do not target a specific substrate, but cleave diverse single-span TM proteins without apparent preference for specific sequence motifs. On the other hand, the proteolytic process appears to be tightly regulated, and a deregulation is associated with diseases. A prime example is the intramembrane proteolysis of the amyloid precursor protein (APP) by γ -secretase (GSEC), where misprocessing of APP is related to onset and progression of Alzheimer's disease.

The study the mechanisms and determinants of substrate recognition and cleavage is the aim of a collaborative research project [1] (FOR2290, DFG Forschergruppe "Understanding intramembrane proteolysis").

Results and Methods

Within the research group, experimental methods such as functional assays and structure determinations are complemented by computational approaches.

The investigations using APP processing by GSEC as a read-out indicate that protease activity can be regulated at several levels (Figure 1): (1) Conformational flexibility of their TM domain confers substrates the ability to adapt to interactions with the enzyme [2-4]. (2) Transfer from membrane to hydrophilic environment induces unwinding of substrate TM domain α -helical hydrogen bonds close to the scissile bond [2]. (3) The lipid composition affects interactions and assembly of the GSEC complex, that in turn modulates cleavage [1]. (4) At least part of the substrate's extracellular domain impacts processing as revealed by conversion of non-substrates into substrates upon domain swaps [1].




Figure 1: Regulation of intramembrane proteolysis. The substrate transmembrane domain (dark yellow) and few juxtamembrane residues (dark grey) are sufficient for high-affinity binding and processing by the protease GSEC. Dynamic interactions between full-length substrate, enzyme, and lipids play a regulatory role for protease activity and specificity.




Figure 2: Capture of substrates (APP, NOTCH1) and the apparent non-substrate integrin β 1 (ITGB1) by GSEC predicted from simulations. The GSEC complex consist of four subunits. Presenilin 1 (PS1) harbors the catalytically active sites (red spheres) between its N-terminal and C-terminal fragment (NTF, CTF). Numbers refer to PS1 TM helices. Darker grey color indicates more frequent and/or longer-lasting contacts. The hydrophilic loops extending from PS1 TM helix 2 towards the extracellular and the intracellular side contribute the main PS1 contacts with substrate juxtamembrane residues.

In order to span a range of time and spatial scales, we employ multi-scale molecular dynamic (MD) approaches using the highly scalable GROMACS 2019.3 code (gromacs.org). Ensemble-based coarse-grained (CG) models (MARTINI 2.2 force field, cgmartini.nl) provide the long simulation times and large number of replicas required for an unbiased prediction of substrate-enzyme contacts in membrane using the DAFT approach (Docking Assay For Transmembrane Components, *cgmartini.nl*). A typical DAFT run consist of 1,000 x 1 μ s replicates of a system which corresponds to ~80,000 heavy atoms. Distributed to 8000 Sandy-Bridge nodes on SuperMUC we reached 850 ns/day. With SuperMUC-NG the performance increased to 3,800 ns/day on 24,000 cores. After mapping the CG system back, atomistic resolution MD simulations (CHARMM36 force field, *mackerell.umaryland.edu*) are used to analyze molecular details of conformational dynamics and interactions in the μ s range (210,000 atoms, 960 cores, 110 ns/day).

Structures of human GSEC in substrate-bound states determined by cryo-electron microscopy revealed some insight into mechanisms of substrate recognition at atomic resolution. However, these structures were engineered by site-specific cross-linking and represent static snapshots.

Experiments [1] have shown that encounter of APP and GSEC is highly dynamic and involves intermediate substrate binding to and dissociation from exosites before translocation to the active site takes place. After validating the *in silico* predicted exosites for APP against these experiments, our computational approach was used to probe the capture of various substrates by GSEC in a single-component POPC (1-palmitoyl-2-oleoyl-sn-glycero-3-phosphocholine) bilayer. We find that substrates sample exosites on the catalytic subunit of GSEC (presenilin 1, PS1) with different preferences (exemplified in Figure 2 for APP and NOTCH1). The contact maps for ~15 different substrates could be classified either as APP-like or as NOTCH-like. The apparent non-substrate integrin β 1 (ITGB1) does not contact PS1 at all but sticks to the PEN2 (presenilin-enhancer-2) subunit. This indicates that PEN2 could act as first checkpoint for sorting non-

substrates from substrates [2]. However, the comparison with other non-substrates was hampered by their elusive repertoire.

Ongoing Research / Outlook

To complete our project we will investigate the regulatory role of membrane composition for substrate recruitment. Lipids might either directly interact with protease and substrate or change membrane properties which affect association pathways such as fluidity, thickness or curvature. The previous studies used a bilayer of POPC, the main component of the synaptic plasma membrane. A cholesterol-rich multi-component bilayer including lipids with polar head groups and different fatty acids' chain length and saturation will be more native like.

Sampling of lipid dynamics and interactions will be a challenging task given the large space of slow conformational re-arrangements such as lipid mixing and desolvation of binding interfaces. The necessary long simulation times for a large number of replicates will greatly benefit from enhanced performance provided by SuperMUC-NG.

References and Links

- [1] <https://www.i-proteolysis.de/>
- [2] M. Hitzberger, A. Götz, S. Menig, B. Brunschweiler, M. Zacharias, Ch. Scharnagl., *Semin Cell Dev Biol*, 105: 86-101 (2020). DOI: 10.1016/j.semcdb.2020.04.008
- [3] A. Götz, N. Mylonas, Ph. Högel, M. Silber, H. Heinel, S. Menig, A. Vogel, H. Feyrer, D. Huster, B. Luy, D. Langosch, Ch. Scharnagl, C. Muhle-Goll, F. Kamp, H. Steiner, *Biophys. J.* 116: 1-18 (2019). DOI: 10.1016/j.bpj.2019.04.030
- [4] A. Götz, Ph. Högel, M. Silber, I. Chaitoglou, B. Luy, C. Muhle-Goll, Ch. Scharnagl, D. Langosch, *Scientific Reports* 9: 3521-3532 (2019). DOI: 10.1038/s41598-019-41766-1

PIP₂: Another player in biased signaling of G protein-coupled receptors?

RESEARCH INSTITUTION

¹Institute of Nanotechnology, KIT, Karlsruhe

PRINCIPAL INVESTIGATOR

Wolfgang Wenzel¹

RESEARCHER

Kristyna Pluhackova²

PROJECT PARTNERS

²Department of Biosystems Science and Engineering, ETH Zürich, Basel

SuperMUC Project ID: pr27wi

Introduction

The ability of a cell to react to extracellular stimuli is enabled by a complex protein machinery attached to the cell membrane. The core of this machinery consists of transmembrane proteins termed G protein-coupled receptors (GPCRs) and their associated G proteins and arrestins localized inside the cell. While G proteins mostly trigger only signaling pathways, arrestin can either desensitize and recycle the receptor or trigger diverse signaling pathways on its own. Due to their important roles in healthy and disease states, GPCRs are targeted by one third of all current drugs.

Over the last years it was discovered that GPCR signaling is modulated by diverse mechanisms. This modularity is termed biased signaling. One modulator of GPCR signaling is the membrane composition. Thereby, the biased signaling can result from different (co)localization of GPCRs, G proteins, and arrestin, from changes in the oligomerization state of the receptor and from differences in the GPCR conformational flexibility. The most important lipidic modulators are cholesterol [1] and acidic lipids, mainly phosphatidylinositol phosphates (PIPs). The latter were found to be essential for complex formation of β_2 -adrenoreceptor (AR) (a prominent GPCR responsive to adrenaline) with the kinase (enzyme responsible for GPCR phosphorylation) GRK5. Moreover, acidic lipids stabilize the active state of the

receptor and can influence signaling pathways by determining which G protein will preferentially bind the given GPCR. Moreover, arrestin is only able to internalize GPCRs in presence of PIPs[2]. Although binding sites on both, arrestins and on a range of GPCRs have been determined in the past, no molecular resolution information is available on PIP binding to the complex of a GPCR with arrestin so far. In order to obtain this information and its possible impacts on the biased signaling of GPCRs by modulation of the relative conformation of the arrestin and the GPCR, here we have performed multiscale molecular dynamics simulations of β -arrestin2 bound to β_2 -AR in membranes with and without PIP₂.

Results and Methods

Sequential multiscale molecular dynamics (MD) simulations spare computational resources by studying time consuming processes (like lipid binding to the protein complex or protein-protein association) using coarse-grained (CG) resolution. After conversion back to all-atom (AA) resolution using the tool *backward* [3], the simulations were continued atomistically, refining the lipid-protein and the specific protein-protein interactions.

In CG MD simulations group of atoms are bundled into so called beads or particles having average properties of the atoms they represent. The resulting speed-up amounts to the factor of 350 for the here applied methodology and simulation system. However, CG simulations cannot be as well parallelized as AA systems of the same size. In detail, we have used 2 SuperMUC-NG nodes (48 cores each) for CG and 13 nodes for AA simulations, producing ~340 ns and ~6.25 ns simulation time per hour, respectively. All CG simulations were run for 10 μ s and all AA simulations for at least 1 μ s. All simulations have been performed using the well-established, high performing simulation engine GROMACS (www.gromacs.org) in single precision. GROMACS has only low requirements concerning I/O. The trajectories and other output files are written regularly (every ~20–30 minutes) by appending data to a few binary and text files. All I/O is done by the first MPI rank only. Typical simulations




Figure 1: GPCRs bind ligands (orange) on the extracellular side (top) and transmit signals over the membrane (blue sticks and balls) to the inside of the cell (bottom) where they bind trimeric G proteins (left). After receptor phosphorylation of the C-terminus (small orange balls) the binding of arrestins competes with G proteins. Once arrestin binds to the receptor intracellular binding pocket (right), the receptor is desensitized, internalized and either degraded or recycled.




Figure 2: PIP₂ binds rapidly and specific to the β_2 -AR/ β -arrestin complex. PIP₂ lipids were attracted to specific binding spots on arrestin and receptor within a microsecond, and stayed stably bound there (almost half of the lipids were stably bound to arrestin and a third to receptor). The temporary interaction of PIP₂ with less specific binding spots lasted 300 ± 35 ns for arrestin and 406 ± 61 ns for receptor.




Figure 3: PIP₂ binding determines arrestin binding depth and orientation. Complexes of β_2 -AR with β -arrestin embedded in membranes without PIP₂ (left) show enhanced variability in orientation and insertion depth of arrestin as compared to membranes with PIP₂. Moreover, in presence of PIP₂ arrestin prefers to bind deeper.

require 2 GB and 12 GB storage, for 10 μ s of a CG and 1 μ s of an AA simulation, respectively. However, the storage requirements can increase significantly if the output is written out more often, or if the storage of forces and velocities is needed. Additionally, we have performed extensive steered-MD simulations in which the arrestin was pulled away from the receptor at 0.2 m/s. The necessary increased box size lead to reduction of production to 3.5 ns per hour on 12 SuperMUC-NG nodes.

Ongoing Research / Outlook

By steered-MD, we have pulled diverse complexes (including different activation states of the receptor and diverse membrane composition) of β_2 -AR and arrestin apart, and recorded force-time curves and combined them with atomic-force-microscopy. The acquired results build a solid basis for further investigations of the **regulation mechanisms of extracellular signal transmission to the cell**. We were able to pinpoint residues that are especially

important for both stabilisation of the bound states and during unbinding. This knowledge is of special **pharmacological importance**, because currently only the activation state of the receptor is targeted by pharmaceuticals. However, arrestin binding to the receptor competes with binding of its cognate G protein and causes receptor desensitization, thus opening new pathways for pharmacological intervention.

References and Links

- [1] Pluhackova, K., et al., PLOS Computational Biology, 2016. 12(11): p. e1005169.
- [2] Gaidarov, I., et al., The EMBO Journal, 1999. 18(4): p. 871-881.
- [3] Wassenaar, T.A., et al., J. Chem. Theory Comput., 2014. 10(2): p. 676-690.

Regulation of proton transfer in respiratory

complex I

RESEARCH INSTITUTION

Technical University of Munich

PRINCIPAL INVESTIGATOR

Ville R. I. Kaila

RESEARCHERS

Max E. Mühlbauer, Patricia Saura, Andrea Di Luca, Daniel Riepl, Michael Röpke

PROJECT PARTNERS

—

SuperMUC Project ID: pr27xu

Introduction

All organisms need energy to power their essential processes like movement, growth, and reproduction. In all eukaryotes and several bacteria, aerobic respiration forms the main energy production system during which electrons, protons, and oxygen are combined to form water. The energy transduced is converted into an electrochemical proton gradient (or proton motive force, pmf) across a biological membrane that further drives the synthesis of adenosine triphosphate (ATP), the principal energy currency of cells.

The respiratory complex I is the entry point for electrons into many aerobic respiratory chains. Complex I takes up electrons from nicotinamide adenine dinucleotide (NADH) and transfers them to ubiquinone (Q10). The free energy is used to build up a pmf by pumping protons (H^+) across a membrane. These processes are separated by large distances of ca. 200 Å (Figure 1), but the coupling principles still remain poorly understood. Single-point mutations in complex I have been linked to several human mitochondrial diseases, why the system is also of great biomedical relevance. In this supercomputing project, we addressed how proton pumping in complex I is activated and regulated, and we designed specific mutations that disturb the function of this complex cellular machinery.

Results and Methods

In this project, we identified several key molecular elements that control proton pumping in complex I. We had previously suggested that charged interfaces between the proton-pumping modules could function as on/off-switches for proton transfer within each subunit. In this project, we calculated free energy profiles for inducing conformational changes in these ion-pairs between closed (“switch off”) and open (“switch on”) states, and how this transition depends on the hydration levels of the protein. We found that the open state is favored at high hydration levels, but kinetically inaccessible in the dry state (Figure 2).

We also probed how the proton transfer reactions are modulated by the proposed ion-pair switching process and the hydration levels. To this end, we used snapshots obtained from our large-scale MD simulations for computing QM/MM free energies for the proton transfer process. Our results suggest that the proton transfer is feasible in the open ion-pair state, but not in the closed ion-pair state (Figure 3).

Interestingly, each proton transfer reaction is coupled by the charged interface to the next subunit that, in turn, provides a key step in the signaling propagation across the membrane domain of complex I. To test the coupling principles further, we mutated identified key residues along the proton transfer wires both in silico and in vitro biochemical experiments. Our simulations suggested that the introduced substitutions indeed destabilized the proton transfer reaction, and experimentally lead to a lowered proton pumping and electron transfer activity. In summary, we demonstrated how hydration, conformational switching in conserved ion-pairs, and proton transfer reactions themselves provide key elements in proton pumping of respiratory complex I.

In this project, we used vast resources provided by SuperMUC Phase 2 and later on by SuperMUC-NG to perform large-scale classical molecular dynamics simulations of respiratory complex I embedded in a lipid membrane/water/ion environment. We employed the NAMD2 molecular dynamics-engine for our simulations with about one million atoms on microsecond timescales. Excellent scaling enabled us to utilize up to 1,536 cores per simulation and to generate about 10 TB of data concentrated in tens of large coordinate trajectories. In subsequent free energy calculations performed using replica exchange umbrella sampling, we performed 20 parallel simulations using up to 11,520 cores. In total, 24 M core-hours were dedicated to this project.

Ongoing Research / Outlook

The resources provided by this SuperMUC project were key to generate the data used in deducing the




Figure 1: Complex I catalyzes electron transfer from NADH to quinone in its hydrophilic domain (in white) and pumps protons across the antiporter-like subunits of the membrane domain (colored). Nqo13 is an antiporter-like subunit, connected to its neighboring subunits by charged interfaces (blue squares), which can form contacts between the subunits. The proton transfer within the subunit is indicated by an arrow.




Figure 2: The hydration level within a proton pumping subunit regulates the dynamics of ion-pair opening. In the “open” state (high R), a charged contact is formed with a neighboring subunit, but not in the “closed” state. The ion-pair opening affects in turn the proton transfer reaction barriers (see Figure 3).





Figure 3: Hybrid QM/MM calculations, where a chemically reactive region described at the QM (DFT/B3LYP) level is embedded in a classical environment, reveal that proton transfer is thermodynamically and kinetically favored in the open conformation (in blue), but unfavored in the closed state (in red).



mechanism presented above, and resulted in one JACS-publications [2]. During this project, our simulations migrated from SuperMUC Phase 2 to SuperMUC-NG seamlessly. We are looking forward to broadening our understanding of biological energy conversion in follow-up projects on the SuperMUC-NG infrastructure.

References and Links

- [1] <https://villekaila.com/>
 [2] Mühlbauer, M. E. , Saura P, Nuber F, Di Luca A, Friedrich T, Kaila VRI, J. Am. Chem. Soc. 142, 13718–13728 (2020).

Advanced sampling of peptide and protein association and protein conformational transitions

RESEARCH INSTITUTION

¹Lehrstuhl für Molekulardynamik, Physik-Department T38, TU München

PRINCIPAL INVESTIGATOR

Martin Zacharias¹

RESEARCHERS

Jonathan Coles¹, Manuel Hitzenberger¹, Maria Reif¹, Danial Pourjafar Dehkordi¹, Julian Hartmann¹, Florian Kandzia¹, Korbinian Liebl¹, Asman Nayis¹, Till Siebenmorgen¹, Paul Westphälinger¹, Aliaksei Krukau², Gerald Matthias²

PROJECT PARTNERS

²Leibniz-Rechenzentrum

SuperMUC Project ID: pr27za

Introduction

Atomistic Molecular Dynamics (MD) simulations are widely used to study the dynamics of proteins and peptides and the association of biomolecules. However, the simulation time scales that can be reached are often still too short to reach biologically relevant conformational changes or to follow association processes. We have developed and improved replica-exchange based advanced sampling methods to specifically tackle large scale conformational changes in biological macro-molecules and to enhanced the sampling of binding events. The techniques are based on the Hamiltonian replica exchange method. In several parallel running simulations specific biasing potentials are added to the force field (Hamiltonian) of the system. The biasing potential promotes transition in a collective degree of freedom or allows for more rapid association/dissociation of biomolecular complexes. Due to possible exchanges between replica simulations including a reference system under the control of the original force field it is possible to specifically accelerate conformational transitions or binding events of interest. We have applied these techniques to study the 90kDa heat-

shock protein (Hsp90) which is an essential molecular chaperon protein that plays a vital role in the folding process of several client proteins [1]. It is found in bacteria as well as eukaryotes and is essential for cell viability and plays a pivotal role in many signaling and regulation pathways. In its active conformation it forms a homodimer and its chaperone activity depends on ATP binding and hydrolysis. During its work cycle Hsp90 can adopt different global conformations covering a range of tightly bound closed structures up to widely open conformations [2,3]. In addition, we developed an efficient REMD method to accelerate the simulation of protein-protein and protein-peptide binding [4]. The approach can be used as efficient approach for predicting biomolecular complex formation. The approaches make efficient use of the SuperMUC parallel computer facilities.

Results and Methods

Hamiltonian replica exchange simulations on yeast Hsp90 middle and C-terminal dimer

In order to better understand the domain dynamics comparative Molecular Dynamics (MD) simulations of a sub-structure of Hsp90, the dimer formed by the middle (M) and C-terminal domain (C), were performed. Since this MC dimer lacks the ATP-binding N-domain it allows studying global motions decoupled from ATP binding and hydrolysis. Conventional (c)MD simulations starting from several different closed and open conformations resulted in only limited sampling of global motions. However, the application of a Hamiltonian Replica exchange (H-REMD) method based on the addition of a biasing potential extracted from a coarse-grained elastic network (ENM) description of the system allowed much broader

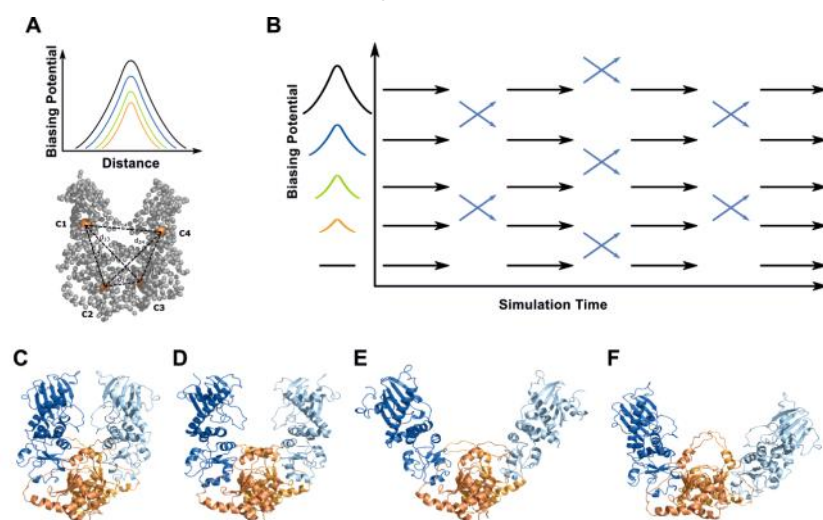


Figure 1: ENM coupled REMD simulations on a truncated yeast Hsp90 (dimer of M- and C-domains). (A) For the ENM-coupled REMD the center-of-mass of M- and C-domains (orange spheres, C1-C4) served as the 4 centers to design a biasing potential (BP) in the replica runs. It acts along the distances between the 4 centers (dashed lines). The BP increases along the replicas (B) and the width is based on the distance fluctuation derived from the ENM calculations. (B) Illustration of ENM coupled REMD. (C-F) Snapshots of global domain motion during the simulations (M: blue, C: orange).




Figure 2: (left panels) Illustration of the biased van der Waals interaction along replicas runs during the RS-REMD approach. (right panels) Broad sampling of peptide (red) placements in highest replica around whole receptor (blue) and in intermediate replica as well as narrow sampling around the native binding site in the reference replica. The peptide start placement is in light blue.

sampling of domain motions than the cMD simulations (termed ENM coupled REMD). The approach forms an effective multi-scale methodology in which directions of large scale global conformational transitions (extracted from a low resolution technique) can guide and enhanced the high-resolution sampling of the multi-domain structure [3] (Fig. 1). With this multiscale approach it was possible to extract the main directions of global motions, the energy landscape and to elucidate the mechanism of the global structural transitions of the MC dimer [3]. Future applications to Hsp90 with the N-domains could be useful to elucidate global motions in the full Hsp90 molecule.

Studying protein-protein and protein-peptide binding

In the addition to studying global protein motions, we developed a H-REMD based method to systematically investigate the formation of protein-protein and protein-peptide complexes and the refinement of predicted complexes. In each replica it includes different levels of a repulsive biasing between partners. The bias acts only on intermolecular interactions based on an increase in effective pairwise van der Waals radii (repulsive scaling (RS)-REMD) without affecting interactions within each protein or with the solvent (Fig. 2). For a set of 5 protein test cases (out of 6) the RS-REMD technique allowed the sampling of near-native complex structures even when starting from the opposite site with respect to the native binding site for one partner. The method allowed also efficient refinement of pre-docked protein-protein and protein-peptide complexes [4,5]. Even in case of relatively large protein-protein complexes the method allows the rapid sampling of near-native complex geometries. This was also possible when starting from the worst possible initial placement on the opposite side (with respect to the native binding site) of the surface of a partner protein (Fig. 3). We have recently further extended the methodology such that a simultaneous refinement and free energy type of scoring of docked complexes is possible [5]. It is based on the observation that the technique samples both near-natively bound structures and dissociated states of a given protein-protein or protein-peptide complex along the replicas. The




Figure 3: (upper panel) Snapshots of sampled starting, intermediate and near native complexes (green: receptor; light grey: ligand protein; dark grey: native bound ligand) during a RS-REMD simulation (approximate simulation stage and sampling times are indicated). (lower panel) Ligand protein sampling density on receptor surface (light grey dots) in the first third, second third and final part of the simulation.

extensive overlap of sampled states in the replica simulations allows one to extract associated free energy differences and ultimately a free energy estimate of binding.

Besides of identifying near-native binding geometries and estimating free energies of binding the methodology is also useful to characterize the full energy landscape for protein-protein binding on the surface of each partner protein and to identify transiently bound states as well as pathways from a nonspecifically bound configuration towards the native complex geometry. It is very demanding but efficient on SuperMUC, and a step forward to a realistic prediction of protein-protein and protein-peptide interactions. An extension for refinement and free energy evaluation of docked drug-receptor complexes is in progress.

Ongoing Research / Outlook

Understanding the mechanism of how peptides or proteins bind to other partner molecules is of fundamental importance for basically all biological processes. With our new RS-REMD method we are in position to rapidly sample protein surfaces for potential binding sites under realistic solvation conditions and including full flexibility of binding partners. The new techniques allow a simultaneous refinement of potential complexes and at the same time a scoring of the complexes based on a free energy estimate. The approach for enhanced global domain motions is currently applied to multi-protein complexes involved in the process of RNA interference that is of major importance for regulatory processes of the cell. Due to the large size of these systems systematic simulations are only possible by using SuperMUC.

References and Links

- [1] Pratt WB, Toft DO, *Exp Biol Med* 228 (2003) 111–33.
- [2] Hellenkamp B, Wortmann P, Kandzia F, Zacharias M, Hugel T., *Nat Methods* 14 (2016) 174–180.
- [3] Kandzia F, Ostermeir K, Zacharias M., *Front Mol Biosci.* 6 (2019) 93.
- [4] Siebenmorgen T, Engelhard M, Zacharias M, *J Comp. Chem.* 10 (2020) 1–12.
- [5] Siebenmorgen T, Zacharias M., *J Chem Inf Model* 60 (2020) 5552–5562

Multi-scale Molecular Simulations and

ATPase Activity of Hsp90

RESEARCH INSTITUTION

Department of Chemistry, Technical University of Munich

PRINCIPAL INVESTIGATOR

Ville R. I. Kaila

RESEARCHERS

Sophie L. Mader, Qi Luo, Ana P. Gamiz-Hernandez

PROJECT PARTNERS

—

SuperMUC Project ID: pr53po

6

Introduction

In proteins, biological signals can be transferred over large molecular distances and can thereby enhance or inhibit the protein activity via long-range coupling mechanisms. Such allosteric effects regulate the function of all kinds of proteins, from molecular chaperones and kinases involved in cell signalling, to redox enzymes and ion pumps [1]. An input signal, e.g., a post-translational modification or binding of a ligand, can be propagated through the protein via both electrostatic and conformational changes of central amino acids. Mutation of the amino acids involved in the signal propagation pathway may change the protein's activity.

The molecular chaperone Hsp90 (Figure 1) is involved in the maturation of hundreds of substrate proteins and regulates important cellular pathways in the eukaryotic cell. Amongst these client proteins of Hsp90 are transcription factors and protein kinases that control, e.g., cell growth. Some of these clients play a role in human diseases, including cancer and Alzheimer's disease. Hsp90 has thus evolved into a promising target for drug development [2].

Hsp90 functions as a homodimer that undergoes large conformational changes from an inactive, open state to an active, closed state, constituting its chaperone cycle that is coupled to a slow ATPase activity. To gain insight into the signaling pathways within the chaperone Hsp90 and their effect on its catalytic activity, we performed both quantum chemical and classical atomistic simulations in combination with biophysical and biochemical experiments to elucidate the molecular principles responsible for the function of Hsp90.

Results and Methods

To investigate the connection between conformation and ATPase activity of Hsp90, we performed classical atomistic molecular dynamics (MD) simulations to explore the conformational dynamics, and we calculated free energy profiles for the ATP hydrolysis reaction in different conformational states by employing hybrid quantum mechanical/classical mechanical (QM/MM) umbrella sampling simulations.

We found that conformational changes of a glutamate/arginine ion-pair (R32/E33) in the ATP binding site of Hsp90 influence the free energy barrier of the ATPase reaction by electrostatic tuning effects (Figure 2a). Substitution of the computationally identified switching residue by a neutral amino acid (R32A) led to an Hsp90 variant that decouples the global conformational state from the catalytic activity (Figure 2b, c) [3].

We also studied a post-translationally modified lysine that controls the function of Hsp90. Methylation of the lysine affects the conformational state and catalytic activity of Hsp90, suggesting signal propagation across large distances. MD simulations showed that methylation of the lysine could induce conformational changes in a neighboring ion-pair (Figure 3a, b), which might propagate through the protein. Our studies indicated that the methyl-lysine can be mimicked by an isoleucine [4].

To computationally study the conformational dynamics of Hsp90, we performed MD simulations on microsecond time scales using the software package NAMD [5]. We employed periodic boundary conditions (PBC)




Figure 1: Closed state of the dimeric molecular chaperone Hsp90, with the two protomers shown in orange and blue. ATP molecules are shown in yellow and Mg²⁺ in green. Identified switch-point residues, Arg-32 and Lys-594, are shown in red and cyan, respectively.

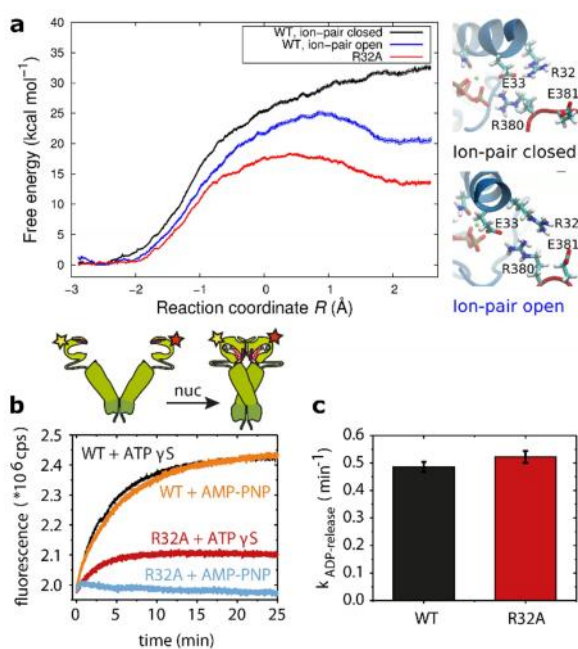


Figure 2: a) Free energy profiles for ATP hydrolysis in Hsp90 obtained from QM/MM umbrella sampling simulations for the wild type (WT) Hsp90 with the R32/E33 ion-pair closed (black curve) and open (blue curve), as well as for the R32A mutant (red curve). b) FRET experiments indicate an impaired closing behavior for the R32A variant compared to WT Hsp90. c) ADP-release assays show a similar global ATPase activity for the wild type Hsp90 and the R32A mutant. The figure is adapted from ref. [3].

and used a 2 fs integration time step. Starting from available crystal structures of the active state of Hsp90, we investigated structural effects of amino acid modifications in monomeric and dimeric models. All structural models were extracted from the Protein Data Bank (PDB) and embedded in a water-ion environment to mimic biological conditions. We typically used 800–1,000 cores per classical MD simulation.

Multiple restrained QM/MM simulations of 1 ps each were performed to estimate free energy profiles for ATP hydrolysis in Hsp90 using umbrella sampling. The simulations were performed in parallel to each other, with each simulation using one node. QM/MM calculations were performed by coupling together the software packages CHARMM (for classical MD simulations) and TURBOMOLE (for quantum chemical calculations) using a Python interface. The simulations were propagated for maximum 48 hours and were restarted from the latest checkpoint until the desired simulation time was reached. All in all, we used 10 M core-hours for this project, requiring 2 TB of storage in the WORK partition of SuperMUC. The work resulted in two publications, both of which were published in Nature Communications [3,4].

Ongoing Research / Outlook

In this project, SuperMUC provided us with the necessary computing resources to study the structure, function, and dynamics of the Hsp90 chaperone. This allowed us to investigate important structural changes in different variants and models of the Hsp90 dimer in atomic detail, and to probe the connection between protein conformation and catalytic activity. Future work will focus on characterizing the structural ensemble of




Figure 3: a) Close-up view of the dimerization interface of Hsp90 after 200 ns MD simulations with a methylated lysine (metK594). b) Methylation of the lysine induces conformational changes in the neighboring E590/R591 ion-pair in both chains of the Hsp90 dimer. The figure is adapted from ref. [4].

Hsp90 in different nucleotide states and with various point mutations of previously identified amino acid residues that stabilize the closed conformation of the dimer. In combination with biophysical experiments, we will investigate structural determinants involved in the large-scale conformational changes leading from the open to closed state of the Hsp90 dimer.

References and Links

- [1] <https://villekaila.com/>
- [2] Biebl MM, Buchner J. Cold Spring Harb. Perspect. Biol. 11, a034017 (2019). DOI: 10.1101/cshperspect.a034017.
- [3] Mader SL, Lopez A, Lawatscheck J, Luo Q, Rutz DA, Gamiz-Hernandez AP, Sattler M, Buchner J, Kaila VRI, Nat. Commun. 11, 1410 (2020). DOI: 10.1038/s41467-020-15050-0.
- [4] Rehn A, Lawatscheck J, Jokisch M-L, Mader SL, Luo Q, Tippel F, Blank B, Richter K, Lang K, Kaila VRI, Buchner J. Nat. Commun. 11, 1219 (2020). DOI: 10.1038/s41467-020-15048-8.
- [5] Phillips JC et al., J. Chem. Phys. 153, 044130 (2020). DOI: 10.1063/5.0014475.

Metadynamics reveals the secrets of G-protein coupled receptors

RESEARCH INSTITUTION

¹Computer-Chemistry Center, Friedrich-Alexander-University Erlangen-Nuremberg

PRINCIPAL INVESTIGATOR

Timothy Clark

RESEARCHERS

Heinrich Sticht², Peter W. Hildebrand³

PROJECT PARTNERS

²Bioinformatics, Institute for Biochemistry, Friedrich-Alexander-University Erlangen-Nuremberg

³Institute for Medical Physics and Biophysics, Leipzig University

SuperMUC Project ID: pr74su (Gauss Large Scale project)

Introduction

G-protein coupled receptors (GPCRs) are membrane proteins that transmit the effects of extracellular ligands to effect changes in the intracellular G-protein signaling system. Approximately 800 GPCRs are encoded in the human genome and approximately half of marketed drugs target GPCRs. It is therefore not surprising that GPCR research was recognized by the award of the 2012 Nobel Prize in Chemistry to Robert Lefkowitz and Brian Kobilka. Crystal structures of only 64 different GPCRs are currently available. Importantly, GPCRs can exist in active or inactive conformations and in binary complexes with ligands or intracellular binding partners (IBPs, G-proteins or β -arrestin) or in ternary complexes with both a ligand and an IBP. Crystal structures often deviate from the natural system: Proteins, especially membrane-bound ones, do not necessarily crystallize in their biologically active structures and the measures needed to obtain suitable GPCR crystals tend to increase the diversity between the natural environment and the crystal. It is within this context that molecular-dynamics simulations play a special role in GPCR research as a full-value complement to experimental studies.

The ternary complex model and experimental findings suggest that both an agonist ligand and a bound G-protein are necessary in order to activate GPCRs. It is therefore significant that the first molecular dynamics (MD) simulations of a ternary GPCR complex were reported only eight years ago. Such simulations are now commonplace and the comparison between binary ligand-receptor and ternary complexes has become a valuable tool in GPCR research.

Very long timescale MD simulations can be performed on specialized hardware but are less effective on more conventional massively parallel supercomputers because the simulations only scale up to a relatively limited number of CPUs or GPUs. Luckily, modern variations of metadynamics can make very effective use of massively parallel CPU-based hardware, as has been shown in this project [1].

Results and Methods

Methodologically, the project uses classical MD simulations over long time scales (microseconds) or enhanced-sampling techniques such as metadynamics. Protocols that are particularly well suited to SuperMUC-NG have been developed to make optimal use of the available hard- and software.

The project is divided into three subprojects as follows:

Class A GPCR activation

One goal of the project is to characterize and simulate receptor activation. An initial study of variants of the histamine 4 receptor (H4R), for which a wide range of basal activities is observed, showed that the observed activity correlates well with stabilities calculated using Constraint Network Analysis.

A general activation index was trained using sixteen multi-microsecond simulations of class A GPCRs in defined activation states. This index, denoted A^{100} , can also be calculated for X-ray crystal structures. A^{100} was calculated for 275 X-ray crystal structures for 114 different receptors as validation. The “active” and “inactive” classes are separated extremely well, whereas the histogram for the “intermediate” structures shows no indication of stable partially active structures, which oscillate between stable active and inactive states [2].

Ligand binding properties of the H_1 receptor

One aim of this project was to study ligand binding to the histamine H_1 receptor and the effect of sequence variation on this process. Unbiased MD simulations on a microsecond time scale were performed to deduce the histamine-binding site. The final histamine-binding pose in the orthosteric pocket (Fig. 2) is characterized by interactions with Asp107, Tyr108, Thr194, Asn198, Trp428, Tyr431, Phe432, and Phe435. The conformational stability of the complex structure obtained was subsequently confirmed in two μ s equilibrium MD simulations, and a metadynamics simulation proved that the detected binding site represents an energy minimum [3].




Figure 1: Calculated A^{100} values for the available GPCR X-ray structures. The bars of the histogram are colored according to the activity class assigned in the experimental work [2].




Figure 2: Structure of the Histamine- H_1 -Receptor (blue ribbon) with the modelled binding site of histamine (space-filled presentation). The lipids of the cellular membrane are depicted as grey/orange lines.




Figure 3: Proposed Process of GPCR-G Protein Complex Formation [4]: (A) Before formation of the $\beta 2AR$ -Gs protein complex, the $\alpha 5$ helix of Gs(GDP) is bent, with H387GAs and Y391GAs interacting with the Ras domain, whereas E392GAs and R389GAs are surface-exposed. (B and C) Straightening of the $\alpha 5$ helix (B) is required for an early interaction with the $\beta 2AR$, which involves E392GAs and R389GAs and may initiate GDP release (C). Intermediate state position of $\alpha 5$. (D) Relatively large conformational changes between GPCR and Gs are required to change from the intermediate state shown in (C) to a nucleotide-free complex (PDB: 3SN6). H387GAs and Y391GAs interact with $\beta 2AR$ in this state (enlarged window on top of D).

Receptor-G protein coupling specificity

In order to investigate how receptors can activate different intracellular downstream signaling proteins such as heterotrimeric G proteins ($G\alpha\beta\gamma$; families G_i , G_s , G_q , $G_{12/13}$) we investigated $G_s(GDP)$ spontaneous binding to the $\beta 2AR$. Despite elucidation of various GPCR-G protein complexes, the structural basis of coupling specificity is still unclear. This may be because all structures have been traced in a nucleotide-free state, under non-physiological conditions i.e. in the absence of nucleotides. We observe spontaneous binding of $G_s(GDP)$ to the activated $\beta 2$ -adrenoceptor ($\beta 2AR$) in agreement with an X-ray structure of a $\beta 2AR$ construct where only the far C-terminus of the α -subunit of Gs was co-crystallized [4]. Our work ($\sim 100 \mu s$ MD) provides a first glimpse on the initial interactions of receptors with G proteins, revealing novel recognition sites controlling coupling specificity.

Ongoing Research / Outlook

New simulation protocols are being developed to investigate the above questions in more detail in a standard fashion. The project is ongoing and follow-up proposals are planned.

References and Links

- [1] P. Ibrahim and T. Clark, *Curr. Opin. Struct. Biol.*, 2019, 55, 129-137. <https://doi.org/10.1016/j.sbi.2019.04.002>
- [2] P. Ibrahim, D. Wifling and T. Clark, *J. Chem. Inf. Comput.* 2019, 59, 3938-3945. <https://doi.org/10.1021/acs.jcim.9b00604>
- [3] C.A. Söldner, A.H.C.Horn and H.Sticht, *J. Mol. Model.* 2018, 24, 346. <https://doi.org/10.1007/s00894-018-3873-7>
- [4] X. Liu et al., P.W. Hildebrand and B. K. Kobilka, *Cell* 2019, 17, 1243-1251.e12. <https://pubmed.ncbi.nlm.nih.gov/31080070/>

Blood flow simulations

RESEARCH INSTITUTION

Universität Bayreuth

PRINCIPAL INVESTIGATOR

Stephan Gekle

RESEARCHERS

Christian Bächer, Achim Guckenbergler

PROJECT PARTNERS

—

SuperMUC Project ID: pr84mo

6

Introduction

The behavior of biological cells in hydrodynamic flows can be incredibly complex. As experimental investigations in this field are notorious difficult, especially due to the restriction of standard microscopy methods to 2D, high-performance computer simulations offer unique insights.

Results and Methods

We highlight here two key papers resulting from the project:

Bistability of red blood cells [2]

Red blood cells flowing through capillaries assume a wide variety of different shapes owing to their high deformability. Predicting the realized shapes is a complex field as they are determined by the intricate interplay between the flow conditions and the membrane mechanics. In this work we construct the shape phase diagram of a single red blood cell with a physiological viscosity ratio flowing in a microchannel. We use both experimental *in vitro* measurements as well as 3D numerical simulations to complement the respective other one. Numerically, we have easy control over the initial starting configuration and natural access to the full 3D shape. With this information we obtain the phase diagram as a function of initial position, starting shape and cell velocity. Experimentally, we measure the occurrence frequency of the different shapes as a function of the cell velocity to construct the experimental diagram which is in good agreement with the numerical observations. Two different major shapes are found, namely croissants and slippers. Notably, both shapes show coexistence at low (1 mm/s) and high velocities (43 mm/s) while in-between only croissants are stable. This pronounced bistability indicates that RBC shapes are not only determined by system parameters such as flow velocity or channel size, but also strongly depend on the initial conditions.

Flow-accelerated platelet biogenesis is due to an elasto-hydrodynamic instability [3]

Blood platelets are formed by fragmentation of long membrane extensions from bone marrow

megakaryocytes in the blood flow. Using lattice-Boltzmann/immersed boundary simulations we propose a biological Rayleigh-Plateau instability as the biophysical mechanism behind this fragmentation process. This instability is akin to the surface tension induced break-up of a liquid jet but is driven by active cortical processes including actomyosin contractility and microtubule sliding. Our fully three-dimensional simulations highlight the crucial role of actomyosin contractility, which is required to trigger the instability, and illustrate how the wavelength of the instability determines the size of the final platelets. The elasto-hydrodynamic origin of the fragmentation explains the strong acceleration of platelet biogenesis in the presence of an external flow, which we observe in agreement with experiments.

Our simulations then allow us to disentangle the influence of specific flow conditions: while a homogeneous flow with uniform velocity leads to the strongest acceleration, a shear flow with a linear velocity gradient can cause fusion events of two developing platelet-sized swellings during fragmentation. A fusion event may lead to the release of larger structures which are observable as preplatelets in experiments.

Together, our findings strongly indicate a mainly physical origin of fragmentation and regulation of platelet size in flow-accelerated platelet biogenesis.

Ongoing Research / Outlook

In the future, we plan to extend the active shell model used in [3,4] towards viscous membranes.

References and Links

- [1] biofluid.physik.uni-bayreuth.de.
- [2] Guckenbergler, A., Kihm, A., John, T., Wagner, C., & Gekle, S., *Soft Matter*, 14 (11), 2032–2043 (2018). <http://doi.org/10.1039/C7SM02272G>.
- [3] Bächer, C., Bender, M., & Gekle, S., *PNAS*, 12, 202002985 (2020). <http://doi.org/10.1073/pnas.2002985117>.
- [4] Bächer, C., & Gekle, S., *Phys. Rev. E*, 99 (6), 062418 (2019). <http://doi.org/10.1103/PhysRevE.99.062418>.




Figure 1: Setup of a red blood cell in a channel, taken from [2].




Figure 2: Phase diagram of a red blood cell in channel flow, taken from [2].





Figure 3: Simulation of a megakaryocyte protrusion fragmenting into platelets, taken from [3].

Appendices



SuperMUC Phase 1 and Phase 2: System Description

SuperMUC was the leadership class supercomputer at the Leibniz-Rechenzentrum (Leibniz Supercomputing Centre, LRZ) in Garching near Munich (the MUC suffix refers to the Munich airport code). With more than 241,000 cores and a peak performance of the two installation phases of more than 6.8 petaFLOP/s ($= 6.8 \times 10^{15}$ floating point operations per second), it was one of the fastest supercomputers in the world. SuperMUC strengthened the position of Germany's Gauss Centre for Supercomputing (GCS) 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 Phase 1 started user operation in July 2012, and the Phase 2 extension in June 2015. The final shutdown for SuperMUC Phase 1 was in January 2019, and for SuperMUC Phase 2 one year later, in January 2020.

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 GBytes (Phase 1) and 64 GBytes (Phase 2), respectively, and a smaller number of fat compute nodes with a main memory of 256 GBytes. The network interconnect between the nodes allowed excellent scaling of parallel applications up to the level of more than 100,000 tasks. SuperMUC consisted 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 com-

pute nodes within an individual Island were 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 enabled a bi-directional bi-section bandwidth ratio of 4:1 (intra-Island / inter-Island). An additional system segment, designed to evaluate future architecture and programming concepts, was called SuperMIC. It contained 32 Intel Ivy Bridge-EP nodes with two Intel Xeon Phi "Knights Corner" accelerator cards each. See the technical data (p. 276) for more details.

SuperMUC Phase1 and Phase2 were loosely coupled through the General Parallel File System (GPFS) and Network Attached Storage (NAS) File systems, jointly used by Phase 1 and Phase 2. Both phases were operated independently, but supplied the same programming environment.

SuperMUC used a new, revolutionary form of warm water cooling developed by IBM. Active components like processors and memory were directly cooled with water that could have an inlet temperature of up to 40 degrees Celsius. This High Temperature Liquid Cooling together with very innovative system software reduced the energy consumption of the system by up to 40%. In addition, LRZ buildings were heated re-using this energy.

Permanent storage for data and programs was provided by a 16-node NAS cluster from NetApp. This primary cluster had a capacity of 3.5 Petabytes and had demonstrated an aggregated throughput of more




Figure 1: SuperMUC Phase 1 on the left and SuperMUC Phase 2 on the right, in the server room.

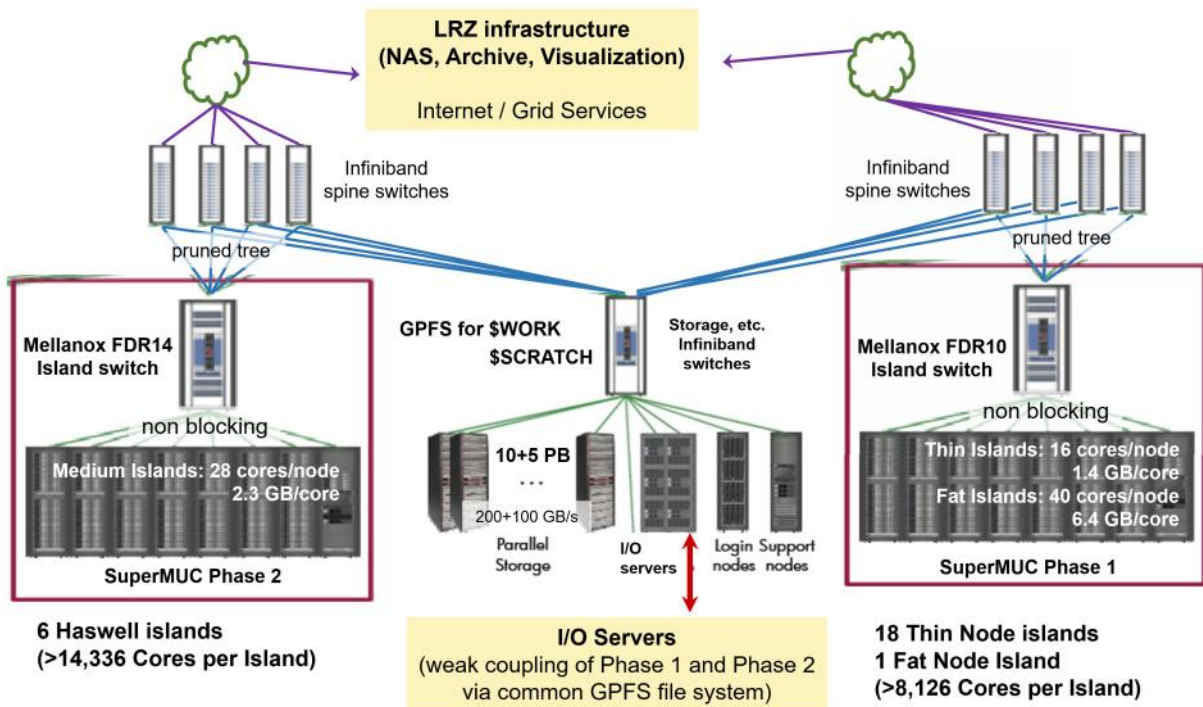


Figure 2: Schematic view of SuperMUC Phase 1 and Phase 2.




Figure 3: SuperMUC Phase 2. Rear of racks with warm water cooling.

than 12 GB/s using NFSv3. NetApp's Ontap 8 Cluster-Mode provided a single namespace for several hundred project volumes on the system. Users could access multiple snapshots of data in their home directories. For additional redundancy, data were regularly replicated to a separate 4-node NetApp cluster with another 3.5 PB of storage for recovery purposes. Replication used Snapmirror-technology and ran 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

was available. Disk storage subsystems were built by DDN (for Phase 1) and Lenovo (for Phase 2). The storage hardware consisted 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 were based on IBM Spectrum Protect, providing more than 30 Petabytes of capacity to the users of SuperMUC. Digital long-term archives help to preserve simulation results. User archives were also mirrored to a remote site for potential disaster recovery.

Technical data

Installation Phase	Phase 1			Phase 2
Installation Date	2011	2012	2013	2015
End of operation	12/2019	12/2019	12/2019	1/2020
Islandtype	Fat Nodes	Thin Nodes	Many Cores Nodes	Haswell Nodes
System	BladeCenter HX5	IBM System x iDataPlex dx360M4	IBM System x iDataPlex dx360M4	Lenovo NeXtScale nx360M5 WCT
Processor Type	Westmere EX Xeon E7-4870	Sandy Bridge EP Xeon E5-2680	Ivy Bridge EP and Xeon Phi 5110P	Haswell EP Xeon E5-2697 v3
Nominal Frequency [GHz]	2.4	2.7	1.05	2.6
Performance per core	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 multiply-adds every cycle using 4 threads	16 DP Flops/cycle = 41.6 DP Flops/s two 4-wide fused multiply-adds
Total Number of nodes	205	9,216	32	3,072
Total Number of cores	8,200	147,456	3,840 (Phi)	86,016
Total Peak Performance [PFlop/s]	0.078	3.2	0.064 (Phi)	3.58
Total LinPack Performance [PFlop/s]	0.065	2.897	N/A	2.814
Total size of memory [TByte]	52	288	2.56	194
Total Number of Islands	1	18	1	6
Typical Power Consumption [MW]	< 2.3			~1.1
Components				
Nodes per Island	205	512	32	512
Processors per Node	4	2	2 Ivy Bridge EP + 2 Phi 5110P	2
Cores per Processor	10	8	8 (Ivy Bridge EP) + 60 (Phi)	14
Cores per Node	40	16	16 (host) + 120 (Phi)	28
Logical CPUs per Node (Hyperthreading)	80	32	32 (host) + 480 (Phi)	56
Login Servers for users	2	7	1	5
Memory and Caches				
Memory per Core (for applications) [GByte]	6.4 (~6.0)	2 (~1.5)	4 (host) + 2 x 0.13 (Phi)	2.3 (2.1)
Size of shared Memory per node [GByte]	256	32	64 (host) + 2 x 8 (Phi)	64
Bandwidth to Memory per node [GByte/s]	136.4	102.4	Phi: 384	137
Interconnect				
Technology	Infiniband QDR	Infiniband FDR10	Infiniband FDR10	Infiniband FDR14
Intra-Island Topology	non-blocking Tree			non-blocking Tree
Inter-Island Topology	Pruned Tree 4:1		N/A	Pruned Tree 4:1
Bisection bandwidth of Interconnect	12.5 TByte/s			5.1 TByte/s
Storage and System Software				
Operating System / Batch System	Suse Linux Enterprise Server (SLES) / IBM Loadleveler			
Parallel Filesystem	15 PByte, IBM GPFS for SCRATCH and WORK 250 GByte/s aggregated bandwidth			
File System for HOME	3,5 PByte, NetApp NAS, 12 GByte/s aggregated bandwidth			
Archive and Backup	> 30 PByte, IBM TSM			
System Management / Monitoring	xCat from IBM / Icinga, Splunk			




Figure 4: 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 5: One of the Infiniband switches.




Figure 6: Several racks of SuperMUC Phase 2.

SuperMUC-NG: System Description

SuperMUC-NG (Next Generation) completed its installation in September 2018, and started general user operation in August 2019. It replaced the predecessor system “SuperMUC” as the leadership class supercomputer at the Leibniz-Rechenzentrum (Leibniz Supercomputing Centre, LRZ) in Garching near Munich (the MUC suffix refers to the Munich airport code). With more than 311,040 cores, a main memory capacity of 719 TByte and a peak performance of 26.9 peta-FLOPS ($= 26.9 \times 10^{15}$ floating point operations per second), it is among the fastest supercomputers in the world. SuperMUC-NG entered the TOP500 list of supercomputers in the world at rank 8 in November 2018.

SuperMUC-NG is tightly integrated into the European high-performance computing ecosystem. LRZ is a European Centre for Supercomputing and a Tier-0 Centre for the Partnership for Advanced Computing in Europe (PRACE).

All compute nodes are equipped with Intel Xeon Skylake Platinum 8174 processors. The internal interconnect is a fast OmniPath network with a link bandwidth of 100 Gbit/s. The compute nodes are bundled into 8 domains (islands). Within one island, the OmniPath network topology is a ‘fat tree’ for highly efficient communication. The OmniPath connection between the islands is pruned with a pruning factor of 1:3.75.

In addition, 65 nodes of the LRZ Compute Cloud were procured with the SuperMUC-NG system to provide services to users that cannot be offered otherwise.

The LRZ Compute Cloud currently consists of 82 Cloud Compute Nodes with 40 cores and 192 GB memory, 32 accelerated Cloud Compute Nodes with 40 cores, 768 GB memory and two Nvidia Tesla V100 GPUs with 16 GB high bandwidth memory, and 1 Cloud Huge Node with 192 cores and 6 TB memory.

SuperMUC-NG integrates Lenovo DSS-G for IBM Spectrum Scale (aka GPFS) as building blocks for the storage WORK (34 PiByte, 300 GB/s aggregated bandwidth) and SCRATCH (16 PiByte, 200 GB/s aggregated bandwidth). The LRZ Data Science Storage offers additional 20 PiByte (with 70 GB/s aggregated bandwidth) for long-term storage of project data that can also be shared with the science community throughout the world. The HOME filesystem has a capacity of 256 TiByte.

LRZ's tape backup and archive systems are operated via the Spectrum Protect (formerly Tivoli Storage Manager) software from IBM, providing 260 Petabytes of capacity to the users of SuperMUC-NG. Digital long-term archives help to preserve simulation results. User archives are also mirrored to a remote site for potential disaster recovery.

Like its predecessor, SuperMUC-NG is cooled with hot water with an inlet temperature of up to 50 degrees Celsius and an outlet temperature of up to 56 degrees Celsius. These high temperatures allow to use the waste heat to generate cold water using a large set of adsorption chillers produced by the company Fahrenheit. The adsorption chillers are based on the physical



Figure 1: SuperMUC-NG.

principle of using solid materials of enhanced water vapor adsorption capacity for cooling via the evaporation of water. The cold water from the chillers is used to remove the waste heat of air-cooled components such as network switches, power supplies and storage. In addition, LRZ's buildings are heated re-using SuperMUC-NGs heat dump.

SuperMUC-NG is available to all German and European researchers to expand the frontiers of science and engineering. 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 45 million core-hours. Smaller proposals by German scientists can be submitted throughout the year directly to LRZ.




Figure 2: Dual node chassis of SuperMUC-NG.

Technical data

Compute Nodes	Thin Nodes	Fat Nodes	Total (Thin + Fat)
Processor Type	Intel Skylake Xeon Platinum 8174	Intel Skylake Xeon Platinum 8174	Intel Skylake Xeon Platinum 8174
Cores per Node	48	48	48
Memory per Node [GByte]	96	768	N/A
Number of Nodes	6,336	144	6,480
Number of Cores	304,128	6,912	311,040
Peak Performance @ nominal [PFlop/s]	26.3	0.6	26.9
Linpack [PFlop/s]	–	–	19.476
Memory [TByte]	608	111	719
Number of Islands	8	1	9
Nodes per Island	792	144	N/A
Filesystems			
High Performance Parallel Filesystem	50 PiB @ 500 GB/s		
Data Science Storage	20 PiB @ 70 GB/s		
Home Filesystem	256 TiB		
Infrastructure			
Cooling	Direct warm water cooling		
Waste Heat Reuse	For producing cold water with adsorption coolers		
Software			
Operating System	Suse Linux Enterprise Server (SLES)		
Batch Scheduling System	SLURM		
High Performance Parallel Filesystem	IBM Spectrum Scale (GPFS)		
Programming Environment	Intel Parallel Studio XE, GNU compilers		
Message Passing	Intel MPI, (OpenMPI)		

More than three decades of super-computing at

LRZ – peak performance and machine balance

The following is a retrospective of high performance computing at the Leibniz Supercomputing Centre, seen through the eyes of our colleague, Dr. Matthias Brehm, who joined LRZ in 1986, working with HPC systems throughout most of his career. He was leading the LRZ HPC user support since 2000, until his retirement in 2020.

Gigascale, Vector computers

At the end of 1988, the first supercomputer was installed at LRZ, a Cray X-MP/24, with 2 processors and 4 megawords (32 Mbyte) of main memory. This interim system had a peak performance of 400 MFlop/s. One year later, it was exchanged for a CRAY Y-MP/432, i.e., a system with 4 processors and 32 megawords (256 Mbyte) of main memory and a peak performance of 1,333 MFlop/s. Finally, the system was upgraded to 8 Processors (2,666 MFlop/s) and 1 GByte.

At that time, the performance advantage of the CRAY Y-MP over the mainframe-like Cyber systems at LRZ was more than one magnitude. It was not only this big jump in compute performance, but also the unprecedented memory bandwidth which made it relatively easy to be used. One processor of the system could perform two loads and one store operation from memory per cycle, and could perform an addition and multiplication in parallel. Neglecting side effects, the system was able to sustain the following loop kernels with near peak performance.

$$A(i) = B(i) + \alpha * C(i) \quad (\text{Eq.1})$$

$$A(i) = A(i) + \alpha * C(i) \quad (\text{Eq.2})$$

Why is this so important? The general triad (Eq.1) is representative for many engineering applications whereas the linked triad (Eq.2) is typical for problems with matrix multiplications. The machine balance for the Cray Y-MP can be calculated as

$$\text{memory bandwidth / peak performance} = \\ (3 \text{ words/cy}) / (2 \text{ Flops/cy}) = 12 \text{ Byte / Flop}$$

The CRAY T90 (4 Procs, 7.2 GFlop/s, 1 GByte) (installed in 1996) had an even slightly better balance of 13.3. It was never reached by any other system at LRZ. The good machine balance and well-performing auto-vectorizing compilers made it relatively easy to be programmed. Background monitoring by hardware performance counters showed a relatively high percentage compared to peak performance what could be gained in everyday operation (see Table 1).

The Cray vector systems had first approaches for parallel processing: macro-tasking (parallelism at the subroutine level with explicit calls), micro-tasking (parallelism at the loop level), and auto-tasking (automatic multitasking). The latter two were directive driven.

Microprocessor, First parallel computers

Microprocessors had triggered the boom in personal computers. The exponential growth in transistor density and compute performance (Moore's law), cheaper and larger memory technology led to the so-called "Attack of the Killer Micros" in the early 1990ies. The small memory on the vector computers, limited possibilities to improve a fixed sized problem by vectorization or parallelization (Amdahl's law) were replaced by expectation for scalable algorithms which can make use of cheap processors as well as of more and cheap memory to solve ever larger problems (Gustafson's law). Instead of costly connections from CPU to memory, fast and low latency caches mitigated the effect of missing memory bandwidth.

The first parallel system KSR1 was installed in 1993 (1.3 GFlop/s, 1 GByte). It was a virtual shared memory system, where data that were not present in the local cache of a node were automatically transferred from another node in the network. Therefore, it had a Non-uniform Memory Access (NUMA). Parallelization was performed in the shared memory by directives similar as today with OpenMP

```
C*KSR* tile (ij) n
DO I=1,Nx ...
```

Two message passing libraries (TCGMSG and PVM) were also available. While KSR1 was running quite stable, the target system for the second phase (KSR2, 110 nodes, 8.8 Gflop/s) had too many technical problems and was not accepted by LRZ. Also, the system vendor, Kendall Square Research, collapsed due to accounting irregularities.

Since LRZ had no financial losses, an IBM SP2 (77 nodes, 21 Glop/s, 17 GB) system could installed as replacement in 1996. The SP2 was equipped with RS/6000 Power2 RISC processors. The first MPI standard was released in June 1994, but was not available from the start of operation of the system, only one year later. First versions of OpenMP followed in 1998.

The SP2 paved the way for parallel computing at LRZ. However, the extensive task of the users was now to

convert their sequential algorithms to message passing and to make efficient use of the cache hierarchy. Also, a first embodiment of parallel file systems was available. PIOFS was the ancestor of GPFS, which was later used on LRZ systems.

A remarkable breakthrough for massively parallel computing was the victory of the IBM SP2 “Deep Blue” over the chess world champion Kasparow in 1996.

Terascale, Combining vector and parallel computing

Fortunately, it was possible to combine the “good” characteristics of a vector computer with the scalability of parallel systems. A Fujitsu VPP 700 with 52 nodes (114 Glop/s, 104 GB) was installed in 1998. However, it was already clear that the machine balance of 8 Byte/Flop was at the limits of what is technically practicable.

The first system that was available for researchers from all over Germany was the Hitachi SR8000 (2 TFlop/s, 1.4 TByte for Phase 2). It was comprised of 168 compute nodes, each with eight compute processors plus one service processor, totaling 1,344 compute processors. At the time of installation, the SR8000 was the fastest civil computer in the world, ranked at No. 4 of the TOP500 list. It was the first computer in Europe to reach more than one TFlop/s.



Figure 1: Hitach SR8000 in the old computer room of LRZ.

The S8000 provided a unique hybrid approach. The processors originated from an IBM PowerPC design with additional hardware support for fast thread synchronization. It had a relatively high memory bandwidth, which was used by specific extensions: 160 floating point registers and up to 16 outstanding prefetch instructions were available to prefetch cache lines in advance of the actual computation to the L1 cache (128 kByte). No other caches were present in the node. It also had preload instructions for loading single data in advance. A sliding window technique was used to breakdown the prefetches from the physical to the logical RISC registers. The software pipelining by the compiler for these features was called pseudo-vectorization. It was complemented by the possibility of auto-parallelization for shared memory parallelization with eight threads, or by the OpenMP-like distribution of large loops over the node-proces-

sors, called Cooperative Microprocessors in a single Address Space (COMPAS). Parallelization across nodes could also be done by directly assigning an MPI process to a single processor (MPP mode) or by assigning the whole node to an MPI process (COMPAS node)

Already during the installation phase of the SR8000, it became evident that the LRZ building in downtown Munich had no perspective for containing even larger systems. A new building in Garching then provided excellent infrastructure and enough floor space, power, cooling, and energy efficiency. The first “Computer Cube” from 2006 was soon extended to a “Double Cube” in 2011.

Large Shared Memory, but the Free Lunch was over

With the general transition to 64-bit architectures Intel tried to establish a new processor line, the Itanium. Very large memory could now be addressed and the new programming paradigm “Explicit parallel instruction computing” should make it possible to execute several instructions in parallel. The parallel instructions were arranged by the compiler, rather than using complex hardware for instruction scheduling. Already the first version of Itanium was capable of execution six instructions per cycle. One big advantage, particularly for HPC applications, was the predictability regarding when data from memory or cache would be available in the execution units [2].

LRZ installed an SGI ALTIX system with the Itanium architecture (IA-64) in its new building in Garching. The system was again ranked among the TOP10 system in the world (62.3 TFlop/s, 39 TByte). The increase in peak performance was very impressive, while the machine balance dropped to 0.56 Byte/Flop. The SGI Altix was a cache coherent non-uniform memory access (ccNUMA) system where the whole main memory of the system could be reached from each processor. In principle, now was working what was promised by the KSR system 13 years ago. The dominant programming model kept to be message passing. Apart from some difficulties with the parallel file systems, the SGI Altix provided many years of good and stable service for the LRZ users. On the software side, some unification of software environment for Intel compilers and libraries could be achieved with the Linux-Cluster of LRZ, even though IA-64 and X86 were not binary compatible.



Figure 2: SGI Altix in the new computer room of LRZ.

The Itanium was a disaster for Intel [2]. Production of successor chips was delayed, AMD successfully approached the market with 64-bit extensions for the x86 line, SGI filed for bankruptcy protection, and, perhaps most importantly, Microsoft and many other providers were not willing to support Itanium for their software (legacy drop). Finally, Intel dismissed the Itanium line and concentrated on its x86-64 Xeons.

A substantial fraction of the performance increase came from the rise of processor frequency of microprocessors. But around the year 2004, this trend stopped at approximately 3.5 GHz. The slogan at this time was: “the free lunch is over” [3]. Any increase of frequency of a single processor would exponentially increase the power (and heat), and would cause severe technical challenges. Instead of increasing the processor frequency, it was more efficient to use two or more cores to gain the same peak performance. The increase in main memory continued to be much slower (“memory wall”), leading to larger caches and deeper and non-uniform memory hierarchies. More intra-node and inter-node parallelism had to be used by the applications. Thus, the physical size of the systems again grew.

Petascale, Warm-Water Cooling, Multi-Cores

Energy consumption is an important fraction of the total cost of ownership for HPC systems. The cost for cooling may sum up to about 30% of the energy bill. Therefore, the LRZ’s building extension in 2011 also featured an additional cooling loop for directly cooling the most power-hungry components with warm water of up to 55°C. This temperature enabled free cooling without the use of additional compressor chillers. LRZ was a pioneer for this technology which is now commonly used all over the world.

In 2012, SuperMUC went into operation. Again, a very substantial increase in peak performance could be reached. With the extension of Super-MUC Phase 2, the system had more than 241,000 cores and a combined peak performance of the two installation phases of more than 6.8 PFlop/s (482 TByte memory) [4]. The machine balance lowered to 0.26 Byte/Flop for Phase 1 and 0.12 Byte/Flop for Phase 2, respectively. A multicore node of the Phase 1 consisted of sixteen Sandy Bridge cores arranged in two sockets, therefore




Figure 3: IBM iDataPlex “SuperMUC Phase 1” in the again extended computer room.

the memory access was non-uniform. For Phase 2, Haswell nodes with 28 cores in two sockets were used (more details on SuperMUC Phase 1 and Phase 2 can be found in the Appendix).

The high peak performance of SuperMUC came from the Advanced Vector Extensions (AVX, AVX2) of the x86 instruction set. AVX introduced a three-operand SIMD instruction format. AVX2 on Haswell could support vector instructions like $a = a + b \cdot c$ on four times three double precision operands simultaneously per cycle. But Data needed to come from the L1 Cache.

With SuperMUC, a complete x86-64 software stack and programming environment, ranging from desktop, over cloud nodes and the Linux-Cluster to the highest end was established, offering an enormous variety of open-source and commercial software. This attracted many new users.

The latest supercomputer at LRZ is SuperMUC-NG, its architecture is described in more detail in the Appendix. The AVX-512 instruction set of the Skylake processors increased the width of the vector register file from 256 to 512 bits, thus doubling the peak vector performance. Already during the procurement of the system, it was clear that conventional microprocessors reached many technological limits, particularly power density and energy consumption. Further substantial increase in performance can only be achieved by the use of accelerator hardware like GPUs. Artificial intelligence and machine learning further drive the application of GPUs. Perhaps for AI a similar milestone as the historic chess match in 1996 was the defeat of a human player by Google DeepMind in a Go tournament. AI, ML and Big Data will shape HPC and add new scientific cases.

Outlook

The next step in HPC, “Exascale” can only be reached with accelerators. Like the transition to parallel programming in the 1990s, accelerated computing needs severe and disruptive changes in software and in the programming models. The performance of SuperMUC-NG Phase 2 will come primarily from accelerated nodes. Also, dedicated file systems will be available for the requirements of AI/ML.

Over the years, we have seen that the performance increase in HPC did not come for free. It goes hand in hand with the use of specialized hardware and features: vectorization, SIMD, parallelization, shared memory, MPI and OpenMP, optimum use of caches, order of data, mixed precision programming, multi-cores, accelerators, off-loading, hyperthreading, use of appropriate files systems and (very!) much more. Any failure in using one of these features may result in a performance drop, sometimes of more than one order of magnitude. Even with newly available high bandwidth memory, the memory wall will remain.

Year	System	Peak (GFlop/s)	Memory (GByte)	Machine Balance (Byte/Flop)	Everyday Percentage of Peak
1990	CRAY YMP/8	2.7	0.5	12.0	~33%
1996	IBM SP2	21	17	2.4	~8%
1998	Fujitsu VPP	114	104	8.0	~22%
2002	Hitachi SR8000 Phase 2	2,000	1,400	2.7	~12%
2008	SGI Altix Phase 2	62,300	39,000	0.6	~6%
2016	IBM SuperMUC Phase 2	6,800,000	534,000	0.1	~5%
2020	Intel/Lenovo SuperMUC-NG	26,900,000	719,000	0.05	~4% (preliminary)


Table 1: Key characteristics of LRZ HPC Systems . The last column "Everyday Percentage of Peak" describes the typical average performance of the application-mix at LRZ, measured by hardware counter monitoring over several years. For specific applications and/or a lot of effort and hand-optimization, a substantial higher fraction of peak is still possible on modern machines, which was demonstrated with several codes on SuperMUC and SuperMUC-NG.

The key performance data of the LRZ high-end systems are summarized in Table 1. It is not surprising that the ratio of peak performance to what could be gained in "every day operation" lowered over the years. The data were taken from LRZ background monitoring. Of course, this is heavily dependent on the actual application mix and large differences exist between scientific areas. We only want to highlight the trend, not the actual values. And please, do not mix these data with the often cited "sustained performance" of a system.

Over three decades we have seen an increase in peak performance of ten million. Even when considering the drop of the achieved performance, we still see roughly one million. However, only those users who are willing to consciously invest in adopting their applications to the ever-changing hardware will remain scientifically competitive and will stay 3-5 years ahead of others, who do not so. As in the past, LRZ will provide courses and workshops for users to manage the challenges.

References and Links

- [1] https://link.springer.com/chapter/10.1007/3-540-44520-X_189
- [2] <https://www.elektronikpraxis.vogel.de/epic-fail-20-jahre-intel-itanium-a-855870/>
- [3] <http://www.gotw.ca/publications/concurrency-ddj.htm>
- [4] <https://doku.lrz.de/display/PUBLIC/Decommissioned+SuperMUC>



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 in the time frame July 2018 – June 2020. The 115 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 was in operation 2012-2019, **SuperMUC Phase 2** from 2015 to 2020, and **SuperMUC-NG** (above) started a friendly user operation in June 2020, with general user operation starting in August 2020. A detailed system description of the three systems can be found in the appendix.

The articles provide an overview of the broad range of applications that use high performance computing to solve 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 978-3-9816675-3-0
<https://www.lrz.de/hpcbooks>