P. Bastian • D. Kranzlmüller • H. Brüchle • G. Mathias
EDITORS

High Performance Computing

in Science and Engineering Garching/Munich 2024

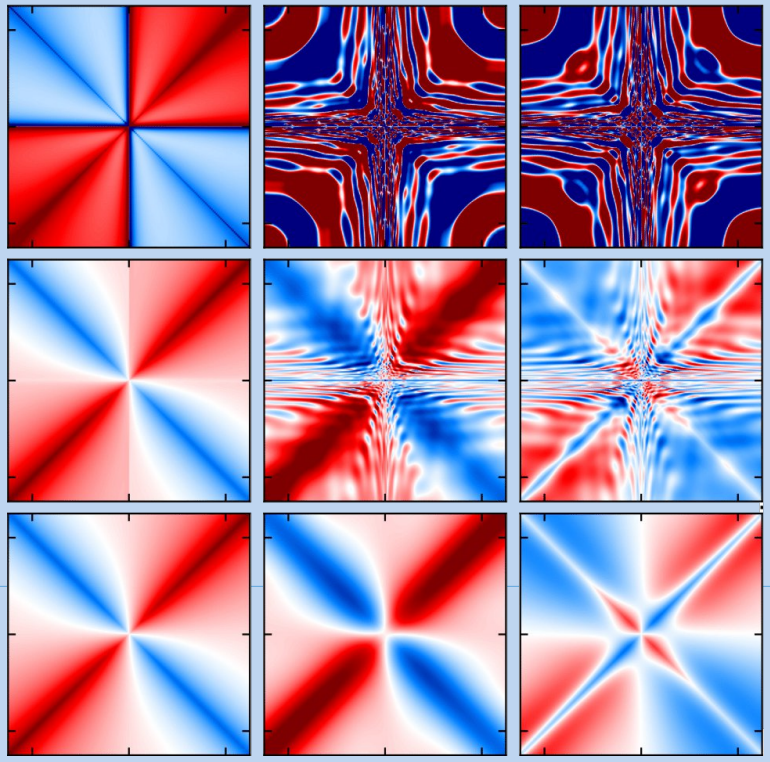

Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

Titelbild entnommen aus dem Artikel von Jan von Delft. Das Bild zeigt eine Darstellung von Vierpunkt-Vertexfunktionen von Quantenverunreinigungsmodellen. Die von seiner Arbeitsgruppe entwickelte numerische Mehrpunkt-Renormierungsgruppe kann viel niedrigere Temperaturen erreichen als die modernsten Quanten-Monte-Carlo-Methoden. Diese Ergebnisse helfen, Quantenphänomene in stark korrelierten Elektronensystemen zu verstehen, wie etwa Hochtemperatur-Supraleitung und Quantenkritikalität. Weitere Informationen im Artikel ab Seite 76.

Impressum:

Leibniz-Rechenzentrum
der Bayerische Akademie der Wissenschaften
Boltzmannstraße 1, D-85748 Garching bei München
lrzpost@lrz.de, www.lrz.de

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

Bezugsadresse:

Leibniz-Rechenzentrum (LRZ)
Boltzmannstraße 1, D-85748 Garching bei München
https://www.lrz.de/hpcbooks

ISBN 978-3-9816675-6-1

P. Bastian • D. Kranzlmüller • H. Brüchle • G. Mathias
EDITORS

High Performance Computing

in Science and Engineering Garching/Munich 2024

Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

Table of Contents

Preface

Chapter 01 – Astrophysics

Chapter 02 – Chemistry and Material Sciences

Chapter 03 – Computational Fluid Dynamics and Engineering

Chapter 04 – Earth, Climate and Environmental Sciences

Chapter 05 – Elementary Particle Physics

Chapter 06 - Life Sciences

Chapter 07 - Plasma Physics

Chapter 08 - Math, Computer Sciences

Appendices

The high-performance computer SuperMUC-NG at the Leibniz Supercomputing Centre is based on Lenovo ThinkSystem SD650 DWC compute nodes and equipped with Intel Skylake Xeon Platinum 8174 processors. The machine hosts 6,336 thin compute nodes, each with 48 cores and 96 GB memory, and 144 fat compute nodes, each with 48 cores and 768 GB memory per node. User operation started in August 2019. (Picture: Axel König).

The Editors

Prof. Dr. Peter Bastian
Chairman of the SuperMUC-NG Steering Committee, Interdisciplinary Center for Scientific Computing, Heidelberg University

Prof. Dr. Dieter Kranzlmüller
Chairman of the Board of Directors Leibniz Supercomputing Centre, Bavarian Academy of Sciences and Humanities

Dr. Helmut Brüchle
Editor-in-Chief Computational X Support, Leibniz Supercomputing Centre, Bavarian Academy of Sciences and Humanities

Dr. habil. Gerald Mathias
Head of Computational X Support Leibniz Supercomputing Centre, Bavarian Academy of Sciences and Humanities

Preface

This book presents a comprehensive compilation of the latest results from simulations performed on the high-performance computer SuperMUC-NG, operated by the Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences and Humanities, a member of the Gauss Centre for Supercomputing (GCS). It covers the years of service between January 2022 and December 2023. SuperMUC-NG remains to be a highly valued resource by its users, enabling and accelerating a broad range of scientific research, as this collection of articles shows.

The book-cover features a representation of fourpoint vertex functions of quantum impurity models. The multipoint numerical renormalization group (mpNRG) developed by Jan van Delft's group can access much lower temperatures than the state-of-the-art quantum Monte Carlo (QMC) methods. These results help to understand quantum phenomena in strongly correlated electron systems, such as high-temperature superconductivity and quantum criticality (p. 76).

Lattice Quantum Chromodynamics test the standard model of particle physics to its limits. Szabo and co-workers reached ever higher precision in the predicted values of certain decay constants (p. 232) and contributed (p. 234) to the ongoing discussions about the myon magnetic moment, which was voted among the '10 Breakthroughs of the Year' by the Science Magazine.

Hydrogen is seen as a key energy carrier for a sustainable economy. Heinz Pitsch and his group use direct numerical simulations to study hydrogen combustion. The transformation of energy systems towards hydrogen poses new challenges that arise from the specific molecular transport and combustion properties of hydrogen, which differ significantly from conventional fossil fuels (p. 168 and p. 176).

Weather and climate simulations are among the most common applications of high-performance computing that make news headlines. Juan Pedro Mellado and co-workers investigated, how stratocumuli change with climate warming, using direct numerical simulations. These low-level clouds at the top of the atmospheric boundary layer extend over thousands of kilometres in the eastern boundaries of the subtropical oceans. They reflect more incoming solar radiation than the underlying surface of the ocean, while they emit similarly in the long-wave range. This combination of large coverage and net cooling effect substantially affects the Earth's radiative budget (p. 200).

The search for new compounds is a fundamental challenge in chemistry and material science. The group of Miguel Marques combines density functional theory and machine learning to speed up the process by orders of magnitude. Their neural network architecture is specialized for predicting new thermodynamically stable materials. They predicted the stability of around 15 billion compounds and investigated the most promising compounds with highly accurate ab-initio calculations. Their results have been added to a materials database (p. 94), which can be accessed publicly by experimentalists for further research.

All the scientific groups that contributed to this book rely heavily on the availability of leadership-class HPC machines such as SuperMUC-NG. We therefore 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 bei München, November 2024

Peter Bastian
Dieter Kranzlmüller
Helmut Brüchle
Gerald Mathias

Astrophysics

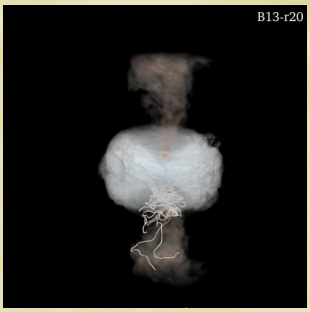

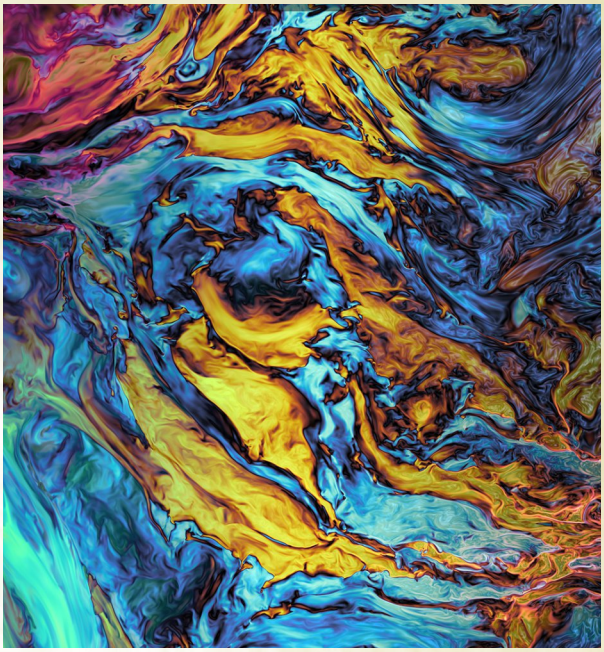

Analyzing and Interpreting Compact Binary Mergers

Research Institution
[1]Institute for Physics and Astronomy, University of Potsdam

Principal Investigator
Tim Dietrich[1]

Researcher
Edoardo Giangrandi[1], Nina Kunert[1], Ivan Markin[1], Anna Neuweiler[1], Federico Schianchi[1], Henrique Leonhard Gieg[1], Maximiliano Ujevic Tonino[2], Peter Tsun Ho Pang[3]

Project Partner
[2]Centro de Ciências Naturais e Humanas, Universidade Federal do ABC, São Paulo
[3]Nikhef, Amsterdam

Funding
Starting Grant of the European Research Council (ERC, SMArt, 101076369)

SuperMUC Project ID: pn29ba

Introduction

Compact objects, such as black holes and neutron stars, emit gravitational waves – tiny ripples in the fabric of spacetime – when they orbit around each other and eventually merge. The era of gravitational-wave astronomy began with their first direct detection of a binary black hole merger in September 2015. Just two years later, the first simultaneous detection of gravitational waves and electromagnetic signals generated by the merger of a binary neutron star system has been made. This multi-messenger event provided unique insights into the physics of compact binary systems, allowed for the testing of theoretical models for the emitted gravitational and electromagnetic waves, and enabled studies covering subatomic to cosmic length scales. To obtain physical information of the observational data, one needs to cross-correlate the observational data with theoretical predictions. For this purpose, we have developed a software framework to extract system parameters from observations using Bayesian statistics. In this regard, our research focuses not only on the analysis of signals but also on the development of accurate theoretical models for compact binary systems. The latter requires numerical-relativity simulations that solve Einstein's Field Equations together with the equations of general relativistic hydrodynamics. The support from LRZ and our granted resources on SuperMUC-NG allowed us to improve our signal analysis infrastructure significantly, enabled us to perform injection studies to assess the accuracy of information extraction from signals, and allowed us to improve our in-house numerical-relativity code.

Results and Methods

Our work is dedicated to the study of neutron star mergers and covers both the analysis of observational data and the performance of accurate simulations. For the former, we have developed a nuclear physics and multi-messenger astrophysics (NMMA) framework [1]. We use Bayes' theorem to analyze gravitational waves and electromagnetic signals, including kilonovae and the afterglow of gamma-ray bursts, simultaneously. The code is MPI-parallelized and publicly available at: https://github.com/nuclear-multimessenger-astronomy.

For our numerical-relativity simulations, we use the BAM code [2] that evolves the gravitational field in time with the methods-of-line approach and applies finite difference stencils for spatial discretization as well as high-resolution shock-capturing schemes to handle shocks in the hydrodynamic variables. BAM adopts an adaptive mesh refinement technique consisting of a hierarchy of refinement levels, necessary to resolve the different length scales (the strong-field region near the compact objects as well as the far-field region where the gravitational wave signals are extracted). The code uses a hybrid OpenMP/MPI parallelization strategy and is continuously upgraded.

In the following paragraphs, we summarize some selected results and highlights of the last year.

Binary Neutron Stars with Microphysics
We implemented in the last years a new hydrodynamics module able to handle three-dimensional nuclear theory-based equation of states [3]. This enables our numerical-relativity code to additionally evolve the temperature and electron fraction of the baryonic matter composing the stars in our simulations and to handle thermal pressure in a more physically consistent way. Even more importantly, this module provides the basis for an accurate microphysical description of baryonic matter, providing all the information needed to compute neutrino-baryon interaction rates. Neutrino interactions are of fundamental importance in determining the electron fraction of the ejecta, on which kilonova light curves and nucleosynthesis yields are strongly dependent. Our first attempt to include neutrino interactions involved implementing a neutrino leakage scheme [3]. Recently, we implemented a more advanced scheme based on the solution of first-order multipolar transport equations of neutrino energy and momentum [4]. This allows for a more accurate estimate of electron fraction and neutrino radiation pressure.

Prompt Black Hole Formation in Binary Neutron Stars Mergers with Spin
One crucial feature in the post-merger dynamics of a binary neutron star system is the fate of the remnant. A short living remnant is believed to be associated to a lower amount of proton-rich ejecta due to wind mecha-

nisms from the disk, leading to a redder kilonova, and the emission of a smaller amount of energy through gravitational waves. Systems undergoing a prompt gravitational collapse should be easily distinguishable and are of particular interest. The estimate of a prompt collapse mass threshold can indeed bring fundamental information on the equation of state of matter at supranuclear densities. We investigated how the spin of the neutron star affects the post-merger dynamics, and particularly the prompt collapse threshold [5]. We tested whether the total spin (or weighted averaged spin) is a sufficient additional parameter to capture the effect of spin on prompt collapse threshold mass. For this purpose, we have simulated a total of 28 binary neutron star configurations employing different equations of state, spin configurations, and mass ratios. We find that the time between merger and black hole formation is longer for aligned spinning configurations and that these systems also lead to the formation of a more massive disk giving potentially rise to brighter electromagnetic counterparts.

Chemical Distribution in the Dynamical Ejecta

In addition to our studies of compact binary systems with the help of numerical-relativity simulations, we have also performed multi-messenger analyses of the kilonova AT2017gfo, the optical and near-optical counterpart to the gravitational wave signal GW170817. Since the kilonova is triggered by the radioactive heating of newly synthesized radioactive elements multi-messenger analyses have the potential to provide new insights on fundamental physics questions such as the formation of heavy elements in our Universe. In addition to these astrochemical investigations, they also allow for a better understanding of supranuclear dense matter, the nature of gravity, and the expansion rate of our Universe. In a recent study [6], we incorporated new constraints on the inclination angle under which AT107gfo was observed. These constraints, arising from Very Long Baseline Interferometry (VLBI), helped to reduce degeneracies between different parameters and allowed for a more precise measurement of the source's properties. We combined this new VLBI information with an updated model for our kilonova, based on radiative transfer simulations that are embarrassingly parallelizable and were run on 49,152 cores on SuperMUC-NG simultaneously. We obtained new insight into the composition of the material that was ejected from the BNS system during and after the collision and found a very strong angular dependence of the ejected material, in contrast to other recent studies performed in the literature.

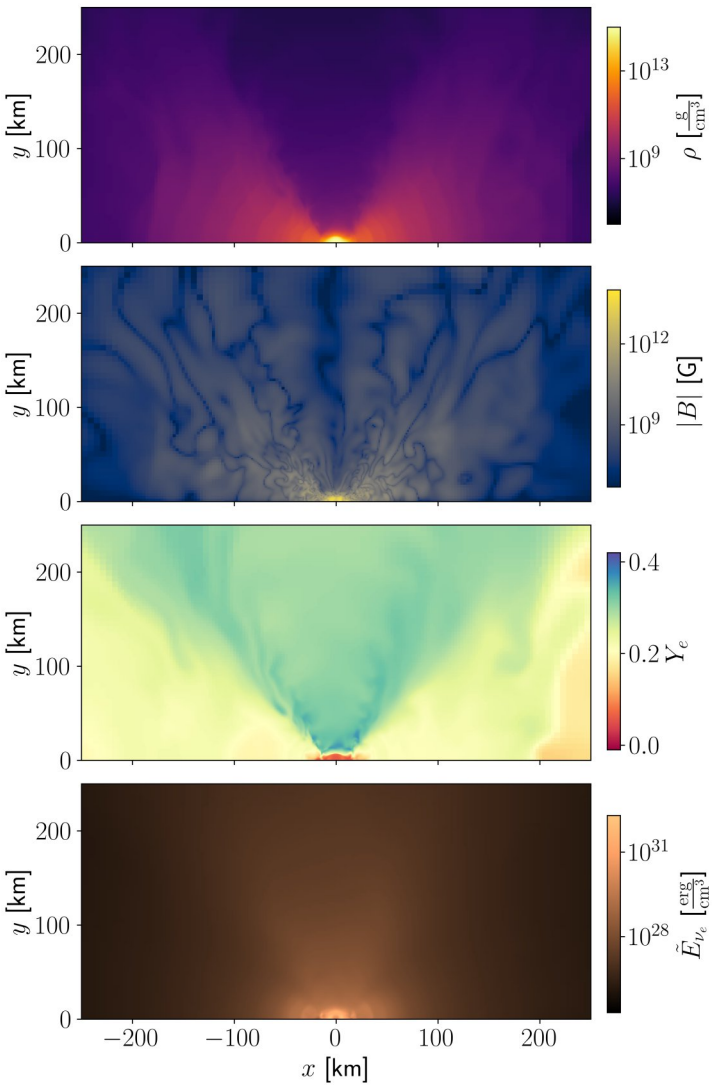

Ongoing Research / Outlook

We are currently upgrading our numerical-relativity code infrastructure to allow for the simultaneous treatment of composition effects, neutrino radiation, and magnetic fields. While the extensions towards the inclusion of magnetic fields increases the complexity of our code, it also enables more realistic simulations in which we can investigate how magnetic-driven winds change the behavior of the outward flowing material and consequently the electromagnetic counterparts connected to the merger of neutron stars. Given the ongoing fourth observing run of the international network of gravitational-wave detectors, we are also eagerly waiting for the next observation of gravitational waves from a binary neutron star mergers such that the newly developed multi-messenger astrophysics tool that we have developed [1] can be employed and unfold its full potential.

References and Links

[1] P. T. H. Pang et al., Nature Commun 14 (2023) 8352.
[2] B. Brügmann et al., Phys Rev D 77 (2008) 024027.
[3] H. Gieg et al., Universe 8 (2022) 370.
[4] F. Schianchi et al., (2023), Phys Rev D 109 (2024) 044012.
[5] F. Schianchi et al., Phys. Rev. D 109 (2024) 123011.
[5] S. Anand et al, (2023), arXiv: 2307.110800.

Figure 1: Snapshot from a binary neutron star simulation performed with neutrino transport scheme and inclusion of magnetic field. From top to bottom, we show the rest-mass density, the magnetic field strength, the electron fraction, and the energy density of electron neutrinos.

Superstars: simulating galaxies in the era of large Galactic surveys

1

Research Institution
[1]Astrophysics Research Institute, Liverpool John Moores University

Principal Investigator
Robert Grand[1]

Researcher
Rüdiger Pakmor[2], Francesca Fragkoudi[3]

Project Partner
[2]Max Planck Insitute for Astrophysics
[3]Durham University

SuperMUC Project ID: pn73wa

Introduction

Most of the visible starlight in the Universe emanates from spiral galaxies like the Milky Way. These beautiful structures are the culmination of a vast range of physical processes that begin propagating from just after the Big Bang. In the standard cosmological model, tiny fluctuations in the matter density of the early Universe grow as the Universe expands and cools until they collapse into halos under the action of gravity. The halos drag cooling gas into their centres, which begins forming a rotationally supported disc within which stars form as further cooling and collapse takes place. These stars hold memory of the conditions of the galaxy at their time of birth, because their positions, velocities, and chemical compositions are inherited from their nascent gas clouds. Their chemistry remains unchanged over Cosmic evolution, however a star's position and velocity may be further modified by dynamical evolution, through e.g. radial migration from bars and spiral arms, and the external influence of mergers and nearby galaxies. Thus, the hundreds of billions of stars contained within galaxies as seen today are fossil records imprinted with the origin story of galaxies. Deciphering these cosmic clues is critical for a complete understanding of our current cosmological model and theories of galaxy evolution.

However, these observables are only snapshots in time, and their interpretation can only be achieved through the lens of theoretical modelling. The highly complex and non-linear nature of the processes involved eludes simple models and makes cosmological hydrodynamic simulations essential. However, as current and future Galactic surveys become increasingly large, simultaneously resolving newly observed features and relevant physical phenomena on a range of scales is beyond current state-of-the-art techniques; we are missing the link between these features and their likely origin.

In our "Superstars" project, we build on the highly successful Auriga cosmological zoom-in simulations of Milky Way-mass galaxies [1,2]. Through a series of novel, computationally efficient numerical techniques designed to push the resolution of the stellar component, we add internal galactic collision-less dynamics at unprecedented fidelity and detail, thereby overcoming the numerical challenge to connect galaxy formation processes on a range of scales to observed Galactic structures. The resulting suite of simulations will provide new insights into a variety of puzzles, such as: the origin and nature of galactic spiral arms; the formation of the prominent "Gaia Snail Shell" feature; and the origin of the myriad stellar streams that surround the Galaxy.

Results and Methods

We have used a modified version of the N-body, magneto-hydrodynamics (MHD) code AREPO. AREPO uses a highly-parallel architecture based on an hybrid MPI+openMP parallelization scheme to evolve the equations of motion of gas cells, star particles, dark matter particles, and black holes. Our simulations use the "zoom-in" technique: a high-resolution region is drawn around an isolated Milky Way like galaxy which is itself surrounded by a low-resolution large-scale Cosmic Web. The setup, code, and galaxy formation model has been used in the past to great effect [1,2]. We have made several novel modifications to the code, including: i) foming a total of 64 star particles per star-forming gas cell, giving about 100,000,000 star particles in the stellar disc (a resolution ~10 times higher than any other cosmological simulation of equivalent type, and matching that of the best idealised simulations); ii) a high-cadence output (5 Myr) for star particles allowing unprecedented deep scrutiny of dynamical evolution (across 2,764 snapshots); iii) a new type of gas "dwarf" refinement which enhances the resolution of gas within the lowest mass galaxies which are otherwise poorly (or not at all) resolved; and iv) the production of r-process (rapid neutron capture) elements from binary neutron star systems to enhance theoretical predictions of chemical abundance patterns.

We performed an initial pilot simulation with the first 2 of these developments, which we used to shed light on the origin of the so-called "Gaia Snail Shell": a feature that we now think indicates the Milky Way was significantly perturbed by its misshapen dark matter halo in the recent

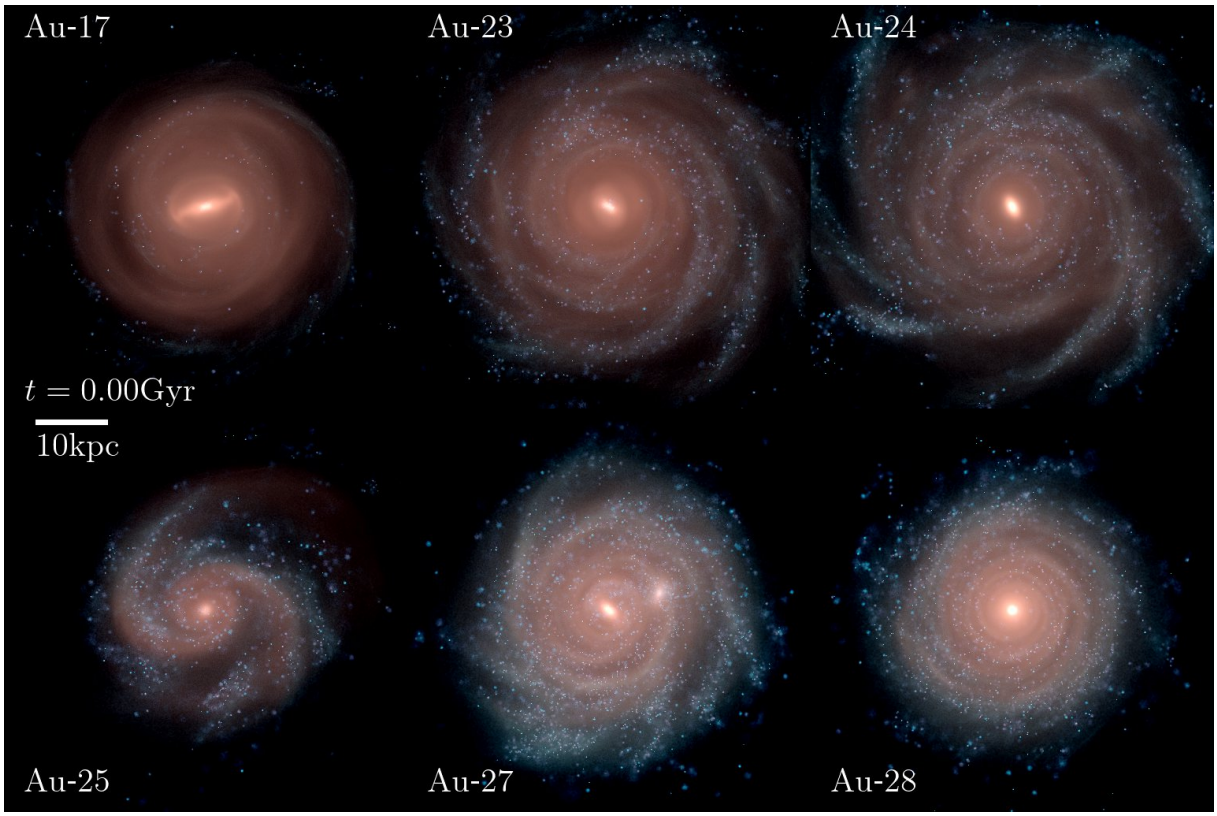

Figure 1:High-resolution stellar light images of each of the 6 simulated Milky Way-like galaxies viewed face-on. Redder light comes from older stars, whereas bluer light comes from younger stars. The latter depict spiral arms particularly clearly. Also visible are: a strong central bar (top-left) and a satellite galaxy hovering above a Milky Way (bottom-centre).

past [3]. These code modifications were crucial to obtain this result, because without them the Snail Shell is simply not resolved. We have since performed 6 simulations with all the above developments, each run using 4,800 cores for a total of 17.9 M core-h and 148 TB of memory. On top of this, we have produced merger trees and further supplementary information in post-processing needed for analyses of these large datasets. Figure 1 shows images of the present-day stellar light distribution of these six galaxies viewed face-on. Figure 2 shows the mass function of luminous subhalos within 1 megaparsec of the main galaxy for a Superstars simulation in comparison to a normal Auriga run.

Ongoing Research / Outlook

The main suite of Superstars simulations is now complete and they are now being analysed. In particular, an overview paper of the simulations and a companion paper presenting a detailed numerical study to demonstrate the robustness of the results are being prepared. Subsequent analyses of this large, rich dataset will yield dozens of publications and significant advances in our understanding of galactic dynamical evolution, the nature of dark matter, and even the galactic birth site of our Sun. We anticipate that the simulation methodology established here will be extended further to enable star-by-star comparisons with observations, which would be another world first.

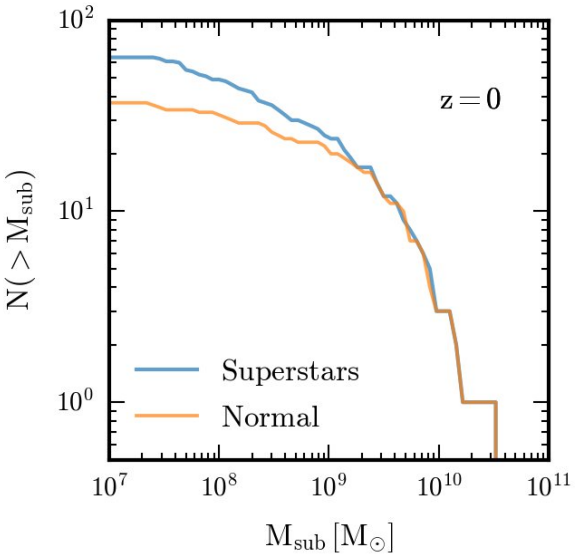

Figure 2: The cumulative subhalo mass function of luminous galaxies within 1 Mpc of the Milky Way-like galaxy at redshift zero. Superstars resolves almost twice as many galaxies with total masses lower than 10^9 Solar masses compared to normal Auriga. Note the excellent convergence for more massive galaxies.

References and Links

[1] https://wwwmpa.mpa-garching.mpg.de/auriga/
[2] Grand R. J. J. et al. 2017, MNRAS, 467, 179.
[3] Grand R. J. J. et al. 2023, MNRAS, 524, 801.

The effects of magnetic fields and cosmic ray feedback on galaxies and the circumgalactic medium

1

Research Institution
[1]Cardiff University, [2]University of Zurich

Principal Investigator
Freeke van de Voort[1]
Rebekka Bieri[2]

Researcher
Rüdiger Pakmor[3]

Project Partner
[3]Max Planck institute for Astrophysics

SuperMUC Project ID: pn68ju

Introduction

Galaxies are beautiful and complex celestial objects. They are unimaginably large, far away, and each one is different. Dwarf galaxies tend to be messy and irregular, whereas galaxies like our own Milky Way are often large star-forming spiral discs. More massive galaxies in general do not form stars very efficiently and are more spherical. How our universe creates this diversity of galaxies is still an open question.

In order for galaxies to grow, they need to accrete gas from their surrounding gaseous haloes, also called the circumgalactic medium (CGM). This provides fresh fuel for star formation and for the central supermassive black hole. Outflows, driven by, for example, supernovae or supermassive black holes, eject gas from the galaxies back out into the CGM. By doing so, they enrich the CGM with heavy elements, magnetize it, and drive turbulence. Galaxies and their CGM are thus intimately connected and need to be studied in tandem if we want to understand the full galaxy ecosystem.

Here, we used SuperMUC-NG to run a large suite of cosmological zoom-in simulations – high-resolution simulations that follow a large region of the universe for the full 13.8 billion years of cosmic history. Each simulation focuses on a specific galaxy and its large-scale environment. Our simulation suite covers the mass range from dwarf galaxies a hundred times smaller than our own Milky Way to galaxies an order of magnitude more massive. We furthermore ran each simulation with different physics models: neglecting magnetic fields, including magnetic fields, and with cosmic ray transport and feedback in addition to magnetic fields. Magnetic fields are important because most of the gas in the universe is ionized and thus flows preferentially along the direction of magnetic field lines. Cosmic rays are relativistic particles, which also follow the magnetic field lines. They could provide an important pressure source, potentially changing the gas flows around galaxies. The new simulation suite enables us to understand the impact these processes have on the CGM and on the embedded galaxies. This is an important step forward, because most state-of-the-art cosmological simulations are do not include magnetic fields or cosmic rays even though they are likely to affect galaxy formation.

We make use of a CGM refinement technique that increases the spatial resolution of the gas in the halo by two orders of magnitude and allows us to study this region in more detail than previously possible. We focus on the morphology of the central galaxies, the properties of the multiphase CGM and how they depend on halo mass, on the impact of the non-thermal physical processes included, and on the numerical resolution. We also use our simulations to compare directly to observations. This tests our theoretical models and guides the interpretation of observations.

Results and Methods

Our simulation suite consists of 56 high-resolution cosmological zoom simulations, run from the very early universe to the present day. We used the code AREPO which solves the equations of magneto-hydrodynamics on an unstructured moving mesh coupled to gravity and a galaxy physics model. To properly model low-mass galaxies, we ran them at even higher resolution than the massive galaxies. We showed that this is vital for resolving the turbulent dynamo and thus for the convergence of the magnetic field strength [2].

We focused on 14 galaxy haloes with different physical processes included. The fiducial model contains magnetohydrodynamics and feedback from supernovae and supermassive black holes, but no cosmic ray feedback. To better understand the impact of different physical processes we add runs without magnetic fields, as well as runs including cosmic ray amplification, transport, and feedback.

We already showed that magnetic fields cannot be neglected if we want to model the cosmic flow of gas and thus the formation and evolution of galaxies [1,2]. The magnetic field strength in the CGM varies by orders of magnitude, but can be as strong as 1 microGauss. We show that the thermal pressure dominates globally, but locally, in the cool gas, the magnetic pressure is dominant. Magnetic fields slow down the flow of gas in the CGM and the total gas fraction and metal mass fraction in the halo are higher, because less gas escapes the halo. Additionally, magnetic fields reduce the scatter in radial velocity, whilst enhancing the scatter in the

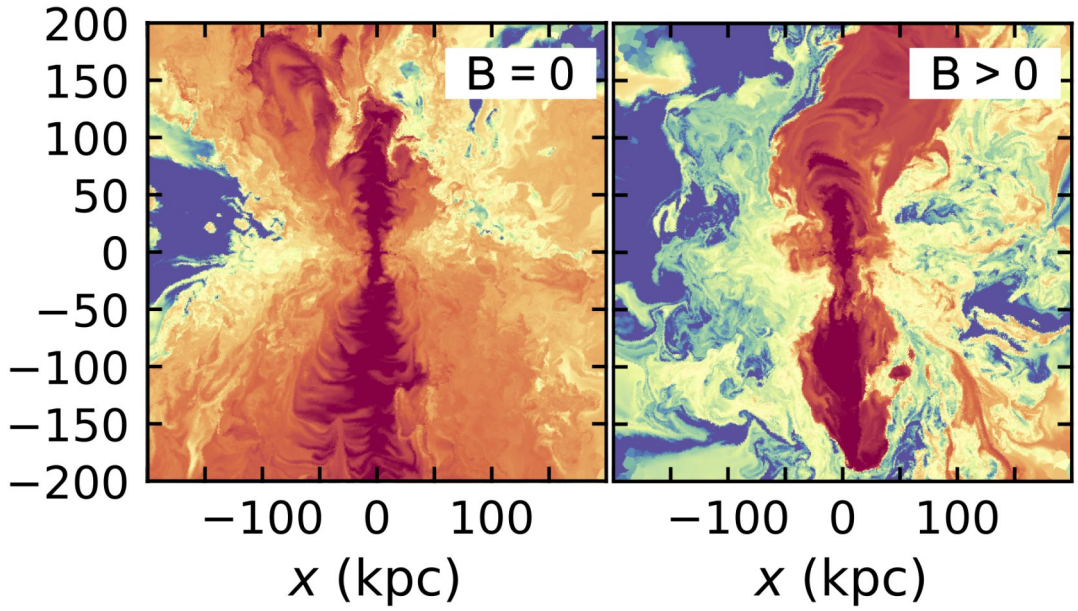

Figure 1: Images of the gas in and around an edge-on Milky Way-mass galaxy in a simulation without and with magnetic fields. This shows the distribution of heavy elements where red regions are the most enriched. Heavy elements are more spread out across the CGM without magnetic fields, showing that the gas mixes more efficiently.

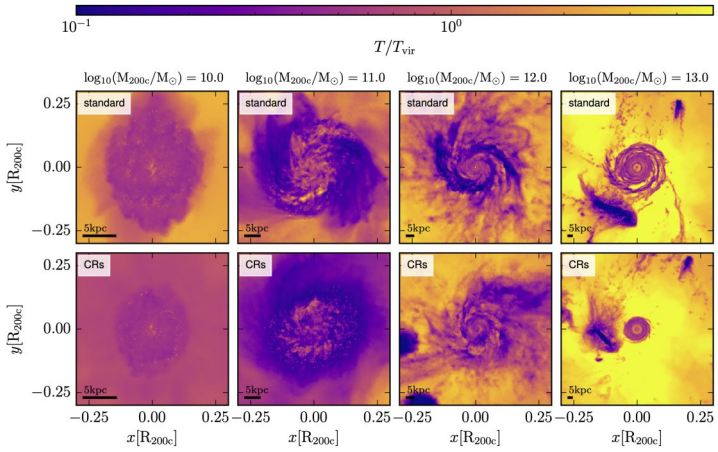

Figure 2: Face-on images of the temperature of four representative galaxies, normalized to the typical temperature of their haloes. The top (bottom) row shows simulations without (with) cosmic ray (CR) feedback. The galaxy discs are typically smaller with cosmic ray feedback. For low-mass galaxies (left panels), the CR simulations reveal cooler surroundings, showing that cosmic rays significantly influence the evolution of these systems.

distribution of heavy elements, which means they are not as well-mixed throughout the halo (see Figure 1). These results show that magnetic fields in the CGM make it more difficult for outflows to mix with the rest of the halo gas and for them to escape the halo. Therefore they are vital for correctly modelling the formation and evolution of galaxies and their large-scale ecosystems. We will study whether these results hold at for galaxies with lower and higher masses as well [3].

In [4] we show that cosmic ray feedback reduces the size of the gas disc on all mass scales (see Figure 2). Cosmic rays enhance the gas temperature of the CGM of low-mass galaxies. Additionally, they enhance the gas flow and the radial velocity of outflows from low-mass galaxies. Similarly, they manage to reduce the star formation rates by up to a factor of 1.5 for the dwarf galaxies. Comparing our simulations with and without cosmic ray feedback, we find that the time at which the star formation histories begin to diverge correlates with the moment the cosmic ray pressure in the disc approximates the thermal pressure of the gas. Investigating different cosmic ray models, we will show in [5] that some cosmic-ray dominated models can lead to extreme changes in the CGM in terms of overall gas structure and temperature. These relative changes persist when the resolution is increased in the CGM.

Ongoing Research / Outlook

We will further analyze our existing simulations [3-5]. We also plan to expand the SURGE simulation suite by including simulations that include thermal conduction, which could change the thermodynamic properties of the gas by exchanging thermal energy along magnetic field lines. We are furthermore interested in the interface between galaxy discs and the CGM, where inflows join the interstellar medium and outflows push gas out into the halo. We will develop a model to increase the gas resolution specifically in this region to study it in more detail than previously possible.

References and Links

[1] van de Voort et al., 2021, MNRAS, 501, 4888.
[2] Pakmor et al., 2024, MNRAS, 528, 2308.
[3] Cook et al., *in preparation*.
[4] Bieri et al., *in preparation*.
[5] Hannington et al., *in preparation*.

1 Unveiling cosmic dawn: the nature of the first galaxies and reionisation

Research Institution
[1]Leibniz Institute for Astrophysics Potsdam

Principal Investigator
Stefan Gottlöber[1]

Researcher
Gustavo Yepes[2], Anne Hutter[3], Pratika Dayal[4]

Project Partner
[2]Universidad Autónoma Madrid
[3]NBI Copenhagen
[4]University of Groningen

Funding
ERC, Villum Fondem (Denmark), MICIN(Spain)

SuperMUC Project ID: pr74no, pn36si

Introduction

About 13.8 billion years ago at the epoch of recombination the charged electrons and protons first became bound to form electrically neutral hydrogen atoms. This recombination occurred about 400,000 years after the Big Bang. At this time our Universe was still almost homogeneous. Gravitational instability forced the matter to clump and to become structured. First objects formed and the first stars started to reionise the Universe. This structure formation can be followed by numerical simulations which take into account the clumping of the dark matter and assume that baryons form stars and galaxies within the dark matter halos. In March 2013 we had started our MultiDark simulation project: a series of cosmological dark matter only simulations where we assumed a spatially flat universe consisting of 30.7 % of dark matter and 69.3 % of dark energy with a present-day expansion rate of 67.8 km/s/Mpc, the Planck cosmology. Between the BigMultiDark (BigMD), the MultiDark (MD), the Small MultiDark (SMD), the Very Small MultiDark (VSMD) and the Extremely Small MultiDark simulations (ESMD) the simulation volume decreases by a factor of ~16. On top of this series we were running also the HugeMultiDark (HMD) simulation and within the new project pn36si we started to rerun the SMD simulation with an eight times higher resolution. All MultiDark simulations are summarized in Table 1.

Many users were interested in using the simulation data directly or analysing the full catalogues by themselves. Therefore, we opened a possibility to download selected simulation outputs (1.7 Tb each) as well as the full ROCKSTAR or galaxy catalogues 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.

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
SMD-HR	400	$7,680^3$	1.2×10^{7}	120
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: The columns give the simulation identifier, the size of the simulated box in Mpc/h, the number of particles, the mass per simulation particle mp (in Msun/h) and the total number of simulation outputs stored Nout .

Results and Methods

The first galaxies emerging from the primordial soup of neutral hydrogen and helium in the first billion years after the Big Bang were the first sources of light and elements heavier than Lithium in our Universe. Their photons ionised the hydrogen in the surrounding intergalactic gas and defined the ionised regions that grew around them. This last major phase transition, the Epoch of Reionisation (EoR), fundamentally governs the formation and evolution of the galaxies we see today. To build a coherent picture of how reionisation happened, we need to understand its sources: How much ionising radiation did the first galaxies produce and how much could escape to ionise the intergalactic hydrogen gas? What underlying physical processes determine their characteristics?

Modelling
Addressing these fundamental questions requires models that can rapidly explore the vast parameter space of galaxy evolution and reionisation physics while resolving the faint low-mass galaxies within a volume large enough for statistics not to be biased by cosmic variance. High resolution dark matter only simulations are best suited

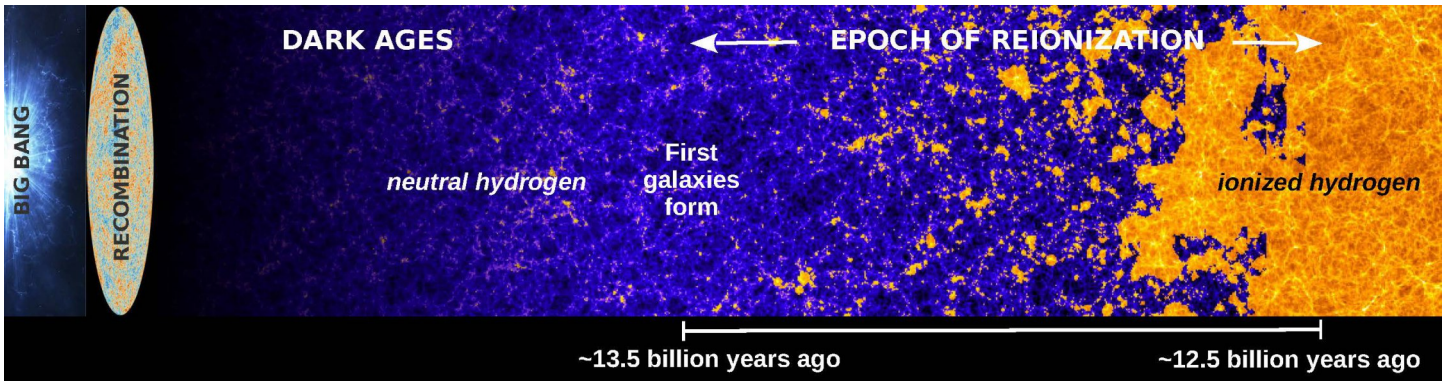

Figure 1: The first billion years of the Universe. The first galaxies form around 200 million years after the Big Bang and start ionising the hydrogen in the intergalactic gas. Ionised regions (yellow) start growing around the galaxies in the neutral (blue) intergalactic gas medium.

because they can be performed in large volumes and allow at the same time high spatial and mass resolution. The outcome of these simulations is the distribution of dark matter halos in which the galaxies are formed. The position, mass and formation history of these halos is known at any desired moment. Based on this information one wishes to model the galaxy evolution and the epoch of reionisation. To this end we have developed the Astraeus framework [2]. It couples self-consistently the outputs of high-resolution, large-volume dark matter simulations with a semi-numerical model of galaxy evolution and reionisation. Astraeus follows the essential physical processes governing early galaxy evolution and reionisation by post-processing a dark matter only simulation's mass assembly histories of dark matter halos and the evolving spatial distribution of the matter density. For each simulated galaxy and at each time, it tracks the amount of gas accreted, the gas and stellar mass merged, the formation of stars and their feedback through supernovae and metal enrichment. It also follows the large-scale reionisation process, tracking the time and spatial evolution of the ionised regions within the intergalactic gas, and its feedback on the gas content in galaxies.

Results

Applying the Astraeus model to the outputs of the VSMD simulation yielded the following significant insights which are also summarized in [5].

The debate over which physical processes were most instrumental in determining which galaxy population contributed most to reionising the Universe—whether numerous faint ones or a few bright ones—is ongoing, and the recent discovery of "little red dots" reignited the discussion whether galaxies with accreting supermassive black holes are significant contributors. Astraeus simulations have shed light on this issue, revealing that the fraction of ionising radiation escaping from galaxies plays a pivotal role, outweighing the suppression of star formation in low-mass galaxies due to radiative and stellar feedback. They also highlighted that the contribution of black hole-dominated galaxies becomes significant only at lower redshifts when reionisation is nearly complete.

As galaxies emit ionising radiation, they leave distinct patterns in the distribution of ionised regions within the intergalactic hydrogen gas, known as ionisation topology (see Fig. 1). Observations that track or map this ionisation topology offer insights into the galaxy population responsible for reionising the Universe. While galaxies with observable Lyman-alpha radiation, sensitive to absorption by neutral hydrogen, have been proposed as a means to map the ionisation topology, Astraeus simulations suggest that the distribution of such Lyman-alpha emitters (LAEs) primarily reflects the overall ionisation state of the intergalactic hydrogen gas rather than the specific ionisation topology [3]. Various studies have shown that cross-correlating the distribution of LAEs with the 21 cm signal from the neutral intergalactic hydrogen gas — measurable with large upcoming radio interferometers like the Square Kilometre Array — traces the global ionisation state of the intergalactic hydrogen gas, but any dependency on the ionisation topology has remained unclear. Using Astraeus simulations, a new analytical model for cross-correlating Lyman-alpha emitters with the 21 cm signal has been developed, uncovering such a dependency and explaining variations among results seen in existing literature [4]. In Fig. 2, we show the large-scale distribution of the 21 cm signal and LAEs — both observable with upcoming near-infrared and radio telescopes. This figure demonstrates that the number of LAEs increases over time as the Universe becomes more ionised while the overall strength of the 21 cm signal weakens.

The James Webb Space Telescope has given us exciting glimpses into the past, unveiling an abundance of bright galaxies in the ultraviolet during early times (~ 200 million years after the Big Bang) that exceed predictions of standard theoretical models. Proposed solutions to this discrepancy include star formation occurring in sporadic bursts in these galaxies or an increased abundance of massive stars in their stellar populations. Analysing the representative galaxy populations in different Astraeus simulations has shown that bursty star formation is predominantly found in lower-mass galaxies, where stellar feedback, rather than gas accretion, drives star formation. Furthermore, extending the Astraeus model to account for a mass distribution of the forming stars favoring the formation of more massive stars in regions with less chemically enriched gas and hotter cosmic microwave background radiation has provided unprecedented insights. These simulations suggest that a higher abundance of massive stars alone is unlikely to explain the high abundance of bright galaxies, as their increased light-to-stellar mass ratio is counteracted by their stronger stellar feedback, suppressing ongoing star formation.

Ongoing Research / Outlook

However, the predictive power of some of these reionisation-focused studies remains limited by the lack of simulated volume. Both, studies of the massive and rare black-hole-dominated galaxies and efforts to constrain the dependency of the ionising emissivity of galaxies on their mass/gravitational potential from the cosmic variance-affected 21cm signal maps, require larger volumes at VSMD-like mass resolution. The larger high-resolution SMD-HR simulation, resolving the low-mass galaxies, will address this issue. It will allow us to study the impact of black-hole-dominated galaxies in the most overdense regions on reionisation and robustly predicting the constraining power of statistical analyses of the 21cm signal. Both will inform the current and upcoming near-infrared, infrared and radio surveys on how to constrain the first galaxies' properties best.

The SMD-HR simulation has been particularly difficult to run in SuperMUC-NG. All the previous simulation boxes reported in Table 1 were run with a special version of the famous TreePM N-body GADGET code. The so-called Lean-Gadget (L-Gadget) code, kindly provided by the code author, Volker Springel, was based on the publicly available Gadget2 version but optimized for an efficient use of memory per MPI task in order to be able to handle large number of particles in large computational boxes. Although there is now a much better version of GADGET, the Gadget4 code, we have decided not to move to the new version to avoid possible differences in the gravitational evolution with respect to the original SMD simulation. Since the L-Gadget2 is only MPI parallel, we cannot make efficient use of all the cores per node in SuperMUC-NG thin nodes, unless we allocate 48 MPI tasks per node. In previous simulations, we could fit almost all these MPI tasks per node, but in the SMD-HR, with 453 billion particles, it was not possible to do it due to memory constraints. With the 1.5 GB memory per core available in the thin nodes of SuperMUC-NG, we could not run this simulation unless we use ~ 50K cores. But, at the same time, having the simulation box split into 50K subdomains produces a large problem of load/balancing, specially at low redshifts when the particle distribution start to cluster into halos. Therefore, to have a compromise between domain decomposition and memory per MPI task, the only possibility is to either allocate few tasks per node in thin nodes, so each MPI task gets enough memory of the node, or else going to the fat nodes and try to allocate as many MPI tasks as possible within each node. The first option is not efficient at all since many cores in each node are not used, so the total cost in core hours is prohibitively. Therefore we have opted for the second option to use the fat nodes, despite the fact that the number of fat nodes are considerably lower than the thin ones, so the queue waiting times are much larger. Nevertheless, given the large amount of memory per fat node (740 GB), it is worth to pay the price for the long waiting times in queues. Moreover, we have done several tests with the fat nodes in order to allocate more MPI tasks than cores available per node using the hyperthreading capacities of the Intel Skylake chips. In our final configuration, we managed to allocate 96 MPI tasks per fat node (twice the number of cores), so we can fit all the SMD-HR simulation in just 96 fat nodes, using a total of 7,680 subdomains, each one assigned to an MPI task. This is our production configuration we are using in SuperMUC-NG to run the SMD-HR simulation. As mentioned above, given the oversubscription of the fat queues, the simulations run is proceeding quite slowly, but we have now managed to run the computational box down to redshift 1.5 and stored more than 100 snapshots from which we have processed them through the ROCKSTAR halo finder in order to get all the halos/subhalos in the simulation box. These catalogs will then be used to construct the merger trees which are the key product for applying the Astraeus model in order to study in much more detail the reionisation process.

References and Links

[1] http://www.cosmosim.org
[2] Hutter et al., MNRAS, 503, p.3698-3723, 2021.
[3] Hutter et al., MNRAS, 524, p.6124-6148, 2023.
[4] Hutter et al., MNRAS, 525, p.1664-1676, 2023.
[5] https://annehutter.github.io/astraeus.html

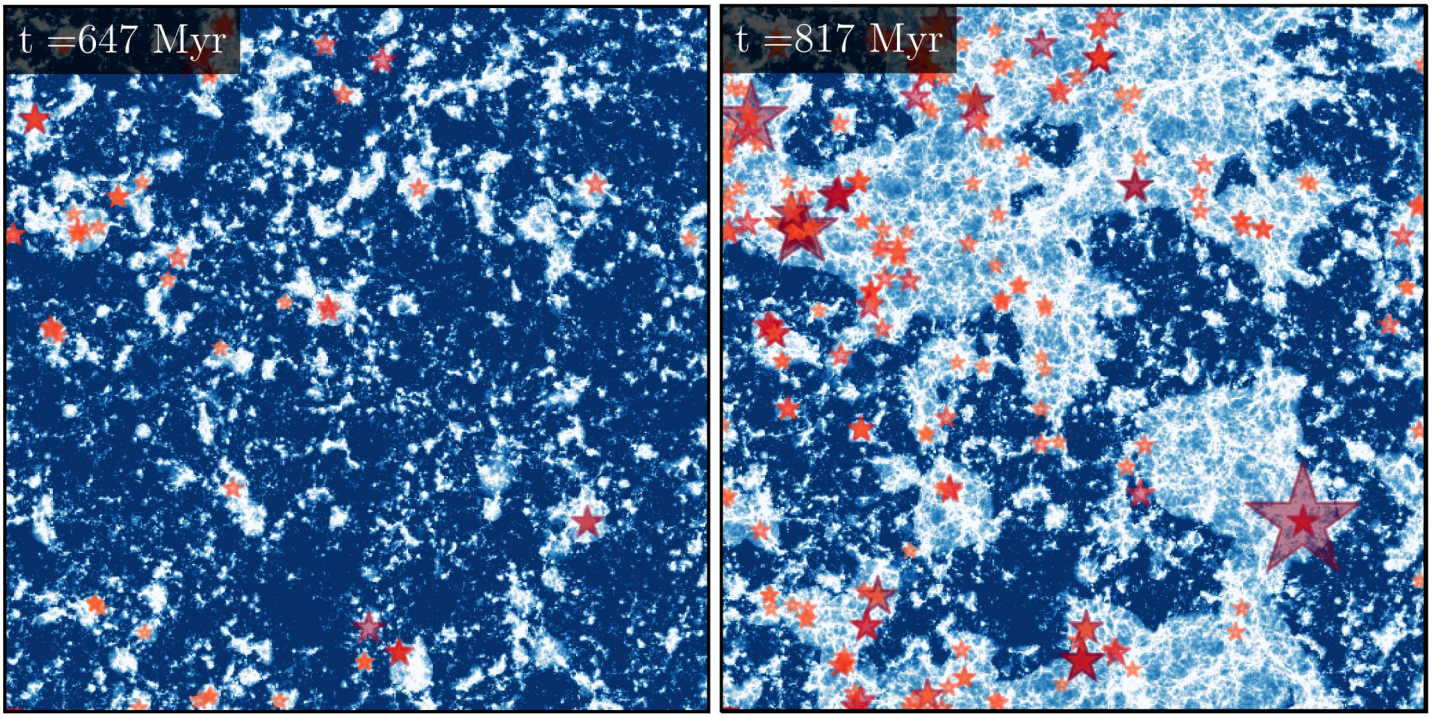

Figure 2: In a slice through the center of the VSMD simulation box the 21cm signal (blue) and galaxies with observable Lyman-α emission (red stars, size scales with brightness in Lyman-α) are shown 647 (left) and 816 million years after the Big Bang (right). Blue indicates a strong 21cm signal and white a weak to no 21cm signal.

GraWindi: Gravitational waves and disc winds from neutron star mergers

1

Research Institution
[1]Friedrich Schiller Universität Jena

Principal Investigator
Sebastiano Bernuzzi[1]

Researcher
William Cook[1], Pedro Espino[1], Rossella Gamba[1], David Radice2, Francesco Zappa[1]

Project Partner
[2]Penn State University, Princeton University, Berkeley University

Funding
ERC 714626 “BinGraSp”, ERC 10104337 “InspiRem”, DFG E 6301/2-1 “MEMI”

SuperMUC Project IDs: pn68wi, pn36ge, pn36jo

Introduction

The scientific breakthrough associated to the LIGO-Virgo observation of gravitational waves (GWs) and electromagnetic (EM) counterparts from a binary neutron star merger (BNSM) has been crucially supported by theoretical predictions provided by simulations in numerical general relativity (NR). Simulating the spacetime and the neutron-star matter fields in 3 spatial dimensions (plus time) is the only way to connect the strong-field dynamics to the observable gravitational and electromagnetic spectra. Crucially, these HPC simulations provide precise calculations for the GWs and for mass outflows of neutron rich material. The former are necessary to detect the signals and identify the properties of the source (masses). The latter are engines for kilonova (Kn) transients which are produced by the radioactive decays of newly formed r-process elements in the outflows.

Jena is at the forefront of BNSM simulations and the related astrophysics modeling of GW and EM signals. The group performs ab-initio global simulations starting from the orbital dynamics and investigating the remnant evolution with sophisticated microphysics and neutrino transport. A particular focus is on producing high-quality data by employing multiple grid resolutions and by performing extensive convergence tests. Simulations data are then employed to develop GWs and light curve models directly employed in observations.

Jena leads the computational relativity (CoRe) collaboration: an international collaborative research effort supporting the emerging fields of gravitational wave and multi-messenger astronomy. Jena and CoRe collaborators continuously develop state-of-art NR codes and release open data for the astrophysics community. A primary example is the largest-to-date BNSM waveform database [1].

The focus of the GraWindi project is to address two outstanding issues in BNSM: (i) compute high-precision waveforms from multi-orbit inspiral-merger simulations, and (ii) characterize disc winds from long-term remnants evolutions. The project is still ongoing; selected published (or submitted for publication) results are presented below.

Results and Methods

This project employed two similar production NR codes developed at Jena (BAM) and PSU (THC). Both codes solve the Z4c formulation of Einstein Equations, previously developed at Jena, coupled to general relativistic hydrodynamics. The equations are discretized with high-order finite differencing methods and explicit time stepping. The spatial domain is covered by a hierarchy of Cartesian grids organized with a box-in-box adaptive mesh. Both codes use a hybrid OpenMP/MPI parallelization scheme and have been shown to efficiently scale up to a few thousands cores on SuperMUC-NG (>70%, strong scaling) in full production runs.

The project leverages on novel algorithms recently implemented in both codes. BAM has been equipped with an entropy-flux-limited scheme that significantly improves waveforms convergence by applying an unlimited high-order scheme on smooth flow regions. THC has been extended with a gray neutrino transport scheme based on the truncated momentum formalism (M1). The scheme improves the robustness of previous general relativistic formulations and implements complete matter-radiation source terms for the first time. THC also features finite-temperature microphysical equations of states for the neutron star matter and a subgrid model (Large Eddy Scheme) for magnetohydrodynamics-induced turbulent viscosity. These new features allow unprecedented simulations in terms of GWs precision (BAM) and long-term evolutions with neutrino radiation (THC).

For waveform modeling purposes we performed 14 novel simulations with BAM of unequal mass, spinning binaries with varying matter equation of state and multiple resolutions. These simulations span more than ten orbits before merger and start from eccentricity reduced, constraint-satisfying data. Building such initial data requires an iterative procedure and several simulations only to prepare the initial configurations. However, eccentricity reduction is a necessary step to achieve higher data quality for waveform modeling.

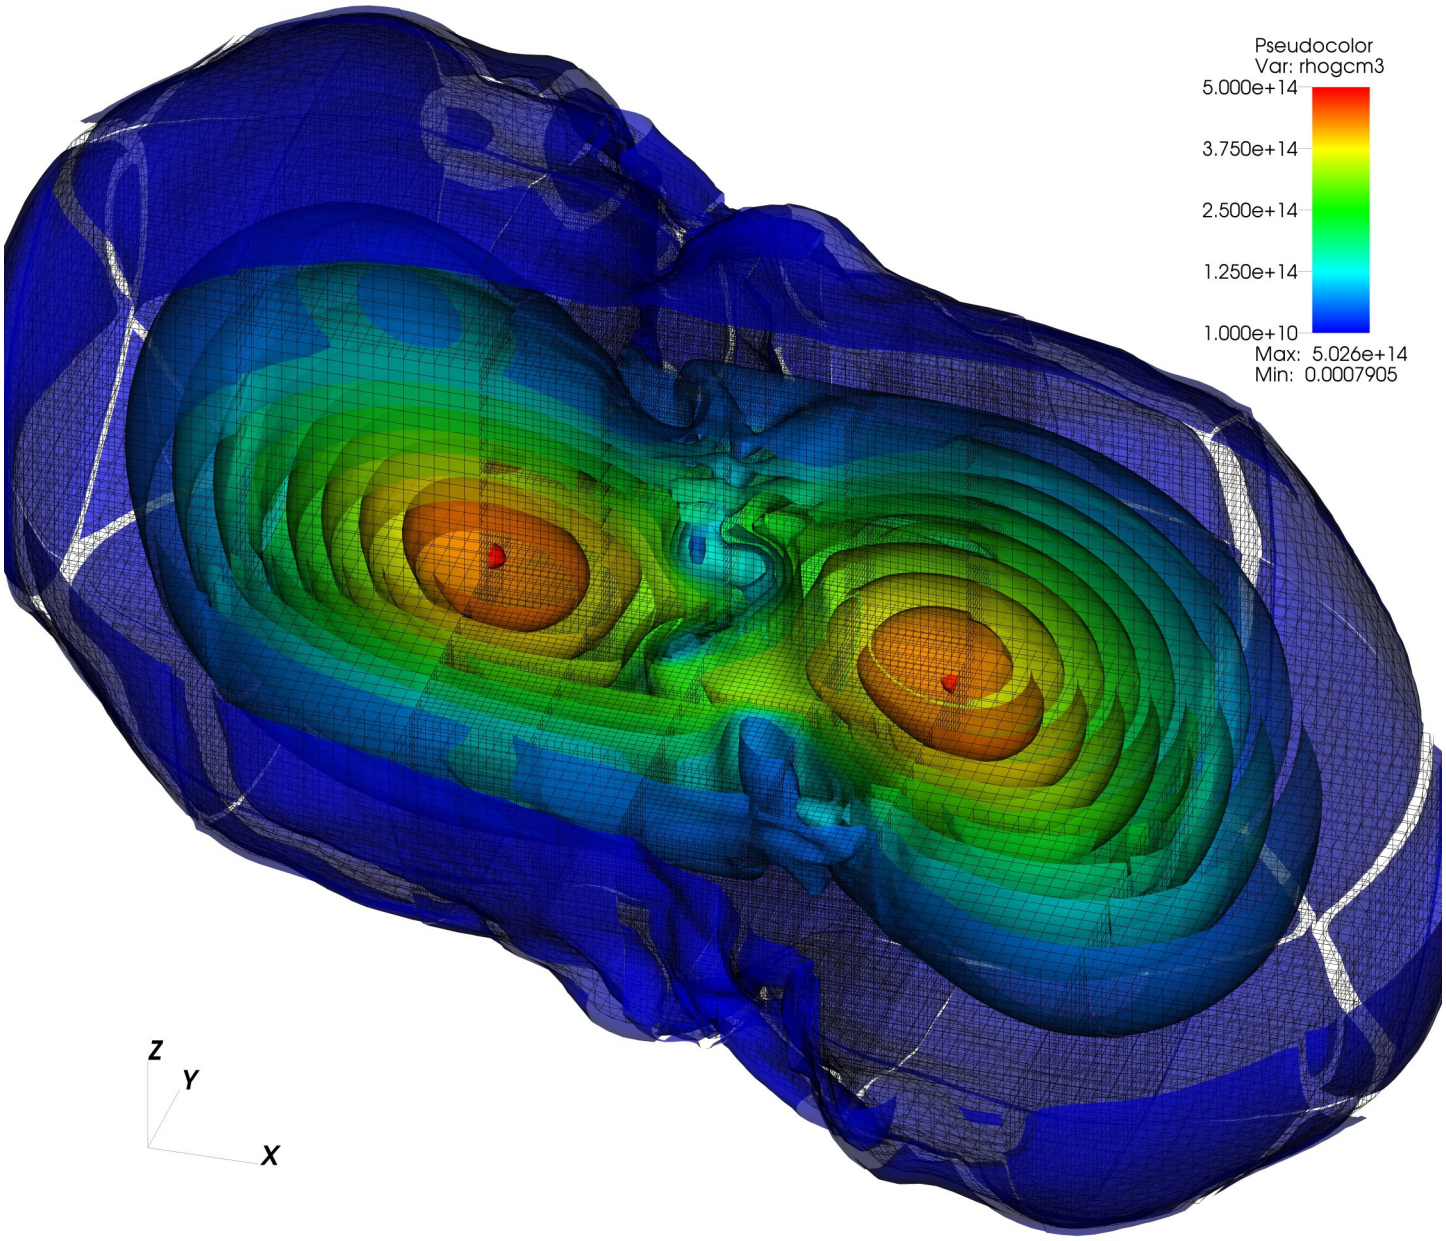

Figure 1: Density isocontours of a BNSM around merger and the GR-Athena++ adaptive mesh refinement structure.

Waveform data have been used to inform a sophisticated effective-one-body model that is employed for the observation of GW signals [2]. The goal is to accurately model the effect of tidal interactions on the waveform phase. These effects are small and difficult to resolve (order of a radian over many orbits) but nonetheless crucial because they carry the imprint of the unknown equation of state of neutron star's matter. The effective-one-body model developed at Jena incorporates all the analytical (perturbative) information about tidal interactions and it is now tested and improved by GraWindi simulations. Future observations of BNSMs, especially by third-generation detectors like the Einstein Telescope, will use such models to clarify the nature of matter at extreme densities of supranuclear densities.

Regarding, BNSM remnants we performed the first 3D ab-initio general-relativistic neutrino-radiation hydrodynamics of a long-lived neutron star merger remnant spanning a fraction of its cooling time scale (~100 ms) with THC [3]. We found that neutrino cooling becomes the dominant energy loss mechanism after the GW-dominated phase of the merger dynamics. A massive accretion disk is formed from the material squeezed out of the collisional interface between the stars. The disk carries a large fraction of the angular momentum of the

system, allowing the remnant massive neutron star to settle to a quasi-steady equilibrium within the region of possible stable rigidly rotating configurations. The remnant remains differentially rotating but it is stable against convection and stably stratified. This implies it is stable against the magneto-rotational instability and that other magnetohydrodynamics mechanisms operating on longer timescales are necessary for the removal of the differential rotation. Our results indicate the remnant massive neutron star is qualitatively different from a proto-neutron stars formed in core-collapse supernovae. Understanding how the magnetic field can emerge from the remnant and create the condition to launch relativistic jets remains a challenge for future simulations.

In a follow-up study [4], we studied out-of-thermodynamic equilibrium effects in BNSM. We found that during merger, the cores of the neutron stars remain cold at temperature of a few MeV; out of thermodynamic equilibrium matter with trapped neutrinos originates from the hot collisional interface between the stars. However, within a few milliseconds matter and neutrinos reach equilibrium everywhere in the remnant. These results show that dissipative effects, such as bulk viscosity, if present, are only active for a short window of time after the merger. Hence, they are likely to be unimportant for GW observations.

Ongoing Research / Outlook

BNSMs are rich astrophysical laboratories that involve all the fundamental interaction at the extreme. Observations of BNSMs have the potential of uniquely informing fundamental physics (including cosmology). However, they require precise theoretical predictions which imply significant theoretical and computational challenges.

In order to tackle these challenges, the Jena group is developing the next generation NR codes. A first step has been taken with the completion of the GR-Athena++, a general-relativistic radiation-magneto-hydrodynamics code for applications to neutron star spacetimes [5]. GR-Athena++ combines a highly scalable oct-tree adaptive mesh refinement a hybdrid parallelism and a task based dynamic scheduling with the state of art NR methods implemented in BAM/THC. The code has been benchmarked against several stringent tests. We showcased, for the first time in BSNM, the use of adaptive mesh refinement to resolve the Kelvin-Helmholtz instability at the collisional interface, Fig.1. GR-Athena++ shows strong scaling efficiencies above 80% in excess of 10^5 cores and excellent weak scaling is shown up to 5×10^5 cores in a realistic production setup. Hence, the code allows for the robust simulation of neutron star spacetimes with exascale computers.

References and Links

[1] http://www.computational-relativity.org
[2] Gamba et al. *submitted to* Phys.Rev.D. https://arxiv.org/abs/2307.15125
[3] Radice & Bernuzzi Astrophys.J. 959 (2023) 1, 46. https://arxiv.org/abs/2306.13709
[4] Espino et al. *submitted to* Phys.Rev.Lett. https://arxiv.org/abs/2311.00031
[5] Cook et al. *submitted to* Astrophysical Journal https://arxiv.org/abs/2311.04989

CLONES: digital twins of the local Universe

Research Institution
[1]Centre de Recherche en Informatique, Signal et Automatique de Lille, Institut d'Astrophysique Spatiale
[2]Leibniz-Institut für Astrophysik

Principal Investigator
Jenny G. Sorce[1,2]

Researcher
Nabila Aghanim[1], Théo Lebeau[1], Salvatore Cielo[3], Margarita Egelhofer[3], Gabriel Jung[4], Klaus Dolag[4], Elena Hernandez-Martinez[4], Benjamin Seidel[4]

Project Partner
[3]Leibniz-Rechenzentrum
[4]Ludwig-Maximilians-Universität München

Funding
ANR/DFG

SuperMUC Project ID: pn57ne, pr74je

Introduction

Two billion light-years wide, the local Universe, is a formidable observational playground for astrophysicists. This tiny bit of the Universe hosts billions of billions of stars, planets gathered in galaxies, including our very own: the Milky Way. We now have a huge amount of data at different wavelengths of tons of galaxies and galaxy clusters in our neighborhood to understand better the Universe, its formation and evolution as well as that of its constituents. However, analyzing and interpreting properly all the observations of our Cosmic Home requires sophisticated cosmological simulations that need to be run on powerful supercomputers.

Nevertheless, typical cosmological simulations reproduce only statistically the large-scale structure of the Universe called the cosmic web. This web defines the distribution of matter, galaxies included, in the Universe into knots, sheets, filaments and voids. In order to refine our conclusions, new simulations designed to resemble the local Universe not only statistically, but also down to the details (galaxies and galaxy clusters), become essential. Thus, CLONES (Constrained LOcal & Nesting Environment Simulations) were born [1]. These digital twins of the local Universe have since been used for several projects. The CLONES initial conditions (a set of dark matter particles with initial positions and velocities) have almost all been run forward in time at SuperMUC-NG to produce various digital twins. Among them, there are those of our own galaxy and its neighbor, Andromeda (HESTIA project), of our closest cluster neighbor, Virgo (Light on the Virgo cluster project) e.g. [2], of another cluster, Coma (Coma Connectivity project) [3], of the inner part of the Universe in its infancy (CoDa project) and of the full local Universe (SLOW and LOCALIZATION projects) e.g. [4].

Still, the first two projects used only zooms of small parts of larger simulation boxes. The third one followed only dark matter particles and the fourth one stopped at high redshifts, namely way before reaching our cosmic time. The fifth project reached our cosmic time in the full box and followed both dark matter and baryonic particles. However, given the scientific ambition of our new LOCALIZATION project, a CLONE with an even higher resolution is required. The new CLONE of the local Universe will thus have a higher resolution than our previous full-box CLONE. In addition, it will include additional physics, in particular non-isotropic AGN jet feedback and black hole spin-dependent feedback. These two additional components permit a more accurate modelling of the physics related to the energetic black holes happening in galaxy clusters. These phenomena are indeed fully involved in galaxy formation and evolution via their effect on the gas (baryonic matter).

All the CLONES, but for our last one, have already been presented in previous editions. Consequently, this paper focuses only on our latest higher resolution full-box CLONE, that includes additional physics. It is currently run for the LOCALIZATION project at LRZ on SuperMUC-NG with the RAMSES code [5].

Results and Methods

Following the methodology we developed and presented in [1] and references therein, the replica of our Cosmic Home is obtained via initial conditions that are run forward in time down to our cosmic time (redshift z=0). In order to properly resolve galaxies of interest for our project, the dark matter particle mass is ~10^9 times the solar mass. There are more than 8.5 billion of them in the simulation box. In addition, the gas distribution and physics properties are represented by cells of less than 10,000 light-years aside in the densest regions. There are an additional more than 2 billion stars and black holes particles.

Already more than sixty million cpu hours were used to simulate the full formation and evolution of our cosmic neighborhood. From 5,280 up to 31,680 cores performed the calculation either from the thin or fat nodes available on SuperMUC-NG. About 200 snapshots were written,

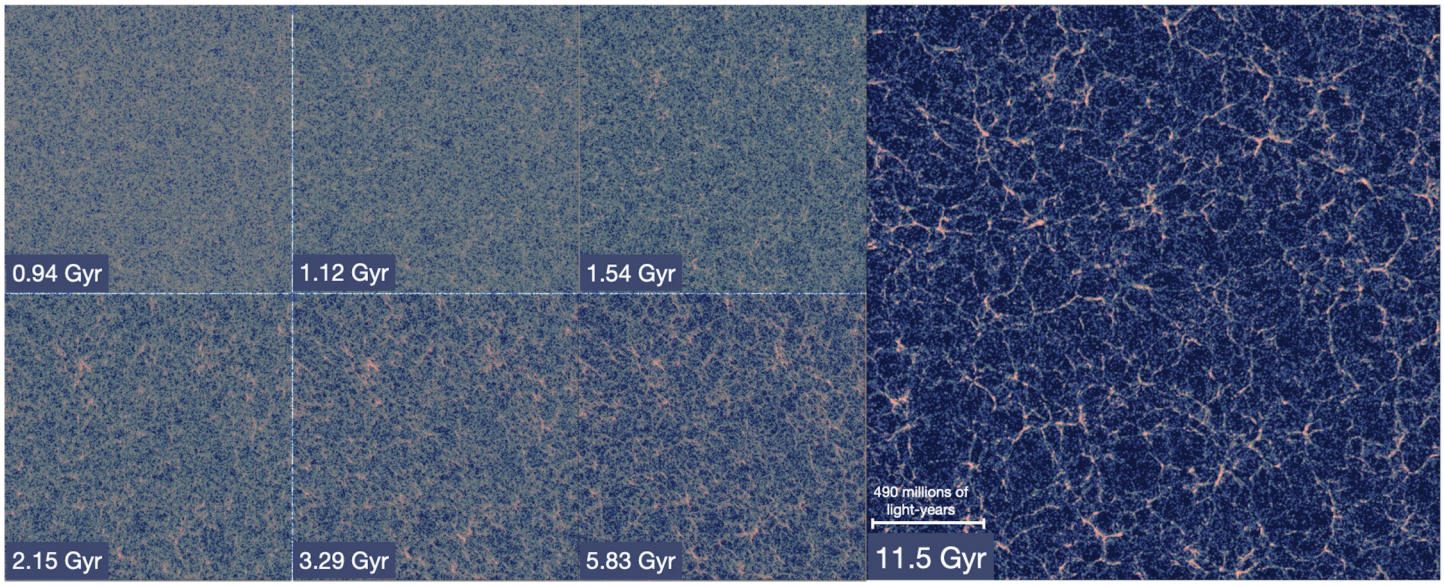

Figure 1: Different time steps showing the formation of the cosmic web of the local Universe digital twin. Each panel shows a slice of the box at a different time with a thickness of about ten million light-years. The age of the local Universe replica at that time is given in Gigayears. The size of the final box is about 2 billion light-years wide. The blueish color shows the dark matter distribution while the reddish color stands for the distribution of baryonic matter.

half were kept for reasonable storage reasons. One snapshot is indeed 8.3 TB of data distributed into 26,411 files. A hundred snapshots means almost a petabyte of data distributed into almost three millions files. The number do not include the pre-computing work required to prepare the initial conditions nor the post-computing time that has been and will be necessary to further analyze the outputs.

In parallel, the RAMSES code set up has been profiled for performance on SuperMUC-NG. A few hotspots were identified and optimized such as a reduction of compilation time and of memory footprint.

Figure 1 shows different steps of the formation of our Cosmic Home. The cosmic web becomes more and more pronounced as structures form and evolve. Filaments become more pronounced. Knots and filaments are filled with matter, galaxies included, while voids, and to a lesser extent sheets, are slowly emptied from their galaxies. Filaments constitute highways for galaxies falling into galaxy clusters that already gather hundreds of them. While the blueish color stands for the distribution of dark matter in the local Universe replica, the reddish color shows the gas.

Ongoing Research / Outlook

The next step is to analyze the simulation as an observer would and to compare it with observations in order to better understand them. To mimic observations, an observer must then be placed at the center of the box. Synthetic observations made by this observer are then derived. With the CLONE, our ultimate goal is to correct observationally derived (cosmological) parameters from any source of observational biases.

References and Links

[1] J.G. Sorce, MNRAS 478 (2018), 5199.
[2] T. Lebeau, J.G. Sorce, N. Aghanim, E. Hernandez-Martinez, K. Dolag, ArXiv231002326L, A&A *in press*
[3] N. Malavasi, J.G. Sorce, K. Dolag, N. Aghanim, A&A 675 (2023), A76.
[4] K. Dolag, J.G. Sorce, S. Pilipenko, E. Hernandez-Martinez, M. Valentini, S. Gottlöber, N. Aghanim, I. Khabibullin, A&A 677 (2023), A169.
[5] Teyssier, R. A&A 385 (2002).

1 Galaxy evolution including multiphase interstellar medium, individual stars and massive black holes

Research Institution
[1]Max Planck Institute for Astrophysics, Garching

Principal Investigator
Natalia Lahén[1]

Researcher
Thorsten Naab[1], Christian Partmann[1], Antti Rantala[1], Peter H. Johansson[2], Stefanie Walch[3], Ulrich Steinwandel[4], Jessica May Hislop[2], Constantina Fotopoulou[1]

Project Partner
[2]University of Helsinki
[3]Physics Institute, University of Cologne
[4]Flatiron Institute, Center for Computational Astrophysics

SuperMUC Project ID: pn49qi

Introduction

Globular clusters (GCs) are massive and ancient star clusters populating practically all present-day galaxies. While GCs are frequently found in the galactic outskirts, the central regions of galaxies are occupied by more massive, aptly named nuclear star clusters (NSCs), and massive black holes (MBHs). Due to their ubiquity, GCs, NSCs and MBHs are thought to originate as natural by-products of the extreme gaseous and stellar densities that occur during the assembly of galaxies. The seeds of MBHs may have formed through runaway collisions in dense proto-GCs, and the spatial coexistence of NSCs and MBHs suggests a common mass-growth scenario. Open questions related to the formation of GCs and MBHs still remain: How were the GCs and MBHs able to form already during the first half a billion years? What role did stellar radiation and metal-enrichment play in the formation of GCs and the growth the MBHs? Which came first, star clusters in the galactic nuclei or the MBHs at the centres of galaxies?

Results and Methods

Hydrodynamical simulations of galaxy evolution
This SuperMUC-NG project is a part of the Galaxy Realizations Including Feedback From Individual stars (GRIFFIN) project [1], an ongoing effort to investigate the cycle of star formation, MBH growth and related feedback in low-metallicity environments. To do so, we run high-resolution hydrodynamical simulations of gas-rich dwarf galaxies (Fig 1) with detailed prescriptions for the cooling of the interstellar medium (ISM), star formation, stellar and MBH dynamics, and feedback from stars and MBHs. Our simulation-code SPHGal [1] models the non-equilibrium cooling and heating mechanisms in the multiphase ISM with a chemistry network coupled with modern smoothed particle hydrodynamics methods. Stars between 0.08-500 solar masses are realised individually and supplemented with evolution models that provide details for the spatially resolved interstellar radiation field, photoionising radiation, stellar winds and supernova (SN) blastwaves. The physics of MBHs is modeled with a sink particle accretion prescription, that explicitly tracks the gas-flow from galactic to sub-parsec scales. Below the resolved scales, the orders of magnitude smaller, unresolved black hole accretion disc is modelled with observationally motivated analytical prescriptions. With this technique, the energy output from the unresolved MBH accretion disc ("MBH feedback") is accounted for.

Gravitational dynamics at high accuracy
For gravitational interactions of individual stars and black holes, we include the publicly available gravitational dynamics module KETJU. KETJU guarantees high accuracy in close-by star-star and star-black hole encounters using the efficient regularised integrator MSTAR [2, 3, Partmann et al. *in prep.*]. The updated MSTAR combines algorithmic regularisation, a minimum spanning tree coordinate system and the Gragg-Bulirsch-Stoer extrapolation technique. Because the integration of high-precision gravitational dynamics of all stars and black holes across the galaxy would be unfeasibly expensive, the module is spatially limited to operate inside star clusters and in the galactic centre to concentrate the computational resources where it matters most. The addition of collisional dynamics provides a natural source of stars that run or walk away from their birth-environments, directly impacting where and when massive stars release their feedback in the galactic ISM.

Tracing where and when chemical enrichment occurs
Our dwarf galaxy models are populated with tens to hundreds of millions of dark matter, gas, and stellar resolution elements. Gas and star particles, that compose the majority of the particles by number, are encoded with the chemical composition of the baryonic matter in the galaxy. In addition to the 6 chemical species traced in the ISM-chemistry, we follow the abundances of 13 individual chemical elements per particle, making the simulations quite memory-intensive. This allows us to trace how, where and when the chemical composition evolves through enrichment by stellar winds and SNe.

Star clusters within their galactic environment
In [4, 5] we followed, for the first time, the formation and evolution of entire galactic populations of star clusters realised with single stars. In addition to a number of

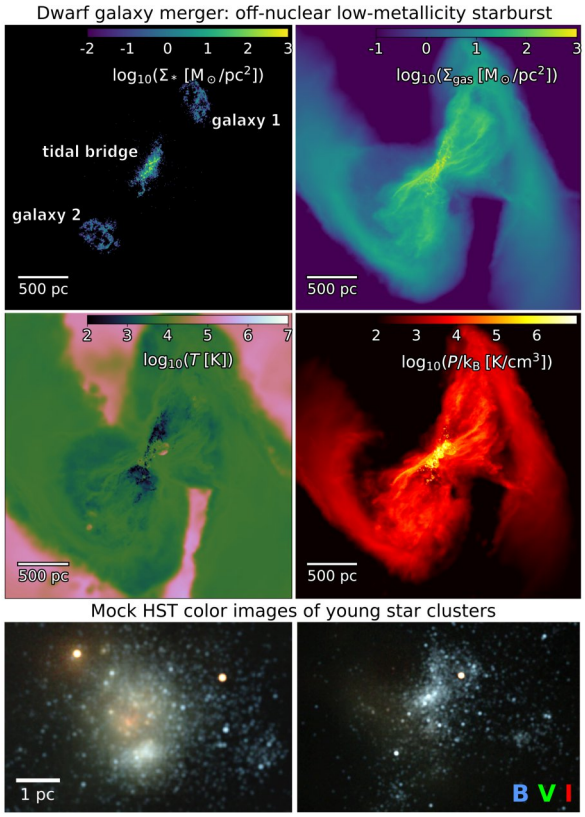

Figure 1: Top: The stellar and gaseous surface density, the temperature and the thermal pressure of the gas in a simulated off-nuclear dwarf galaxy starburst. The simulation follows 20 million gas particles, 8 million old star particles and 3.5 million new single stars. In the tidal bridge star clusters form with masses in excess of 10^5 solar masses. Bottom: Mock Hubble Space Telescope images of young massive star clusters forming in the tidal bridge.

isolated galaxy models, this SuperMUC-NG allocation has been used to simulate a galaxy merger where intense star formation takes place in the tidal bridge between the interacting galaxies (Fig. 1). In this simulation, star clusters with masses up to hundreds of thousands of solar masses form rapidly on timescales of less than 5 million years. Because massive stars live for at least a few million years, this timescale is too short for SNe to significantly enrich their nearby regions in heavy elements. Meanwhile, the formation of the most massive clusters takes long enough for the light-metal-enriched wind-material ejected by high-mass stars to be mixed with the ISM and recycled into new stars. We have thus demonstrated how stellar populations with variations in light-element abundances may emerge in nearly uniform-age star clusters, as is often observed in GCs.

Feedback-regulated growth of black holes

Recent observations have established that many dwarf galaxies host MBHs, albeit with lower black hole mass compared to more massive galaxies. In Partmann et al. in prep., we explore to which extent MBHs in the expected mass range can grow and how they coevolve with their galactic environment. We find that the growth of MBHs in dwarf galaxies is significantly limited by several physical mechanisms. Because gas accretion onto the black hole and star formation compete over the same supply of dense gas, accretion events (as shown in Fig. 2) typically lead to an increase in star formation around the MBH. The most massive among these young stars emit radiation and explode as SNe at the end of their short lifetime, which can expel gas from the galactic

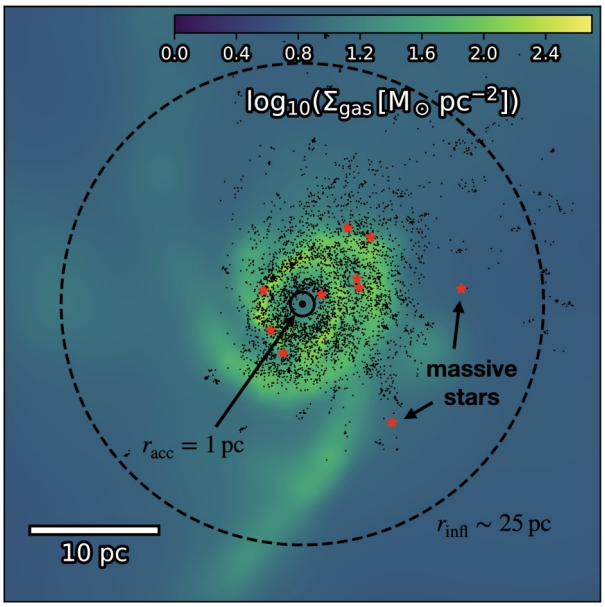

Figure 2: Gas accretion flow inside the central ~25 pc of the galaxy, where the MBH (central black dot) dominates the gravitational potential. Gravitationally trapped gas forms a disc that feeds the MBH and reaches densities that allow for the formation of stars. As soon as the first massive star (red symbols) explodes as a SN, most of the gas will be expelled from the galactic center and the MBH growth stops. The simulation explicitly resolves the accretion flow down to the accretion radius, that is set to 1 pc here.

center. Hence, dense gas around the MBH is short lived and the mass growth of the MBH is significantly limited by star formation and the associated stellar feedback. Repeated accretion and central star formation cycles lead to the formation of a small NSC. In addition to stellar feedback, the feedback from the MBH accretion disc is strong enough to stall MBH growth almost completely – a result that is consistent with observations, indicating that more extreme environments than isolated dwarf galaxies are required to grow black holes. Suitable conditions with higher gas densities and more frequent galaxy mergers might have existed in the early Universe, which we will explore in a future project.

Ongoing Research / Outlook

The SPHGal-simulations offer a large dynamic range in hydrodynamical quantities and gravitational interactions. The methods outlined above make the simulations uniquely suited for studies of star formation and chemical enrichment, as well as gaseous inflows and outflows in star clusters and in the vicinity of MBHs together with the evolving galactic environment. The unique combination of non-equilibrium ISM chemistry, chemical element tracing and collisional dynamics make the simulations computationally demanding and SuperMUC-NG has provided an invaluable resource for running the largest simulations in the ongoing GRIFFIN project. A number of simulations are still in preparation, and the data products of the simulations run so far will be used in follow-up studies of the baryon-cycle and ISM-structure in and around massive star clusters and the galactic centre.

References and Links

[1] https://wwwmpa.mpa-garching.mpg.de/~naab/griffin-project/
[2] Rantala, A. et al., 2017, ApJ, 840, 1, 53.
[3] Partmann, C. et al., *submitted to* MNRAS, arXiv:2310.08079.
[4] Lahén, N. et al., 2023, MNRAS, 522, 2, 3092-3116.
[5] Lahén, N., Naab, T., Szécsi, D., 2024, MNRAS, 530, 1, 645-667, arXiv:2402.09518.

Pulsar radio emissions by kinetic instabilities in electron-positron plasmas

Research Institution
[1]Institute of Physics and Astronomy, Univerity of Potsdam

Principal Investigator
Jan Benáček[1]

Researcher
Jörg Büchner[2,3], Patricio Muñoz[3,2]

Project Partner
[2]Technical University of Berlin, Berlin
[3]Max Planck Institute for Solar System Research, Göttingen

Funding
DFG projects BU 777-17-1 and BE 7886/2-1

SuperMUC Project ID: pn73ne

Introduction

Pulsars are neutron stars that emit radio beams out of their magnetic poles. The radiation is generated in strongly magnetized (10^4-10^{11} T) relativistic electron−positron plasma in their magnetospheres. However, the exact origin of their coherent radiation has yet to be understood. Though several radio emission models have been proposed to interpret the observed radiation properties, the models lack the inclusion of strongly nonlinear effects, resolution of kinetic micro-scales, and self-consistent particle-wave evolution, which obstructs our deeper understanding.

The pulsar radiation is driven by pair-creation events in their gap regions where electric fields can accelerate particles to ultrarelativistic energies. The particles radiate gamma-ray photons that interact with the magnetic field and create electron-positron bunches. It was theoretically proposed and shown by spark simulations that the bunches propagate along the open magnetic field lines and form a “train” of bunches as the spark event quasi-periodically repeats. In the train, they may mutually interact, expand, and can produce electromagnetic radiation.

Results and Methods

We studied the instabilities using kinetic particle-in-cell (PIC) simulations because they can resolve the nonlinear processes behind the coherent radio waves at plasma kinetic micro-scales and describe the particle-wave interaction self-consistently. We studied three promising coherent emission mechanisms and the instabilities of pulsars. For the first time, we studies obtained radiation properties (spectrum, total power, angular dependence) using the PIC simulations.

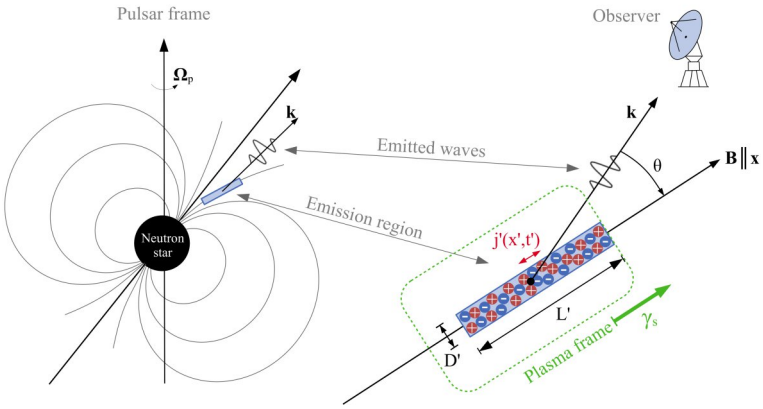

Figure 1: A scheme of the considered emission region for the linear acceleration emission. Adapted from [2].

Interacting plasma bunches forming streaming instabilities and solitons
We studied how the relativistic streaming instability by plasma bunches develops, saturates, and forms the electrostatic solitons. Because the solitons contain a large amount of electric charge, they were proposed to strongly radiate by coherent curvature radiation as they propagate at relativistic velocities along curved magnetic field lines. We showed for the first time that the theoretically predicted solitons can be formed in the plasma in 1D simulations; nonetheless, it was uncertain how the solitons evolve in more dimensional space and what the radiation properties are in the perpendicular direction.

Therefore, we run a few 2D simulations to study the soliton properties in the perpendicular direction to the magnetic field. We found that the solitons evolve very differently in 2D and 1D cases (both studied for the same plasma parameters) [1]. While in the 1D case, the solitons are formed as the instability saturates and survive until the simulation end, in the 2D case, the solitons start to dissipate after the instability saturation, and no solitons are present at the simulation end. This surprising finding can significantly affect how we think pulsar radio emissions are generated because the assumed solitons dissipate in a few microseconds.

Linear acceleration emission of the streaming instability and bunches
Linear acceleration emission (LAE) is produced by charged particles oscillating along magnetic field lines in strong electric fields of plasma instabilities. The radio emission properties of this mechanism were unknown, mostly because of the large complexity of the system, which requires superposing the emitted electromagnetic waves of all individual plasma particles according to their oscillation phase.

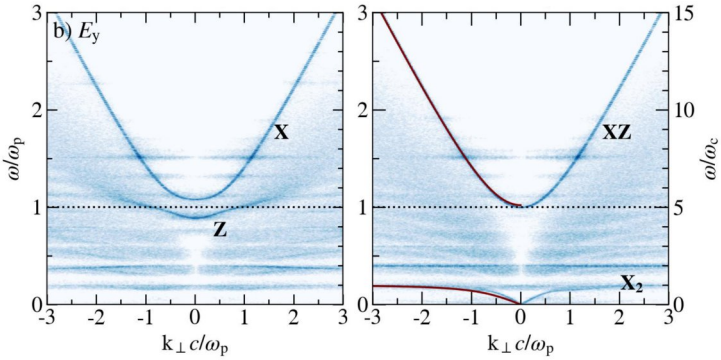

Figure 2: Comparison between dispersion properties of electron-proton (left) and electron-position (right) electromagnetic waves for electron cyclotron maser. Adapted from [3].

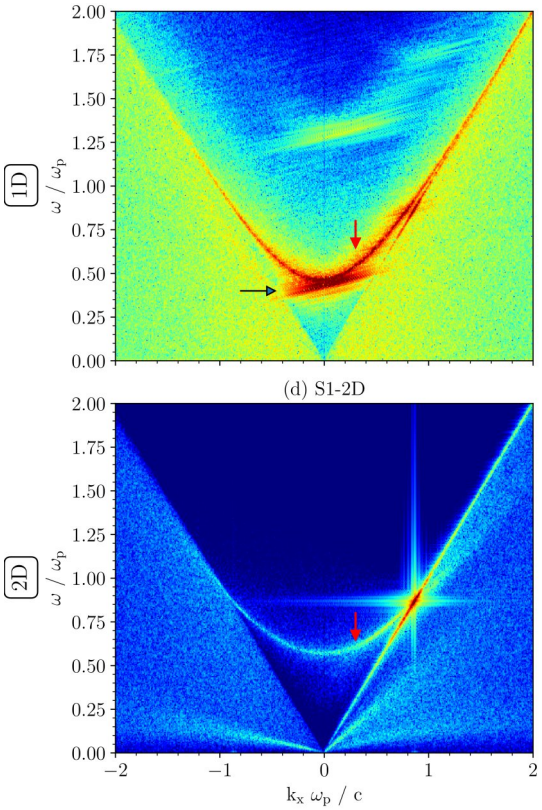

(d) S1-2D

Figure 3: Comparison between dispersion properties of 1D simulations showing solitons (black arrow, top) and 2D simulations (bottom) having only the Langmuir waves. Adapted from [1].

To study the coherent waves, we developed a new numerical approach how to calculate the radio emissions of the plasma from the PIC code by a novel approach that can utilize information on all particles in the simulation and fully considers wave phase when the emitted waves are added up [2]. In contrast to the previous approaches, our approach did not require storing all or selected particles on disks every time step for tens of thousands of time steps but rather used electric currents, which are intrinsically produced by particle motion and stored in grid cells. This approach led to a reduction of computational and storage requirements by a factor of 10^3-10^4. Application of this approach for all particles in our simulation allowed the first detailed studies of the LAE properties.

We found that the total radio emission power by the constrained plasma bunches exceeds by eight orders of magnitude the power by the expanding plasma bunch because the constrained bunches create stronger particle oscillations and currents.

Therefore, the constrained bunches are more likely to be responsible for the observed radiation than the expanding bunches if both cases exist in the pulsar magnetosphere simultaneously. We also analyzed the properties of the produced radiation for both cases and found that the radiation of plasma bunches has a few similar features to the observed pulsar radiation. The radio spectrum contains a relatively flat low-frequency part and the power-law profile for frequencies above ~1 GHz. Also, the radiation power oscillates in micro-second scales and is beamed along magnetic field lines, forming a typical pulsar radio beam.

Electron-positron cyclotron maser creating radio zebras during pulsar interpulse

Between one of the pulsar radio emission features belongs radio zebras - narrowband stripes detected in frequency-time diagrams of pulsar interpulse. The cyclotron mechanism assumes that each emission stripe is generated when the plasma frequency (dependent on the plasma density) is a resonant harmonic of the electron-cyclotron frequency (dependent on the magnetic field) by the presence of an unstable, loss-cone particle distribution.

In this sub-project, we carried out two kinds of simulations. First, a few series of 1D small-scale simulations to cover how the maser evolves for a broad range of initial parameters (thermal parameter of loss-cone, plasma-cyclotron frequency ratio, loss-cone to background density ratio) and 2) three large-scale 2D simulations that covered the directivity of the radiation, i.e., the angle at which the radio waves are emitted to the direction of the magnetic field.

We found that the initial evolution and saturation of the instability are similar to the classical electron-ion cyclotron maser. Moreover, both instabilities differ from analytical predictions, which were made only for the case in which the loss-cone density is negligible in comparison with the background plasma density. While analytical predictions showed that the most intensive electrostatic Bernstein waves have frequencies slightly above the plasma frequency and the electromagnetic X-mode waves occur at harmonic (double) plasma frequency, the waves are generated not only at these frequencies but also in a broader frequency range.

Ongoing Research / Outlook

The studied instabilities strongly depend on the properties of plasma that is released from the pulsar polar cap regions. Therefore, in our follow-up research, we will focus mostly on the dynamics of pair creation spark events and their direct relation to the emitted coherent radio waves and properties of the plasma bunches.

References and Links

[1] Benáček, J., Muñoz, P.A., Büchner, J., and Jessner, A., Astronomy & Astrophysics, forthcoming in 2024, DOI:10.1051/0004-6361/202348087.
[2] Benáček, J., Muñoz, P.A., Büchner, J., and Jessner, A., Astronomy & Astrophysics, 675, A42, 2023.
[3] Labaj, M., Benáček, J., and Karlický, M., Astronomy & Astrophysics, 681, A113, 2024.

1 Numerically probing the co-evolution of supermassive black holes and their galaxies

Research Institution
[1]Interdisciplinary Center for Scientific Computing at Heidelberg University

Principal Investigator
Tobias Buck[1]

Researcher
Aura Obreja[1], Fabrizio Arrigoni Battaia[2], Andrea V. Macciò[3,4]

Project Partner
[2]Max-Planck-Institute for Astrophysics, Garching
[3]Max-Planck-Institute for Astronomy, Heidelberg
[4]New York University Abu Dhabi, United Arab Emirates

Funding
DFG-443044596, Carl-Zeiss-Stiftung NEXUS program

SuperMUC Project ID: pn29mo

Introduction

Galaxies are the lighthouses in the vast ocean of the Universe. Their structure and appearance is shaped by a range of physical processes probing some of the most pressing questions in astrophysics. At the largest scales of megaparsecs, galaxies grow through hierarchical accretion of smaller systems with a sensitive dependence on the cosmological model. At intermediate scales of kiloparsecs, collisionless dynamics give rise to complex morphologies, offering a window to study a still unknown type of matter: dark matter. At the smallest scales of molecular clouds and star clusters (parsecs), galaxies represent laboratories for understanding the complex physics of gas cooling/heating and star formation. Understanding the formation of spiral galaxies (including our own Galaxy, the Milky Way) has been at the forefront of theoretical astrophysics for decades.

Current and up-coming observational facilities like e.g. the James Webb Space Telescope are focusing on the early epoch of galaxy formation to investigate the properties and formation sites of supper-massive black holes (SMBHs). This will not only provide unique insights into the black hole-galaxy connection, but the number of SMBHs at high redshifts can actually constrain the cosmological model.

A comprehensive theoretical picture for the co-evolution of SMBH and galaxies as well as their interactions is still missing, though it is generally recognised that BH feedback is an important factor in regulating star formation in massive galaxies. Yet, the details of the processes are still elusive.

Results and Methods

The project is based on the modern smoothed-particle-hydrodynamics (SPH) code Gasoline2 that overcomes some inaccuracies of the standard SPH formulation [1].

The code incorporates a detailed modelling of all physical processes important for galaxy formation: gas cooling including metal lines, metal production and diffusion, star formation and stellar feedback from supernovae and massive stars as well as several updates to the physics modules. In the last years, we have actively worked in improving our version of Gasoline2 by revisiting the parametrisation of some physical processes and by adding various new physics modules. Specifically we have introduced the following new modules to the Gasoline code:

- the formation, evolution and energy input from SMBHs,
- the effect of the local photoionization field (LPF) on the gas cooling [2],
- an updated prescription for the release of chemical elements synthesized inside stars [3],
- and updated numerical parameters for the star formation prescriptions.

SMBH feedback
Our code follows the creation and evolution (via gas accretion) of the central SMBH and also its energy input (via thermal feedback) into the surrounding gas. We mainly use the Bondi-Hoyle-Lyttleton parametrization to model the accretion of gas by the BH. However, we have also implemented additional accretion models such as the Gravitational Torque Accretion Model or the Viscous-Disk Accretion Model. In the former, matter loses angular momentum due to gravitational instabilities in the accretion disk that surrounds the BH with the accretion rate mainly depending on the mass and size of the

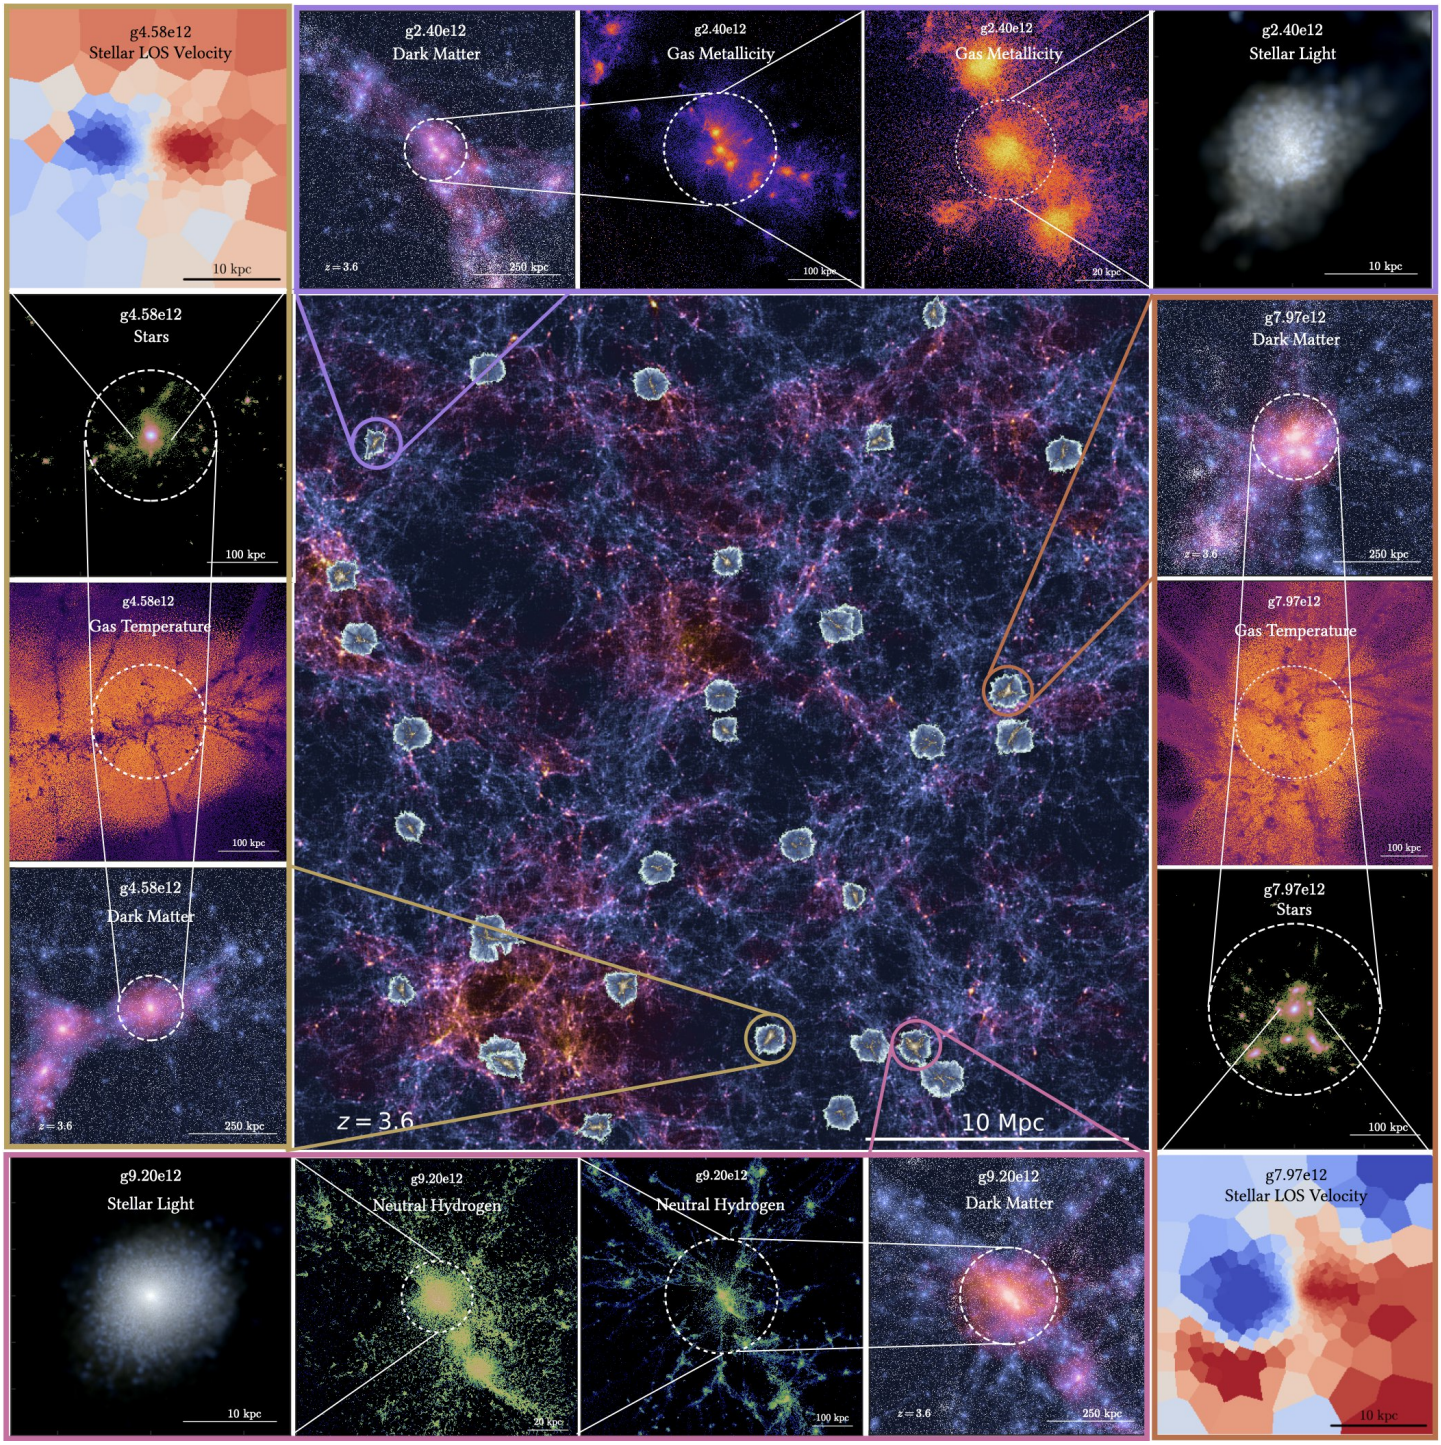

Figure 1: Highlight of the physical and numerical complexity of the "High-z Exploration of Large and Luminous Objects" (HELLO) project. The central panel shows a rendering of the dark matter distribution with the 32 zoom regions overlayed on top. Four representative zoom regions are circled and the smaller panels surrounding the central panel show continuous zooms onto the central galaxy imaging various physical quantities of the simulations thereby highlighting the large dynamical range and physical complexity of the simulation suite.

accretion disk. In the Viscous-Disk Accretion Model, the gas losses angular momentum due to viscous forces inside the accretion disk, making the accretion rate a function of the disk's surface density and its viscosity. The viscosity is calculated following the alpha-disk model, and thus depends on the speed of sound and the height of the accretion disk.

Local photoionization field (LPF)
While on large cosmological scales, the Ultra-Violet Background (UVB) can be approximated as homogeneous, on galaxy scales the high energy photon sources (stars, hot gas, BHs) are distributed highly non-uniformly. These local sources of ionisation and heating can increase significantly the cooling time of the gas, and therefore decrease the star formation rate by partially starving the supply of cold and dense gas. This feedback channel can have a large impact on how gas is distributed across the various phases, both within the body of galaxies (i.e. the morphology and kinematics of neutral hydrogen disks), as well as on their CGM.

Chemical enrichment
Galaxies and their stars are the production sides of all chemical elements more massive than beryllium in the Universe and thus a prerequisite for the formation of rocky planets and life. Accurately capturing and explaining the detailed features in the observational data is notoriously difficult for state-of-the-art numerical models. We have incorporated improved prescriptions for time resolved chemical enrichment into the simulations. The novel chemical enrichment prescription offers several new pathways for studying galaxy formation across cosmic time. Especially, observations of emission and absorption lines of the gaseous components of galaxies and their surroundings at different cosmic times will offer new insights into the physical mechanisms of feedback and galactic winds enriching the CGM and ISM.

Results
Within our LRZ project of 40 Mio. CPUh we have simulated 32 AGN host galaxies to be published under the name "High-z Exploration of Large and Luminous Objects" (HELLO). Each simulation evolves a total of about 107 particles from the Big Bang to about half the age of the Universe following all relevant physics on scales of entire galaxies (mega parsecs) down to scales of star formation (several tens of parsecs), covering ~5 orders of magnitude in dynamical range. Our simulations are in excellent agreement with various scaling relations observed with the Hubble Space Telescope and the James Webb Space Telescope. HELLO galaxies not only capture the heightened star formation rate characteristic of early star-forming galaxies, but also elucidate the role of active galactic nuclei feedback in shaping the galaxy evolution trajectory. In one publication [4], we have used our new simulations to explore the direct effects that AGN radiation has on gas heating and cooling. Assuming the AGN radiation to impact the CGM anisotropically, within a bi-cone of angle α, we found that even a relatively weak AGN can significantly lower the fraction of halo gas that is catastrophically cooling compared to the case of gas photoionized only by the UVB. We have further used our model in accordance with high redshift observations to determine free model parameters and among others found that the bi-cone opening angle takes most likely values around $\alpha = 60$ degrees. For our identified combination of parameters, the CGM mass catastrophically cooling is reduced by half with respect to the UVB-only case, with roughly same mass out of hydrostatic equilibrium heating up and cooling down, hinting to the importance of self-regulation around AGNs. This study showcases how CGM observations can constrain not only the properties of the CGM itself, but also those of the AGN engine.

Parts of the simulation suite have further been used to infer the spin parameter, or normalized angular momentum λ, of the dark matter halo of the Milky Way from observational data. We have presented a first attempt to measure λ starting from the dynamically distinct disks and stellar halos identified in high-resolution cosmological simulations with our Machine Learning tool Galactic Structure Finder. Additionally, we have theoretically investigated the impact of early massive mergers on the chemical evolution of Milky Way-like galaxies [5]. Recent observations of the Milky Way found an unexpected steepening of the star-forming gas metallicity gradient around the time of the Gaia-Sausage-Enceladus merger event. We found a strong steepening of the gas metallicity gradient at early times for all four galaxies in our sample which is caused by a sudden increase in the cold gas disk size (up to a factor of 2) in combination with the supply of un-enriched gas by the merging dwarf galaxies. We find that such mergers can contribute significantly to the formation of an observable bimodality in stellar disk chemistry as it is observed in the Milky Way.

Ongoing Research / Outlook

As modern telescopes begin to probe ancient galaxies right up to the frontier of reionisation, it is critical that numerical simulations similarly sharpen their gaze at the early epochs of our Universe. As such, HELLO is a welcome endeavor to better inform our knowledge of the dispositions of the galactic ancestors of extant galaxies like our Milky Way (MW). The HELLO galaxy simulations achieve a desired resolution of roughly twice that of its hydrodynamical code's predecessor during the evolution of MW-like galaxies. With increasing numbers of seemingly overmassive black holes being discovered at extreme redshifts, it is crucial to further investigate the characteristics of galaxies during the early Universe to help understand the nature of coevolution between black holes and their host galaxies. Armed with a firm test bed of realistic simulations of our Galaxy at early times, HELLO provides a path to improved insights into the environments of ancient massive black holes. Thus, the HELLO suite will serve for various scientific studies over the next years and we anticipate that it will produce a similar scientific impact as its predecessor the NIHAO suite.

References and Links

[1] Wadsley, J. W., et al., MNRAS, 471 (2), OUP, 2357–2369, 2017.
[2] Obreja, A., et al., MNRAS, 490 (2), OUP, 1518–1538, 2019.
[3] Buck, T., et al., MNRAS, 508 (3), OUP, 3365–3387, 2021.
[4] Obreja, A., et al., MNRAS, 527 (3), OUP, 8078–8102, 2024.
[5] Buck, T., et al., MNRAS, 523 (1), OUP, 1565–1576, 2023.

Internal Gravity Waves Excited by Hydrogen Burning

1

Research Institution
Heidelberg Institute for Theoretical Studies

Principal Investigator
Friedrich Röpke

Researcher
Johann Higl,Giovanni Leidi

Project Partner
—

Funding
Klaus Tschira Foundation ; DFG Grant RO 3676/3-1

SuperMUC Project ID: pn25bo

Introduction

Stars are the building blocks of the visible Universe, which drive the galactic chemical evolution and act as observational tracers of the evolution of the cosmos as a whole. Yet models of stellar structure and evolution rely on parametric models of multi-dimensional phenomena, because many multi-dimensional processes operate on timescales that are up to 10 orders of magnitude smaller than the nuclear timescale, which is dominating stellar evolution during most of a stars' lifetime. The parametrizations of multi-dimensional processes are often based on simplistic assumptions and include free parameters that are adjusted to match observational properties of stars. In this project, we performed three-dimensional simulations of convective core hydrogen burning in main sequence stars, which allow us to test parametrized models and assumptions of two important stellar phenomena.

The most stringent tests of solar models are derived from helioseismology, which has been used to investigate the internal structure of the Sun for decades by observing tiny brightness fluctuations on the stellar surface. Therefore, the free parameters of convection are rather well known for the Sun. This, however, is not the case for most other stars. While the advancement in space-born observations has allowed us to use similar techniques to investigate the structure of distant stars, the seismological data is missing many details due to the larger distance.

Recent systematic observations, e.g., show a power excess in low frequency oscillations (red noise) of intermediate to high mass stars on the main sequence. This feature cannot be explained by coherent oscillations, which indicates that it is driven by a stochastic excitation mechanism. The fact that red noise is seen in stars with and without convective envelopes strongly suggests that it is caused by internal waves created at the boundary of the convective core. This is an intrinsically multi-dimensional process, which cannot be captured by standard 1D stellar evolution models. The shape of the wave excitation spectrum is, hence, highly uncertain. Theoretical predictions of the excitation spectrum from convective regions predict a steep power law. Other simulations using 3D spherical geometry, on the other hand, find excitation spectra significantly shallower than predicted. However, the results of these simulations are likely affected by the high viscosity and artificially increased luminosity of the numerical setup in addition to the anelastic approximation.

We also found shallow excitation spectra in our fully compressible 2D simulations of a star with three times the mass of the Sun without any explicit viscosity but with an increased luminosity [2]. The simulations performed in this project aim to confirm this result in a fully compressible 3D simulation at the nominal stellar luminosity.

The other aspect we were investigating is the mixing of material across the convective boundary of convective cores, which is arguably the free parameter with the largest uncertainty in 1D stellar models. The evolution of stars more massive than the Sun is largely determined by the mass of their convectively mixed core, such that a more massive star with a smaller mixing parameter can have the same observable surface properties as a less massive star with larger mixing values. However, the respective age of these stars will be different. The degeneracy of mass and age has far-reaching consequences for other fields like the characterization of exoplanets and galactic archaeology.

Furthermore the mixing during the early stages of stellar evolution, like hydrogen burning on the main sequence, has a huge impact throughout the remaining evolution up until the death of the star. The question if a star will eventually explode as a supernova is to a large extent determined by the size of the helium core at the end of the hydrogen burning phase.

Results and Methods

For our simulations we use the Seven-League hydrodynamics (SLH) code, which is specifically tailored for low Mach number hydrodynamics. This is necessary since the flow velocity during core hydrogen burning is several orders of magnitude smaller than the local sound

speed. To ensure stability and accuracy in our simulations, we also had to utilize well-balancing techniques that stabilize the background state, effectively mitigating spurious oscillations and maintaining numerical robustness. Moreover, SLH relies on an implicit timestepping algorithm, which enhances the efficiency and stability of our simulations, particularly in low Mach number flows.

Due to the stochastic nature of convective turbulence it was necessary to follow the flow for several characteristic timescales in order to extract a steady excitation spectrum of internal waves. Another issue with internal wave simulations at the nominal luminosity is the resolution of low frequency waves, because the radial wavelength of this type of waves becomes very short and therefore a high resolution in the radial direction is needed. We hence used a spherical grid with 1,152×128×256 cells. Our simulations were highly parallelized and utilized 4,608 cores in pure MPI mode. The final simulation covers a timescale spanning approximately 300 days and used over 14 million core-hours. The output was stored on 7 TiB of disk space in approximately 8,000 output files spread over the simulation time. The long simulation time in combination with the large amount of output files allows the detailed analysis of both high- and low frequency internal waves.

Initial results indicate that a combination of at least two theoretical models is needed in order to reproduce the simulated excitation spectrum.

Previous studies have analyzed the mixing behavior at convective boundaries by performing a series of simulations with artificially increased luminosity, which can then be used to extrapolate to the nominal luminosity case [e.g. 3]. Even though we could directly use our nominal luminosity simulation to get the mixing rate, we decided to also perform simulations with a 1,000 and 10,000 fold increased luminosity. The increased luminosity leads to a flow with a larger Mach number, but does not reach the trans-sonic regime. Therefore, the computational cost for these simulations is significantly reduced. These runs are not only important for comparability reasons with previous studies, but also served as our final preparation step before we started the nominal luminosity run.

Ongoing Research / Outlook

We are currently in the process of comparing excitation spectra generated from both two- and three-dimensional simulations. By analyzing these spectra and interpreting them based on theoretical models, we aim to make predictions for the excitation spectrum of stars with different masses, which can then be tested with further simulations. We are also working on a paper draft about these results, which we plan to submitt to a peer-reviewed journal.

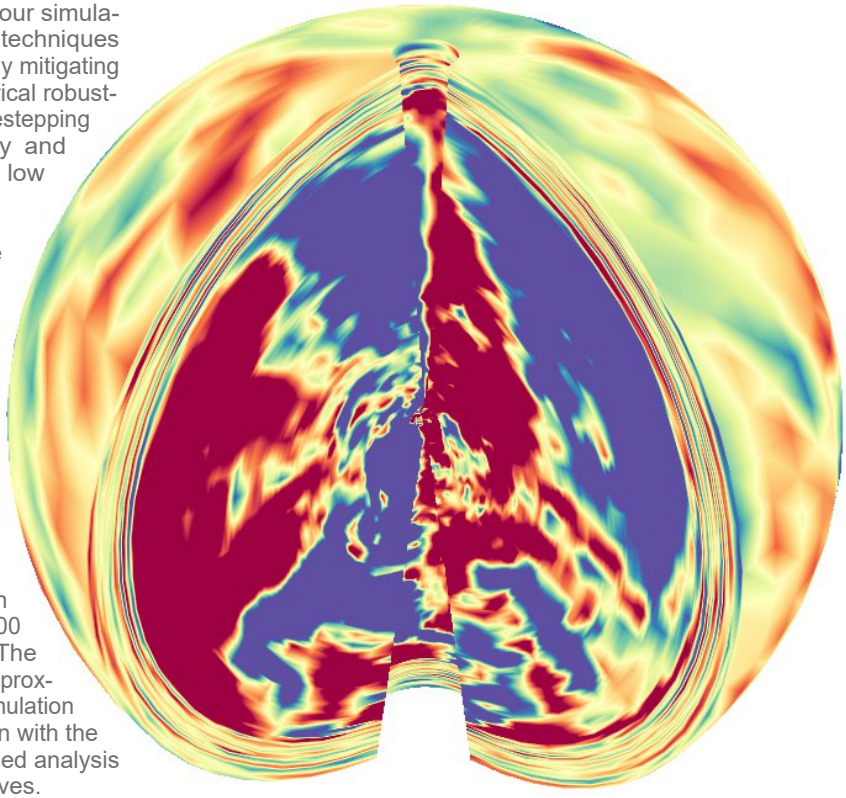

Figure 1: Visualization of the tangential velocity component in the nominal luminosity case. The central region shows the convective motion of the stellar core. The region around is dynamically stable and allows internal waves to propagate. The surface of the sphere shows the typical velocity pattern of propagating waves.

In the future, we would also like to asses the observability of certain features in the spectra. This requires an equilibrium state between the excitation and dissipation of the internal waves. Due to the long travel times of internal waves this would take 10 to 100 times longer simulations than we performed here, which is not achievable with current computing systems. However, studies showed that non-linear wave features do not play a significant role in the observed spectra, i.e. the current practice of using linear stellar oscillation codes to extract oscillation frequencies from one-dimensional stellar models can hardly be improved by multi-dimensional simulations. However, oscillation codes cannot predict the expected amplitudes of standing waves without additional information about the driving mechanism. The excitation spectra we extract from our simulations can be used exactly for that purpose, which will give us more insight into ways of detecting signatures of convection in asteroseismic observations.

References and Links

[1] www.h-its.org
[2] Horst, L., et al. 2020, A&A, 641, 18.
[3] Horst, L., et al. 2021, A&A, 653, 55.

The smallest galaxies providing big insights into dark matter

1

Research Institution
[1]Institute for Astronomy, University of Hawaii
[2]Max-Planck-Institut für Astrophysik

Principal Investigator
Thales Gutcke[1], Natalia Lahen[2]

Researcher
Volker Springel[2], Rüdiger Pakmor[2], Thorsten Naab[2], Azadeh Fattahi[3], Finn Giddings[1], Jessica Hislop[4]

Project Partner
[3]Institute for Computational Cosmology, University of Durham
[4]Department of Physics, University of Helsinki

Funding
NHFP #HF2-51480, STScI, DFG, EXC-2094 – 390783311, ORIGINS

SuperMUC Project ID: pn73we

Introduction

This project targets various central questions in modern astrophysics, including “What is the nature of dark matter?”, “How do galaxies form and evolve?” and “Do we understand the extremes of the universe?”. Dwarf galaxies provide a natural laboratory for confronting these questions as we will explain below. The formation of dwarf galaxies tracks an extreme situation in various ways. Dwarf galaxies are assumed to be the very first type of galaxy to form in the earliest Universe. This means they form from gas clouds with the simplest chemical composition and begin forming stars at the first moment when this becomes physically possible. This makes them the smallest entities in the Universe that are capable of forming stars. It also means that studying the smallest dwarf galaxies can tell us about the conditions and physical processes at work in the very beginning of the Universe. In the structural build-up of galaxies, it is assumed that galaxies merge and form hierarchically. This means that small galaxies form first and then merge to create larger galaxies. Thus, dwarf galaxies are the building blocks of all other galaxies. They are also the most abundant type of galaxy in the Universe. All of these characteristics of dwarf galaxies make them highly sensitive to the nature of dark matter, too. The aim of this project is to create a highly accurate cosmological, hydrodynamical simulation model that can produce extremely realistic representations of dwarf galaxies right into the centers, where dark matter models can be tested.

Results and Methods

The first result of this project is a suite of six highly realistic dwarf galaxy simulations. The computational requirements are significant due to the precision gas physics calculations and the fact that every star is created and tracked individually. Of special note are the individual supernova explosions of dying stars that cause atomic bomb-like explosions, heating and moving gas clouds around at high speeds (see Fig. 1). These features make the newly developed code [2] state-of-the-art, since it employs various numerical and astrophysical modeling techniques that have never before been used in cosmological simulations. The precision we aim for requires that the model is able to resolve individual supernova blast waves. Preliminary work has shown that this necessitates a mass resolution of 4 solar masses, meaning that a single dwarf galaxy is resolved with between 0.5-2 billion resolution elements. Thus, we are constrained to run all simulations at this resolution or better if we want to retain this vital feature of the model. We have carried out six ultra-high resolution cosmological zoom-in simulations of dwarf galaxies in that have masses of a few billion solar masses. These simulations constitute the highest resolution dwarf galaxies run fully cosmologically across the age of the Universe to date. The galaxies were selected to form in isolation and distanced from larger neighbors, the simulations probe universal aspects of cosmological evolution and of the effects of an external radiation field, without being influenced by effects from the environment or infall into larger halos. In the recently published study, we strove to validate the model. We studied the present-day properties of the simulated galaxies and compare the global stellar properties with measurements from Local Group dwarf galaxies. We show that the stellar mass, stellar size, stellar chemical composition, kinematics and morphology are well matched by our simulations [3]. These results provide the basis on which we can proceed to study alternative dark matter in these galaxies.

Ongoing Research / Outlook

There exists a discrepancy between the inferred dark matter density profiles for observed dwarf galaxies and simulated ones (see Fig. 2). Observations seem to indicate a flatter inner profile that standard simulations produce. This is a an ongoing area of active research, and various solutions have been suggested. To retain the current standard cosmological model, internal energetic processes such a supernova and other such processes have been proposed to play an important role in the flattening of the profile. As well, the formation channels for nuclear stellar clusters in the centers of galaxies are still under discussion, with the infall of

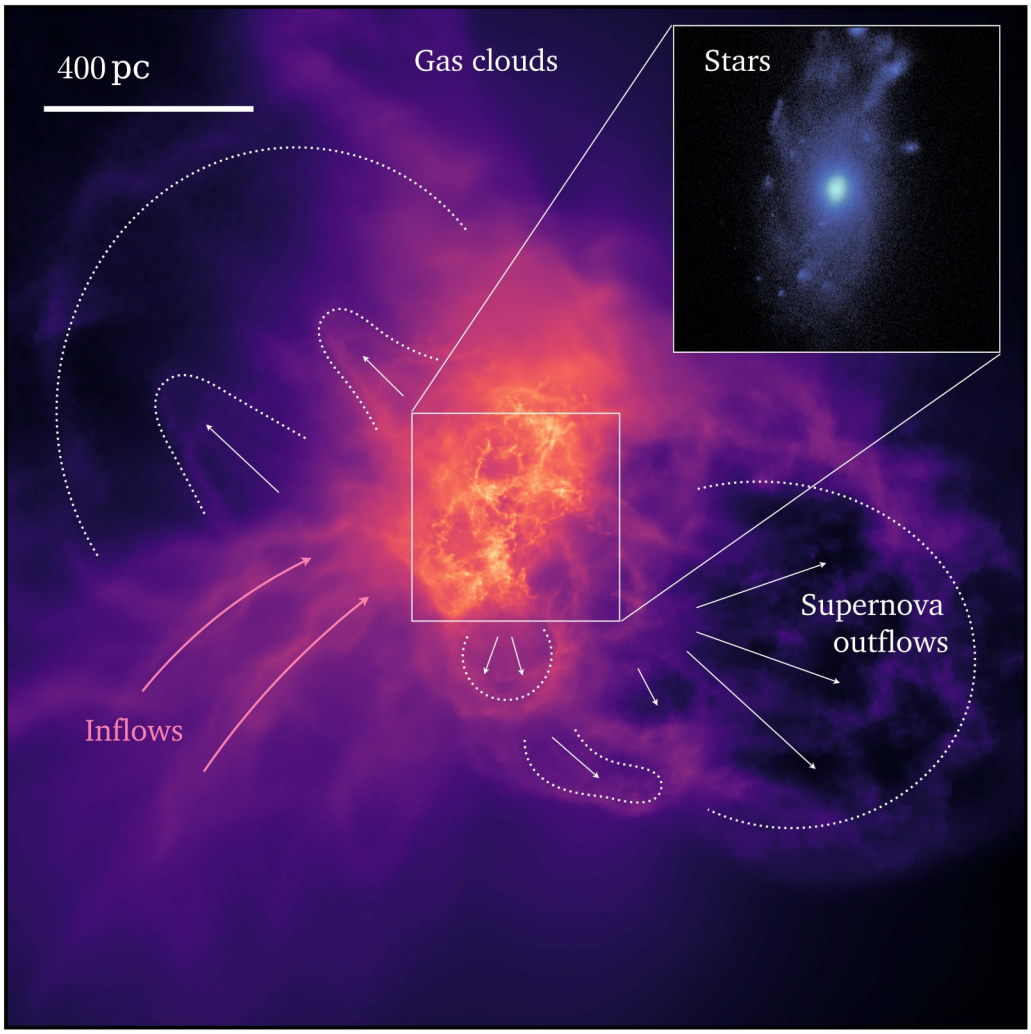

Figure 1: Gas clouds around a dwarf galaxy forming in the early Universe. Due to supernova explosions in the central region, gas is being accelerated and pushed outwards from the center at high speeds of hundreds of km per second.

globular clusters or the in-situ star formation being the main options discussed. In our ongoing work, we study the formation of dark matter density profiles in our simulated dwarf galaxies. Our findings so far show no signs of the formation of a core in the sample of simulated galaxies. However, we see the impact of the addition of gas and supernova in the dark matter profiles in these simulations. For one of the halos, it is possible to see an increase in the dark matter density in the inner regions. We argue that this is explained by a significant star formation event that might also be linked to the subsequent formation of a nuclear stellar cluster in the central region. While the flatter profiles we see in standard cosmology may be still reconcilable with better understanding the gas physics, self-interacting dark matter (SIDM) is a favored family of dark matter models that may provide an alternative explanation for inconsistencies. However, since gas physics is equally active in an SIDM universe, it remains crucial to test these models with the same set of detailed baryonic processes. Self-interacting dark matter (SIDM) postulates new gauge bosons, which opens up a hidden sector of dark matter models. The boson facilitates annihilation and self-scattering between dark matter particles, altering the distribution. Our model is the ideal testbed for this study, since the small-scale gas physics is highly resolved. We have already performed the dark matter-only simulations and there is a clear flattening of the dark matter profile in the SIDM simulation (se Fig. 2, red region). We expect to present the results including gas physics that predict the dark matter profile shape for galaxies 4 orders of magnitude smaller than recent studies [4].

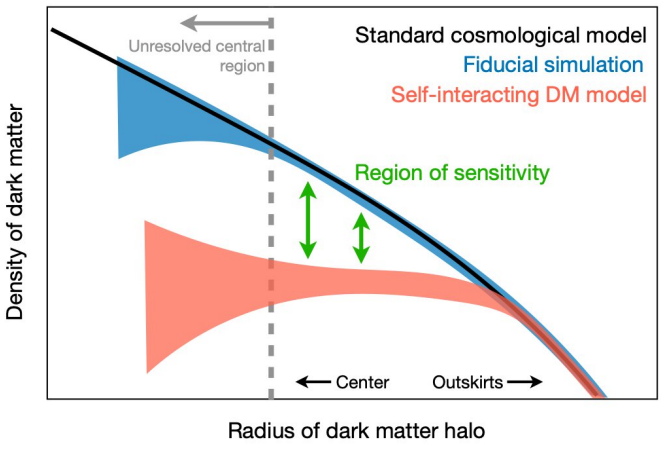

Figure 2: The centers of dwarf galaxies are a sensitive probe of different dark matter models. Our preliminary results show a lower central density for an alternative model: self-interacting dark matter.

References and Links

[1] http://thalesada.github.io/lyra
[2] Gutcke T. A., Pakmor R., Naab T., Springel V., 2021, MNRAS, 501, 5597.
[3] Gutcke T. A., Pakmor R., Naab T., Springel V., 2022, MNRAS, 513, 1372.
[4] Correa C. A., Schaller M., Ploeckinger S., Anau Montel N., Weniger C., Ando S., 2022, MNRAS, 517,3045.

1 Insights on the origin of ORCs from cosmological simulations

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

Principal Investigator
Klaus Dolag[1]

Researcher
Ludwig M. Böss[1], Bärbel S. Koribalski[3], Ulrich P. Steinwandel[4], Milena Valentini[1]

Project Partner
[2]Australia Telescope National Facility, CSIRO
[3]Center for Computational Astrophysics, New York

Funding
DFG EXC-2094 – 390783311 (ORIGINS), ERC-2019-AdG 882679 (COMPLEX)

SuperMUC Project ID: pr86re (Gauss Large-Scale project)

Introduction

Odd Radio Circles (ORCs) are a newly discovered class of radio sources (see Figure 1 and references in [1]), showing large scale, ring-like diffuse radio emission with a diameter of hundreds of kpc, around central, elliptical galaxies, without any detected counterparts at non-radio wavelength. One far-reaching possibility is that these are merger-driven shocks, similar to what is observed as radio relics in galaxy clusters, propagating away into the outer circum-galactic and intergalactic medium [1].

Results and Methods

We performed a high-resolution, non-radiative simulation of a galaxy with a virial mass of $M \sim 1.2 \times 10^{12} M_{\odot}$, taken from the COMPASS [2] simulation set, which realizes zoomed initial conditions for various halos from very massive galaxy clusters down to normal galaxies selected within a 1 Gpc parent simulation. Here we used the asin region at 25,000x resolution level, resulting in a gas particle mass of $M_{gas} = 6.3 \times 10^3\ h^{-1}\ M_{\odot}$ and a softening of $\epsilon = 110\ h^{-1}$ pc so that the halo is resolved with 20 million gas particles within the virial radius at redshift $z = 0$. Resolving such shocks at large distances from the center of the halo either needs some shock refinement, or, as in our case, employing a very high number of resolution elements to increase the number of particles tracing the low density regions around galactic halos. The simulation has been performed as non-radiative versions using OpenGadget3, which is an advanced version of P-Gadget3, featuring an updated SPH scheme. The high number of particles used allows us to focus on the thermodynamic structures in the outskirts of the halo with very high resolution. Here the simulation contains almost 40 million gas particles within the accretion shock region and due to the adaptive nature of the simulation, the resolution at the accretion radius varies between 10 and 30 h^{-1} kpc and the spatial resolution at the radius of the observed radio ring is typically a factor 3 better. To be able to track the shock structures (see Figure 2) within the halo, we applied an on-the-fly shock finder [4]. To model synchrotron emission we use Crescendo, a Fokker-Planck solver for spectral cosmic rays in cosmological simulations [5], which we applied in a post-processing mode.

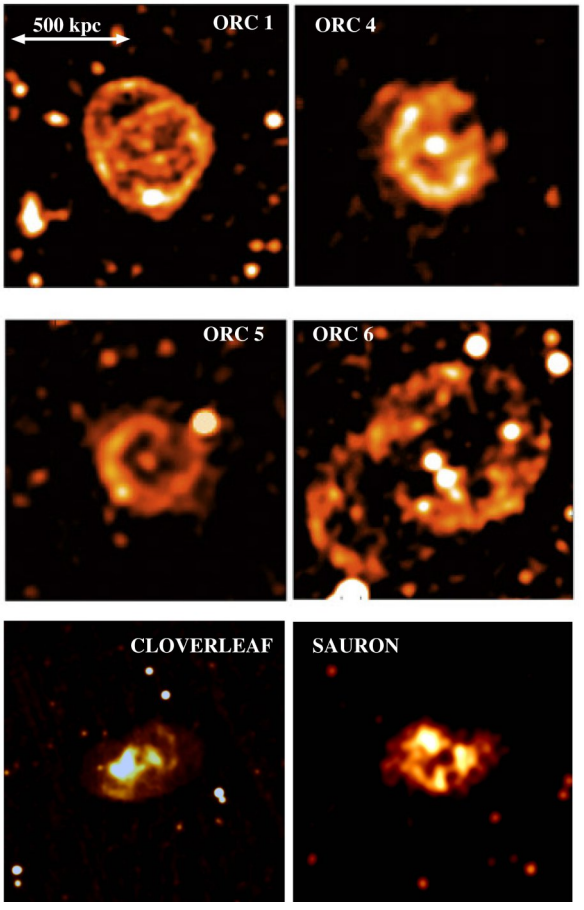

Figure 1: Collection of observed odd radio circles through different radio telescopes. From top left to bottom right: ORC 1 (MeerKAT), ORC 4 (GMRT), ORC 5 (ASKAP), ORC 6 (ASKAP), CLOVERLEAF (ASKAP) and SAURON (MeerKAT). See [1] for details and references.

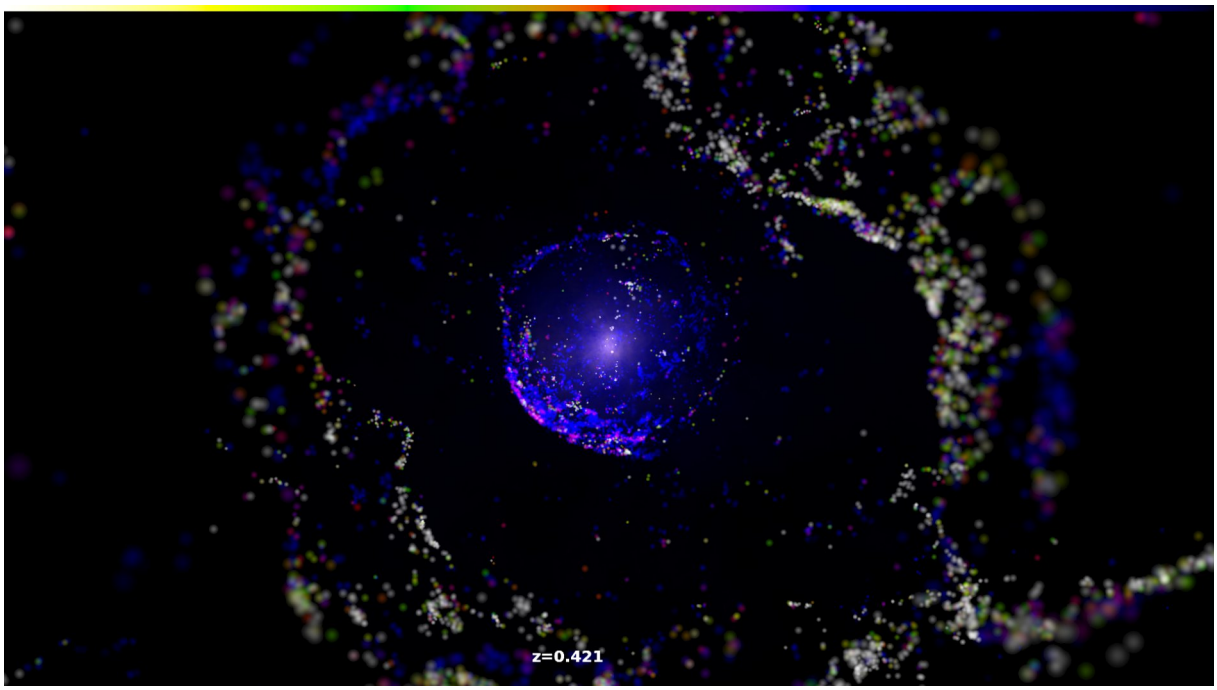

Figure 2: Ray tracing Image (Movie [3] available at [2]) of a 1.3 h^{-1} cMpc region using Splotch. Shown is a composition of the hot gas (mostly visible in the central part, color-coded by temperature) with the individual, shocked resolution elements color-coded by their Mach number $\mathcal{M}$ (where blue, red, yellow and white correspond to $\mathcal{M}$ ~ 2, 3, 4 and > 5, respectively).

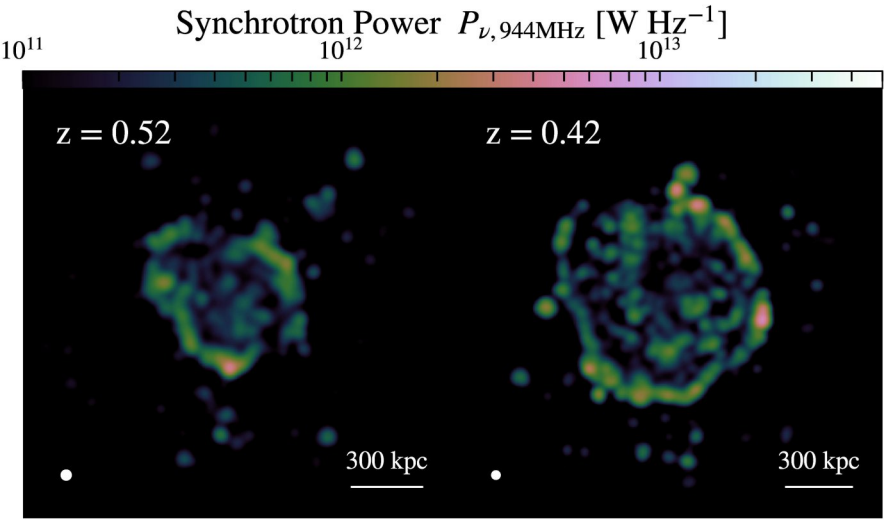

Figure 3: Two examples for the radio synchrotron emission at 944 MHz of the simulated system at z = 0.52 and z = 0.42. The images are smoothed with a beam, corresponding to the observations.

These structures obtained could indeed be the origin of a newly detected radio emission feature (see Figure 3). The best matching morphology found in our galaxy simulation originates from a very-extreme merger event, where the final halo mass of the galaxy increases by a factor of three due to multiple merger components. By identifying similar extreme merger cases in a cosmological simulation, we deduced that only 5% of galaxies undergo such dramatic events, which would be generally in line with the rare number of observations of such ORCs. This scenario also predicts that the galaxies showing these radio emission phenomena should be predominately quiescent galaxies. Furthermore, our scenario of internal shocks also agrees well with the radio Mach number of $\mathcal{M}$ ~ 2.1 – 2.4 inferred from observations of ORC 1.

Ongoing Research / Outlook

Recently, one new ORC system was discovered at (cosmologically speaking) very close distance to us. This opens the window to utilize X-Ray telescopes to try to detect the hot gas around these galaxies (observational campaigns are currently ongoing). On the other hand, improving the theoretical understanding of the processes leading to these systems need even higher resolution and the inclusion of many more physical processes directly within the simulations, making this to a challenging endeavor for next generation of HPC systems.

References and Links

[1] Dolag et al., ApJ 945 (2023) 182
[2] http://www.magneticum.org/complements.html#Compass
[3] https://wwwmpa.mpa-garching.mpg.de-HydroSims/Magneticum/Images_Movies/asin_25000_igm_shocks_combined.avi
[3] Beck et al., MNRAS 458 (2016) 2080.
[4] Böss et al., MNRAS 519 (2023) 548.

The Young and the Wild: What happens to Protoclusters forming at z ~ 4?

1

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

Principal Investigator
Klaus Dolag[1]

Researcher
Rhea-Silvia Remus[1], Helmut Dannerbauer[2]

Project Partner
[2]Instituto de Astrofísica de Canarias, Tenerife

Funding
DFG EXC-2094 – 390783311 (ORIGINS), ERC-2019-AdG 882679 (COMPLEX)

SuperMUC Project ID: pr83li

Introduction

Overdensities of galaxies at very high redshifts have been observed in increasingly large amounts in the last few years, reaching redshifts as high as z = 6 and more. Assuming that those massive agglomerations of galaxies are the cores of structures that will collapse into very massive galaxy clusters at present day, these structures have been named protoclusters. Some of these observed protocluster reach masses high enough to challenge predictions from ΛCDM cosmological simulations. For example the two massive protoclusters observed at z ~ 4, namely SPT2349-56 at z = 4.3 with a total mass of more than 1 x $10^{13}M_{\odot}$ and the even more massive protocluster reported at z = 4.0 with a total mass above 4 x $10^{13}M_{\odot}$. Both of these protocluster have large numbers of member galaxies with extremely high total star formation rates of more than 6,000 $M_{\odot}$/yr.

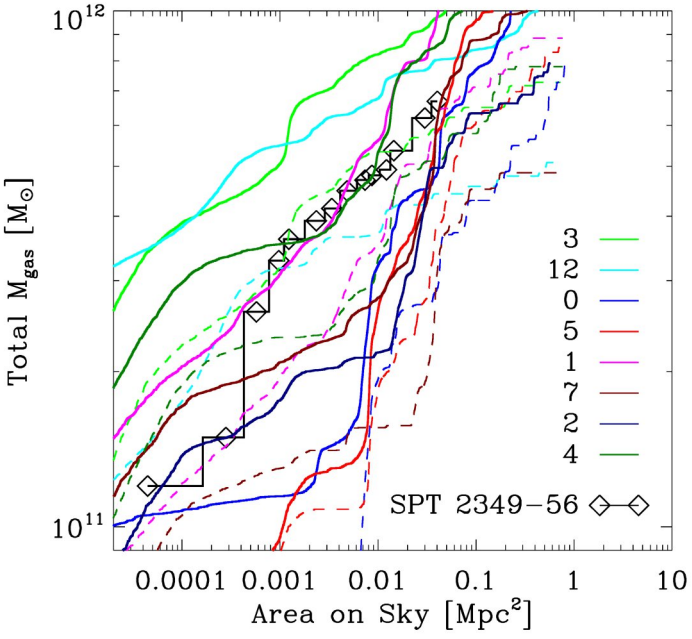

Results and Methods

To find protoclusters at high redshift comparable in mass to those observed recently, a large simulation volume is required which also follows galaxy formation physics. For this study (for details see [1]) we use one of the largest volumes (e.g. *Box2/hr*) from the cosmological, hydro-dynamical simulation suite *Magneticum Pathfinder* [2]. Here, the resolution is high enough to resolve galaxies down to baryonic masses of $M_{bar} > 10^{10}M_{\odot}$, e.g. high enough to well represent the galaxies currently observed in protoclusters, and the simulations box-size of $(909\ Mpc)^3$ is large enough to have such massive structures already formed within the young universe. This simulation is resolved with 2 x $2{,}880^3$ particles, giving a particle mass resolution of mDM = 9.8 x $10^8M_{\odot}$ and mGas = 2.0 x $10^8M_{\odot}$ for dark matter and gas, respectively. Since each gas particle can spawn up to four stellar particles during its lifetime, the mass of a stellar particle is $m_{star} \sim 5 \times 10^7M_{\odot}$. For dark matter and gas particles the same softening is used, with $\epsilon_{DM} = \epsilon_{Gas} = 5.0$ kpc, while for the stars a softening of $\epsilon_{star} = 3$ kpc was adopted. We searched for protocluster at z = 4.2 based on identified bound structures within our simulation, ranked according to their total mass. This provides us with 42 structures with total masses above $M_{tot} = 1 \times 10^{13}M_{\odot}$, from which we selected a set of 8 optimal objects (see Figure 2) which showed extreme values in total mass, number of member galaxies or star-formation rate. Analyzing these structures in detail we could show tha we find several virialized structures at z = 4.2 with a similar number of member galaxies (richness) and the same dynamical properties of the member galaxies as the observed protocluster SPT2349-56 at z = 4.3. The simulations also predict that several of the member galaxies of these structures are fast rotating systems, in agreement with the observed findings for member galaxies of SPT2349-56. However, while the general stellar and total masses of the simulated protoclusters and the number and dynamics of member galaxies within resemble the observations closely, the instantaneous integrated starformation rate within such simulated protoclusters is a factor ~ 2 … 3 smaller than the observed values, although the simulations are reproducing

Figure 1: Total gas mass versus area-on-sky around the most massive member galaxy for the eight example protoclusters. The solid lines show the gas content within the individual galaxies, while the dashed lines show the total gas mass, including the diffuse gas between the galaxies. For comparison, the values for the protocluster SPT2349-56 at z = 4.3 are shown as black diamonds and solid line.

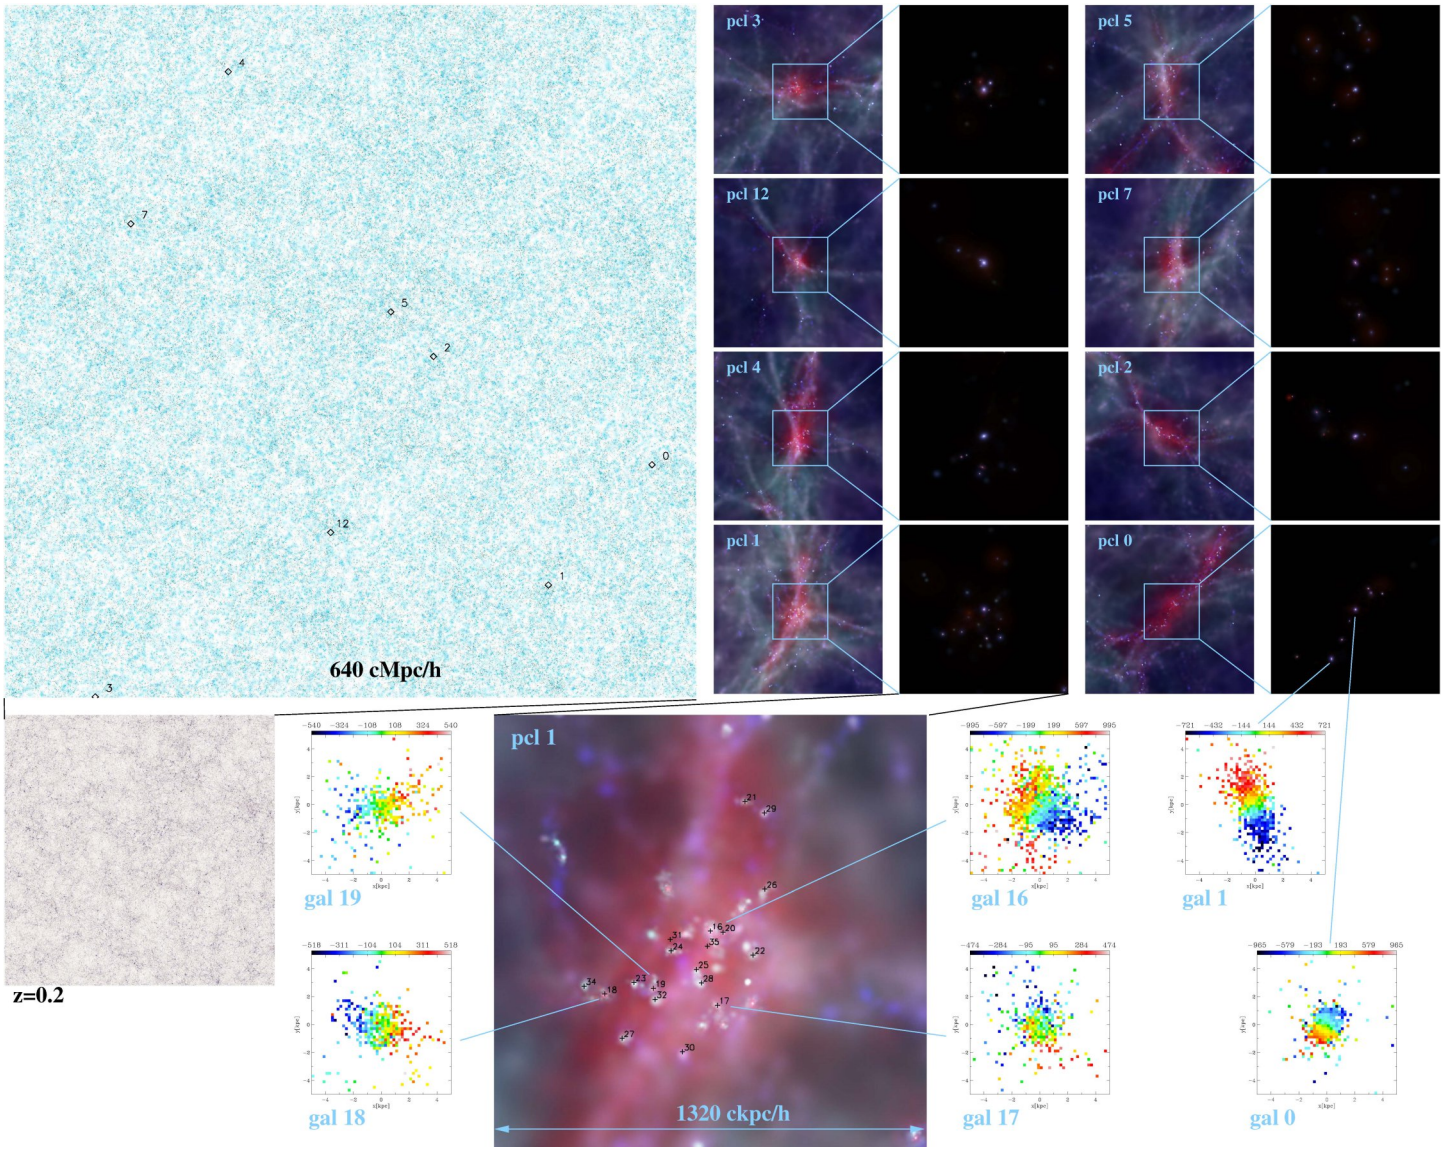

Figure 2: Upper left: Stellar distribution of Magneticum Box 2b at z = 4:2, with the labeled locations of the protoclusters PCl 3/12/0/5/1/7/2/4. Upper right: Zoom-in on the 8 selected protoclusters, with the left panels showing the gas content, from cold (blue) to hot (red), and the right panels showing a zoom on the central stellar components within 1,875 kpc co-moving, which corresponds to a physical box length of 353.77 kpc, with the colors marking the age of the stars (from young (blue) to old (red)). Lower right: Larger version of the central area of PCl 1 with both gas and stars shown in the same color scheme as in the small panels, with a box-length of 353.77 kpc. The galaxies are labeled according to their stellar mass. In addition, the rotation maps of the gas are shown for six example galaxies, four from PCl 1 and two from PCl 0. Lower left: Stellar distribution of the Magneticum Box 2b volume at z = 0.

the observed integrated gas mass within the protoclusters. As the simulations also reasonably well reproduce the main sequence of star-forming galaxies as well as the observed stellar mass function up to this redshift, this indicates that the star-formation in the simulations lacks the ability to reproduce higher star-formation efficiency (or accordingly lower depletion timescales) at least in certain environments. The simulations predict that the star-forming gas in protoclusters at redshifts of z ~ 4 is already enriched to roughly solar values. Part of this cold gas is expected to be subsequently heated by feedback and become part of the intra-cluster medium of the forming galaxy cluster. About half of the gas within these structures at z = 4.2 is already significantly heated to temperatures around 1keV, and a very small fraction (~ 2%) of this hot gas is already enriched to one tenth of the solar value. Using the full power of the simulation, we traced the evolution of these protoclusters and refuted the hypothesis that the extremely massive structures found at high redshifts really are the progenitors of the most massive galaxy clusters at present day. In fact, from our 8 examples chosen to be a among the top in these measures at z = 4.2, none is among the 10 most massive clusters at z = 0 and one of them evolves barely into a very low mass cluster.

Ongoing Research / Outlook

We currently extend this study to the a new cosmological simulation which is constrained to resembles the Local Universe (SLOW, [3]), where we can predict what type of protocluster many of the well-known galaxy clusters in our local neighborhood have been at early times.

References and Links

[1] Remus et al., ApJ 950 (2023) 191.
[2] http://www.magneticum.org
[3] Dolag et al., A&A 677 (2023) 169.

Simulating the LOcal Web (SLOW) – III: Synchrotron

1

Emission from the Local Cosmic Web

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

Principal Investigator
Klaus Dolag[1]

Researcher
Ludwig M. Böss[1], Ulrich P. Steinwandel[2], Elena Hernández-Martínez[1], Benjamin Seidel[1], Jenny G. Sorce[1]

Project Partner
[2]ICenter for Computational Astrophysics, New York
[3]Centre de Recherce en Informatique, Signal de Automatique de Lille, France

Funding
ANR-21-CE31-0019 / DFG 490702358 (LOCALIZATION), ERC-2019-AdG 882679 (COMPLEX)

SuperMUC Project ID: pn68na (Gauss Large-Scale project)

Introduction

The observations of radio relics and radio haloes within galaxy clusters unveil that efficient acceleration of Cosmic Ray (CR) electrons provide powerful tracers of the intra-cluster medium's magnetic field. The canonical process of CR acceleration in galaxy clusters is diffusive shock acceleration, where the electrons or protons are accelerated from the thermal pool by scattering off magneto-hydrodynamic (MHD) turbulence up- and downstream of shock fronts, gaining energy at every crossing from up- to downstream. Such processes also promise the possibility to trace the even more dilute medium within cosmic filaments, which are expected to be embedded within cosmological accretion shocks.

We used a constrained, cosmological simulation of the Local Universe [1,2] to study such potential emission from CRs accelerated at shocks in merger and accretion-shocks of galaxy clusters and accretion-shocks around cosmic web filaments. For this, we realized the largest, cosmological MHD simulation with an on-the-fly Fokker-Planck solver to model the spectra of the CR electron and proton population and its evolution within the cosmic structures [2]. Our $(750\ \mathrm{Mpc})^3$ large simulation volume of the Local Universe was sampled with $2\mathrm{x}3{,}072^3$ particles.

Results and Methods

To model CRs we employ the on-the-fly Fokker-Planck solver Crescendo [4], which attaches population of CR electrons and protons to all resolution elements of the simulation, using 24 bins for electrons and 8 bins for protons to sample the momentum space of their distribution function. These populations are then evolved for every of our 30 billion resolution elements in time by solving the diffusion-advection equation in the two-moment approach. Taking advantage of an on-the-fly shock detection scheme we account for the injection of CRs at shocks via diffusive shock acceleration (DSA), implying a sonic Mach number - dependent acceleration model which takes into account the different acceleration efficiencies for CR electrons and protons depending on the angle between shock and magnetic field. In addition, CRs evolve following adiabatic changes due to density changes of the surrounding thermal gas and energy losses of electrons due to synchrotron emission and inverse Compton scattering of cosmic microwave background photons.

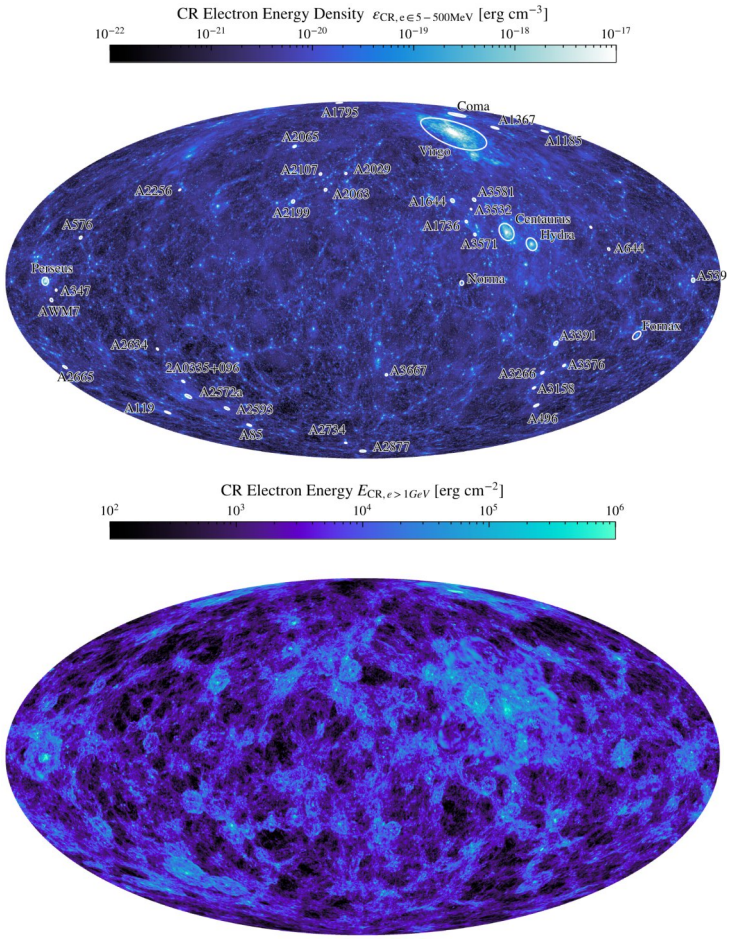

Figure 1: Full-Sky projections of the CR electron component, showing the long-lived electrons in the range E ∈ [5 − 500] MeV (top) and the fast cooling ones with energies E > 1 GeV (bottom). Taken from [3].

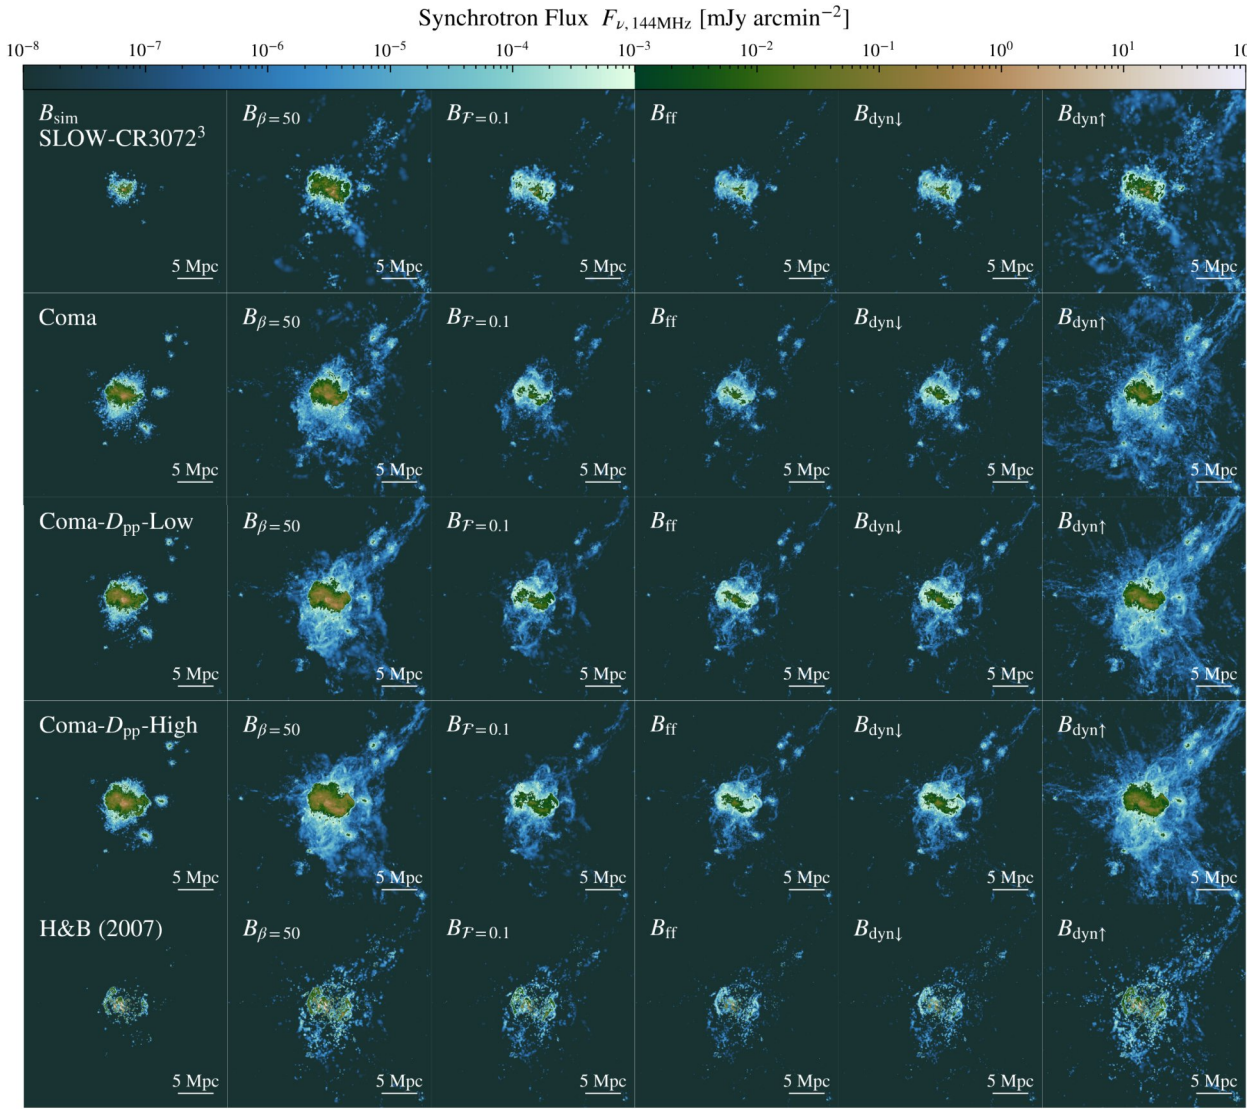

Figure 2: Synchrotron flux of the Coma cluster region at 144 MHz (LOFAR). From top to bottom different CR electron acceleration simulations are shown, while from left to right the magnetic field models are varied. The lowest row shows the result for painting the synchrotron emission following a simplified model in a post-processing approach. Taken from [3].

Figure 1 shows full-sky projections of the predicted CR electrons, where the simulation counterparts of well-known galaxy clusters within the Local Universe are labeled. In addition, we performed a higher resolution, zoom-in simulation of the Coma galaxy clusters, where we also followed turbulent re-acceleration of CR electrons. Since the synchrotron emissivity scales strongly with the magnetic field strength, the observability of synchrotron emission of the cosmic web is tightly linked to the magnetic field strength in filaments. We tested six different models for the magnetic field configurations within the parts of the cosmic structures where due to resolution limits we did not fully resolve the magnetic field amplification through to the dynamo effect. This allowed us to study the threshold for obtaining considerable diffuse emission from cosmic web filaments and we showed that with our choice of CR injection parameters we require fields of the order B ~ 100 nG to obtain considerable emission which can be detected by the next generation of radio instruments. Figure 2 shows the impact of the choice of parameters within our acceleration model as well as the magnetic field models on the emission predicted for the zoom-in simulations of our Coma cluster replica.

Ongoing Research / Outlook

We are currently investigating the radio emission predicted for a set of individual, prominent galaxy clusters within the Local Universe which have counterparts within our simulation to shed light on their actual, cosmological formation history. In addition, we are also investigating the predicted gamma ray signal of the CR proton component. As this expected signal is not yet detected, it is vital to compare the predictions from the simulation to the current, upper limits from the observations, as this can give further hints on the yet not fully understood processes of accelerating CR electrons and protons through cosmological shocks within galaxy clusters and the large scale structures.

References and Links

[1] Sorce et al., MNRAS 478 (2018) 5199.
[2] Dolag et al., A&A 677 (2023) 169.
[3] Böss et al, arXiv231013734B (2023).
[4] Böss et al., MNRAS 519 (2023) 548.

Zooming in on Galaxy Clusters

Research Institution
[1]Leibniz-Institut für Astrophysik Potsdam
[2]Niels Bohr Institute, Copenhagen

Principal Investigator
Thomas Berlok[1,2]

Researcher
Christoph Pfrommer[1], Ewald Puchwein[1], Rainer Weinberger[1], Joseph Whittingham[1], Martin Sparre[1], Larissa Tevlin[1], Lorenzo Perrone[1], Rüdiger Pakmor[3], Rosie Talbot[3]

Project Partner
[3]Max-Planck Institute for Astrophysics, Garching

Funding
ERC-ADV PICOGAL 101019746, DFG FOR 5195, DFG 444932369, MSCA-HE-2022 101106080

SuperMUC Project ID: pn68cu

Introduction

Galaxy clusters are the largest gravitationally bound objects in the Universe and provide the opportunity to study cosmology, structure formation, and plasma astrophysics. The intracluster medium is an ionized gas which permeates the space between galaxies in galaxy clusters. It has a very low density but due to its sheer volume it nevertheless contains more mass than all the stars found in the galaxies themselves. The intracluster medium has a very high temperature and consequently emits copious amounts of X-rays that transport energy out of the system. This cooling appears in many clusters to be offset by injection of energy by a super-massive black hole which resides at the center of the cluster, a so-called active galactic nucleus (AGN). How the energy in relativistic jets emitted by the AGN is transported from the center and ultimately transferred to the ambient intracluster medium is not yet well understood. Answering this question is important for understanding why the intracluster medium does not collapse and why some clusters are colder at their centers. The role of mergers, i.e., galaxy clusters colliding due to their gravitational attraction, also needs to be understood. Galaxy cluster mergers are the most energetic events in the Universe and can violently disturb the gas. A fraction of this energy is used for accelerating relativistic particles (so-called cosmic rays), injecting turbulence and amplifying magnetic fields. These are all topics that we are addressing with our SuperMuc-NG allocation. Modern cosmological simulations self-consistently include mergers (by evolving the hierarchical collapse of dark matter) and include increasingly sophisticated models for the AGN feedback. They are therefore the natural tool for understanding the combined influence of the above-mentioned processes. A limitation of cosmological simulations is, however, that they treat the plasma as a non-relativistic, collisional fluid, even though the intracluster medium is so hot that particles rarely collide and AGN jets contain magnetized and relativistic particle populations. Using our SuperMuc-NG allocation, we are improving on the state-of-the-art using novel implementations of weakly collisional Braginskii viscosity [1] (which reduces gas mixing), cosmic ray protons [2] (which exert a pressure on and heat the ICM), and cosmic ray electrons [3] (which emit synchrotron radio radiation) in the moving mesh code AREPO [4].

Results and Methods

Using the SuperMUC-NG allocation, we have performed a cosmological simulation that models a large volume containing 272 massive galaxy clusters (see Figure 1). We have then selected 25 clusters for re-simulation using the zoom-in technique (see Figure 1 for an illustration of this technique and a look at the cosmic web in our parent simulation). We are using the AREPO code for these simulations [4]. AREPO is a moving, adaptive mesh code for modeling cosmological fluids. The moving mesh allows AREPO to perform especially well if a system is moving through space with a large bulk velocity. AREPO includes algorithms for solving the magnetohydrodynamic equations which govern the cosmological fluid (coupled to the gravitational dynamics of cold dark matter) and the astrophysics of galaxy formation is included via the IllustrisTNG model [5]. Our large parent simulation contains roughly two billion resolution elements, and our highest resolution zoom-in simulations have 13,824 times better mass resolution than the parent simulation. This means that each of our high-resolution zoom-ins contain more than a billion resolution elements. These simulations probe the regime of very massive galaxy clusters at a resolution not previously achieved with the AREPO code. An example of one of these simulations is shown in Figure 2. The data for this figure is from a simulation that used 8,448 cores and required 3.3 M core-hours. In addition to the significant computational cost, these simulations also require significant storage capabilities, as we use frequent outputs so that the time evolution of the merger seen in Figure 2 can be studied in detail (e.g. 30 TB for this simulation). The simulations would thus not have been feasible without our SuperMUC-NG allocation. We are also using these new simulations to study turbulence generation, amplification of magnetic fields and the role of galaxy feedback for these processes. We illustrate this part of the project in Figure 3.

Ongoing Research / Outlook

The first set of production simulations, which use ideal MHD and the IllustrisTNG galaxy formation model [5], are currently completing on SuperMuc-NG. These simulations form the foundation on which we are building

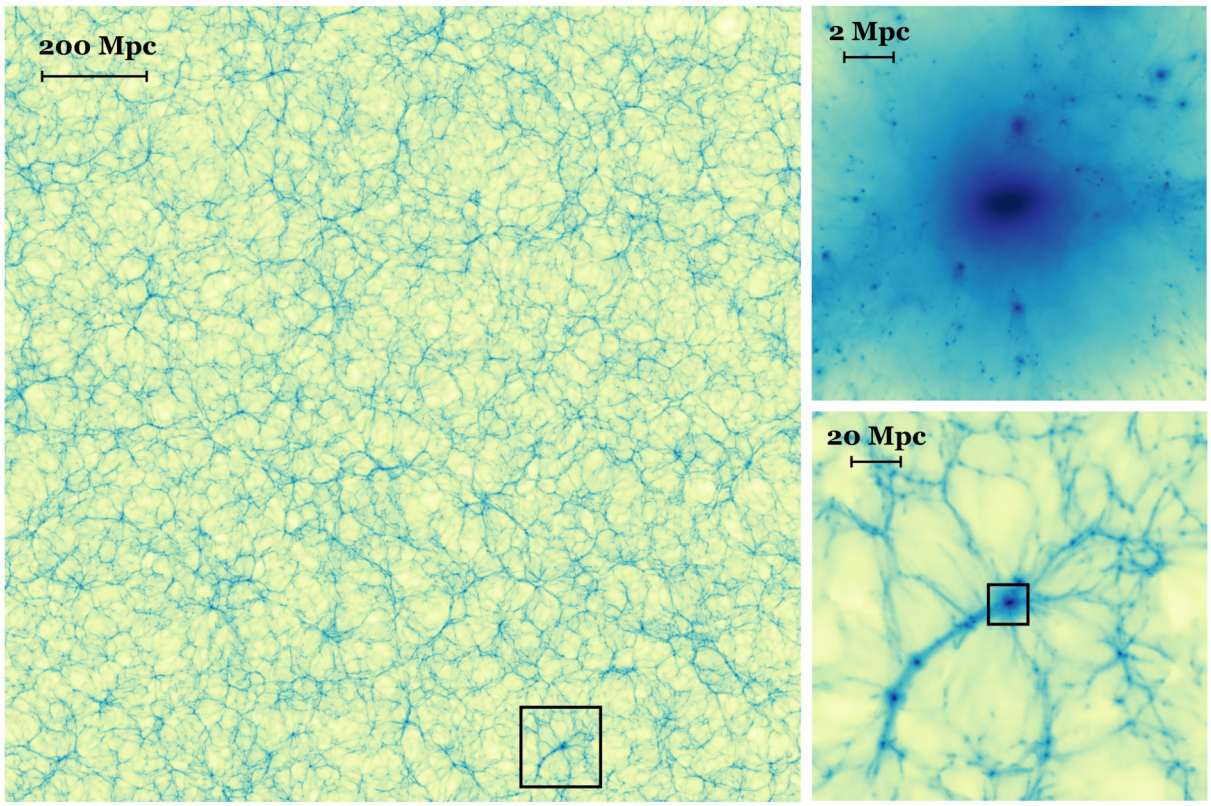

Figure 1: Illustration of the zoom-in technique for cosmological simulations. A large volume (left), containing 272 massive galaxy clusters, is simulated. Subsequent zoom-in simulations focus on single objects (upper right), which allows the simulations to reach a much better spatial resolution.

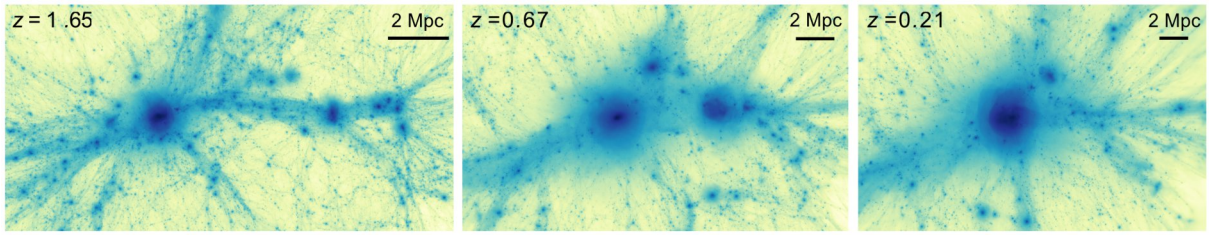

Figure 2: A major merger of two galaxy clusters from one of our high-resolution, non-radiative cosmological simulations (shown here in projected gas density). The merger results in shocks moving through the intracluster medium, which are visible as sharp jumps in density at z = 0.21. The shocks power acceleration of electrons, which leads to radio synchrotron emission co-located at the shocks (known as radio relics). The merger also injects additional gas turbulence, which drives a magnetic dynamo.

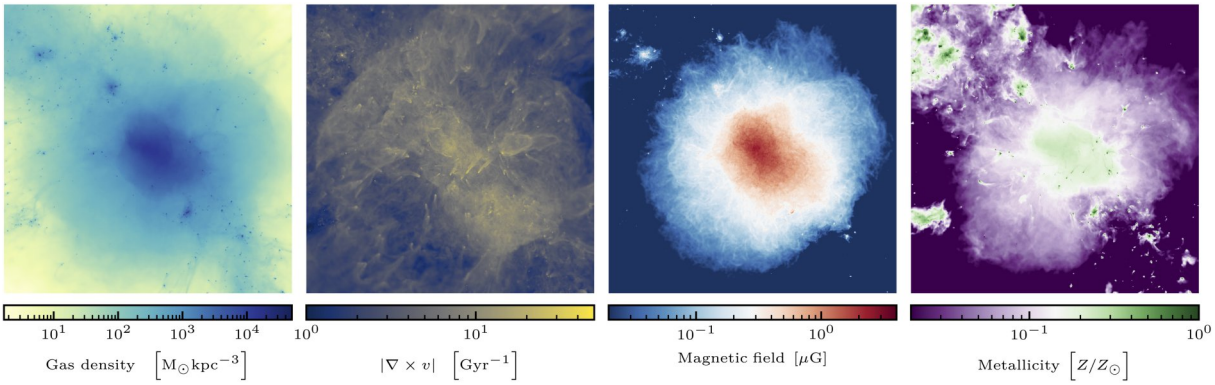

Figure 3: A massive galaxy cluster today (redshift zero) in a simulation with galaxy formation physics turned on. The small dense clumps in the density panel are individual galaxies, which drive turbulent wakes visible in a projection of the vorticity (a proxy for turbulence, second panel). The turbulence amplifies the magnetic field strength to observed microgauss levels (third panel). Finally, the galaxies distribute metals via winds and stripping, building up a global radial gradient in metallicity. The panels are 15 Mpc wide.

novel simulations that include Braginskii viscosity, cosmic ray protons and spectral electrons and a sophisticated AGN-jet. These improvements to the physical models will allow a more direct comparison with radio and X-ray observations.

References and Links

[1] Berlok, T., Pakmor, R. & Pfrommer, C. MNRAS, p. 2732, 2019.
[2] Pfrommer, C., Pakmor, R., Schaal, K., Simpson, C. M. et al. MNRAS, 465:4500, 2017.
[3] Winner, G., Pfrommer, C., Girichidis, P. & Pakmor, R. MNRAS, 488(2):2235, 2019.
[4] Springel, V. MNRAS, 401(2):791, 2010.
[5] https://www.tng-project.org/

Unravelling the Energy Cascade in Supersonic Magnetohydrodynamic Turbulence Simulation

1

Research Institution
[1]Research School of Astronomy and Astrophysics, Australian National University
[2]Department of Astrophysical Sciences, Princeton University
[3]Canadian Institute for Theoretical Astrophysics, University of Toronto

Principal Investigator
James R. Beattie[1,2,3], Christoph Federrath[1]

Researcher
Ralf S. Klessen[4,5], Salvatore Cielo[6], Amitava Bhattacharjee[2]

Project Partner
[4]Universität Heidelberg, Zentrum für Astronomie, Institut für Theoretische Astrophysik
[5]Universität Heidelberg, Interdisziplinäres Zentrum für Wissenschaftliches Rechnen
[6]Leibniz-Rechenzentrum

SuperMUC Project ID: pn73fi

Introduction

Compressible, trans-to-supersonic, turbulence is ubiquitous across many scales in our Galaxy. Unlike subsonic turbulence, supersonic turbulence is not described by a single energy cascade. The scale that marks the transition from supersonic to subsonic turbulence, the sonic scale, splits the energy spectrum into two separate power laws. This is important in the context of turbulence-regulated star formation models [1], where the sonic scale defines the critical value for which regions inside of molecular gas clouds collapse under their own gravity to form stars. The sonic scale was recently measured for the first time within a hydrodynamical supersonic turbulence simulation using a grid resolution of $10{,}080^3$ in [2]. This is the only calculation with sufficient resolution to separate the energy injection scale, supersonic cascade, sonic scale, subsonic cascade and dissipation scale of the turbulence. Whilst this calculation brought great insight into the exact position of the sonic scale and the nature of supersonic turbulence, the influence that magnetic fields have on not only the sonic scale, but the two cascades, is more or less unknown, yet is of tantamount importance for understanding the nature of supersonic, magnetized turbulence in our Galaxy.

In this LRZ project, we run and process the first supersonic, magnetised turbulence simulation that is sufficiently resolved to measure the exact position of the (magneto)-sonic scale, a number of other important scales, and the supersonic and subsonic energy cascades, simultaneously. We do this by running a magneto-hydrodynamical (MHD) turbulence simulation with $10{,}080^3$ grid cells, distributed over almost 140,000 compute cores and run for over 80 million compute-core hours on SuperMUC-NG.

Results and Methods

Numerical and high-performance computing methods
We use a modified version of the MHD code flash [3]. Our code uses a highly-optimised, hybrid-precision, second-order, positivity-preserving MUSCL-Hancock HLL5R Riemann scheme [4] to solve the ideal,

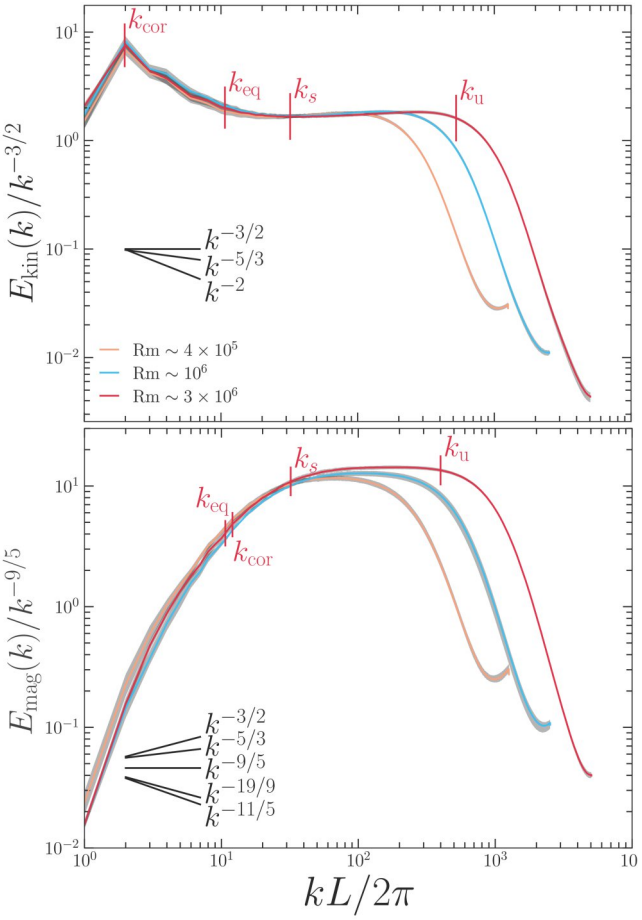

Figure 1: The kinetic (top) and magnetic (bottom) energy spectra for three runs at increasing resolution: $2{,}520^3$ (yellow), $5{,}040^3$ (blue) and $10{,}080^3$ (red), labelled by their estimated magnetic Reynolds numbers (Rm). The kinetic energy spectrum shows a dichotomy: $\propto k^{-2}$ (Burgers-type turbulence) on supersonic scales and $\propto k^{-3/2}$ (Iroshnikov-Kraichnan-type turbulence) on subsonic scales. The magnetic spectrum exhibits a single power law developing only at very high resolutions, $\propto k^{-9/5}$, which is currently unexplained by any theories of high Rm turbulence. The outer scale k_{cor}, magnetic-kinetic equipartition scale k_{eq}, sonic scale k_{sonic} and microscale k_u are annotated on the spectra, highlighting the large range of scales that are captured by these extremely high-resolution simulations.

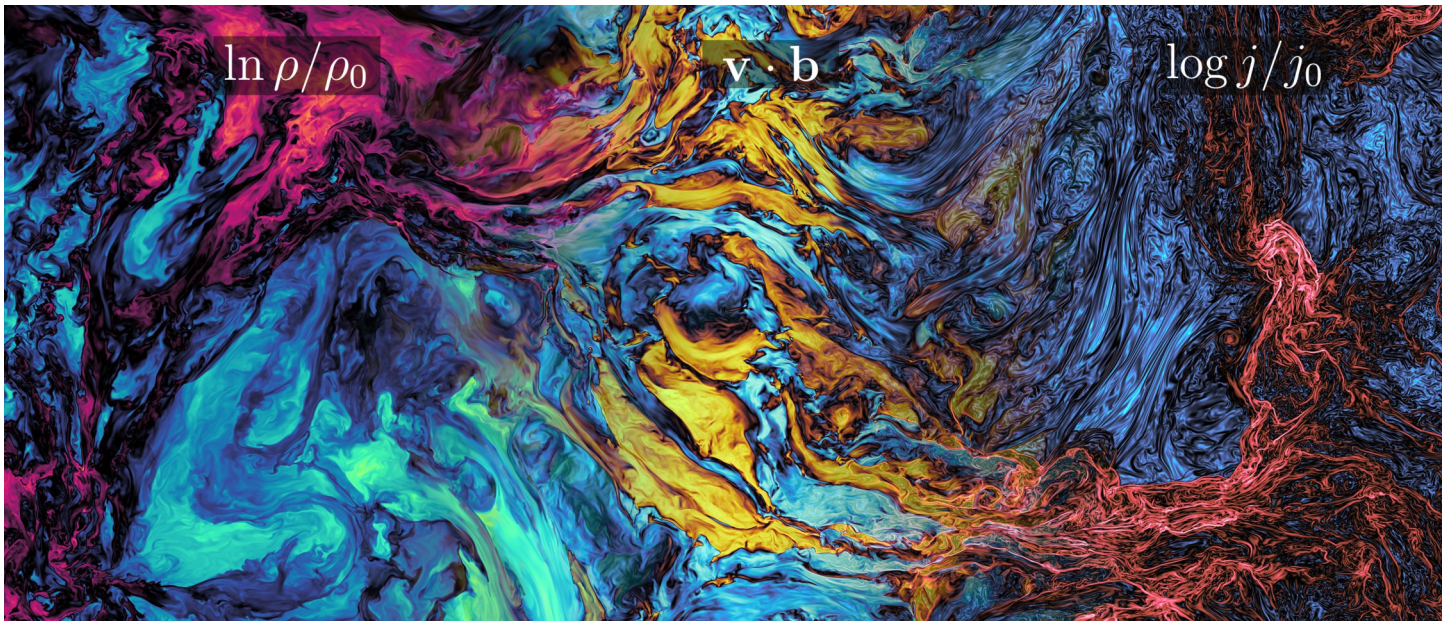

Figure 2: A two-dimensional slice of the gas density $ln\,\rho/\rho_0$ (left; high, red; low, green), point-wise cross helicity $\vec{v}\cdot\vec{b}$ (center; parallel, yellow; antiparallel, blue) and the logarithmic current density $log|\vec{j}|/j_0$ (right; high, red; low, blue), $\vec{j}=(1/4\pi)\nabla\times\vec{b}$, with the mean of each quantity shown in black. This image was selected for the DFG 2023 Christmas Card, and DFG 2024 Research Calendar. It shows not only the intricate structure of the gas and current density, the former setting the initial conditions for star formation in interstellar turbulent gas and the latter responsible for magnetic dissipation through intense current sheets in the plasma, but also the exceptional alignment between $\vec{v}$ and $\vec{b}$, which in turn reduces the strength of the nonlinearity in the turbulence and gives rise to a shallow slope in the kinetic energy cascade (as shown in Figure 1).

isothermal, compressible magnetohydrodynamic equations in three dimensions with stochastic, large-scale forcing to drive finite time-correlated supersonic turbulence.

We use a block-structured parallelization, with each 3D computational block distributed onto one single compute core. Each block contains 168×210×210 computational cells, which we pass to each of the 138,240 compute cores (2,880 compute nodes on SuperMUC-NG) to give a total grid resolution of 10,080 grid cells in each of the three dimensions, and 3.0 PB of data products in total. To ensure numerical convergence of the plasma statistics, such as the turbulent energy spectra, we also perform simulations at $1{,}152^3$, $2{,}520^3$ and $5{,}040^3$ grid resolutions. We use the latest methods from [5] to estimate the numerical viscosity and resistivity, showing that we are able to probe (plasma and magnetic) Reynolds numbers between ~10^5-10^6.The data exploration is being complemented with 3D interactive visualisations, on SuperMUC-NG itself, of gas and magnetic field properties at both large and small scales (e.g. individual plasmoids).

Energy spectra results

In Figure 1 we show the kinetic $E_{kin}(k)$ and magnetic $E_{mag}(k)$ turbulent energy spectra, where $k=2\pi/l$ is the wavenumber, at the three different resolutions (or equivalently, magnetic Reynolds numbers) as indicated by the line colour. From the kinetic energy, we see two distinct regimes of turbulence emerging in the high-resolution simulations. On small k (large spatial scales) the turbulence is dominated by networks of interacting shocks, and has an energy spectrum consistent with Burgers turbulence, $E_{kin}(k) \propto k^{-2}$. On large k (small spatial scales, coloured yellow) the energy cascade resembles the incompressible magnetised turbulence Iroshnikov-Kraichnan spectrum, $E_{kin}(k) \propto k^{-3/2}$ corresponding to magnetised vortices, which exchange energy and cascade the energy to smaller and smaller spatial scales. This cascade is weaker (shallower) than the classical Kolmogorov cascade, $E_{kin}(k) \propto k^{-5/2}$ and hence there must be a mechanism for weakening the nonlinearities, in turn making the cascade shallower.

In Figure 2 we show a two-dimensional slice of the gas density $ln\,\rho/\rho_0$ point-wise cross helicity $\vec{v}\cdot\vec{b}$ and current density $log|\vec{j}|/j_0$. The point-wise cross helicity reveals intensely aligned structures between the gas velocity and magnetic field. Alignment reduces the nonlinearity in the induction equation, which weakens the cascade. However, by comparing to alignment theories [6], we find that the turbulence data is inconsistent with these theories, and hence a new theory that aligns the turbulence is required. As we show for the $E_{mag}(k)$, this theory will also need to describe the $E_{mag}(k) \propto k^{-9/5}$ magnetic energy spectrum, which is currently unexplained by all turbulent magnetic energy spectral models (indicated in the lower left of Figure 1; ranging between dynamical alignment [6] to tearing and reconnection-related instabilities, e.g., [7]).

Ongoing Research / Outlook

We are currently writing-up the final results for the first release of this simulation data, exploring the nature of the energy spectra, the length scales of the turbulence, and the process that gives rise to the weakening of the subsonic cascade. This will be the first in a number of studies that seek to unravel this regime of magnetised turbulence.

References and Links

[1] Federrath & Klessen, Astrophys. J., 761, (2012).
[2] Federrath et al., Nat. Astr., 5 (2021) 365-371.
[3] Fryxell et al., Astrophys. J. Supplement, 131, (2009) 273-334.
[4] Waagan, Federrath & Klingenberg, J. Comp. Phys., 230, (2011) 3331-3351.
[5] Malvadi Shivakumar & Federrath, (2024), arXiv:2311.10350.
[6] Boldyrev, PRL, 96, 11, (2006), 115002.
[7] Dong et al., Science Advances, 8, 49, (2022).

1 Radiation-Hydrodynamical Simulations of Early Galaxies

Research Institution
[1]Leibniz Institute for Astrophysics Potsdam (AIP)

Principal Investigator
Ewald Puchwein[1]

Researcher
Rahul Kannan[2], Aaron Smith[3], Josh Borrow[4], Enrico Garaldi[5], Laura Keating[6], Nele Stachlys[1], Federico Marinacci[7], Rüdiger Pakmor[5], Martin Sparre[1], Mark Volgelsberger[8]8, Lutz Wisotzki[1]

Project Partner
[2]York University, [3]University of Texas at Dallas, [4]University of Pennsylvania, [5]Max-Planck Institute for Astrophysics (MPA), [6]University of Edinburgh, [7]University of Bologna, [8]Massachusetts Institute of Technology (MIT)

SuperMUC Project IDs: pn29we, pn68sa

Introduction

Understanding the formation of first galaxies and their interactions with the surrounding gas filling the Universe represents one of the main frontiers of research in modern astrophysics. The advent of telescopes and instruments such as the James Webb Space Telescope (JWST), the Atacama Large Millimeter/submillimeter Array (ALMA), and the Multi Unit Spectroscopic Explorer (MUSE) has only recently made the observation of these processes possible.

Our program leverages cosmological radiation-hydrodynamical simulations to numerically model these phenomena. Our goal is to simulate a diverse sample of galaxies, spanning a wide range in mass, with a state-of-the-art modelling of the physics of galaxy formation and including effects like fully-coupled radiative transfer. The latter allows us to investigate how the (UV) light emitted by these galaxies affects their environment, and how it heats and ionizes the Universe in a process called cosmic reionization. We also explore how well our current models of galaxy formation agree with the latest observational data sets, which probe significantly deeper into the early Universe than previously possible.

Results and Methods

Modelling the formation and evolution of galaxies is numerically very challenging due to the wide range of scales involved, from a representative cosmological volume, hundreds of millions of lightyears across, to individual molecular clouds within galaxies. Our codes use adaptive spatial and time resolution for this, which concentrates the numerical power in those dense, highly dynamical regions within galaxies where most of the action is happening, while maintaining a well-balanced distribution of the computational workload across the CPU cores (and MPI tasks) which are employed in the calculation.

Another very challenging aspect of our simulations is following the radiation/light emitted by our simulated galaxies as it travels through our simulated Universe. As light is literally the fastest thing we know, following its propagation requires exquisite time resolution, resulting in a large number of time steps. Even when using various algorithmic optimizations (a moment based radiative transfer scheme, sub-cycling, a reduced effective propagation speed), the computational cost is quite substantial.

Overall, despite the adaptive resolution employed and the many algorithmic optimizations, the amount of computing resources needed, and the amount of data produced make using a machine like SuperMUC-NG indispensable. For example, the largest run in our set of simulations requires almost 30,000 CPU cores and 40 TB of RAM.

This project builds upon the THESAN simulation project [1] (also performed on SuperMUC-NG, pn56ku), which simulated a larger volume. Here, we selected a subset of galaxies from THESAN and simulated them at higher resolution and with an even more sophisticated treatment of the relevant physical processes.

Ongoing Research / Outlook

We are currently analyzing our sample of simulated galaxies and comparing their properties to observations of real galaxies in the very distant Universe. First results include:

- Many of the main galaxy properties (like mass in stars, dust content and temperature) are in good agreement with observations of real galaxies, suggesting our galaxy formation model is capturing the relevant processes well.
- The escape of far UV radiation from

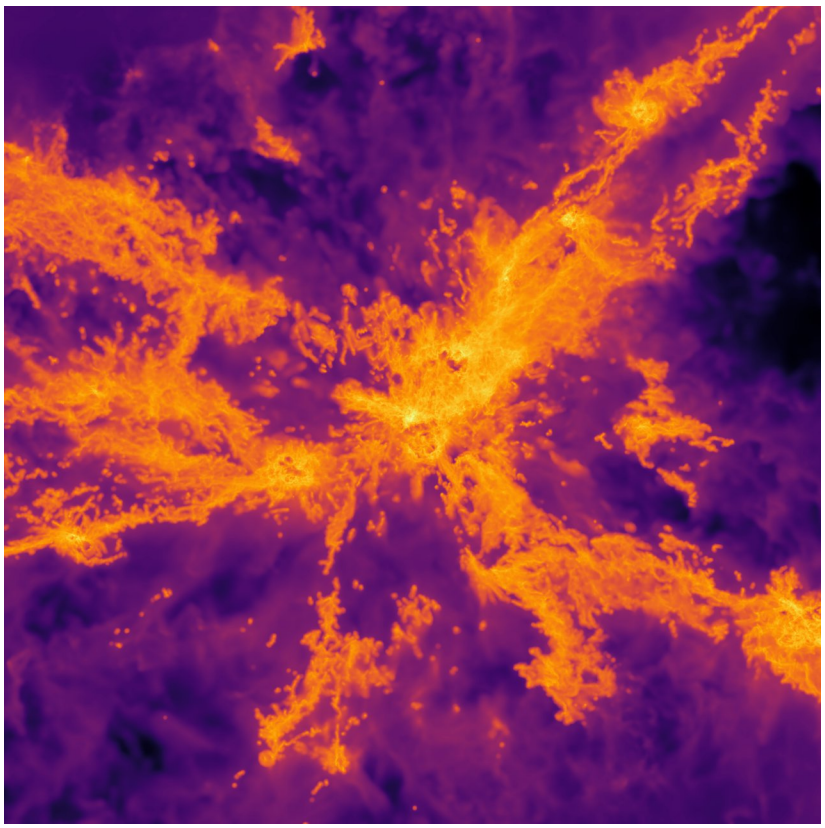

Figure 1: Neutral hydrogen (yellow) around one of our simulated early galaxies (roughly 2 billion years after the big bang). Prominent clouds and filaments that self-shield from ionizing radiation are visible. These structures significantly influence the escape of ionizing radiation and distinctively shape observable signatures such as emission in the Lyman-alpha and H-alpha lines of hydrogen.

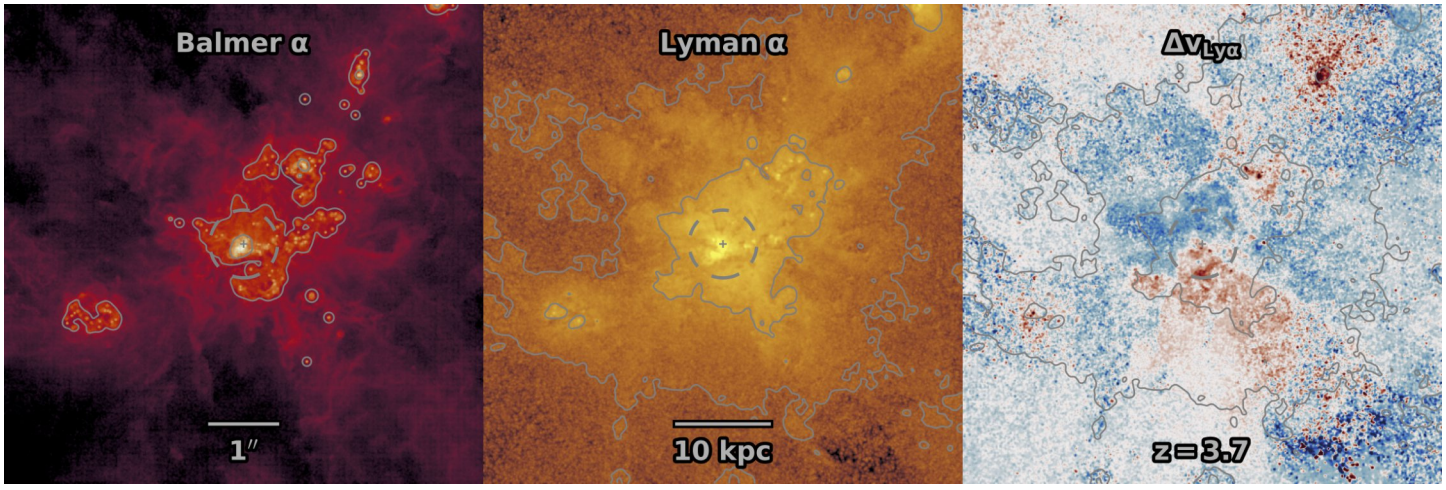

Figure 2: What would our simulated galaxies look like in state-of-the-art observations? Synthetic observations in the Balmer-alpha/H-alpha line (left panel) and Lyman-alpha line (middle panel) of hydrogen are shown. The right panel shows the Doppler shift (velocity) of the emitted Lyman-alpha radiation, an important diagnostic to understand gas distributions and motions in galaxies.

galaxies that then heats and ionizes the Universe is mediated by energy and momentum injection from massive stars. In particular, when they explode as supernovae at the end of their life, channels of low density are created out of the interstellar medium through which radiation can efficiently escape.

A vast amount of further analysis is currently carried out, including:

- Investigating neutral hydrogen in galaxies and how it shields itself against ionizing radiation (see Fig. 1).
- Investigating the detailed observational properties of the simulated galaxies, e.g., their emission in the Lyman-alpha and H-alpha/Balmer-alpha lines of neutral hydrogen (see Fig. 2), as well as in metal lines. What information about the physics of galaxy formation is encoded in it?

References and Links

[1] https://www.thesan-project.com/

Galactic Superwinds and the Galactic Baryon Cycle: Cloudy or Foggy?

Research Institution
[1]Max Planck Institute for Astrophysics, Garching
[2]Purdue University Fort Wayne

Principal Investigator
Ryan Jeffrey Farber[1,2]

Researcher
Max Grönke[1], Fernando Hidalgo-Pineda[2], Meg Blackburn[3]

Project Partner
[2]University of Glasgow

SuperMUC Project ID: pn49ye (Gauss Large-Scale project)

Introduction

On Earth, the water cycle proceeds by three processes. First, water evaporates from the ocean and lakes. Second, water vapor condenses to form clouds. Third, clouds precipitate droplets of water which fall to the Earth as rain. After the rain falls to the Earth, the water droplets collect and eventually flow to lakes and oceans to be evaporated again. Galaxies cycle their star-forming fuel (called "baryons" in distinction to dark matter) in a similar three-step process (see Figure 1). First, massive stars catastrophically end their lives in explosions, referred to as supernovae. The explosions launch remnant material from the stars into the interstellar and circumgalactic media. Second, gravity congeals the ejected material into clouds. Third, clouds cool, contract, and fragment forming a new generation of stars. Galactic outflows play a crucial role in regulating star-formation in the galactic baryon cycle; otherwise, galaxies would stop forming stars much earlier. The most extreme form of galactic outflow is the superwind, whereby material is launched out of the galaxy. Therefore, superwinds may help to explain the enrichment of metals and magnetic fields in the intergalactic medium. Furthermore, superwinds may extend the star-forming lifetime of galaxies; superwinds may prevent galaxies from rapidly accreting cosmologically pristine gas from the intergalactic medium. Despite the importance of the galactic baryon cycle and superwinds in prolonging the star-forming lifetime of galaxies, **the physical mechanism launching star-forming gas into the circumgalactic and intergalactic medium remains unclear**. Moreover, the physical characteristics of the outflows remain to be determined. Originally, work suggested galactic outflows consist of purely hot gas, as supernovae should catastrophically heat any cold clouds during the explosive acceleration process. However, recent work [2] has shown that sufficiently dense cold clouds near the base of the outflows may survive the acceleration, and even grow by cooling any colliding hot gas particles.

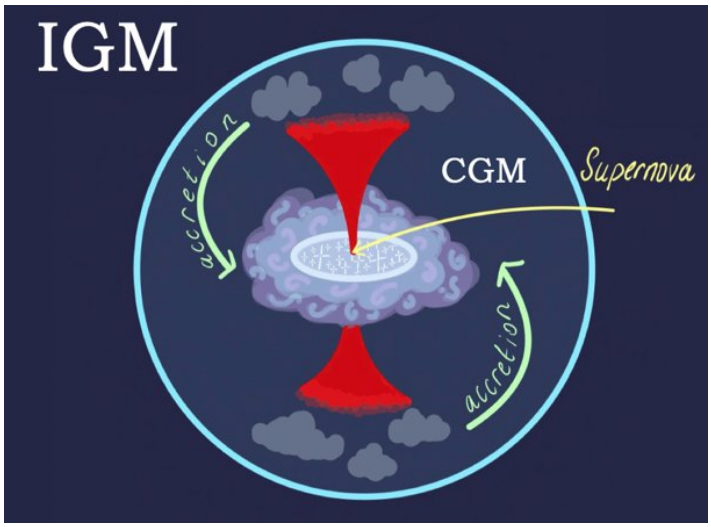

Figure 1: Cartoon of the galactic baryon cycle. Supernovae in the galaxy launch outflows. Cold clouds in the precipitate onto the galaxy to form new stars.

Results and Methods

We have explored the physical mechanisms powering galactic superwinds, considering purely thermal, magnetized, and cosmic-ray accelerated superwinds. SuperMUC-NG has allowed us to perform our studies at unprecedented high resolutions from the galactic disk deep into the circumgalactic medium. We have found a shocking dependence on the phase of the medium with respect to the mechanism driving the outflows. Specifically, we utilized the massively parallelized, Eulerian grid-code FLASH [1] to solve the (two-fluid) (magneto-) hydrodynamic equations, evolving Lagrangian tracer particles to identify the fate of star-forming gas in galactic superwinds. We employ FLASH's unsplit staggered mesh, a finite-volume, high order Godonuv scheme with constrained transport to ensure divergence-free magnetic fields. We include star formation and feedback using active particles, self-gravity of the gas with a multigrid method to solve the Poisson equation, and the Townsend exact integration scheme for radiative cooling. As such, our simulations performed on SuperMUC-NG constitute the physically most sophisticated set of simulations to probe the fate of cold gas clouds deep into the circumgalactic medium to date. As one would expect, magnetized winds are much more filamentary than purely thermal winds. One would expect

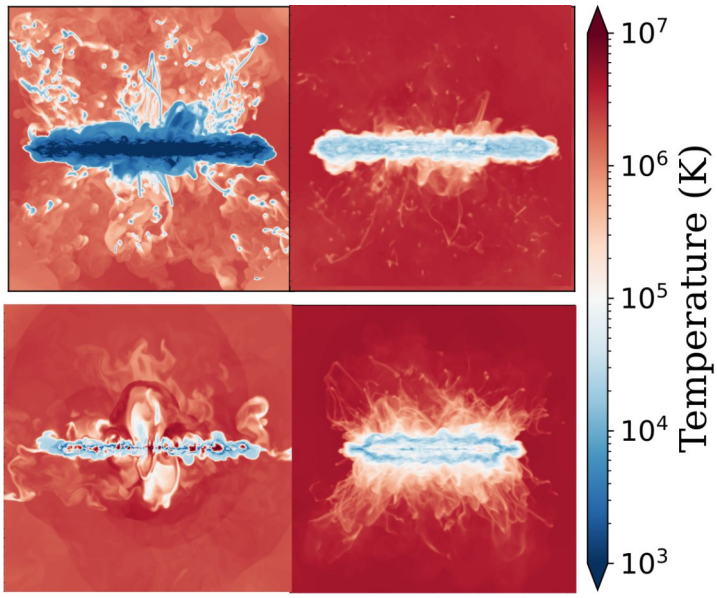

Figure 2: Preliminary results from the Superwind simulations on SuperMUC-NG. Top-left: fast transport cosmic ray simulation (coldest disk and outflows). Top-right: slow transport cosmic ray simulation (much weaker effect). Bottom-left: magnetic fields. Bottom-right: purely thermal (garden-variety galaxy formation simulation).

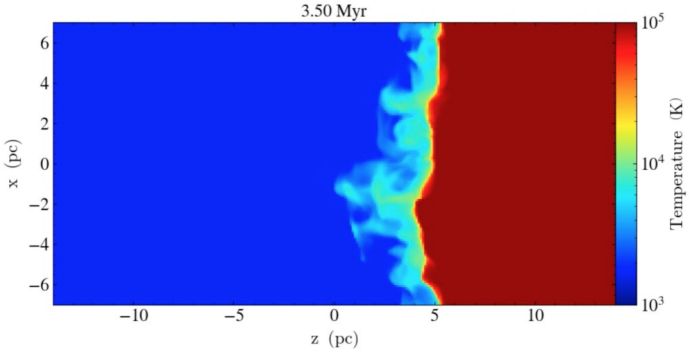

Figure 4: Zoom-in simulation of the interface region of an individual cloud, investigating the mixing rate dependence on cloud temperature. Dark blue indicates cloud material at 1,000 K. Dark red indicates one million Kelvin gas. Intermediate gas around 10,000 K begins to dominate by volume the interface region.

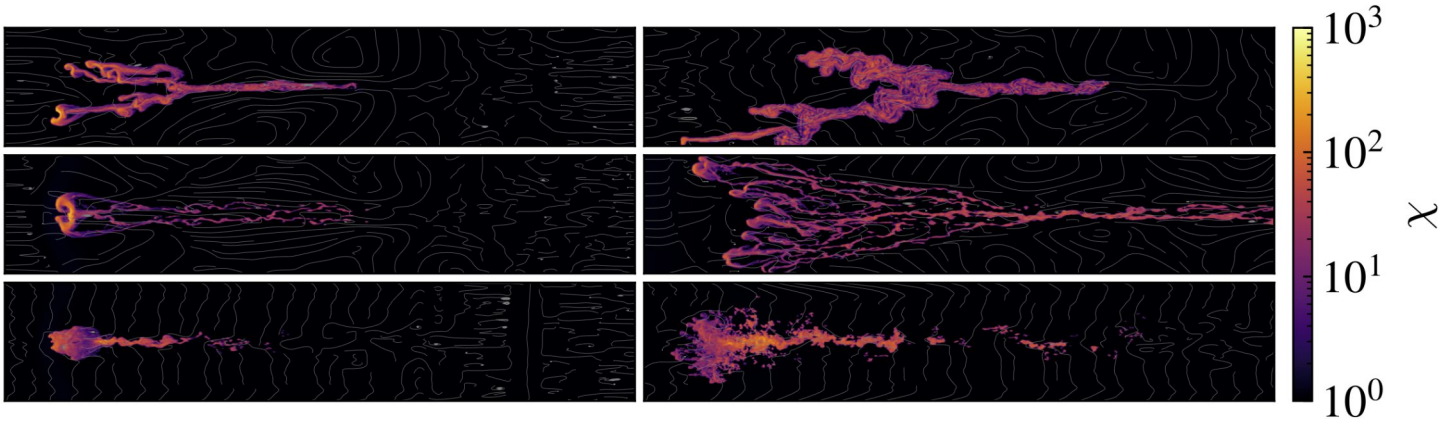

Figure 3: Wind-tunnel simulations of individual clouds in hot winds. The top row shows the evolution of a non-thermal pressure dominated cloud, the middle row indicates the evolution for a weakly nonthermal cloud, and the bottom row shows the evolution for a thermal cloud.

cosmic rays, which contribute non-thermal pressure support to gas in a similar manner to magnetic fields. Shockingly, we find that superwinds are composed of a much larger number of tiny cold clouds than even the purely thermal winds! Moreover, we find an increasing number of clouds the faster the cosmic rays are transported along magnetic field lines, even though cosmic rays with fast transport significantly reduce the star formation (and hence supernova) rate in galaxies (see Figure 2). We have performed follow-up work zooming in for ultra-high resolution to better understand the fate of nonthermally supported clouds. Studying individual clouds in wind-tunnel simulations (see Figure 3), we have found that nonthermally supported clouds survive much more readily than garden-variety thermal plasma, with the survival markedly increased the larger the degree of nonthermal pressure support [3].

Ongoing Research / Outlook

To better understand the results of our superwind simulations, we are presently performing even higher resolution simulations than the wind-tunnel simulations. Namely, we take an individual cloud from the wind-tunnel simulation and zoom in on the interface region to probe the smallest possible scales relevant to our superwind simulations: individual mixing layers. We have found that the clouds outflowing with cosmic rays have cold ~molecular temperatures whereas the thermal winds accelerate predominantly warm atomic clouds. We have found that molecular clouds mix much more efficiently than atomic clouds ([4], Figure 4), providing further understanding of the physical mechanisms supporting enhanced cloud acceleration for the cosmic ray runs than thermal runs in our superwind simulations.

References and Links

[1] B. Fryxell et al., ApJSS, 131, 2000, 273.
[2] M. Gronke, & P. Oh, MNRAS 480, 2018 L111-115.
[3] F. Hidalgo-Pineda, R.J. Farber & M. Gronke, MNRAS 527, 2024, 135-149.
[4] M.G. Blackburn & R.J. Farber, *to be submitted to* RNAAS.

1 Multimessenger signatures of neutron-star merger remnants

Research Institution
[1]GRAPPA, Universiteit van Amsterdam

Principal Investigator
Philipp Mösta[1]

Researcher
Pablo Bosch[1], Sanjana Curtis[1], Sebastiaan de Haas[1], Roland Haas[2], Erik Schnetter[3]

Project Partner
[2]National Center for Supercomputing Applications, University of Illinois, USA
[3]Perimeter Institute, Ontario

Funding
NWO ENW-XL QCD

SuperMUC Project ID: pn49ju (PRACE project ID 2021250017)

Introduction

The era of multimessenger astronomy has been inaugurated with the extraordinary detection of gravitational waves from two merging neutron star by advanced LIGO in combination with optical, infrared and gamma-ray observations of the same event. These observations have started to revolutionize the understanding of short gamma-ray bursts, of the physics of neutron stars, and of the origin of the heavy elements, like silver, gold, and platinum. However, they also pose many pressing questions. What was the fate of the binary after coalescence? Which mechanisms generated the observed electromagnetic signals? For this PRACE project we have used SuperMUC-NG to address some of these questions.

The remnant evolution depends on an interplay between rotation, nuclear forces, magnetic fields, and turbulence. Exactly how those elements combine determines the amount and composition of the elements produced in these events, which in turn sets the electromagnetic signatures that can be observed with current time-domain telescopes. To predict detailed electromagnetic signals expected from these remnants, we have performed state-of-the-art general-relativistic neutron-star merger simulations.

Results and Methods

The simulations include all four fundamental forces in a fully-coupled way: general relativity for gravity, magneto-hydrodynamics to model the plasma effects, tabulated nuclear equations of state to capture the effects of the strong force, and an neutrino transport to model the emission and absorption of neutrinos via weak interactions. The simulations are the first to include all such elements and at the same time employ the very high resolution needed to resolve the small-scale turbulent motions in the remnant. The codebase we have employed is GRHydro [1], an open-source GRMHD code part of the Einstein Toolkit [2] that we develop in the group at GRAPPA in Amsterdam. GRHydro is built on the Cactus framework, a portable framework for the development of partial differential equation solvers, and the Carpet AMR driver, implementing MPI/OpenMP parallelization. Our largest simulations have scaled up to ~50,000 cores on SuperMUC-NG.

We have used this codebase to produce full 3D GR radiation-MHD simulations of NS merger remnants. This first set of simulations have included a Leakage neutrino transport and studied the effect of different magnetic field configurations on the postmerger remnant. We have performed 5 simulations with varying magnetic field strength and geometry. These simulations allowed us to study how sensitively the outflow properties, nucleosynthetic yields and the kilonova lightcurves depend on the magnetic field. We have used our full postprocessing pipeline and produced detailed nucleosynthetic yields and kilonova lightcurves for comparison with observations. With these simulations we have been able to quantify how much variety in the multimessenger signals from different merger remnants we can expect [3]. In addition, the simulations will help us explore the variety of realized magnetic field configurations at the end of the simulations when a black-hole - accretion torus system has formed. This will serve as initial conditions for other groups in the community that perform long-term black-hole - accretion torus system simulations with a fixed background spacetime. This work is currently ongoing and we are collaborating with multiple research groups to make full use of the dataset we created with this PRACE allocation.

In a second step, we have studied how the detailed turbulence and inclusion of M1 neutrino transport affect the kilonova signals from the remnant left behind by the merger of two neutron stars. The evolution of the remnant depends sensitively on the developing turbulence and associated angular momentum redistribution. This is set by the magnetorotational instability, a plasma instability that produces extremely strong magnetic fields by

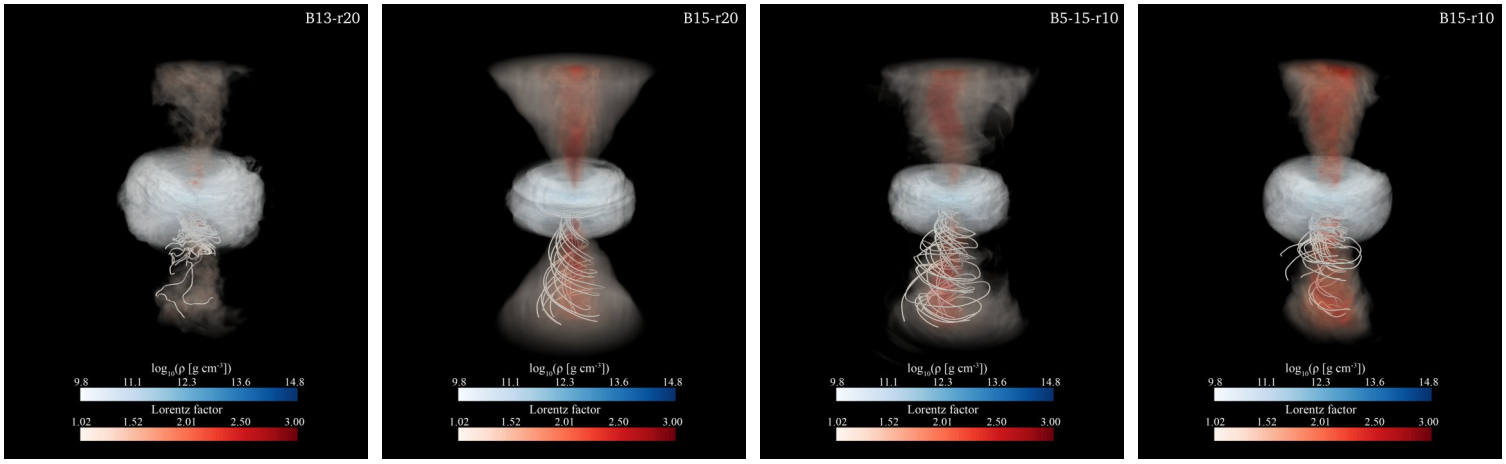

Figure 1: Volume renderings of the outflows (white-red colormap) and density for the accretion torus (white-blue colormap) for different simulations of neutron-star merger remnants presented in [3,4]. The magnetic field lines are also shown in the lower plane ($z < 0$, where z is the vertical axis) in white.

extracting rotational energy and redistributing angular momentum. This reduces rotational support from the inner regions of the remnant, which can cause it to collapse to a black hole. We have utilized the tier-0 computational power of SuperMUC-NG to carry out high-resolution radiation-magnetohydrodynamics simulations with

an M1 neutrino transport scheme. This simulation has allowed us to study to predict more accurate kilonova lightcurves from the remnant. With this simulation, we have been able to successfully explain the blue kilonova component observed for GW170817 [4], a component of the observed signal that simulations have struggled to reproduce so far. The results will help maximize the science return from observations of GW170817 and similar future events.

Ongoing Research / Outlook

We are currently building on the successful simulations performed on SuperMUC-NG that are presented here. As a follow-up, we are studying full binary neutron-star merger simulations to predict the complete multimessenger signatures expected from such events. These simulations will complement the work presented here as they will provide a complete picture of how the different outflow components from the merger and post-merger phase combine. In addition, we are also using the datasets created on SuperMUC-NG to connect to perform long-term follow-up simulations in collaboration with other research groups.

References and Links

[1] P. Mösta et al., Class. and Quant. Grav., 31, 1 (2014).
[2] www.einsteintoolkit.org
[3] S. de Haas et al., MNRAS 527, 2 (2024).
[4] S. Curtis et al., ApJL 961 L26 (2024).

The electron acceleration in astrophysical shocks

1

Research Institution
Leibniz-Institut für Astrophysik Potsdam

Principal Investigator
Mohamad Shalaby

Researcher
–

Project Partner
–

SuperMUC Project ID: pn29wa

Introduction

We observe the Universe primarily by examining the different kinds of radiation emitted by various objects. One type, called thermal radiation, comes from gases that are in a balanced state, while another type, called non-thermal radiation, comes from energetic particles like electrons and protons. These particles, also known as cosmic rays, don't settle into a balanced state and can maintain their high energy levels for a long time as they travel through space. Non-thermal radiation from cosmic rays covers a wide range of energies and can be found in various celestial objects, from large galaxies to smaller astronomical scales. Understanding how these energetic particles, especially electrons, are accelerated and move around in space is crucial for interpreting what we observe in our universe. In our galaxy, cosmic rays are mainly produced by shock waves. While we have a good understanding of how protons and heavier ions are accelerated by these shocks, we're still figuring out how electrons, which are much lighter, get their energy boost. Unlike heavier particles, electrons can't directly interact with shock waves to gain speed. This poses what scientists call the "electron injection problem." Researchers have proposed a two-step solution to this problem: first, the electrons need to be heated up significantly, and then they need to interact with specific small-scale magnetic waves to get the necessary kick in energy. Recent studies have suggested that certain types of waves, like whistler waves, might do the trick. However, these waves might not be powerful enough, especially in some types of shock waves. Recently, a new discovery called the "intermediate-scale instability" has shed light on this problem. This instability generates specific types of waves near shock waves, and our simulations have confirmed that these waves are indeed effective in accelerating electrons [1]. This finding opens up new possibilities for understanding how shock waves work in space. Here, we investigate the dependence of the efficiency of this mechanism on the angle (θ) between the large-scale magnetic field and the direction of shock propagation, as illustrated in Figure 1 which shows a simplified illustration of how these shock waves are simulated.

Results and Methods

We run our numerical particle-in-cell simulations using the SHARP code [2,3]. As shown in Figure 1, the shock is formed by sending a supersonic isotropic electron-ion plasma towards a reflecting wall on the left. The interaction of the reflected and the incoming particles leads to shock formation. Our code has an unprecedented ability to avoid numerical heating and thus, enables more faithfully to explore the underlying physical process in our simulations. The memory management and access in our code is optimized to greatly reduce the computational cost for our simulations. For the fist time, this enabled simulations with a realistic ion-to-electron mass ratio and as such, enabled us to discover the physical mechanism responsible for accelerating electrons at electron-ion shocks. The dependence of the accelerated particle spectra on the angle between the upstream magnetic field and the shock propagation direction is shown in Figure 2. This demonstrates that ion and electron acceleration are most efficient in parallel and quasi-parallel shocks when simulated in 1D3V (one spatial dimension and three velocity dimensions). This conclusion regarding ion acceleration aligns with previous simulations in 1D3V, 2D3V, and 3D3V using the hybrid-PIC technique with simulated only the ion dynamics [4].

Ongoing Research / Outlook

In all these simulations, we have used parameters relevant for probing the electron acceleration at supernova driven shocks. These results are in good agreement with global, magneto-hydrodynamic modeling of the multi-wavelength emission from the supernova remnant SN 1006 which required more efficient electron acceleration at parallel shocks in comparison to perpendicular shock configurations. We plan to further investigate the impact of the new instability on the acceleration of electrons at shocks with parameters relevant for merger shocks in the intracluster medium which gives rise to the puzzling class of radio shocks.

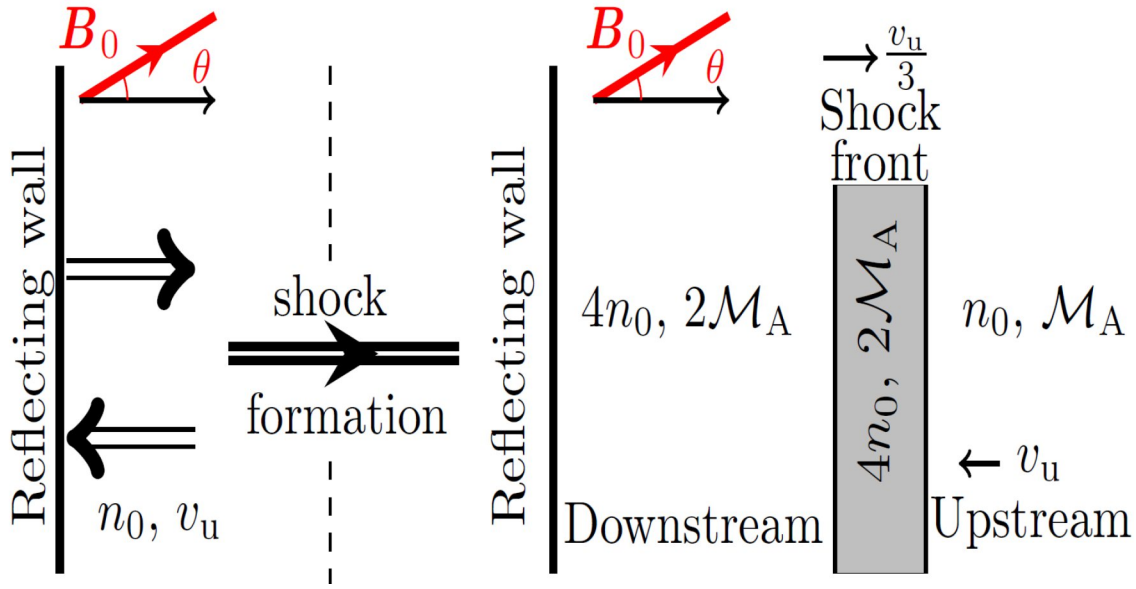

Figure 1: A Schematic representation of how shock formation and evolution are simulated.

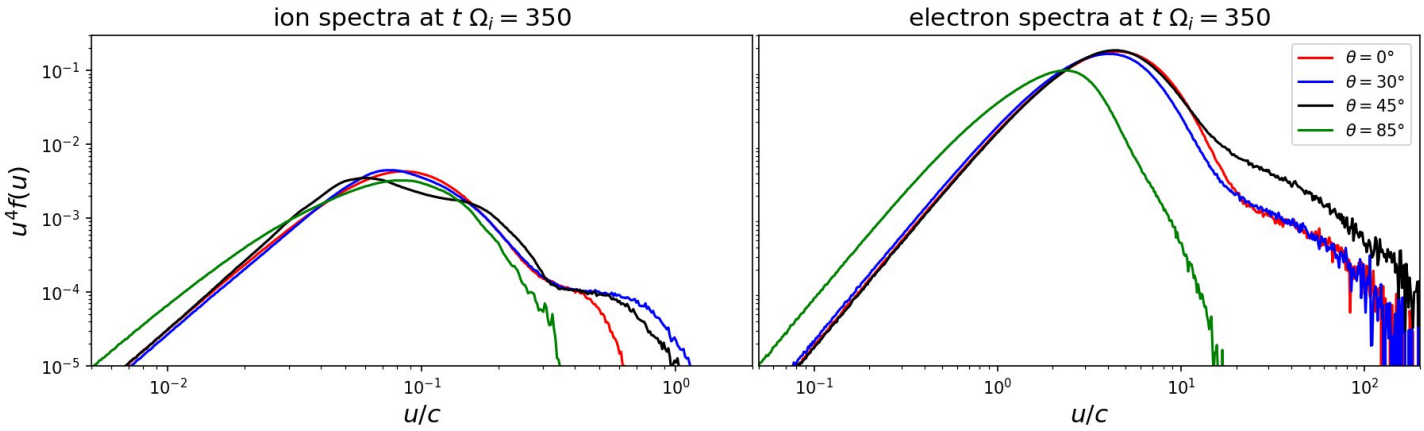

Figure 2: Downstream momentum spectra for ion (left) and electron (right) in various simulations. This shows that ion and electron acceleration are most efficient in parallel and quasi-parallel shocks when simulated in 1D3V (one-spatial dimension and 3-velocity dimensions).

References and Links

[1] Shalaby, M., Lemmerz, R, Thomas, T., & Pfrommer, C. 2022, ApJ, 932, 86.
[2] Shalaby, M., Broderick, A. E., Chang, P., et al. 2017, ApJ, 841, 52.
[3] Shalaby, M., Thomas, T., & Pfrommer, C. 2021, ApJ, 908, 206.
[4] Caprioli, D., & Spitkovsky, A. 2014, ApJ, 783, 91.

Neutron Star Formation from Collapsing White Dwarfs

Research Institution
[1]Max Planck Institute for Astrophysics, Garching

Principal Investigator
Hans-Thomas Janka[1]

Researcher
Eirini Batziou[1,2], Robert Glas[1], Jakob Ehring[1,2,3], Oliver Just[4], Ernazar Abdikamalov[5]

Project Partner
[2]UM School of Natural Sciences, Garching
[3]Max Planck Institute for Physics, Garching
[4]GSI Darmstadt
[5]Nazarbayev University, Astana, Kazakhstan

Funding
SFB 1258, EXC 2094

SuperMUC Project ID: pn49sa

Introduction

Neutron stars are the most compact objects in the Universe with typically 1.5 times the mass of our Sun compressed into a sphere of just about 25 km in diameter, implying central densities higher than those in atomic nuclei. Most neutron stars are formed as remnants of massive stars when the degenerate core of these stars becomes gravitationally unstable and collapses, while most of the stellar matter is ejected in a violent supernova explosion with velocities up to 10,000 km/s. Two such neutron stars in a binary system can collide in a violent merger event after having approached each other on a spiral orbit over hundred of millions to billions of years, driven by the continuous emission of gravitational waves.

Such signals from the late inspiral phase shortly before the final merger have been captured already by the laser interferometer detectors of LIGO and Virgo. The most prominent case is the first-ever detected event of August 17, 2017, termed GW170817, which was accompanied by a short gamma-ray burst, GRB170817A, and an astronomical transient, AT2017gfo. This so-called kilonova emitted ultraviolet, optical, and infrared light over a period of about 10 days, powered by the energy released in the radioactive decay of very neutron-rich, trans-iron nuclei formed in the so-called rapid neutron-capture process. Binary neutron star mergers are thus the first observationally confirmed cosmic sources of super-heavy chemical elements including silver, cadmium, and caesium. According to numerical models also ultra-heavy species such as gold, platinum, thorium, and uranium could be produced, but it is still unclear whether these chemical elements were indeed made by GW170817.

Besides being the remnants of massive stars, neutron stars might alternatively form by the gravitational collapse of white dwarfs. White dwarfs are degenerate remnants of stars with masses lower than those collapsing to neutron stars. Living in close binary systems, they can accrete matter from a normal companion star or they can collide with a second white dwarf after a long inspiral evolution due to gravitational-wave emission. Both scenarios are expected to lead to explosions called thermonuclear ("Type Ia") supernovae. However, under special circumstances also an accretion-induced collapse (AIC) or a merger-induced collapse (MIC) of the white dwarfs to neutron stars might result. Such events have been proposed decades ago, but they have not been studied in much detail and have not been discovered so far. The emerging neutron stars may have masses, recoil velocities, and spins different from neutron stars born in collapsing massive stars. Moreover, mass ejection by AIC and MIC events could contribute to the enrichment of galaxies with heavy chemical elements, and predicting the ejecta properties is crucial for identifying these collapse events in observations. Our project aims to explore these questions by developing a first set of long-term neutrino-hydrodynamical simulations of collapsing white dwarfs for a representative set of initial conditions.

Results and Methods

The PI's team performed two-dimensional (axi-symmetric; 2D) simulations with the Alcar code, which integrates the Euler equations for the conservation of mass, momentum, energy, and electron-lepton number of the stellar plasma, coupled with an energy and velocity dependent transport solver for neutrinos and antineutrinos of all three lepton flavors (electron, muon, and tau neutrinos). The Newtonian hydrodynamics equations include source terms for gravity (with general relativistic corrections) as well as for the exchange of electron-lepton number, momentum, and energy between neutrinos and the stellar plasma. The neutrino source terms account for all relevant interactions (beta processes, pair-neutrino creation and annihilation, and scattering reactions). The transport treatment is based on a two-moment scheme for the coupled neutrino momentum and energy

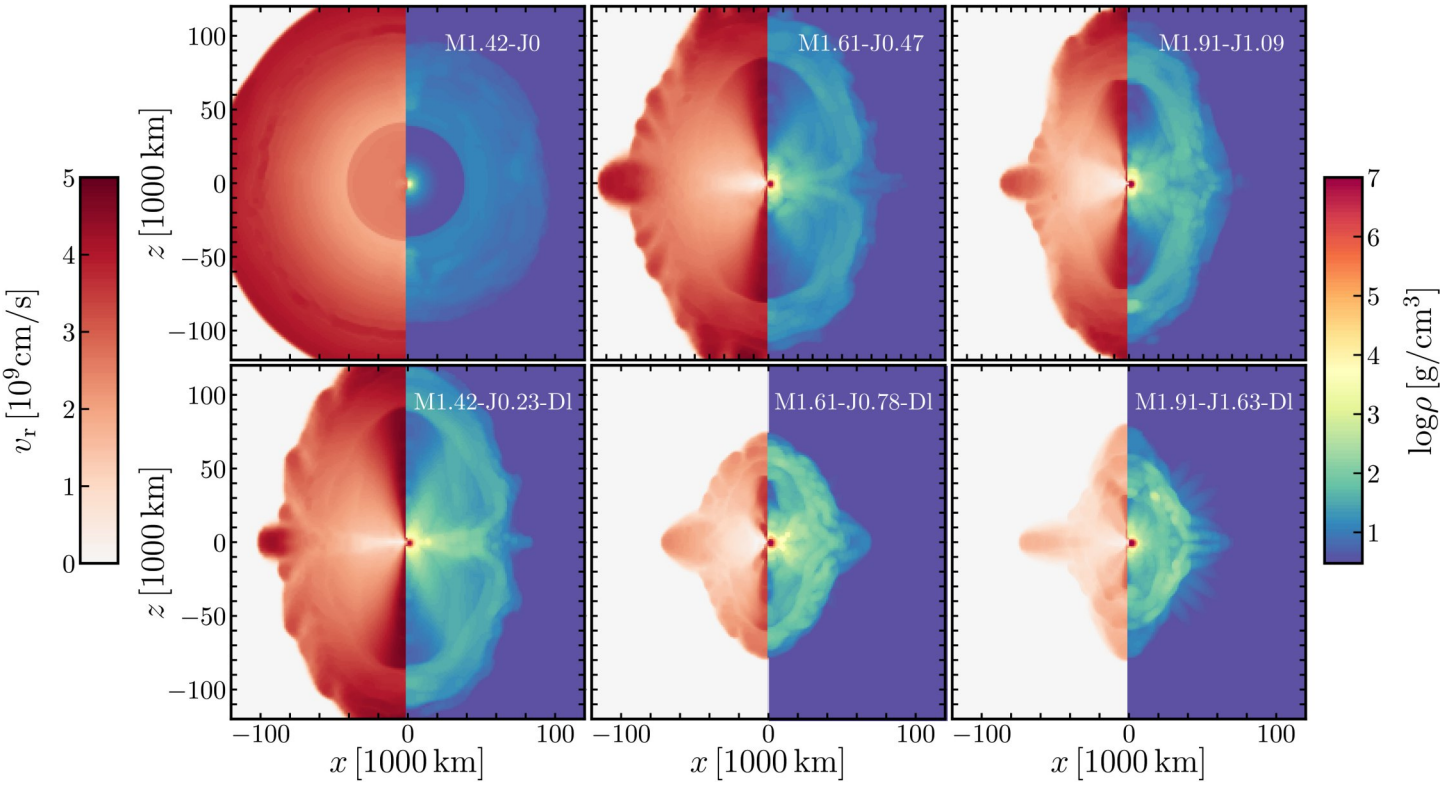

Figure 1: Snapshots of expanding ejecta in our six simulations at 3 seconds after the neutron star formation by white-dwarf collapse. The left half-panels display the radial velocity, the right half-panels the logarithm of the matter density. The upper left model is non-rotating, the models in the right column rotate most rapidly. The masses of the initial white dwarfs are 1.42, 1.61, and 1.91 solar masses (from left to right).

equations with an analytic closure relation for the neutrino pressure tensor ("M1 scheme"). The partial differential equations are discretized by a conservative finite-volume method (Godunov scheme) with higher-order accuracy in space and time and using an approximate Riemann solver. The numerical integration is performed by explicit time stepping and implicit treatment of most source terms. The pressure of the stellar plasma (composed of atomic nuclei, nucleons, electrons, positrons, photons) as a function of the thermodynamic variables is described by a modern equation of state ranging from sun-like conditions to the super-nuclear densities and several 100 billion Kelvin characteristic of newly formed neutron stars.

The simulations were initiated from six white dwarf models in rotational equilibrium for different masses, central densities, rotation rates, and rotation profiles, supported by degenerate-electron pressure and representing various conditions that may result from accretion and merger evolution. High spatial resolution (1.875 degrees in lateral angle and up to 1,600 radial zones) was essential to resolve the steep density gradient near the surface of the forming and contracting neutron stars. With an allocated 18.31 million core-h and 960 cores used on SuperMUC-NG per run, nearly 9 seconds of evolution could be covered for all models.

These simulations show that rotation is a crucial ingredient that determines the properties of the neutron stars and ejecta. Ejecta masses and energies are tightly correlated and can vary from about 0.006 to over 0.05 solar masses and from 10^{49} erg to 3×10^{50} erg, respectively. Rapidly spinning white dwarfs form neutron stars enveloped by thick equatorial belts of centrifugally supported matter that may strip off more mass when evolving by angular momentum loss over secular time scales. The mass ejection occurs predominantly to the poles such that the expanding clouds of matter display prominent bipolar structures, bounded by shock fronts. The spatial extension of the spherical or prolate shocks at a given time (3 seconds after the neutron star started to form in Figure 1) scales directly with the ejecta energy.

Interestingly, there are distinctive differences in the dynamics and composition of ejecta from non-rotating and rotating collapsing white dwarfs. In the nearly spherical ejecta of non-rotating models, moderately neutron-rich matter is expelled with highest velocities first, whereas later ejecta evolve to increasingly proton-rich conditions, analogous to what was found when non-rotating cores of stars near the low-mass end of core-collapse progenitors form neutron stars in so-called electron-capture-like supernovae. Rotation reverses the picture such that the ejecta are increasingly more neutron-rich the later they get expelled (Figure 2). In the most rapidly spinning cases the earliest ejecta are proton-rich. These differences can be explained by the dynamics of the expanding matter and the impact of neutrino emission and absorption in the outflows. Proton-rich and neutron-rich conditions lead to characteristic differences in the chemical elements that are assembled when the initially hot matter cools by expansion and neutrons and protons recombine to alpha particles and heavy nuclei. A rapid

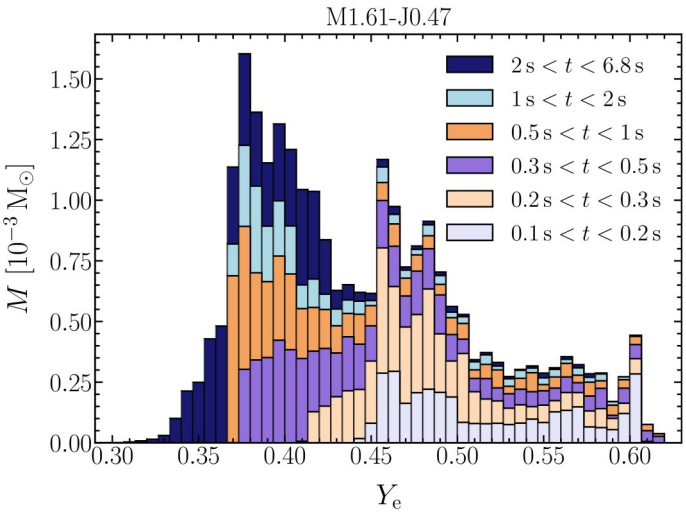

Figure 2: Ejecta mass versus "electron fraction", which is the ratio of electron (or proton) number density to the total number density of protons plus neutrons. In this moderately fast rotating case of a collapsing white dwarf with 1.61 solar masses, the early ejecta are moderately neutron-rich (Y_e ~ 0.45) to significantly proton-rich (Y_e ~ 0.6), and increasingly neutron-rich (Y_e < 0.5) ejecta follow as time progresses.

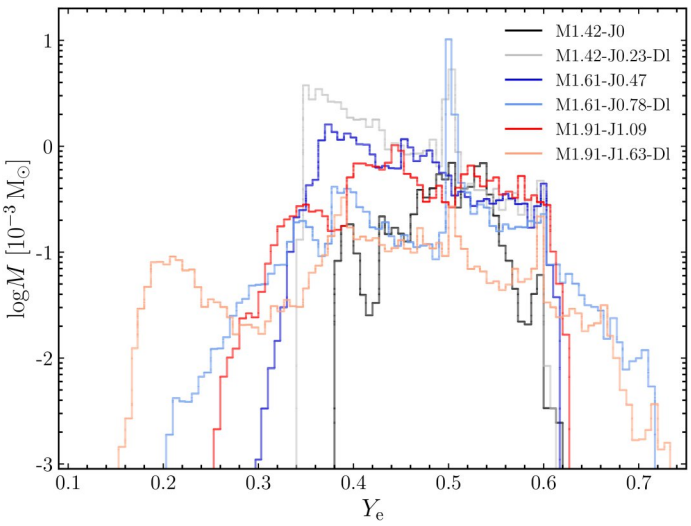

Figure 3: Ejecta mass distributions versus electron fraction for our six simulations. Collapsing white dwarfs with fast rotation eject highly neutron-rich matter (Y_e much below 0.5). Lanthanides will form for Y_e < 0.3, actinides for Y_e < 0.25.

neutron-capture process can occur for neutron excess, possibly even forming the heaviest chemical elements from gold to uranium in the cases of the fastest spinning white dwarfs, where some ejecta are extremely neutron-rich (Figure 3).

Ongoing Research / Outlook

The simulations carried out in this project are the first of their kind in following the formation process of spinning neutron stars created by white-dwarf collapse over several seconds. Since some of them have been finished only recently and the rest will reach their targeted end soon, publications of the results are in preparation at the time of writing this report.

Besides the ejecta conditions, the models also predict the associated neutrino signals and will be evaluated for their gravitational-wave emission. Moreover, our hydrodynamic results will be analyzed for the nucleosynthesis of chemical elements by post-processing them with nuclear reaction networks. Ultimately, this will facilitate detailed predictions of characteristic properties of observable electromagnetic emission connected with the mass ejection during AIC and MIC events. This will guide astronomical searches for such kilonova-like transients.

In a next step of improving the theoretical models, more realistic initial configurations, once available, will have to replace our constructed white dwarfs. On a longer perspective, 3D simulations are needed because triaxial instabilities during the collapse of rapidly and differentially rotating white dwarfs may lead to interesting phenomena, e.g., gravitational-wave features, enhanced mass loss, changes of the ejecta geometry, though the basic picture emerging from our benchmark models is unlikely to change. Carrying 3D simulations over similarly long time scales with sufficient resolution will be challenging. This holds, in particular, for 3D simulations including magnetic fields, which will unavoidably be amplified to dynamically relevant strengths in the differentially spinning, new-born neutron stars. The high spatial resolution needed to follow magnetohydrodynamic instabilities on short length scales will pose even greater demands on computational resources. Our models therefore define only the starting point of much more work to be done in the future.

References and Links

[1] J.C. Ehring, "Accretion-Induced Collapse of White Dwarfs", Master Thesis, LMU (2019).
[2] E. Batziou, "Neutron Star Formation in Accretion-Induced Collapse of White Dwarfs", PhD Thesis, TUM (2023).

Explaining Supernova Properties by Neutrino-driven Explosions

1

Research Institution
[1]Max Planck Institute for Astrophysics, Garching

Principal Investigator
Hans-Thomas Janka[1]

Researcher
Daniel Kresse[1,2], Robert Bollig[1], Robert Glas[1], Malte Heinlein[1,2]

Project Partner
[2]TUM School of Natural Sciences, Garching

Funding
SFB 1258, EXC 2094

SuperMUC Project IDs: pr48ra, pr74de, pr53yi (Gauss Large-Scale projects), pn25me, pn69ho

Introduction

Stars are cosmic fusion reactors, which gain energy by nuclear reactions of light atomic nuclei to heavier ones. In stars of more than about nine solar masses, a sequence of burning phases thus assembles successively heavier chemical elements, starting from hydrogen fusion to helium as in the Sun, and continuing with helium, carbon, neon, oxygen, and silicon burning until a core of iron builds up at the center of the star. Iron as the atomic nucleus with the highest binding energy per nucleon cannot produce energy by further burning, and thus the growing iron core cannot escape a catastrophic end.

When its mass approaches a critical limit, particle reactions including neutrino production set in and trigger the collapse of the iron core to a neutron star within less than a second. In the formation of such a compact remnant with a diameter of ultimately only 20–30 kilometers, huge amounts of gravitational binding energy are released mostly through neutrinos. Less than one percent of this energy is sufficient to power a supernova explosion that expels the stellar layers around the new-born neutron star into circumstellar space with velocities up to more than 10,000 kilometers per second.

The supernova ejecta do not only contain the heavy chemical elements that the massive progenitor star has bred by nuclear fusion during millions of years of its life, but they also contain new heavy nuclei created by nuclear reactions during the explosion. A fair fraction of these products of explosive nucleosynthesis are trans-iron elements that cannot be made by stellar fusion reactions. Also radioactive isotopes of titanium, aluminium, and of iron-group nuclei, for example of nickel, are ejected. Their energy release by radioactive decays powers the bright electromagnetic emission of the supernova phenomenon for years. Due to radioactivity of the longer-living species, they also serve for supernova diagnostics, especially as important probes of the physical processes that take place at the origin of the explosion.

Supernovae play a crucial role for the enrichment of the universe with heavy elements in the so-called baryon cycle, where new generations of stars are born from the supernova-inseminated gas of the preceding stellar generations. Energy, momentum, and shock-wave accelerated cosmic rays released by stellar explosions have an important influence on the dynamical evolution of galaxies. Supernovae as the violent end stages of stellar evolution are the birth sites of neutron stars and stellar-mass black holes, whose terminal collisions in binaries were the first-ever detected sources of gravitational waves. Moreover, due to the extreme conditions of density and temperature at their centers, supernovae are unique cosmic laboratories for nuclear and particle physics, which permit to study the properties of ultra-dense matter and processes involving neutrinos and non-standard particles at conditions that cannot be reached in any terrestrial experiment.

For all these reasons it is an important task of supernova theory to better understand the physics of this fascinating phenomenon and to predict (a) which stars explode or collapse to black holes without concomitant explosion, and (b) the associated consequences in order to interpret the observed properties of supernovae and their compact and gaseous remnants. Because of the enormous complexity of the physical processes, theoretical predictions are mainly based on detailed numerical simulations. The simulations need to self-consistently model the hydrodynamics of the stellar plasma in three spatial dimensions including the crucial effects of neutrinos, which transport energy, momentum as well as electron and muon lepton number out of the hot, newly forming neutron star. The simulations must also take into account all relevant nuclear and neutrino reactions.

Results and Methods

Such 3D simulations are among the most challenging and computing-time consuming numerical calculations performed in contemporary theoretical astrophysics. They are only feasible with the power of massively paral-

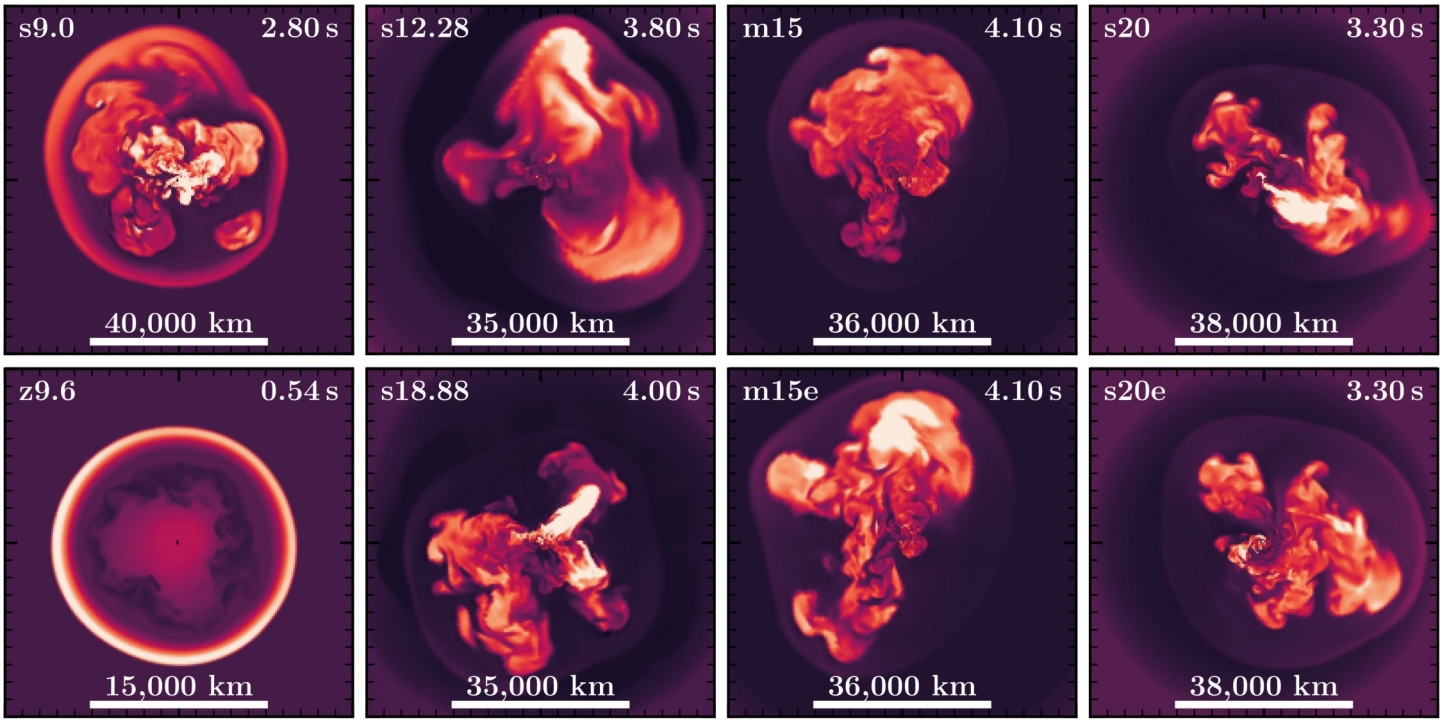

Figure 1: Overview of our current set of self-consistent 3D supernova models for progenitor stars of different masses as indicated by the model names in the left upper corner of each panel; the numbers are the stars' birth masses in solar masses. The plots show vertical cuts of the color coded entropy distribution typically at several seconds after the neutron star (invisibly small at the center) has started to form (time given in the top right corner of each panel) and close to the stage when the explosion energy has saturated (Figure 3).

lel supercomputers, and they are the task of the scientific project described here, carried out on SuperMUC-NG and partially based on results previously obtained with computing resources on SuperMUC. Because the progenitor's mass and pre-collapse structure have an important influence on the outcome, core-collapse simulations must be conducted for a wide spectrum of evolutionary stellar models, and they have to the continued over long evolution periods to determine the observable properties of supernovae. Hydrodynamics and neutrino transport have to be applied for time scales of several seconds to cover the neutrino-emission phase of the hot neutron star, and thereafter pure hydrodynamics for many more minutes to hours until the supernova shock finally breaks out from the star's surface.

We performed our simulations with the Prometheus-Vertex code, which is a state-of-the-art neutrino-hydrodynamics program for supernova simulations. The Prometheus module for Newtonian hydrodynamics is based on a higher-order finite-volume (Godunov-type) scheme for conservative, time-explicit integration of the Euler equations for mass, momentum, energy, electron-lepton number, and a chosen set of chemical elements, using an exact Riemann solver. The Vertex module handles the neutrino-energy and fluid-velocity dependent neutrino transport and solves the coupled energy and momentum equations of all neutrino species, i.e., electron neutrinos, antineutrinos, and one, two, or four separately treated species of heavy-lepton neutrinos. The variable Eddington tensor for closing the two-moment set of equations is computed from the Boltzmann equation, whose solution is iterated for convergence with the moment equations. The transport solver works with implicit time integration, which permits time steps up to about 100 times longer than the hydro steps. A ray-by-ray method allows for extremely efficient parallelization of the transport in multi-D applications. General relativistic corrections are taken into account in the transport and in the source terms for gravity applied in the hydrodynamics equations. The code contains the world-wide most comprehensive and detailed treatment of all relevant neutrino interactions including those with muons. It is also supplemented with a small nuclear network to follow the chemical composition changes due to nuclear reactions. A set of modern equations of state is available to describe the thermodynamics of the stellar plasma for densities ranging from the super-nuclear interior of neutron stars to the conditions in the low-density circumstellar medium.

For more efficient long-term 3D simulations with neutrino physics, we developed the NEMESIS scheme (Neutrino-Extrapolation Method for Efficient SImulations of Supernova explosions; [1]). This new scheme is computationally less expensive (saving up to a factor 10 in computing time) and permits larger time steps, because the evolution inside the new-born neutron star is not followed in detail, whereas the treatment still takes into account the contraction of the compact remnant, its radiation of neutrinos (using results from 1D neutron-star cooling simulations), and the neutrino interactions in the surrounding medium, which power the blast wave.

Our set of supernova simulations spans a mass range from 9 to 20 solar masses (Figure 1). It includes several cases where the stellar collapse and early explosion phase had been computed with resources from previous LRZ and Gauss projects [2,3] and that could now be extended to later evolution times by the application of the NEMESIS scheme [1,4].

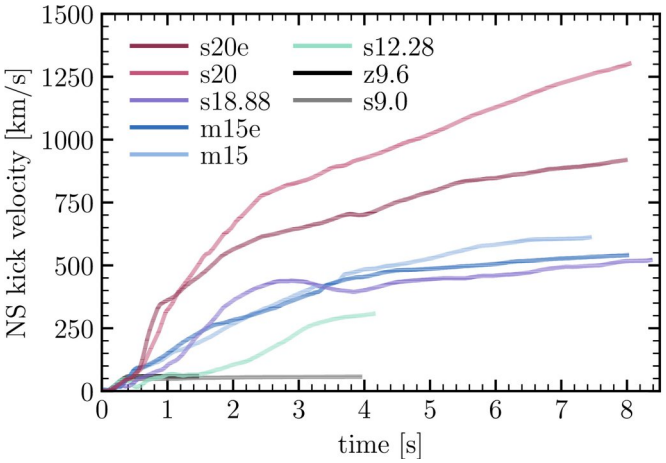

Figure 2: Kick velocities of the neutron stars as functions of time. They range from several 10 km/s to more than 1,000 km/s, explaining the measured space velocities of pulsars. The numbers in the model names mean the star's birth mass in solar masses.

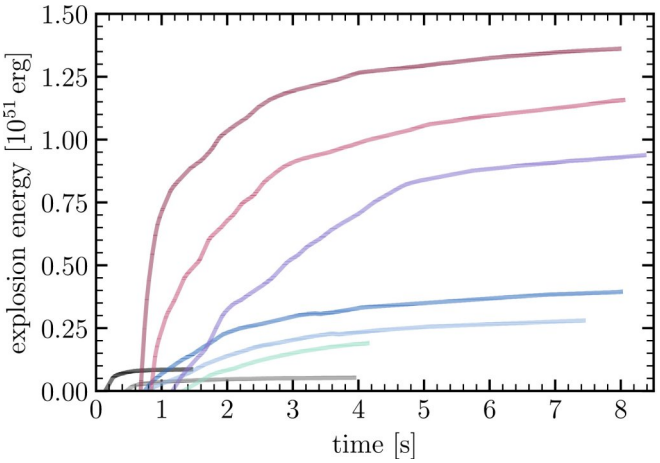

Figure 3: Explosion energies of our supernova models as functions of time. They vary between about $5{\cdot}10^{49}$ erg and nearly $1.3{\cdot}10^{51}$ erg (1 erg = 10^{-7} J), well covering the range of most observed supernovae. For model names and masses, see Figure 2.

The simulations are able to explain the observed kick velocities that neutron stars receive by the recoil associated with highly asymmetric matter ejection in the explosion and by asymmetric neutrino emission [4]. We obtained velocities between a few 10 km/s for neutron stars born in supernovae of low-mass progenitors and more than 1,000 km/s for heavier neutron stars formed during the explosions of massive progenitors (Figure 2). The lowest velocities are mainly caused by asymmetric neutrino emission, because low-mass progenitors explode with low energies or nearly spherically or both. In contrast, the higher velocities are a consequence of asymmetric mass ejection in the more energetic supernova explosions of higher-mass progenitors. The explosion energies in our set of models saturate partially only after several seconds of neutrino heating of ejected matter and range from $5{\cdot}10^{49}$ erg to nearly $1.3{\cdot}10^{51}$ erg (Figure 3). Both the neutron-star kick velocities and the explosion energies are in good agreement with observations of pulsars and the majority of supernovae, respectively. Since our long-term simulations cover the evolution phase that is relevant for nuclear reactions and for neutrino-induced nucleosynthesis, we have started to analyze our explosion models for the chemical-element production in ejected matter, applying a large nuclear reaction network in post-processing the hydrodynamic results. A first piece of work focused on the formation of radioactive ^{56}Ni and ^{44}Ti, both of which are crucial for supernova diagnostics [5]. Again, the results are in auspicious agreement with observations and offer a solution of the long-standing problem that spherically symmetric (1D) explosion models underproduce ^{44}Ti compared to the measured values.

Ongoing Research / Outlook

Our project defines only a starting point for exploring the vast diversity of supernova explosions connected with progenitor stars of different birth masses. Also the stellar metallicity (i.e., the amount of chemical elements heavier than hydrogen and helium in the gas of the star at its formation) and the stellar rotation are important degrees of freedom to be varied. Moreover, the mass-loss or mass-gain histories by stellar winds and binary interaction play important roles in defining the pre-collapse properties of supernova progenitors. Future studies therefore need to investigate much larger sets of models. This will require resources that can be provided only by the next generation of massively parallel supercomputers. Meanwhile we will post-process our entire set of explosion models for the chemical element formation due to nuclear and neutrino reactions. And we will extend the hydrodynamical explosion modeling until the shock penetrates the stellar surface in order to facilitate the computation of supernova light curves and spectra.

References and Links

[1] D. Kresse, "Towards Energy Saturation in Three-dimensional Simulations of Core-collapse Supernova Explosions", PhD Thesis, TUM (2023).
[2] G. Stockinger et al., MNRAS 496 (2020) 2039–2084.
[3] R. Bollig et al., Astrophys. J. 915 (2021) 28.
[4] H.-Th. Janka & D. Kresse, eprint arXiv:2401.13817.
[5] A. Sieverding et al., Astrophys. J. Lett. 957 (2023) L25.
[6] https://wwwmpa.mpa-garching.mpg.de/ccsnarchive/

Chemistry and Material Sciences

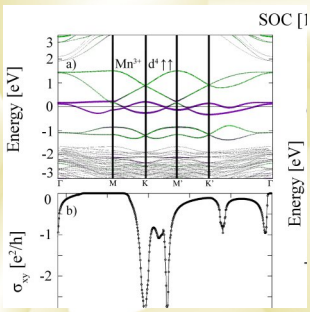

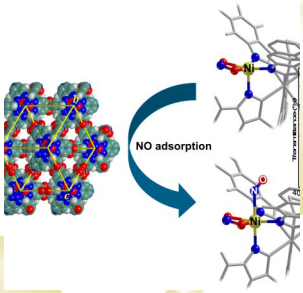

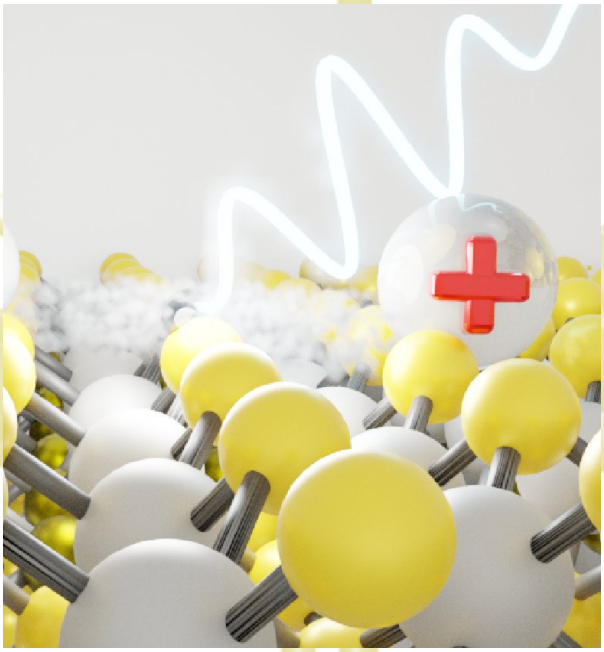

Light-matter interactions on the nanoscale

Research Institution
[1]Institute of Theoretical Physics, University of Regensburg

Principal Investigator
Jan Wilhelm[1]

Researcher
Carlo A. Pignedoli[2]

Project Partner
[2]Swiss Federal Laboratories for Materials Science and Technology (Empa)

SuperMUC Project ID: pn72pa

Introduction

Electronic excitations in matter play a pivotal role in various physical phenomena, including light absorption and transport. The characteristics of these excitations are strongly influenced by the host material. Excitons, which are bound electron-hole pairs (sketch in Fig. 1), exhibit a remarkable and unusually strong electron-hole binding in low-dimensional semiconductors that have emerged in the last decade. When stacking two atomically thin semiconductors on top of each other, the atomic alignment between the layers can exhibit periodic variations, leading to a new type of in-plane superlattice known as the moire superlattice. Excitons in moiré structures have gained enormous attention recently thanks to their highly unusual exciton properties which include spatial confinement due to the moire potential, interlayer, and intralayer charge transfer. Furthermore, electronic properties of moire lattices can be tuned by the band alignment and the twist angle between the layers such that moire structures hold great promise as an exciting platform for probing quantum phenomena over the next decade.

Results and Methods

Gaining insights into excitons in moire structures can be achieved through a combination of experiments, theoretical models, and computations. In Ref. [2], we focus on the GW method from many-body-perturbation theory which is an approximation for the electronic self-energy that allows for computing the electronic band structure of a given material. Importantly, GW accounts for the nonlocal, frequency-dependent screening of the interaction between electrons which is crucial in moire bilayers. The GW band structure is then the basis for the description of excitons via the Bethe-Salpeter equation. Currently available plane-wave-based GW algorithms are however incapable of treating low-angle moiré cells that contain thousands of atoms, despite their computational scalability to the largest supercomputers.

We have developed a low-scaling GW algorithm in the CP2K package [2] based on atomic-orbital basis function. We have extensively benchmarked our new 2D-periodic GW algorithm on Supermuc-NG, and we could deal with a $MoSe_2/WS_2$ heterostructure with up to 984 atoms in the unit cell, see Fig. 2. The size of this unit cell is 8 nm^2, so reaching the nanoscale. This calculation took 42 hours on 128 fat nodes on Supermuc-NG.

Large-memory nodes were required because intermediate three-index tensors were stored in the GW algorithm. After a rewrite of the GW code in summer and autumn 2023, we eliminated this memory bottleneck such that the code also can run efficiently on thin nodes. The trick was to recompute intermediate three-index tensors instead of storing them in memory. The fat nodes have been crucial to benchmark the first version of the GW algorithm.

Ongoing Research / Outlook

We are working to push the low-scaling GW algorithm [2] to even larger unit cells, which is interesting because experimentally relevant moiré structures with low twist angles can contain more than 10,000 atoms in the unit cell. We are optimizing CP2K in collaboration with LRZ and Intel to achieve best performance and low-memory requirements on state-of-the-art supercomputers. We are also using the numerical tricks of Ref. [2] to accelerate the Bethe-Salpeter equation, such that the study of exciton wavefunctions and exciton binding energies will be accessible for moiré structures.

References and Links

[1] https://www.ur.de/physics/wilhelm
https://www.empa.ch/web/s205/home
[2] M. Graml, K. Zollner, D. Hernangomez-Pérez, P. E. Faria Junior, J.Wilhelm, arXiv preprint 2306.16066 (2023), *accepted in* J. Chem. Theory Comput. (2024).

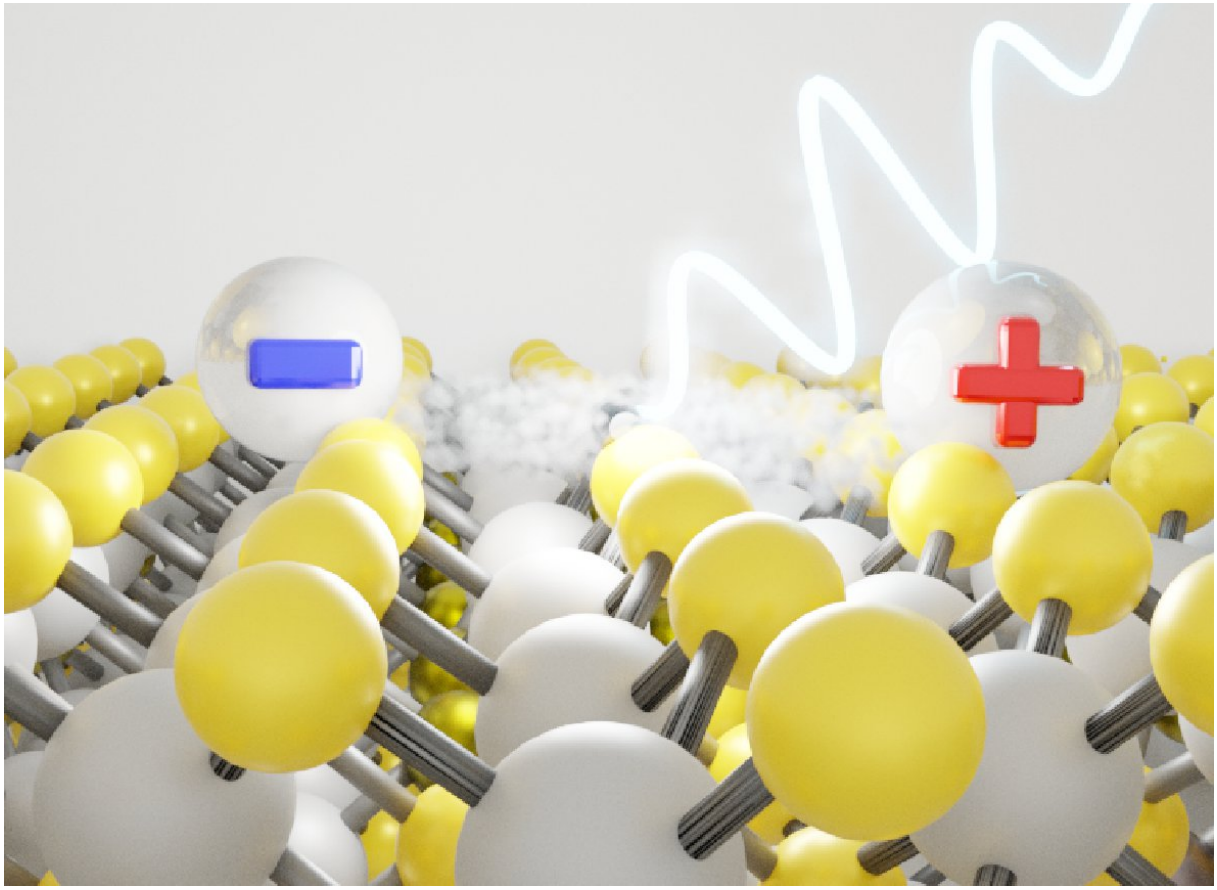

Figure 1: Artistic sketch: A material is irradiated by light and an electron-hole pair is formed.

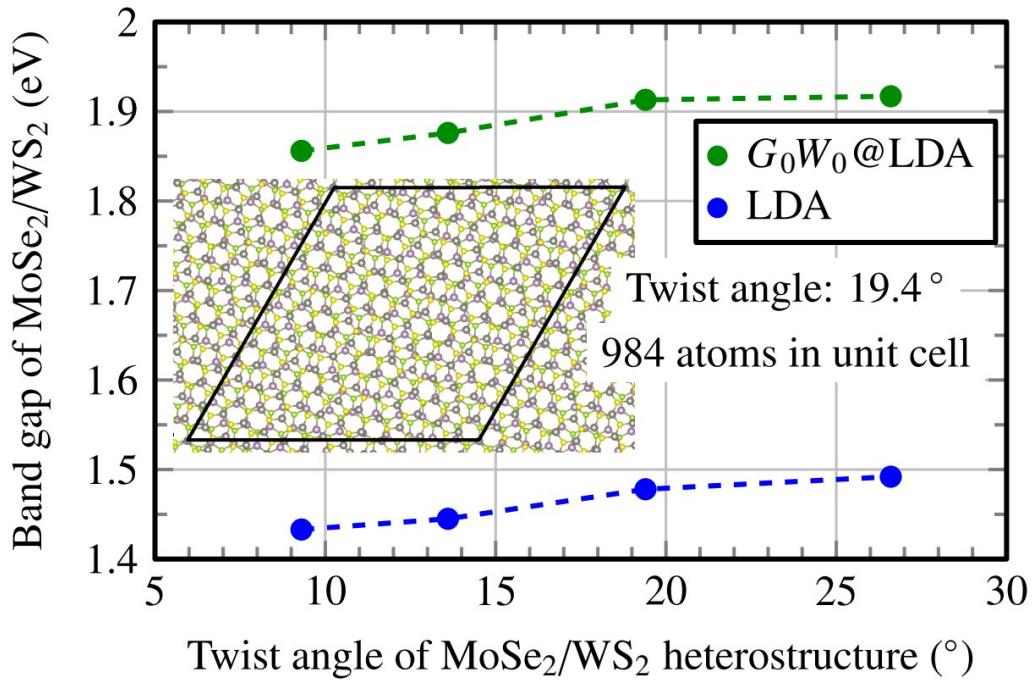

Figure 2: Band gap of a $MoSe_2/WS_2$ heterostructure as function of the twist angle. Inset: Unit cell (black rhomboid) for 19.4° twist angle contains 984 atoms [2].

Accelerated Materials Discovery with Automation and Machine-Learned Chemical Knowledge

Research Institution
[1]Friedrich-Schiller-Universität Jena
[2]Bundesanstalt für Materialforschung und -prüfung

Principal Investigator
Janine George[1,2]

Researcher
Philipp Benner[1], Aakash Naik[1,2], Joana Bustamante[1,2], Dr. Christina Ertural[1]

Project Partner
[3]University of Oxford
[4]Lawrence Livermore National Laboratory

Funding
BAM Junior Group

SuperMUC Project ID: pn73da

Introduction

This project aims to accelerate the search for new materials (e.g., for thermoelectric applications, battery materials, magnets, and other materials classes) based on ab initio high-throughput studies. High-throughput searches are typically restricted to known materials. This project explores strategies (data-driven chemical heuristics in subproject 1 and machine-learned interatomic potentials in subproject 2) to go beyond current database entries and include such computationally demanding properties in high-throughput searches. To accomplish each subproject, we develop automated workflows for high-throughput computations and provide large open databases of computed materials properties to the research community.

Results and Methods

Training data is needed to develop data-driven chemical heuristics and machine-learned interatomic potentials. In computational materials science, we typically rely on density functional theory-based (DFT) training data. Specifically, we use the density-functional theory package VASP [6], very often in combination with the PBE exchange-correlation functional embedded in the PAW method. We typically run these computations on 2-3 nodes with 96-144 CPUs and multiples thereof in the micro or general queue of the cluster.

To compute complex properties such as bonding-based descriptors or vibrational properties of materials, we typically have to combine multiple DFT runs (sometimes up to several hundred) and pre- and post-processing steps (e.g., projection of a plane-wave based basis set on atomic orbitals for bonding properties as implemented in LOBSTER and subsequent analysis of the data with LobsterPy [1,7], or computations of vibrational properties with the finite displacement method as implemented in Phonopy).

Performing high-throughput computations of such complex properties would not be possible in a manual way. Therefore, we develop workflow systems, specifically allowing for the demands of computational materials science projects (including options to use different Slurm submission scripts for different projects or jobs, allowing for dynamic generation of new jobs). One of the developments that we are involved in is jobflow[2] – a software for workflow development – and subsequent specific computational materials science workflows as implemented in the software atomate2 [3]. All the mentioned developments are open-source projects.

In this way, we have now generated an open quantum-chemical bonding database [4]. We are currently working on the integration of the software in a larger materials properties database that would allow for retrieving the data via an API and, therefore, allow for easier reuse of the data. The workflows and codes to accomplish this are all clearly documented with tutorials.

These bonding properties serve as new descriptors for materials science properties. We then retrieve the relationships between quantum-chemical bonding descriptors and materials properties, making the acceleration of the prediction of materials properties possible in the long run. Based on simple random forest models, we have already shown that the quantum-chemical bonding descriptors indeed have predictive power for vibrational properties [4].

Besides automating DFT runs and pre- and postprocessing steps, machine-learned interatomic potentials can be used within such workflow systems. This allows the automation of their training and benchmarking against DFT data. Our workflow developments for computed vibrational properties with DFT have been extended in such a way that they can now also be used with machine-learned interatomic potentials. We have used this workflow to benchmark a new foundation model for the predic-

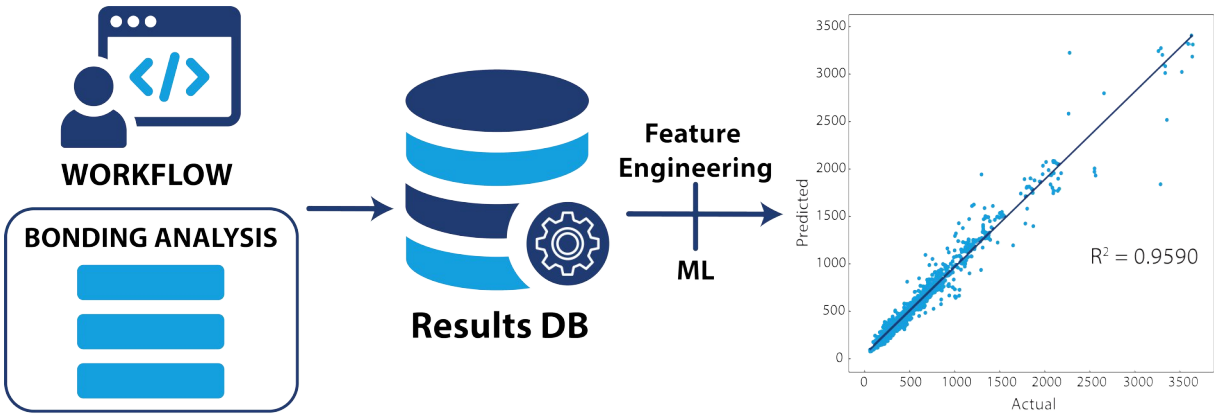

Figure: Schematic representation of the data generation workflow to analysis procedure. Adapted from [4]. Creative Commons CC-BY 4.0.

tion of phonon properties. While the overall phonon band structures cannot be exactly reproduced, overall bandwidths agree to an acceptable degree, potentially allowing for efficient pre-screening of materials' properties such as (low) thermal conductivity [5].

Ongoing Research / Outlook

Future work will be devoted to providing workflow solutions for the automatic development and benchmark of machine-learned interatomic potentials. This potentially includes efficiently and automatically combining heterogeneous computing resources (CPUs and potentially GPUs).

Furthermore, our dataset of quantum-chemical bonding descriptors is further extended to allow for the development of more deep-learning models, requiring more data.

References and Links

[1] J. George et al., ChemPlusChem 2022, 87, e202200123.
[2] A. S. Rosen et al., J. Open Source Software 2024, 9, 5995.
[3] A. Ganose, et al., https://github.com/materialsproject/atomate2, 2023.
[4] Naik, A.A., Ertural, C., Dhamrait, N. et al. A Quantum-Chemical Bonding Database for Solid-State Materials. Sci Data 10, 610 (2023). https://doi.org/10.1038/s41597-023-02477-5
[5] I. Batatiaet al., DOI 10.48550/arXiv.2401.00096.
[6] www.vasp.at
[7] https://github.com/jageo/lobsterpy

Large-scale phase-field simulations of solid-state sintering processes

Research Institution
[1]Helmholtz-Zentrum Hereon

Principal Investigator
Vladimir Ivannikov[1]

Researcher
Peter Munch[2,3], Martin Kronbichler[2,4], Christian Cyron[1,5], Thomas Ebel[1], Magdalena Schreter-Fleischhacker[6]

Project Partner
[2]University of Augsburg, [3]Uppsala University, [4]Ruhr University Bochum
[5]Technical University of Hamburg, [6]Technical University Munich

Funding
Helmholtz-Zentrum Hereon (I2B MgSinter project)

SuperMUC Project ID: pn36li (Gauss Large-Scale project)

Introduction

Sintering is a physically complex process that includes various mechanisms interacting and competing with each other. The obtained densification and microstructure of the sintered packing are of key interest. The accurate prediction of the powder coalescence for a given material and heating profile is a challenging multiphysics problem that couples mass transport and mechanics and involves multiple distinct stages: “early stage” vs. “later stage” (see Fig. 1). These rheological differences justify the application of specialized numerical models and methods with different computational costs for each of the stages.

At Helmholtz-Zentrum Hereon, we conduct experiments and develop numerical tools for simulating sintering processes. These include DEM models [1], which are limited to early-stage simulations but allow to simulate thousands of particles on a simple laptop, and phase-field-based models [2], which are inherently expensive but allow for more accurate predictions. Recent development efforts target the optimizations of our phase-field solver so that tens of thousands of grains can be simulated as in the DEM case. This is necessary to obtain statistically meaningful representative results as well as necessitates the use of supercomputers and the optimization of the code for them.

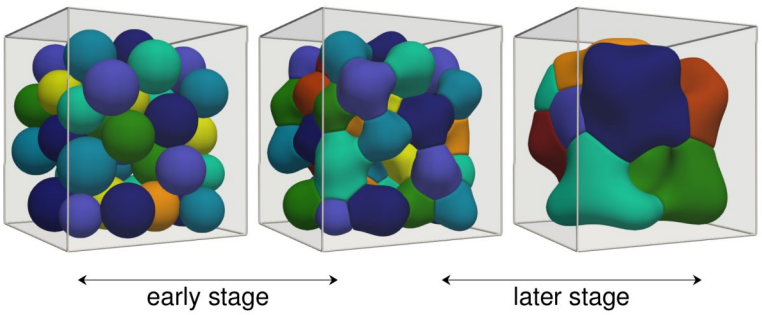

Figure 1: Visualization of different phases of sintering for a 51-particle packing in 3D.

Results and Methods

Numerics, solution process, and challenges
In the context of phase-field simulations, sintering is described by one Cahn-Hilliard equation and by one Allen-Cahn equation per grain. In practice, non-neighboring grains are merged to order parameters, reducing the number of Allen-Cahn equations needed and the computational costs. The resulting time-dependent non-linear partial differential equations are discretized in space with the finite element method and in time with the backward differentiation formula (BDF2). The nonlinear system is solved with a Newton–Krylov approach, in which we either evaluate exactly the Jacobian or approximate it via finite difference (Jacobian-free Newton–Krylov).

Methods
In order to solve the resulting nonlinear system efficiently, we have proposed different ingredients previously. In the following, we only shortly summarize the most important aspects and refer the interested readers to [3].

We apply matrix-free operator evaluation for fast Jacobian and residual evaluation [4]. Material kernels that are evaluated often have been heavily optimized to use SIMD vectorization and to use well the cache. We have developed block-Jacobi preconditioners that are cheap to set up and to apply, since we apply the same block to multiple grains, allowing to work with multivectors. We proposed a novel distributed grain-tracking and remapping algorithm, which replaces all-to-all communication patterns used in the literature to collect mesh/grain information from all processes. Furthermore, we heavily rely on adaptive mesh refinement, which is crucial to minimize the number of degrees of freedom and concentrate the work at interesting regions (interfaces of grains).

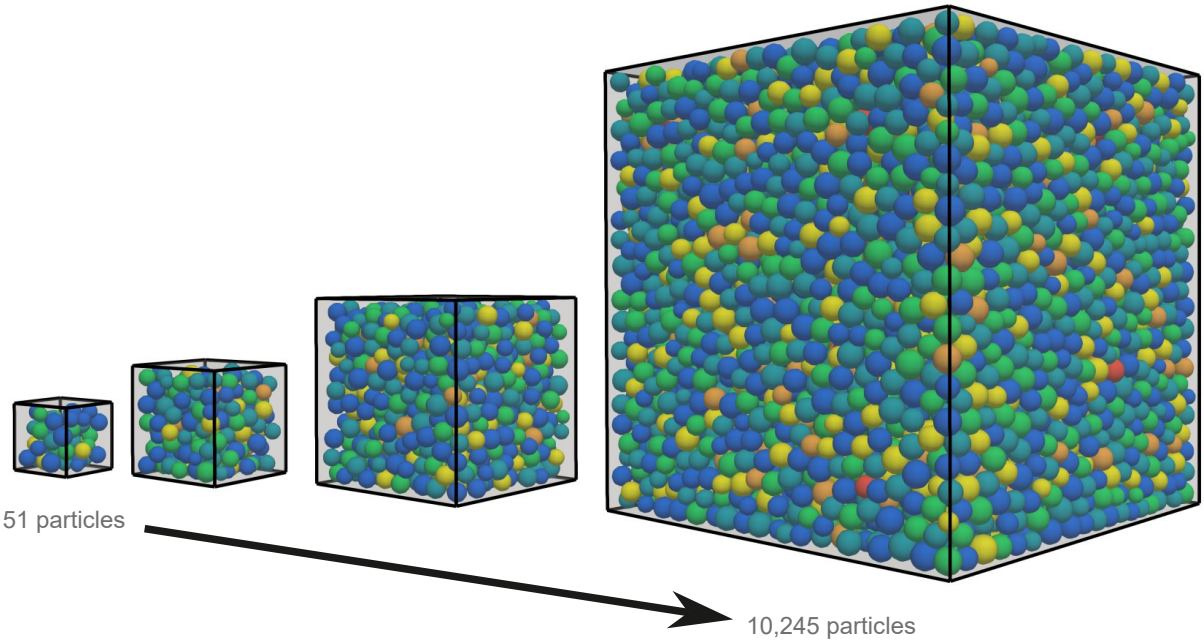

Figure 2: Particle packings considered for the scaling study: 51 - 10,245 particles [3].

While porting the code to SuperMUC-NG, which allowed us to increase the number of simulated grains from hundreds to tens of thousands, we observed non-negligible setup costs, which were related to unfavorable complexities of the algorithm that did not hurt on small scales. We eliminated these and now are able to scale simulations up to 1k nodes. We are in the progress to use these capabilities to quickly run sensitivity studies for different materials, which are investigated by the experimental group at Helmholtz-Zentrum Hereon.

Results

Fig. 2 shows some grain packings we considered during our scaling studies. 51 grains (left panel) is a configuration that we mostly consider during code development, since simulations can be quickly run within a few hours on a few nodes of our in-house cluster. 10k grains, on the other hand, is a configuration crucial to obtain statically relevant measurements and requires supercomputing facilities. Fig. 3 indicates that one needs at least 128 nodes for such a configuration due to memory requirements. Furthermore, Fig. 3 contains scaling results of the linear solver for a representative simulation. The timings of the linear solver, which is the most time-consuming part of the nonlinear solution procedure, contain both the evaluation of the Jacobian, which uses the efficient and optimized matrix-free kernels, and the developed block preconditioner. One can observe that, for sufficiently large problem sizes per process, it is possible to increase the number of processes by a factor of 16–32 with only a loss of 25% in parallel efficiency. Furthermore, the lowest times to solution are reached for about 10k DoFs per process, which is in good agreement with the scalability of FEM solvers from other application fields.

Software

All software developments are available within the open-source project hpsint [5]. It targets HPC systems for

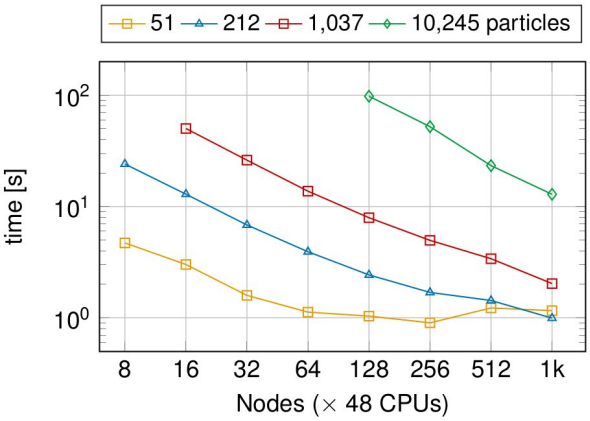

Figure 3: Scaling of linear solver for 51 - 10,245 particles [3]. Times are accumulated over 10 time steps.

sintering applications and is based on deal.II, p4est and Trilinos. The implementation has been verified against results in the literature and shows a clear speedup compared to similar FEM-based solvers.

Ongoing Research / Outlook

Ongoing research adds further physical models that target limitations regarding rigid-body motions and enable liquid-phase sintering relevant in the context of selective laser-melting additive manufacturing processes. Furthermore, we are implementing optimizations regarding data structures, cache locality and preconditioning, which we have identified in [3] and promise a further speedup of 2x.

References and Links

[1] Ivannikov et al. 2021. MSMSE.
[2] Ivannikov et al. 2023. CPM.
[3] Munch et al. 2024. CMS.
[4] Kronbichler, Kormann. 2012. C&F.
[5] https://github.com/hpsint/hpsint

Multipoint Vertex Functions in Strongly Correlated Electron Systems

Research Institution
[1]Arnold Sommerfeld Center for Theoretical Physics, Ludwig-Maximilians-Universität München

Principal Investigator
Jan von Delft[1]

Researcher
Seung-Sup Lee[2], Andreas Gleis[1], Mathias Pelz[1], Johannes Halbinger[1], Jeongmin Shim[1]

Project Partner
[2]Seoul National University, South Korea

Funding
DFG EXC-2111 (No. 390814868), DFG LE 3883 (No. 2-2(678243)), Munich Quantum Valley, International Max Planck Research School for Quantum Science and Technology (IMPRS-QST), National Research Foundation of Korea (NRF), Institute for Basic Science (IBS)

SuperMUC Project ID: pn25ze (Gauss Large-Scale project)

Introduction

In quantum many-body physics, correlation functions, usually abbreviated to correlators, are the quantum expectation values of operators acting on different space-time points. When a correlator involves ℓ space-time points, it is called an ℓ-point correlator, and when ℓ is larger than two, it is a multipoint correlator. When the system of interest is electrons in a solid and the operators are electron creation and annihilation operators, the correlators are also called electron Green's functions. (Here "creation" means the introduction of an electron to a solid from the outside; "annihilation" is the opposite operation.) The Green's functions are important since they determine various dynamical responses and spectral properties of the solid. For example, the single-electron function, which is a two-point correlator defined by a pair of annihilation and creation operators, is probed by angle-resolved photoemission spectroscopy (ARPES) spectra. On the other hand, the two-electron function, a four-point correlator defined by two pairs of annihilation and creation operators, is relevant to the structure factor in inelastic neutron scattering as well as the optical conductivity. Four-point vertex functions are key ingredients of two-electron Green's functions and describe the effective interaction between two electrons in the presence of a many-electron background. Hence vertex functions are important in quantum field theory approaches to quantum many-body problems. Enigmatic quantum phenomena in strongly correlated electron systems, such as high-temperature superconductivity and quantum criticality, typically result from strongly interacting electrons at low temperatures. In such cases, the precise computation of the Green's and vertex functions is notoriously hard; indeed, it constitutes one of the holy grails in condensed matter physics. In this project, we have focused on computing four-point vertex functions in a class of quantum many-body systems, called quantum impurity models, as the first steps towards the holy grail. An impurity model consists of a subsystem in which electrons are interacting ("impurity") and the rest in which electrons are not interacting ("bath"). This setting of the model may sound specific, but according to the dynamical mean-field theory (DMFT) and its extensions, impurity models are useful simplifications of the lattice description of strongly correlated electrons. Once the four-point vertex function at the impurity is computed, it can be used to derive reliable estimates of Green's functions of the original lattice system by solving a set of equations from quantum field theory (see Outlook for detail).

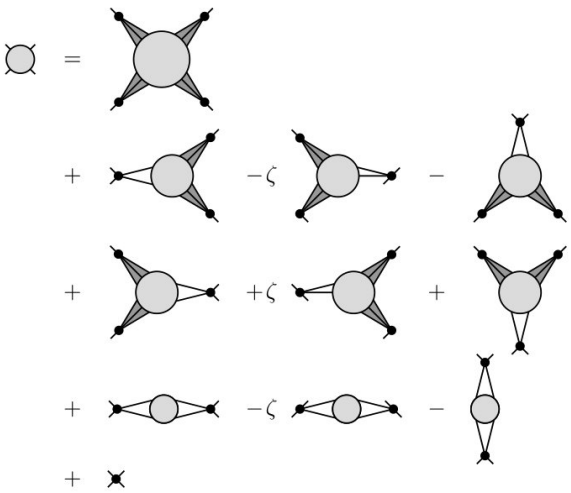

Figure 1: Feynman diagrams of the four-point symmetric improved estimator (sIE). The four-point vertex function on the left-hand side is decomposed into the diagrams on the right-hand side under the sIE, and each diagram corresponds to a different asymptotic frequency dependence of the vertex.

Results and Methods

To compute the three- and four-point vertex functions of quantum impurity models, we use the multipoint numerical renormalization group (mpNRG) method [1,2] developed by some of us. It is a tensor network method specialized for solving impurity models; here the many-body energy eigenstates of the Hamiltonian are represented by matrix product states (MPSs) that are one-dimensional tensor network states. Our previous works [1,2] show that the mpNRG gives accurate results of the four-point vertex functions in the Matsubara formalism (MF), where the functions' arguments are imaginary frequencies. It can access much lower temper-

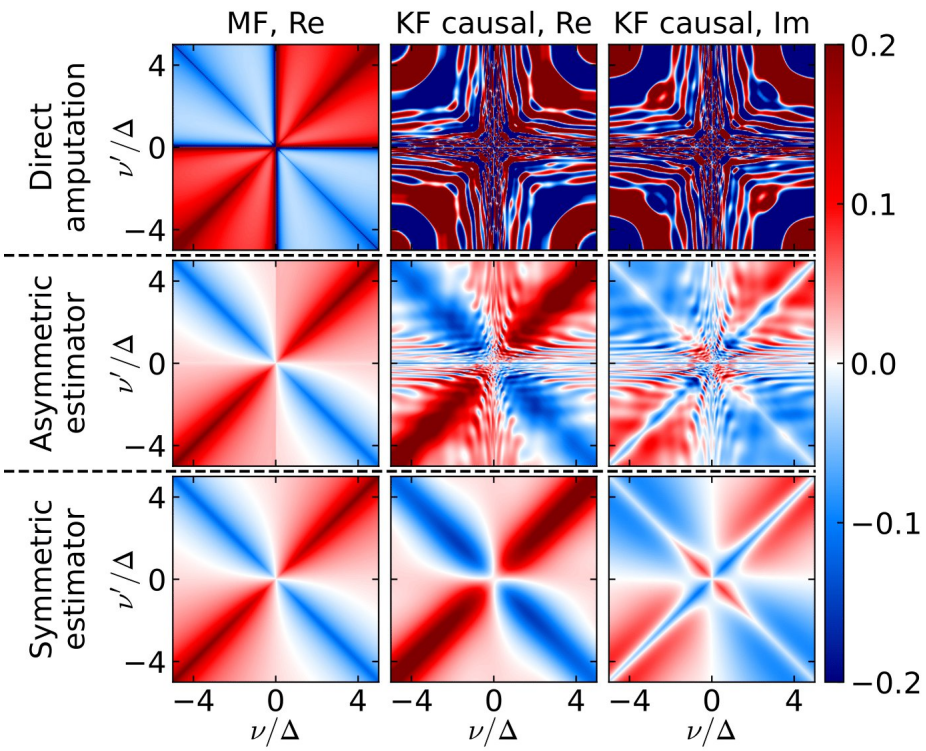

Figure 2: The local four-point vertex function of the single-impurity Anderson model in the Matsubara formalism (MF) and Keldysh formalism (KF). The three rows show the results obtained by direct amputation of the correlator, by using the asymmetric improved estimator (aIE), and by using the symmetric improved estimator (sIE) derived in our project. The first, second, and third columns show the MF vertex (which is purely real), the real part of the causal (c = ----) component of the KF vertex, and its imaginary part, respectively.

atures than the state-of-the-art quantum Monte Carlo (QMC) methods do. Of course, the imaginary frequencies are mere theoretical tools, so to compare with experimental measurements, one should be able to compute the vertex functions in the Keldysh formalism (KF), where the functions' arguments are real frequencies. To obtain reliable results of KF vertex functions, the use of an improved estimator (IE) is crucial. An IE is an analytically exact expression of the quantity of interest (e.g., vertex functions) in terms of various correlation functions. In its numerical evaluation, numerical artifacts in the correlation functions (e.g., discretization artifacts in the NRG method) are largely canceled out, leading to higher accuracy. However, the previously proposed IEs for vertex functions, including the asymmetric IE (aIE) used in Refs. [1,2], were not adequate for our purposes, especially for the real-frequency KF. For example, the aIE improves, due to the asymmetric nature of its formulation, the accuracy only for some components of KF vertex functions. In this project, we developed symmetric improved estimators (sIEs) [3] that yield accurate estimates of the multipoint vertex functions in KF. Our sIEs are derived by applying equations of motion recursively, leading to lower-point correlators. As a result, sIEs for a four-point vertex are given in terms of four-, three-, and two-point functions (see Fig. 1). We demonstrate the advantages of the sIE by computing the KF four-point vertex of the single-impurity Anderson model, which is a paradigmatic impurity model. Figure 2 shows how the sIE improves the quality of the vertex function by comparing the vertex functions obtained from the direct amputation, the aIE, and the sIE. The numerical artifacts are adequately eliminated only for the sIE. While the use of the sIEs for vertex functions has been found to give numerically much more accurate results than the aIEs, this gain in accuracy comes with increased computational costs. Fortunately, the evaluation of the sIEs can be parallelized over a large number of threads and nodes, which made SuperMUC-NG an ideal platform for our project. The most CPU- and memory-intensive part in calculating a four-point vertex is the computation of four-point partial spectral functions (PSFs), which are three-dimensional arrays that contain the core information of the four-point correlators appearing in the sIE. There are 120 four-point PSFs for different tuples of operators. The PSFs are independent, so they can be computed on different nodes with small communication overhead. Moreover, by slicing the PSFs, the computation can be further parallelized so that different slices are treated by different workers. Within each worker, the most of CPU time is spent on contracting tensors that represent the operators and the energy eigenstates. To obtain the sIE results shown in Fig. 2, we used roughly 1,000 nodes for 60 hours on SuperMUC-NG, which is hard to accomplish on the other HPC systems.

Ongoing Research / Outlook

Within our theoretical framework, the four-point vertex of the effective impurity model, which is the output of mpNRG and sIE calculations, will be fed into a method for solving field-theoretical equations as input. We are currently aiming at solving the parquet equations, a set of self-consistent equations that involve four-point vertex functions. A related SuperMUC-NG project, coined multiloop functional renormalization group (mfRG), is led by other members within the group of Prof. Jan von Delft. The solution of the parquet equations with the impurity vertex input amounts to the "parquet dynamical vertex approximation (DΓA)", which is believed to be a powerful method for studying strongly correlated electrons in two dimensions. We plan to investigate various open problems for future work within this field: strange metallicity, quantum criticality, and high-temperature superconductivity, to name a few.

References and Links

[1] F. B. Kugler et al., Phys. Rev. X 11, 041006 (2021).
[2] S.-S. B. Lee et al., Phys. Rev. X 11, 041007 (2021).
[3] J.-M. Lihm et al., Phys. Rev. B 109, 125138 (2024).

The kagome lattice antiferromagnet: tuning frustration and quantum fluctuations

Research Institution
[1]Bielefeld University, Faculty of Physics

Principal Investigator
Jürgen Schnack[1]

Researcher
Jörg Schulenburg[2], Johannes Richter[2,3]

Project Partner
[2]Otto-Von-Guericke University, Magdeburg
[3]Max Planck Institute for the Physics of complex Systems, Dresden

SuperMUC Project ID: pr62to

Introduction

The spin-1/2 kagome Heisenberg antiferromagnet (KHAF) and its relatives such as the square-kagome Heisenberg antiferromagnet (SKHAF) are at the same time very prominent as well as challenging spin models in the field of frustrated quantum magnetism. The first challenge concerns the nature of the ground state on which a plethora of studies exist. The second challenge concerns the thermodynamic properties 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, i.e., by evaluating the magne-tization *M* as a function of the external magnetic field *B*.

In our project we address the second as well as the third point. After having dealt with the phenomenon of magnon crystallization successfully, we focused on the magnetization process and in particular on the asymmetric melting of magnetization plateaus in this phase of the project. This kind of research requires a very high level of quantum physics as well as supercomputing experience, so that we are practically the only group in the world working at this level [1]. Competitors such as Tôru Sakai (Japan) treat only much smaller spin systems (e.g. $N = 36$ compared to $N = 42$) on the Japanese super computers.

Results and Methods

The investigated spin systems are modeled by a spin-1/2 Heisenberg Hamiltonian augmented with a Zeeman term,

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

where J is the antiferromagnetic nearest-neighbor exchange coupling and B the external magnetic field. The complete eigenvalue spectrum of a spin system composed of spins $s = 1/2$ 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

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

In the spirit of trace estimators, the trace is in a Monte-Carlo fashion replaced by a much smaller sum over R random vectors $|\nu>$ for each symmetry-related orthogonal subspace $\mathcal{H}(\gamma)$ 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 $|\nu>$. The vector $|n(\nu)>$ is the n-th eigenvector of the Hamiltonian in this Krylov space. This allows to evaluate typical observables such as magnetization and specific heat.

We could produce two major results in this project:
(1) We discussed the phenomenon of asymmetric melting in the KHAF and traced it back to a combined effect of unbalanced magnetization steps on either side of the investigated plateau as well as on the behavior of the density of states across the plateau [2]. We also compared our findings to the SKHAF that behaves similarly at low temperatures at zero field, but as we could demonstrate differently at 1/3 of the saturation magnetization. Both systems possess a flat one-magnon band and therefore share with the class of flat-band systems the general property that the plateau that precedes the jump to saturation melts asymmetrically. But now the minimal susceptibility bends towards lower fields with increasing temperature.

Figure 1 graphically summarizes the melting of magnetization plateaus by plotting the differential magnetic susceptibility for the KHAF with $N = 42$ sites (left). A flat magnetization plateau corresponds to zero susceptibility, melting increases the susceptibility, and an asymmetric increase expresses itself as a banana-shaped feature. This behavior is clearly visible in the region around

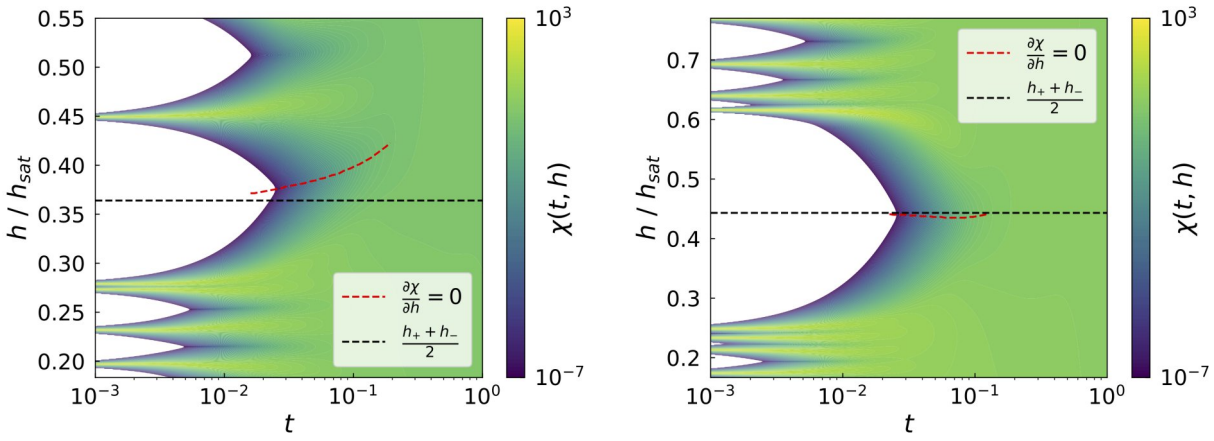

Figure 1: The differential magnetic susceptibility for the KHAF (left) and SKHAF (right) with *N = 42* sites [2].

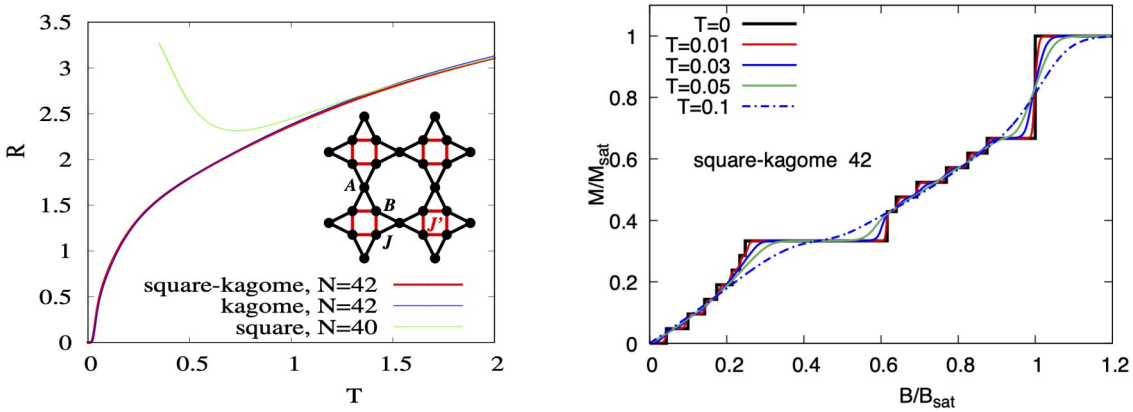

Figure 2: Left, modified Wilson ratio *R* of the spin-half SKHAF (*N = 42*) compared with respective data of the kagome HAF (*N = 42*) and the square-lattice HAF (*N = 40*); inset: Sketch of the square-kagome lattice. Right, magnetization of the SKHAF for *N = 42* and *T ≥ 0* [3,4].

h ~ 0.28 - 0.45 and additionally highlighted by the red dashed curve (minimum of the susceptibility) bending towards higher fields compared to a symmetric behavior shown by the black dashed curve. However, for the SKHAF (right) the 1/3-plateau melts symmetrically, see region around *h ~ 0.27 - 0.61*. The dashed red curve does not deviate from the symmetric black dashed line.

(2) In two extensive studies we approached the physics of the spin-half SKHAF [3,4], that is related to the KHAF, but naturally also allows for two different exchange constants. The SKHAF is a non-Archimedean lattice without magnetic long-range order and sometimes also called shuriken lattice. We filled the gap of missing finite-temperature studies of the spin-1/2 SKHAF by large-scale numerical simulations of finite lattices of up to N = 54 sites.

Figure 2 shows the Wilson ratio for the *N = 42* SKHAF in comparison with the *N = 42* KHAF (both with non-magnetic spin liquid ground state) as well as for the *N = 40* square-lattice HAF (with a magnetically ordered ground state). The striking accordance of the kagome and square-kagome data is obvious signaling the dominance of singlets as T → 0. This behavior agrees with earlier findings. Other interesting properties of the SKHAF are related to the magnetization process in an applied magnetic field. There are two well pronounced plateaus at 1/3 and 2/3 of the saturation magnetization and a jump from the 2/3 plateau directly to saturation caused by the flat one- magnon band. The melting of the plateaus with growing temperature is faster for the upper 2/3 plateau and it melts asymmetrically.

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-NG. 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. One possible direction to follow in the future is the evaluation of needed matrix elements *on the fly*.

References and Links

[1] http://obelix.physik.uni-bielefeld.de/~schnack, https://wasd.urz.uni-magdeburg.de/jschulen/spin/
[2] H. Schlüter, J. Richter, J. Schnack, J. Phys. Soc. Jpn. 91 (2022) 094711.
[3] J. Richter, O. Derzhko, J. Schnack, Phys. Rev. B 105 (2022) 144427.
[4] J. Richter, J. Schnack, Phys. Rev. B 107 (2023) 245115.

Understanding the origin of the resistive switching in oxide perovskites

2

Research Institution
The NOMAD (Novel Materials Discovery) Laboratory at the Fritz-Haber-Institut der Max-Planck Gesellschaft and IRIS Adlershof of the Humboldt-Universität zu Berlin

Principal Investigator
Wahib Aggoune, Matthias Scheffler

Researcher
–

Project Partner
–

SuperMUC Project ID: pn49he

Introduction

This project is funded by the Leibniz association [5] and it is a collaboration with experimental partners at the Leibniz Institute for Crystal Growth (IKZ-Berlin).

Titanate-based oxide perovskites, with a general formula $ATiO_3$ (A=Sr, Ba, etc), emerged as highly promising materials for the next generation of electronic devices because of their variety of outstanding physical properties such as structural phase transitions, ferroelectricity, insulator-metal transition, photocatalysis, superconductivity, etc. Among them, cubic strontium titanate ($SrTiO_3$) is considered the ideal prototype for exploring these fundamental properties and their potential applications. At room temperature, $SrTiO_3$ exhibits a paraelectric state with a high static dielectric constant. It is found that it may approach a ferroelectric phase at low temperature [1]. This state is also reported in nanometer-thick films of $SrTiO_3$ and attributed to the formation of Ti antisite defects (Ti occupying Sr site) [2]. The off-displacement of Ti atoms is assumed to induce local polarization.

This characteristic motivated researchers to explore the physical properties of the defected $SrTiO_3$ samples. Recently, pronounced switching of the resistivity has been observed in thin films, depending upon the voltage history [3]. This property became the key for developing nonvolatile memories [4]. Such a property is achieved in films grown under Sr-*poor* and Ti-*rich* conditions [3]. An electronic mechanism based on switching the defects polarization which, in turn, allows the charging and decharging of the defects trap site, is proposed to explain the resistive switching behavior. However, a convincing proof is still messing and the full physical mechanism is unclear.

To understand the microscopic mechanism behind the resistive switching, we employ density functional theory (DFT) to explore the stability of these defects, under different growth conditions, and estimate the induced polarization. Employing the nudged-elastic band (NEB) method, we investigate the energy barrier for defects diffusion as well as for switching their polarization.

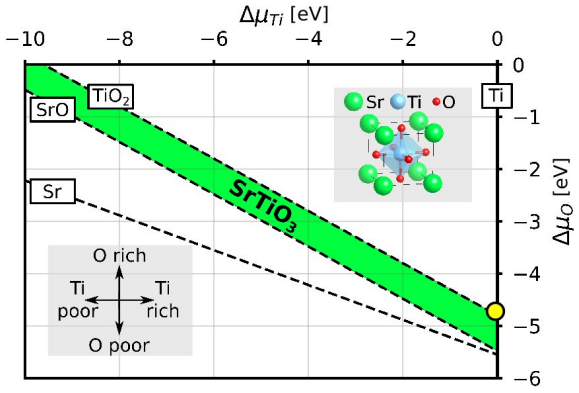

Figure 1: Phase diagram of bulk $SrTiO_3$ as a function of the chemical potential of Ti and O.

Results and Methods

All calculations were performed using FHI-aims [6] which is an all-electron full-potential package implementing DFT and beyond DFT methodologies. It uses numeric atom-centered orbitals and is a highly efficient package also with hybrid functionals. Low- communication parallelization of all real-space grid-based algorithms, that was recently advanced towards exascale hardware, a ScaLapack-based, and an eigenvalue solver (ELPA), customized handling of the linear algebra for all matrix operations are possible, guaranteeing efficient scaling (CPU time and memory) up to massively parallel computer systems with thousands of cores [6]. To converge the defects formation energy, a large supercell containing 320 atoms is required. PBEsol exchange-correlation functional is used for the structural optimization. The hybrid functional HSE06 is used to compute the energies and the electronic properties. Also this software is being developed toward the exascale [7]. We simulated the different systems using up to 20 full nodes (960 cores). Such large calculations were only possible on the SUPERMUC-NG architecture.

First, we explored the formation of the $SrTiO_3$ phase under different growth conditions. Considering the boundaries related to the formation of Sr, SrO, TiO_3 and Ti, the $SrTiO_3$ crystal is found to appear within the narrow area highlighted in green (see Fig. 1).

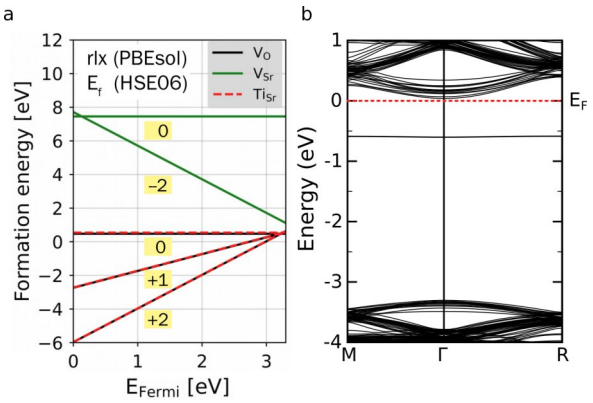

Figure 2: DFT calculations with the HSE06 functional: a) Defect formation energy as function of electronic chemical potential. b) Electronic band structure of the Ti antisite defect.

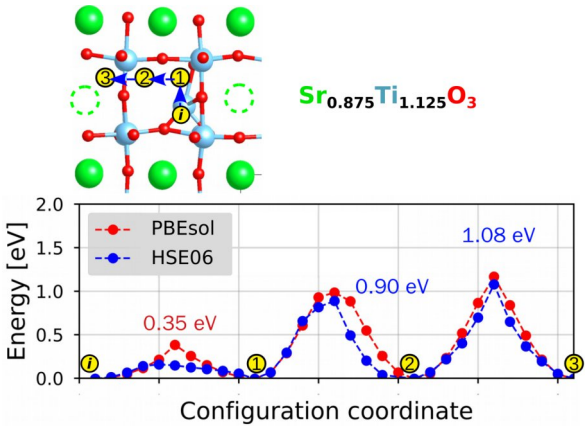

Figure 3: Sketch of the atomic structure of defected $SrTiO_3$ (top) and the energy barriers for Ti diffusion along positions (1), (2), and (3) (bottom).

As the experimental thin films are grown under Sr-*poor* and Ti-*rich* conditions, we considered the chemical potentials highlighted with yellow spot in Fig. 1 and computed the formation energy of different defects such as O vacancy (V_O), Sr vacancy (V_{Sr}) and Ti antisite (Ti_{Sr}).

As depicted in Fig. 2a, our calculations revel that Ti-anstisite is energetically more favorable than O vacancy. For a Fermi level raging between the valence band maximum (VBM) and conduction band minimum (CBM), we found that this defect is stable in the +2 charge state. It becomes neutral for EF slightly below the CBM.

In experiment, the applied voltages are expected to rise the Fermi level, thus we analyze the electronic structure of the neutral Ti antisite defect. Other defects such as O vacancies are ignored as the resistive switching is observed before and after annealing the samples with oxygen.

As depicted in Fig. 2b, the Ti antisite defect has an electronic state at mid-gap. This result supports the trap assisted mechanism of the resisitive switching. Employing the Berry-Phase method, we found that this point defect induces a polarization. We then computed the energy barrier for switching the defect position and also for a possible diffusion process outside its pseudo-cube. For the latter case, we also considered Sr vacancies. As shown in Fig. 3, the energy barrier for switching the position of the Ti atoms to position (1), i.e. also switching its polarization, is low (about 0.35 eV when using PBEsol). Using the HSE06 functional, which describes better the Ti-d states, this barrier decreases to about 0.2 eV.

We found that the diffusion of the Ti antisite outside its peudo-cube is high being about 1 eV. This indicates that the formation of a conductive filament, which can be achieved by diffusion and alignment of these defects, can be excluded. This is inline with experimental observation [3]. It also indicates that the resisitive switching is mainly governed by the Ti antisite defect and its dynamic within the pseudo-cube. Overall, moving from the initial position to position (1) is energetically possible and, importantly, allows to switch the polarization direction. However, the full RS mechanism requires additional calculations to cover other defects and distributions as well as their switching/diffusion barriers.

Ongoing Research / Outlook

This work needs a few final calculations and will then be published. The study will be extended to consider other complex defects depending on the experimental observations and outputs. We will then explore these quantities in selected titanates-based perovskites and their alloys, using SUPERMUC-NG. The resulting database will be used to build a machine learning model for predicting the resistive switching ability of the full space of titanates-based perovskites.

References and Links

[1] T. F. Nova et al., Science 364 (2019) 1075.
[2] D. Lee et al., Science 349 (2015) 1314.
[3] A. Baki et al., Sci Rep 11 (2021) 7497.
[4] M. A. Zidan et al., Nat Electron 1 (2018) 22.
[5] https://www.leibniz-gemeinschaft.de/
[6 V. Blum et al., Comput. Phys. Commun 180 (2009) 2175.
[7] https://www.nomad-coe.eu

The Solvation Properties of Gold at Ore-forming Supercritical Conditions

Research Institution
[1]Lehrstuhl für Theoretische Chemie, Ruhr-Universität Bochum

Principal Investigator
Dominik Marx[1]

Researcher
Jan Noetzel[1], Philipp Schienbein[1], Harald Forbert[2]

Project Partner
[2]Center of Solvation Science ZEMOS, Ruhr-Universität Bochum

SuperMUC Project ID: pr86fo

Introduction

Supercritical water (SCW) and hydrothermal fluids have been focused on a broad range of scientific disciplines, in particular, the geology and geochemistry of the deep seafloor and the earth mantel's crust are intertwined with the solvation properties of water. Pure water reaches its supercritical phase at pressures and temperatures above 221 bar and 647 K, respectively, which defines the supercritical point of bulk water together with a density of 0.322 kg/L. Water is mostly present as SCW in the Earth's mantle and crust where it dissolves minerals and can ejects their solutions through so-called hydrothermal vents located in the deep sea [1].

Remarkably, the solvation properties of SCW can be systematically tuned by the external pressure, which ultimately changes the density of the fluid at a given temperature. This can be roughly understood in terms of the relative permittivity of water which drops from 78 at ambient conditions to 10–25 in the supercritical phase. The latter is in the order of typical apolar organic solvents and thus very different from normal water. Therefore, the solubility of neutral hydrophobic species is typically increased in SCW compared to water at ambient conditions, where they need to be incorporated in void-like structures available in the three-dimensional H-bond network [2]. In contrast, electrolytes do not show such simple trends: From phase diagrams of several binary electrolyte-water mixtures, one can classify electrolytes into soluble and generally insoluble in SCW following the trends in temperature. In the "gas-like" regime (say below roughly 0.2 kg/L), even the soluble electrolytes get expelled from the homogeneous fluid phase, resulting in so-called "supercritical phase separation" [3]. In such aqueous hydrothermal fluids, solvation and mass transport processes have been studied since long for many inorganic compounds like silica, sulfites and simples gases. Remarkably, it has been discovered recently that neutral metal clusters – in particular gold nanoparticles – are dissolved in the first place and subsequently transported in hydrothermal fluids, which provides an interesting explanation on how ore deposits are formed in Earth's crust [3].

Using ab initio molecular dynamics (AIMD), we simulate the neutral gold dimer and tetramer in SCW at a temperature of 750 K and a mass density of the solution below 0.3 kg/L. The fundamental aim of the present investigation is to find out how these neutral metal particles integrate themselves into the highly dynamical fluid environment offered by supercritical water.

In the solvation interactions between water and gold, we found a strong polarization in the molecular dipole moment of both, solvated gold clusters and solvating water molecule due to its interactions [4]. In strong contrast to the gas-like SCW molecules in the systems's background, the solvating water molecule is polarized to dipole moment values (up to 3.06 D) even exceeding average dipole moment at room temperature conditions (RTW), which is 2.86 D for RPBE-D3 water.

Now we are going a step further by analyzing the IR-response of our system in the THz region, which gives us direct access to an observable measure for molecular interactions.

Results and Methods

The present AIMD simulations rely on the same validated electron structure protocol as our previous simulations of pure SCW [5] using RPBE-D3 based density functional theory with a TZVPP basis at a cutoff of 500 Ry. In particular, all simulations are carried out using the CP2k software suite [6] for Born-Oppenheimer propagation of the electronic structure concurrently with the moving classical nuclei. Core electrons are implicitly considered using norm-conserving relativistic GTH pseudopotentials (leaving 11 valence electrons for Au).

The real-time dipole dynamics of the systems, using Wannier-localized electron density, was generated in a set of 51 Newtonian energy conserving simulations (NVE ensemble), where their starting conditions were drawn from the temperature conserving ensemble of the system (NVT ensemble). The electron density was localized every femtosecond for 40 picoseconds in the NVE simulations, which covers fast H-O vibrations and several changes in the configuration of the solvation system within the same trajectory.

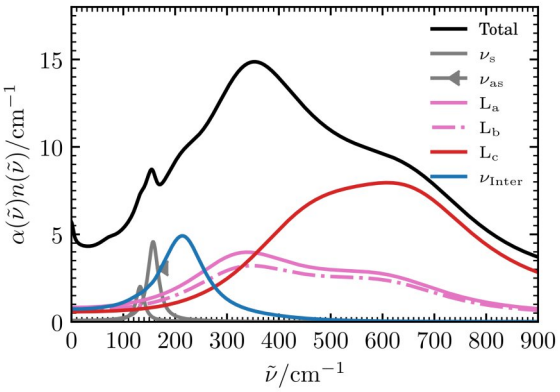

Figure 1: The IR-spectrum of the SSC between the gold tetramer and solvation water (black) dissected in parts: The gold stretch modes (gray), the water-gold interaction (blue) and the solvation water librational motion (red and pink).

The Newtonian trajectories of the nuclei and the pseudo-dynamics of the localized electron density using Wannier centers allows us to calculate the IR-spectra using the time-correlation of the systems dipole moments. In addition, treating the electron dynamics as classical trajectories and the nuclei dynamics as well, allows us to express the systems IR-spectrum as the correlation between well defined entities, molecules or atoms and Wannier centers. In the supramolecular solvation cluster approach (SSC), we define the central gold molecule and its solvating water molecules as SSC, transforming their trajectories in the coordinate system spanned by it's average structure and analyzing the resulting IR-spectrum in terms of eigenvectors of the underlying correlation matrix. These eigenvectors can be interpreted in a similar way as the harmonic eigenvectors of much simpler calculations but assigning a broad spectral lineshape to a single vibrational motion.

For the SSC of the gold tetramer and a single solvation water molecule (where the solvating water is interacting electron density donating along the Au-O axis), shown in Figure 1, we can clearly separate the IR active vibrational stretch motion of the gold-cluster (gray lines at 134 cm^{-1} and 160 cm^{-1}) from the water-gold interaction at 214 cm^{-1} in blue and other water modes in the range of ambient-water librational motion. In the SSC coordinate system the light H-atoms pinning the librational motion in space are degraded in their position in respect to the heavy gold cluster which prevents a clear identification. However, if we remove the gold cluster from the SSC, looking at the SSC of the solvation water only, as shown in Figure 2.A, the eigenvectors of the librational motion can be assigned to L_1 (pink) and L_2 (red). Moreover, the solvation water-only spectrum allows us a direct comparison of the solvation water (sold black line) with SCW (dashed black), bulk water at the same conditions as the solvation water, and water at room temperatures (RTW, dotted). It reveals that the strongly polarized solvating water molecule shows a IR-spectrum in the THz-range much closer to ambient water than to water molecules in the same system not interacting with the gold cluster.

The spectrum of RTW is composed of the well-known H-bond stretch feature at around 200 cm^{-1} and the broad librational signal around 600 cm^{-1}, both hallmarks of the water's H-bond. In the gold-water SSC, the interaction feature at 213 cm^{-1} resulting in a very similar mode, while the libration, not present in SCW due to the lack of H-bonds, is restored by the gold-water interaction.

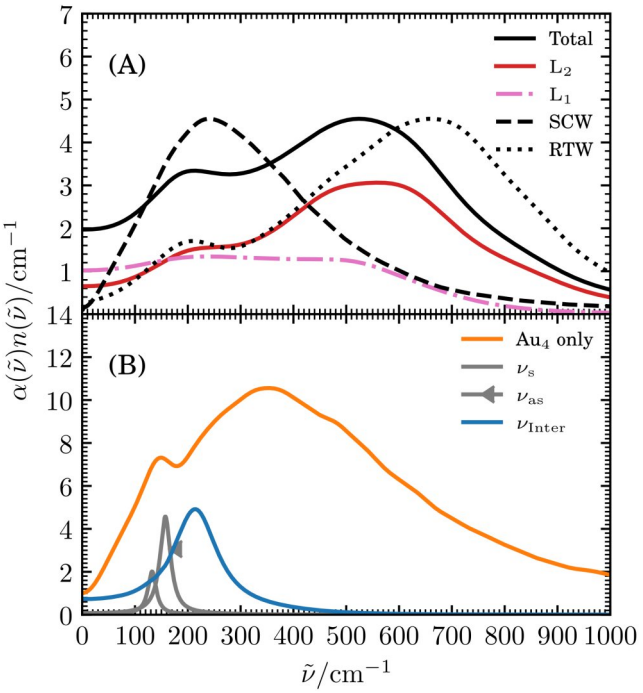

Figure .A: The IR-spectrum of the solvation water (solid black), dissected into the librations (L1, L2) and compared to the IR spectrum of water at room temperatures (black dotted, RTW) and supercritical conditions (black dashed, SCW).
Figure 2.B: The IR spectrum of the gold in SCW only (orange) and its IR active modes (gray lines). The interaction-band with water is shown as blue line.

In summary, the THz-spectrum shows that the solvation interaction of gold in SCW is more than a simple embedding, populating void space in the gas-like phase of SCW. Moreover, the THz-spectrum shows that the well-defined gold-water clusters are an example of water-metal interactions and suggest further investigation and the use of vibrational spectroscopy in the discussion of water-solute interaction in SCW, where the water-water H-bond is absent.

Ongoing Research / Outlook

The presented results are suggesting that interactions, surviving the harsh supercritical conditions, are standing out in the vibrational spectrum. The H-bond network of water is gradually fading along isochoric heating, as it was shown in the transition from RTW to SCW [5]. But how is the solvation interaction changing in between a H-bonded environment and supercritical conditions? In upcoming publications, we will discuss supercritical solvation effects along the transition of RTW to SCW using the well established gold systems. In addition, we will include more types of solute-water interactions using ions, which interact much stronger with water than gold.

References and Links

[1] W. Martin, J. Baross, D. Kelley, M. J. Russel, Nat. Rev. Microbiol. 6, 805–814 (2008); A. Koschinsky, D. Garbe-Schönberg, S. Sander, K. Schmidt, H. Gennerich, H. Strauss, Geology 36, 615 (2008).
[2] Y. Marcus, Supercritical Water, Wiley-Blackwell 2012.
[3] C.A. Heinrich, Miner. Depos. 39, 864--889 (2005)
[4] J. Noetzel, P. Schienbein, H. Forbert, D. Marx, J. Mol. Liq. 362, 119715 (2022).
[5] P. Schienbein, D. Marx: Phys. Rev. E 98, 022104 (2018); J. Phys. Chem. B. 122, 3318–3329 (2018); Angew. Chem. Int. Ed. 59, 18578–18585 (2020); Phys. Chem. Chem. Phys. 22, 10462–10479 (2020).
[6] https://www.cp2k.org/

Antiferroelectric ZrO_2 for piezoelectric devices

Research Institution
Modeling and Simulation Lab, Department of Applied Sciences and Mechatronics, Munich University of Applied Sciences

Principal Investigator
Alfred Kersch

Researcher
Richard Ganser, Luis Azevedo Antunes

Project Partner
–

Funding
DFG project ZEPPELIN

SuperMUC Project ID: pn73hi

Introduction

Reverse piezoelectricity describes the deformation of materials in the presence of an external electric field. To act piezoelectric, the material has to belong to the ferroelectrics, which have non-centrosymmetric crystal phases. The missing inversion symmetry allows that a stable net dipole moment or remanent polarization can form. Well known is PZT (Lead Zirconate Titanate), which is used in a broad range of actuator and sensor applications. But for most ferroelectrics the remanent polarization becomes lost in small structures and thin films. Therefore, the discovery of ferroelectric crystal phases in mixed HfO_2 and ZrO_2 thin film capacitors a few years ago was a breakthrough in micro- and nanoelectronics.

While the material as such was already present in form of a high-k gate and capacitor dielectric in transistor and DRAM, the question of silicon compatibility was proven. The additional nonvolatile functionality of the stabilized ferroelectric phase promises the next big step in chip technology as non-volatile gate or capacitor would combine the fast-operating speeds of transistor and DRAM with the non-volatility of solid-state drives. Whereas the optimum remanent polarization is obtained for mixture of Hf and Zr in equal proportions, with increasing Zr-content the material properties shift from ferroelectric to antiferroelectric, allowing for the exploitation of the giant piezoelectric effects during phase-transitions.

HfO_2, ZrO_2 and the mixtures $Hf_{1-x}ZrxO_2$ stabilize in different crystal phases depending on various effects such as temperature, crystal size and doping. The main crystal phases are:

- monoclinic, at low temperatures in bulk
- cubic, at high temperature in bulk
- orthorhombic ferroelectric, in thin films
- tetragonal dielectric, in thin films

Exploiting the relation between the tetragonal and the ferroelectric phases is a current research topic, showing that the transition can be achieved with an electric field.

The main issue from simulation side is the misfit between the energy landscape derived from 0 Kelvin ab initio calculations and the experimental observations. In these simulations the ferroelectric phase is, due to its lower symmetry, more stable than the tetragonal phase, contradicting the experimental observation of antiferroelectricity (AFE) at room temperature, which is based on the tetragonal phase being more stable. Because experiments show a loss of AFE at thicker film sizes, the suspected reason for this deviation are size and entropy effects. Therefore, the simulation methods needed to be adjusted.

Results and Methods

Since ab initio simulations are based on the solution of the Kohn-Sham-Equation for each structural calculation or time step in molecular dynamics, the system shows a N^3 scaling for calculation time versus the number of atoms. This limits the simulation time to a few picoseconds and system size to few 100 atoms.

An emerging approach to overcome this problem is to utilize machine learning. Using the computing resource of the SuperMUC-NG we were able to calculate the energies of more than 200,000 distorted crystal structures, which were subsequently used to train a neural net. This provides a mapping of the atomic coordinates to the energy and forces, which is the energy landscape. After the considerable investment in the creation of the training data for the neural network, the calculation of the energy of a structure is now 10^4 times faster and almost as accurate as with an ab initio calculation. Furthermore, because modern high-performance GPUs are tailored to the evaluation of neural networks, the calculation of crystal structures of several nanometers size is possible with up to 100,000 atoms, matching realistic grain sizes in thin films.

Experimental measurements discovered an interlayer region between the electrodes and the ZrO_2, which interpolates between the dielectric and the ferroelectric phase. The properties of this interlayer and its stoichiometry somewhat depend on the used electrode which is typically tungsten or titanium nitride. But in a first ap-

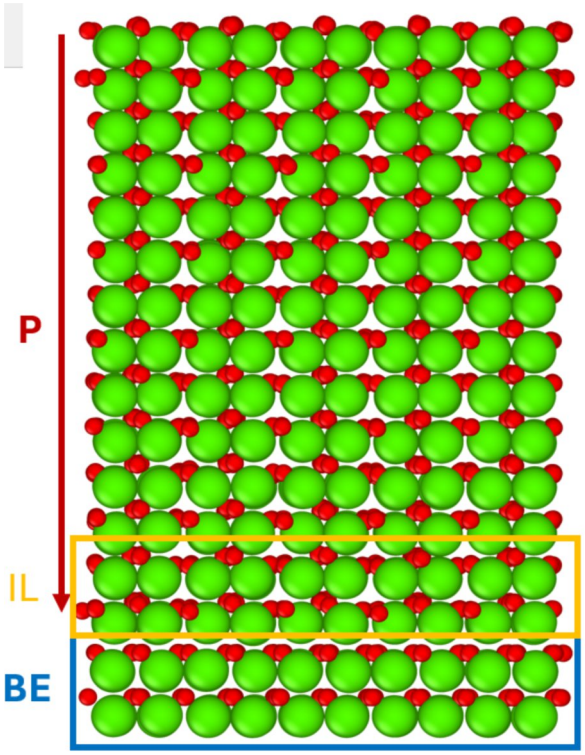

Figure 1: Simulation setup of a part of the ferroelectric grain near the bottom electrode. The interlayer (IL) is created with the bottom electrode (BE) boundary condition fixing the oxide in the polar direction.

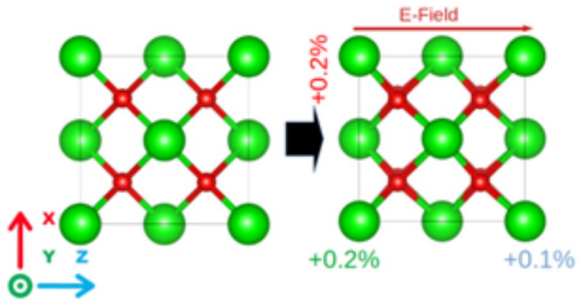

Figure 2: Application of an electric field to the tetragonal phase leads to a volume expansion of the high-dielectric material due to electrostriction. The structural data show a temporal and spatial average of the molecular dynamics results for the grain at room temperature.

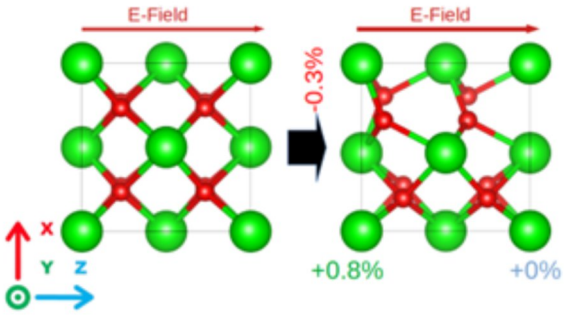

Figure 3: Increasing the electric field further leads to the phase-transition into the ferroelectric phase, accompanied by a giant piezoelectric coefficient due to increasing cell size. A further electric field increase leads to a small contraction due to the negative piezoelectric behavior of the ferroelectric phase, which has been predicted from ab initio calculations.

proach we modeled in molecular dynamics calculations the crystal boundary condition as a one-unit cell wide tetragonal phase layer with oxygen atoms fixed in the polarization direction. The effect on the oxygen atoms is a kind of oxygen stress extending into the ferroelectric, which leads to a shift in the Curie-temperature from 1,000 K for the periodic boundary condition infinite ferroelectric phase to 600 K for 10 nm thick films and below room temperature for thin films of 5nm thickness. To obtain tetragonal phase for very thin films is a nice match, but there remains a discrepancy for slightly thicker film, because polycrystalline ZrO_2 film retaining their antiferroelectricity up to more than 10nm thickness. By applying not only a BE-like boundary condition in the thickness direction but also in the transversal directions, modeling polycrystalline grain boundary instead of epitaxial grain boundary, the Curie-temperature for films thinner than 14 nm is shifted below room temperature. The detailed insight how to achieve tetragonal thin ZrO_2-films at room temperature is a progress and brings experiments and simulation closer together.

Piezo-force-microscopy (PFM) realizes the detailed experimental observation of structural change under an electric field applied on the local scale of a single grain. This allows the reconstruction of field induced phase-transitions and gives insight into the energy landscape. By applying electric fields to our simulation of initially tetragonal grains at room temperature, we were able to reproduce these measurements and draw an atomistic picture of the processes during phase-transitions.

At low electric fields the dielectric tetragonal phase start to deform due to electro-strictive effects. With increasing electric field, the energy-barrier between the tetragonal phase and the ferroelectric phase can be overcome, leading to phase-transition. This effect becomes visible in PFM due to the different unit cell sizes, with the ferroelectric phase being larger than the tetragonal phase, as a giant piezoelectric coefficient, and which is reproduced in the simulation. A further increase of the electric field shows the negative piezoelectric coefficient of the ferroelectric phase in experiment which is matched by the calculated value from molecular dynamics simulations very well.

Ongoing Research / Outlook

We aim to fully reconstruct the energy-landscape of ZrO_2 and to include the detailed effects of different BE boundary conditions representative for electrodes. We will extend the analysis to $Hf_{1-x}ZrxO_2$ mixtures and to slightly oxygen deficit mixtures which are relevant in microelectronics. Further miniaturization of $Hf_{1-x}ZrxO_2$-based devices depends on the final length scales that still allow the ferroelectric functionality to be switched on and off. We find that at a few nm the ferroelectric phase becomes more difficult to stabilize. We will investigate in further research how this can be counteracted.

References and Links

[1] M. Falkowski et al., Appl. Phys. Lett. 2021, https://doi.org/10.1063/5.0029610
[2] R. Ganser, Phys. Rev. Applied 2022 https://doi.org/10.1103/PhysRevApplied.18.054066
[3] P. Lomenzo et al., Adv. Funct. Mat. 2023, https://doi.org/10.1002/adfm.202303636

How do Actinide and Iron Ions Interact with Cement?

Research Institution
Fachgebiet Theoretische Chemie, School of Natural Sciences, TU München

Principal Investigator
S. Krüger

Researcher
I.Chiorescu, A. Kremleva

Project Partner
–

Funding
BMUV 02E11860E

SuperMUC Project ID: pr94je

Introduction

Radioactive waste from nuclear power generation and other industrial and scientific sources needs to be stored safely for a long time due to its radiation and toxicity. Deep underground geological repositories are intended and planned for that purpose in granite, salt, or clay rock formations by various countries, among them Germany. In these underground repositories, cement and concrete will be used as construction material and as engineered barrier against distribution of the waste. Cement is also used to solidify liquid radioactive waste. Thus, the chemical interaction of various elements of radioactive waste and of its containment materials with cement needs to be investigated to determine the retention potential of cement and the risk of migration to the environment. These aspects are of interest for environmental science even beyond the specific application of a geological repository. Uranium is the predominant element in highly radioactive waste because of its use as fuel in nuclear power stations. At environmental conditions it forms the uranyl ion, UO_2^{2+}, in water with U in the oxidation state VI (U(VI)). In a geological repository, anoxic reducing conditions are to be expected in the long term. Under these conditions iron of steel containments and construction materials corrodes to Fe^{2+} ions and hydrogen gas is formed. This process supports reducing conditions. Thus, uranium may be reduced to U^{4+} ions, U(IV), which are less soluble and mobile than U(VI). Cement consists to a considerable extent of calcium silicate hydrate (CSH) phases, which are the main product of cement hydration and essential for its solidification. These phases are layered minerals, showing a central calcium oxide layer, which carries on both sides chains of silica tetrahedra, SiO_4 (Fig. 1). Sorption of uranium and other metal ions may take place at the surface of CSH (adsorption) or between the layers (absorption). Some metal ions may even replace calcium ions of the central oxide layer (incorporation) (Fig. 1). Experiments have shown that uranium and other actinides preferentially sorb at CSH in cement. However, the sorption mode as well as the nature of the sorption complexes at the atomic scale are largely uncertain.

For Ca to Si ratios (C/S) below 1.5, CSH phases correspond to the mineral tobermorite. We have chosen tobermorite with a layer distance of 14 Å as a model for aged CSH to inspect computationally the sorption of uranium, thorium, as well as of iron (Fe^{2+}), which may be a competing metal ion for the actinide sorption under reducing conditions.

Results and Methods

We use quantum mechanical density functional theory to calculate the electronic structure (band structure) of sorbed metal ions in tobermorite. The tobermorite model is treated as an ordered, periodical crystal and the valence electron band structure is determined by the projector augmented wave method (PAW). Valence electron bands are represented by a plane wave basis set and potentials are applied to avoid an explicit treatment of the large number of core electrons of heavy elements like uranium. This method is implemented in the parallel program VASP via MPI [1]. The necessary relativistic treatment of core electrons is commonly restricted to the core potentials. For open electronic shells, spin dependent bands have to be taken into account, which doubles the computational effort. Typical models of a tobermorite crystal comprise about 400 atoms. Surfaces are represented as single, periodically repeated tobermorite slabs with a separating vacuum layer between them. 2-3 water layers are added to the surfaces to model the contact to the surrounding aqueous solution (Fig. 1). These models include about 250 atoms. Straight-forward geometry optimization of structures of this size is precluded due to a large number of local minima on the potential energy surface, mostly caused by weakly interacting water in the interlayer and at the surface. To achieve representative structures, we applied a combination of first-principles molecular dynamics at room temperature to equilibrate the soft degrees of freedom and structure optimization [2]. Equilibration runs of 1 ps followed by geometry optimizations are repeated until energies vary less than about 10 kJ/mol. A typical structure determination of this type needs 10-20 repetitions of these two steps. For a bulk system up to three days on SuperMUC-NG using 144 to 240 cores are spent per repetition. The need for high performance

computing arises not only due to real-time requirements, but more because up to 10 structures (sorption sites) have to be compared for each metal and each sorption mode. Resulting interatomic bond lengths and other geometry parameters are compared to experimental measurements exploiting the fine structure of X-ray spectra (EXAFS). Energies of creation of sorption complexes offer an alternative way for determining preferred structures.

For U(VI) we compared sorption in the interlayer and at the (001) surface of tobermorite [2]. Although sorption energies are comparable at the surface and in the interlayer, structural characteristics of complexes in the interlayer are more similar to measured results. None of the sorption complexes shows all features of the measurements, thus several of them are necessary to rationalize the EXAFS results [2]. Incorporation into the CaO layer has been excluded due to the size of the UO_2^{2+} ion.

For U(IV) we inspected sorption at the (001) surface, in the interlayer and also incorporation into the CaO layer of tobermorite [3]. In agreement with the higher charge of the U(IV) ion, stronger sorption compared to U(VI) has been determined in agreement with experiment. A previous experiment for Np(IV) and a current one for Pu(IV) allow a comparison of structures, as these ions of same charge are only a little smaller (up to 3 %) than U(IV) [3]. Average U-O bond lengths are somewhat overestimated compared to experiment and allow no differentiation between the sorption modes. Distances U-Si and U-Ca and their number are only in good agreement with experiment when U(IV) is incorporated into the CaO layer. This sorption mode is also supported by energy considerations.

To verify the assumption that all actinide elements in the oxidation state IV should follow the same sorption mechanism, we currently inspect the sorption of Th(IV) in comparison to U(IV). The Th^{4+} ion is by 5% larger than U(IV). Indeed, the same sorption complexes have been obtained at most sites and also incorporation into the CaO layer is favored for Th(IV). Aqua ions of Th(IV) sorb as expected weaker than U(IV). The opposite is found for most sites for tetrahydroxo complexes, which are present in solution at the basic conditions prevailing in cement.

The sorption of Fe(II) is inspected for tobermorite models with C/S ratios similar to experiment. The only available EXAFS experiment claims for Fe(II) a sixfold oxygen coordination and distances of Fe to Ca atoms which are shorter than Fe-Si distances [4]. As an effect of Fe sorption, a widening of the interlayer distance of CSH is observed. These results are interpreted to show Fe(II) sorption at surfaces and in the interlayer [4]. For C/S = 1 we obtained in our calculations Fe(II) hydroxo complexes as absorbates in the interlayer. Also two sites in the CaO layer can be occupied by Fe(II). Energies of adsorption yield Fe(II) in the CaO layer as the most stable species, but some of the sorption complexes in the interlayer are nearly as stable. The most stable incorporated species shows 5 neighboring oxygens with an average bond length exceeding the experimental result. Average Fe-Ca distances were determined to be longer than Fe-Si

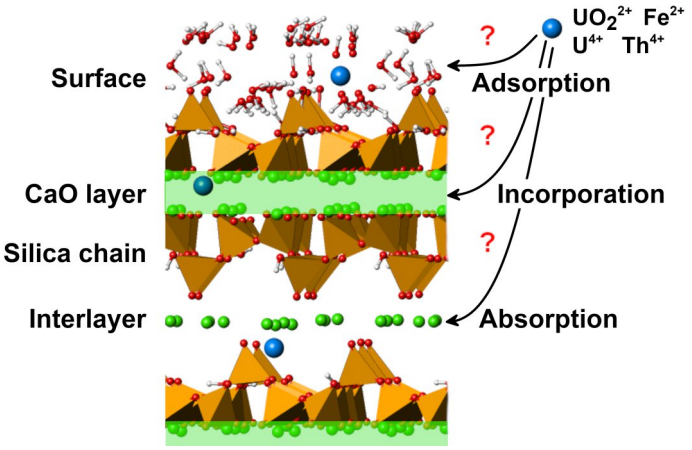

Figure 1: Schematic structure of a tobermorite model for calcium silicate hydrate (CSH). The CaO layer is indicated as green bar and silica rows are shown in orange. Possible sorption modes for actinides and iron ions are indicated by arrows. Blue balls represent sorbed ions at examplary sites at the surface, in the CaO layer, and in the interlayer of CSH. Green balls indicate Ca. Interlayer water has been omitted for clarity.

for all C/S values, in agreement with experimental and computational results for actinides [2,3], but in contrast to experiment for Fe(II) [4]. For absorption into the interlayer the number of oxygens around Fe(II) was determined to 4-5, which is smaller than in experiment, but the average Fe-O bond lengths are in agreement. Again, distances Fe-Si are for all complexes shorter than Fe-Ca distances. Even considering several Fe(II) ions in the interlayer did not yield the experimentally claimed widening of the inter-layer distance. Thus, comparison of geometry characteristics shows that the experimental results cannot be explained by the presence of sorbed Fe(II) ions in 14 Å tobermorite. This is remarkable, as we obtained for actinides good qualitative agreement with EXAFS results [2,3]. Other types of sorption complexes, a different CSH phase, or the formation of Fe-Ca-(O,OH) particles in experiment might be invoked to solve this puzzle. Comparison of sorption energies of Fe(II) and U(VI) ions, which have the same charge, for different hydroxo species in solution as reactants shows that under basic conditions in cement Fe(II) should sorb better. This indicates that Fe(II) released by steel corrosion might compete with uranium for favorable sorption sites and thus may reduce the retarding effect of cement.

Ongoing Research / Outlook

In co-operation with experimental partners we will pursue the open question of the nature of Fe(II) sorbates in CSH by considering Fe(II) adsorption at surfaces of CSH and absorption into tobermorites with narrower interlayer. Further simulations will be devoted to the important question whether U(VI) can be reduced to U(IV) by Fe(II). The lesser mobility of U(IV) compared to U(VI) makes the reduction process particularly interesting, as it hinders the migration of uranium in a geological repository for radioactive waste.

References and Links

[1] www.vasp.at
[2] A. Kremleva, S. Krüger, N. Rösch, Appl. Geochem. 113 (2020) 104463.
[3] I. Chiorescu, A. Kremleva, S. Krüger, Minerals 12 (2022) 1541.
[4] A. Mancini, E. Wieland, G. Geng, B. Lothenbach, B. Wehrli, R. Dähn, J. Colloid Interface Sci. 588 (2021) 692-704.

Ab-initio modelling of paramagnetic sites in zeolites and metal-organic frameworks

Research Institution
[1]Felix Bloch Institute for Solid State Physics, Universität Leipzig

Principal Investigator
Andreas Pöppl[1], Paolo Cleto Bruzzese[1]

Researcher
Lorenzo Donà[2], Bartolomeo Civalleri[2], Prof. Mario Chiesa[2]

Project Partner
[2]Department of Chemistry and NIS Centre of Excellence, University of Turin

Funding
Marie Skłodowska-Curie Grant (agreement no. 813209)

SuperMUC Project ID: pn73va (Gauss Large-Scale project)

Introduction

The implementation of microporous materials with paramagnetic (open-shell) sites is a topic of great interest in several fields of technology, especially in catalysis and quantum sensing. Zeolites and metal-organic frameworks (MOFs) represent ideal cases where single transition metal atoms or ions (TMIs) with one or more unpaired electrons are docked at specific and well-defined sites. Contrary to molecular complexes, the porous network is able to accommodate a variety of TMIs within a three-dimensional array in an ordered fashion. Elucidating the role of the electron spin in such materials is presently one of the most challenging endeavours in the research field. Indeed, the notorious complex nature of zeolites and MOFs systems often prevents mapping their active sites with atomic-scale precision, which is crucial for the enhancements and improvements of their unique properties. In this project we model the paramagnetic active centers in zeolites and MOFs materials providing unique details on their electronic, geometric and vibrational structures. The simulations involved hundreds of atoms – a huge amount in the world of accurate Density Functional Theory (DFT). Due to the large size of the investigated systems and the high level of accuracy of calculations, the feasibility of simulations was possible only thanks to the capabilities of SuperMUC-NG. The results are expected to give an important contribution in the rationalization of the relation between the structure/function properties of the paramagnetic sites [1]. Achievement of this goal would allow improving the design of these materials, as well as reducing their health and environmental impact, with a potential positive outcome also from an industrial perspective.

Results and Methods

The results of the project are divided into two sub-projects: one related to zeolite systems and the second one to MOFs. In both the cases, a full periodic approach based on the exploitation of the CRYSTAL code was adopted. CRYSTAL allows to compute the total energy and related properties of solids (e.g. structural, electronic, vibrational, mechanical, magnetic, dielectric, etc.) at the Hartree-Fock and DFT level. The massive parallel (MPP) version of CRYSTAL exploits MPI libraries and it has been specifically tailored to achieve an excellent scalability in terms of both speed-up and memory usage on SuperMUC-NG. An excellent scalability in terms of both speed-up and memory usage has been obtained up to more than 30,000 processors [2]. In the first sub-project we modelled the location and vibrational structures of copper-, vandayl-, zinc- and cadmium-doped zeolites. For each metal site, several aluminium (Al) distributions were taken into account. By calculating the relative electronic energy of several structures, we determined the most stable location of the TMIs in the zeolitic framework (Figure 1) [3]. The computation of the vibrational frequencies for the four different metal-doped systems unveiled correlations between the relaxation processes of the electron spin and the vibration of nearby nuclei. The study of such phenomenological phenomenon (spin-phonon interaction) is particularly important for the development of quantum computing. The screening of the Al distributions and the consequent calculation of the phonon structures required the running of multiple jobs each using 1,440 cores. Thanks to the large processor counts available in SuperMUC-NG, we performed simultaneously energy plus gradient (E+G) calculations. This allowed to shorten considerably the computational times, especially for the case of frequencies computations where the number of E+G per run is equal to 1,734. In the second part of the project, we focused on modelling the active sites of MOFs containing paramagnetic centers. The geometric and electronic structures of novel nickel-based-MOFs determined by calculations show an outstanding agreement with experimental findings proving that the predicted structures can be considered a reliable model of crystal structures. Moreover, the calculations correctly predicted the change of the structure in presence of solvent. The binding sites of small molecules frequently involved in catalytic reactions (e.g. NO, N_2O and CO_2)

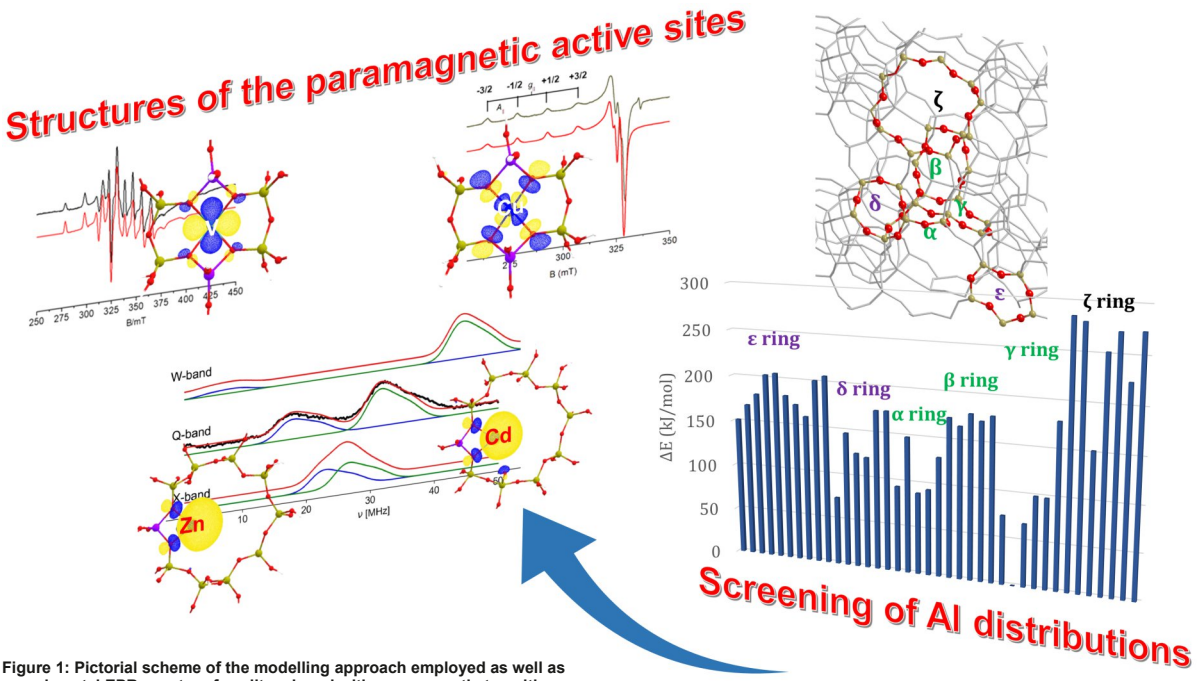

Figure 1: Pictorial scheme of the modelling approach employed as well as experimental EPR spectra of zeolites doped with paramagnetic transition metal ions.

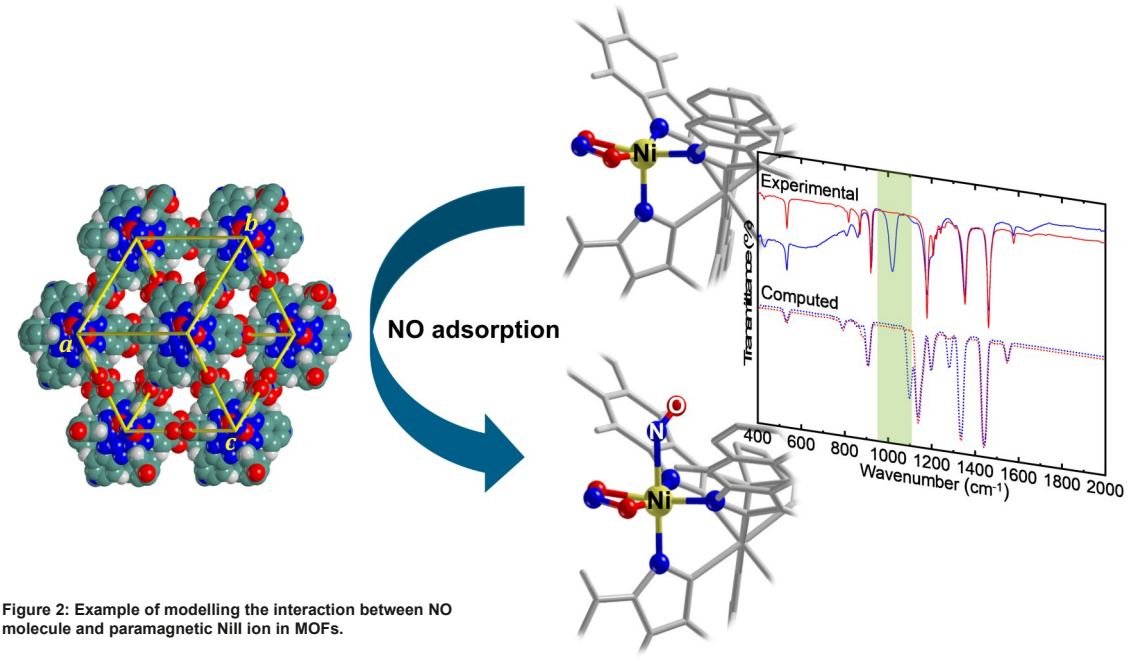

Figure 2: Example of modelling the interaction between NO molecule and paramagnetic NiII ion in MOFs.

were modelled to understand the nature of the chemical bonding between the paramagnetic TMI and the adsorbate (Figure 2). The predicted structures revealed to be fundamental for interpreting the complex spectroscopic data and understanding the relation between structure and reactivity [4]. Due to the larger number of atoms in MOFs unit cells, the typical number of cores used in this case was 2,880.

Ongoing Research / Outlook

The team is currently investigating the impact of the phonon structure on the relaxation properties of zeolites and MOFs containing paramagnetic metal ions. Systematic studies will enable the assessment of the role of the solid support as well as of the spin-orbit effects for the development of solid-state qubit platforms. This might open up the road to the development of novel materials in the quantum computing field.

References and Links

[1] Bruzzese, P. C. st al. J. Am. Chem. Soc., 144, 13079−13083 (2022).
[2] Erba, A. et al. J. Chem. Theory Comput. 13 (10), 5019–5027 (2017).
[3] Liao, Y., Bruzzese, P. C. et al. J. Magn. Reson., 16–17, 100101 (2023).
[4] Thangavel, K., Bruzzese, P. C. et al. Phys. Chem. Chem. Phys., 25, 15702 (2023).
[5] Salvadori, E., Bruzzese, P. C. et al. Acc. Chem. Res., 55(24), 3706–3715 (2022).

Uncovering the physics of heavy fermions with spin chains on metallic surfaces

Research Institution
[1]Institute for Theoretical Physics, University of Würzburg

Principal Investigator
Fakher F. Assaad[1]

Researcher
B. Danu[1], Z. Liu[1,2], B. Frank[2], M. Vojta[2], L. Janssen[2], M. Raczkowski[1], J. Schwab[1], F. Mila[3], T. Grover[4]

Project Partner
[2]TU-Dresden
[3]École Polytechnique Fédérale de Lausanne
[4]University of California at San Diego

Funding
ct.qmat cluster of excellence, KONWIHR, DFG

SuperMUC Project ID: pn73xu (Gauss Large-Scale project)

Introduction

The idea of using nature to investigate the physics of correlated quantum systems dates back to R. Feynman. Quantum simulators refer to controlled quantum systems that can be used to simulate models of correlated electrons. The difficulty to control quantum systems with respect to external perturbations and decoherence, requires benchmarking of quantum simulators with existing algorithms on classical computers. In this short report, we will highlight aspects of our research in this domain. The models that we simulate are so called dimensional mismatch Kondo systems. Consider a d-dimensional array of spin-½ degrees of freedom, embedded in a d + 1 dimensional metallic host. For d = 1 this would, for example, correspond to a spin chain deposited on a metallic surface (see later). The questions we address are fundamental, and pertain to the physics of so called heavy fermion materials. i) In the presence of magnetic impurities how many charge carriers should one count? $N_{Metal} + N_{Spin}$, corresponding to a heavy fermion phase or only N_{Metal} corresponding to a Kondo breakdown phase. Here N_{Metal} counts the number of electrons in the conducting material, and N_{spin} the number of spin-½ degrees of freedom ii) What is the dynamics of the spin system once it is coupled to the metallic host? iii) The model will contain various phases and we are interested in understanding the nature of the corresponding zero temperature phase transitions.

For a special set of parameters, dimensional mismatch Kondo systems can be investigated numerically with the so called auxiliary field quantum Monte Carlo method. The Algorithms for Lattice Fermions, ALF [1,6], package that we are maintaining and further developing provides a very versatile openmp/mpi implementation of this algorithm adequate for massively parallel architectures. The generality of the formulation allows for the simulations of e.g. correlated electrons coupled to phonons, of mixed spin-electron systems, as well as (frustrated) spin systems coupled to phonon degrees of freedom.

Quantum simulators of dimensional mismatch Kondo systems can be realized by manipulating magnetic adatoms on surfaces. In particular, consider the metallic surface of copper, Cu(100), covered by a so called decoupling layer as realized by an insulating film of Cu_2N. Using scanning tunnelling microscopy the authors of Ref. [2] are able to construct and investigate the physics chains of Co adatoms on this surface. Using classical computers [3] and comparing model calculations to the experimental results, we can show that the authors of Ref. [2] have actually created a quantum simulator to investigate the physics of heavy fermions. Cobalt is a spin-3/2 atom. Since on a surface inversion symmetry is broken, spin-orbit coupling has to be taken into account in the modelling. It turns out that the sign of the easy axis anisotropy favors a spin doublet with z-component of spin ½. Importantly, time reversal symmetry is not broken thereby ensuring Kramers degeneracy. The cobalt spin ½ degrees of freedom can interact with each other via a Heisenberg interaction of magnitude J_H and with the spin of the electrons of the metallic copper surface via a so called Kondo coupling of magnitude J_K. The fact that the spin ½ degree of freedom couples to both the conduction electrons and to the other spin ½ degrees of freedom provides the desired competition of entanglement that will ultimately lead to the realization of phases and phase transitions of heavy fermion materials.

Results and Methods

In Fig. 1 we depict the dimensional mismatch systems that we have considered in Ref. [3]. As mentioned in the introduction, this work is strongly inspired from Ref. [2], and hopefully paves the way to quantum simulations. Depending upon the nature of the host metal, we have learned that various quantum critical points (QCP) can occur. While at strong coupling we observe a heavy fermion metal, the weak coupling phase depends on the metallic host. For the semi-metal, the spin and conduction electrons decouple (Kondo breakdown), while for the metallic host, we observe so called dissipation induced long ranged magnetic ordering [5].

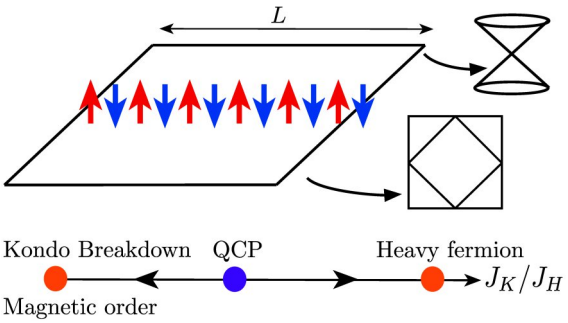

Figure 1: Spin chain on a semi-metallic or metallic surface. For large values of JK/JH, we observe in both cases a heavy fermion phase. In contrast at weak coupling the results depend upon the metallic host. Kondo breakdown occurs for a semi-metal whereas a magnetically ordered heavy fermion state emerges for the metal surface. See Ref. [3] for a detailed account of this research.

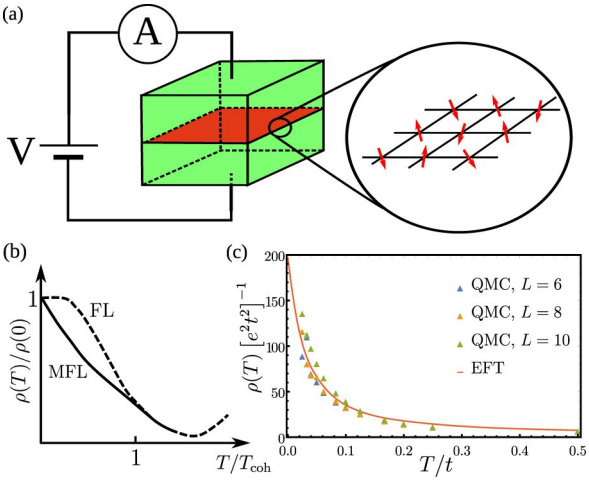

Figure 2: Magnetic layer embedded in a three dimensional metallic host. At the quantum critical point we observe non-fermi liquid behavior signalized by a linear in T saturation of the resistivity. See Ref. [3] for a detailed account of this research.

In Fig. 2 we show our results for a three dimensional metallic host. In this case, we can analyze analytically and numerically the physics of the quantum critical point, by investigating the resistivity across the magnetic layer.

As briefly mentioned in the introduction, the simulations carried out here are based on the ALF-library [1] implementation of the auxiliary field quantum Monte Carlo method. Generically, simulations of quantum systems scale exponentially with systems size. This scaling reflects the dimension of the Hilbert space and is a consequence of the so called negative sign problem. This problem can be avoided in some special cases, and the computational cost turns out to be proportional to the cubed of the volume times the inverse temperature. Here we do not consider issues related to long autocorrelation times. The systems considered here do not suffer from the negative sign problem. They are non-trivial, model realistic systems, and host a number of quantum phases of matter and quantum phase transitions that we will continue to study. Access to supercomputing resources is imperative since the competition between energy scales that drive systems through quantum critical points is inherently of non-perturbative origin. Numerical simulations hence play an important role in guiding analytical work.

Ongoing Research / Outlook

Our ongoing research has many facets and the above report only covers aspects of our work. We have developed novel formulations to tackle the electron-phonon problem within the realm of the auxiliary field quantum Monte Carlo algorithm. We are presently using these algorithms to investigate the physics of the Su-Schrieffer-Heeger model. We have equally concentrated on frustrated spin systems pertaining to Kitaev materials such as $RuCl_3$. Here we confront the negative sign problem and remarkably, with our simulations we can reach temperatures down to 40 K, corresponding to experimentally relevant energy scales. We are equally working on Dirac systems that host a number of exotic quantum phase transitions. Using notions of entanglement one of our aims is to probe the nature of these quantum phase transition.

All our algorithmic development is included and documented in the ALF-library. Our ultimate aim is to render simulations of correlated fermion systems user friendly. At present, our code complies to Fortran 2008 standards. We are actively working on a python interface to the ALF code, pyALF [7]. At present pyALF can only run models that have been hard coded in Fortran. One of our aims, is to alleviate this constraint and provide the possibility to specify the model Hamiltonian directly in python. This will allow to combine ALF with the wide range of libraries available in python. Over the years we have greatly profited from KONWIRH for optimization and algorithmic development of the package.

References and Links

[1] F. F. Assaad, M. Bercx, F. Goth, A. Götz, J. S. Hofmann, E. Huffman, Z. Liu, F. Parisen Toldin, J. S. E. Portela, and J. Schwab, SciPost Phys. Codebases (2022), 1.
[2] R. Toskovic, R. van den Berg, A. Spinelli, I. S. Eliens, B. van den Toorn, B. Bryant, J. S. Caux, and A. F. Otte, Nature Physics 12 (2016), 656.
[3] B. Danu, F. F. Assaad, F. Mila, T. Grover and M. Vojta. Phys. Rev. Lett. 123 (2019), 176601. Phys. Rev. Lett. 125 (2020), 206602. Phys. Rev. B 106 (2022), L161103.
[4] Z. Liu, B. Frank, L. Janssen, M. Vojta, and F. F. Assaad, Phys. Rev. B 107 (2023), 165104. Phys. Rev. B 108 (2023), L100405.
[5] M. Weber, D. J. Luitz, and F. F. Assaad, Phys. Rev. Lett. 129 (2022), 056402.
[6] https://gitpages.physik.uni-wuerzburg.de/ALF/ALF_Webpage
[7] https://git.physik.uni-wuerzburg.de/ALF/pyALF

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

Research Institution
Department of Physics, University of Duisburg-Essen

Principal Investigator
Rossitza Pentcheva

Researcher
Achim Füngerlings, Yuman Peng, Okan Köksal, Manish Verma

Project Partner
–

SuperMUC Project ID: pr87ro

Introduction

Besides novel interface-induced phenomena in oxide superlattices, oxygen defects play an important role as they are ubiquitous in transition metal oxides and determine their functional properties. Moreover, under operation conditions, e.g. under applied voltage in a solvent, many oxides undergo a surface transformation. Understanding how these phenomena affect the electronic, magnetic and catalytic properties of the material is essential for technological applications. During the reporting period several of these aspects were explored within project pr87ro.Within the DFG-funded CRC/TRR 247 (project B04), we gained insight into the role of surface orientation [2,3] and surface transformation [3] on the oxygen evolution reaction (OER), the bottleneck half-reaction of electrochemical water splitting, which allows the storage of green energy in the form of H_2-bonds. Within CRC/TRR 80 (project G3 and G8), we explored how the electronic, magnetic and thermoelectric properties of $LaNiO_3$ and its superlattices can be tuned by oxygen vacancies [4]. Last but not least, density functional theory (DFT) results with a Hubbard *U* term and spin-orbit coupling (SOC) reveal intriguing electronic and magnetic reconstructions in artificial superlattices containing the so-called dice lattice [5].

Results and Methods

Facet-Dependent Intrinsic Activity of Single Co_3O_4 Nanoparticles for Oxygen Evolution Reaction
Motivated by nanoimpact experiments that found higher OER activity of cube- vs. sphere-shaped Co_3O_4 nanoparticles [2], we investigated the OER activity of (001)-and (111)-oriented Co_3O_4 surfaces. The DFT+*U* results not only confirm the higher OER activity of the (001) orientation but trace it back to the highly active octahedrally coordinated Co site, whereas for (111), the reaction site is tetrahedrally coordinated Co. Moreover, a different potential-determining step was identified for the two facets: *OH for (001) vs. *O for (111) [2].

Crystal-facet-dependent surface transformation dictates the oxygen evolution reaction activity in lanthanum nickelate
A combined experimental and theoretical study revealed that the three most prominent facets of the perovskite $LaNiO_3$ undergo a transformation into NiOO(H) before the onset of OER. Despite the similarity in the transformed layer, the catalytic activity is substantially different for the different orientations. Our DFT+*U* simulations showed that the higher OER activity of the transformed (111)-facet can be rationalized by the thicker NiOO(H) layer and the structural compatibility with the underlying hexagonal $LaNiO_3$(111). In contrast, the transformed (001) and (110) facets undergo a substantial lateral distortion due to the symmetry and structural mismatch to the underlying perovskite. The coupling to the substrate also influences the ratio of high to low spin Ni ions and by this the OER activity [3].

Tuning of the carrier localization, magnetic and thermoelectric properties in $(LaNiO_{3-\delta})_1/(LaAlO_3)_1$(001) superlattices by oxygen vacancies
In this project using DFT+*U* calculations and Boltzmann transport theory within the constant relaxation time approximation, we explored the effect of oxygen vacancies on the electronic, magnetic, and thermoelectric properties in ultrathin $(LaNiO_{3-\delta})_1/(LaAlO_3)_1$(001) superlattices (SLs). For pristine SLs, we find the stabilization of an antiferromagnetic charge-disproportionated (AFM-CD) $(d^8L^2)_{S=0}$ $(d^8)_{S=1}$ phase, irrespective of strain. At $\delta = 0.125$ and 0.25, the localization of electrons released from oxygen defects in the NiO_2 plane triggers a charge-disproportionation, resulting in a ferrimagnetic insulating behavior. At $\delta = 0.5$, an insulating phase emerges with alternating stripes of Ni^{2+} (high-spin) in octahedral and Ni^{2+} (low-spin) in planar coordination along the [110] direction (S-AFM), irrespective of strain. This results in a robust *n*-type in-plane power factor of 24 μW/K2 cm^{-1} at a_{STO} and 14 μW/K2 cm^{-1} at a_{LSAO} at 300 K (assuming relaxation time $\tau = 4$ fs), shown in Fig. 3. Additionally, the pristine and $\delta = 0.5$ SLs were found to be dynamically stable [4].

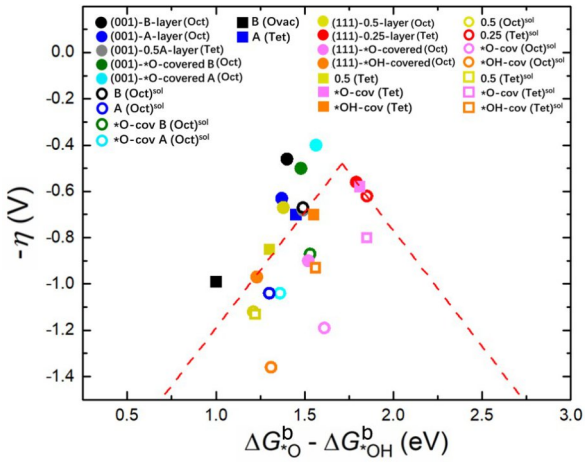

Figure 2: Volcano plot of the overpotential for both Co_3O_4(001) and (111) as a function of the binding energy difference between *O and *OH intermediates. The dashed line is the ideal volcano from the scaling relationship of *OOH and *OH. Modified from Ref. [2].

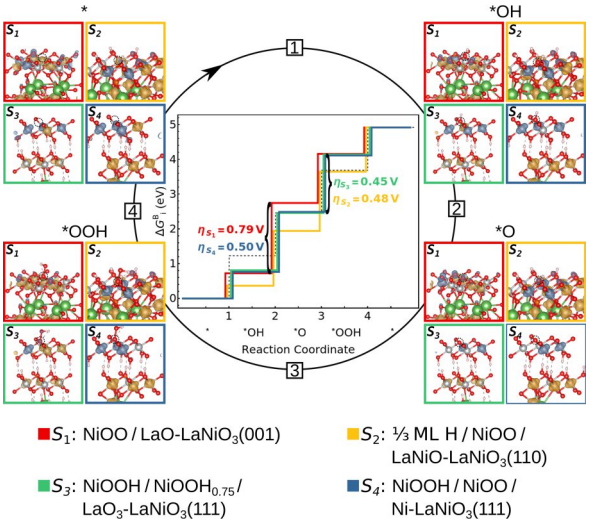

Figure 1: Reaction-free energies for the most favorable reaction site at transformed $LaNiO_3$(001), (110) and (111) surfaces. Side views of the intermediates, including spin density, are framed in the corresponding colors. Modified from Ref. [3].

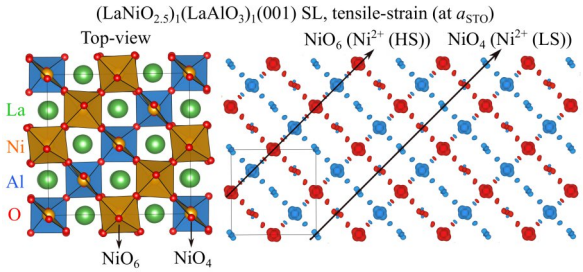

Figure 3: (a) Optimized $(LaNiO_{2.5})_1/(LaAlO_3)_1$(001) SL at a_{STO} showing alternating stripes Ni^{2+} in octahedral/planar configuration, (b) corresponding spin density with stripe-like antiferromagnetic ordering of Ni^{2+}, high spin. Blue and red colors denote to positive and negative spin densities, respectively. Adapted from Ref. [4].

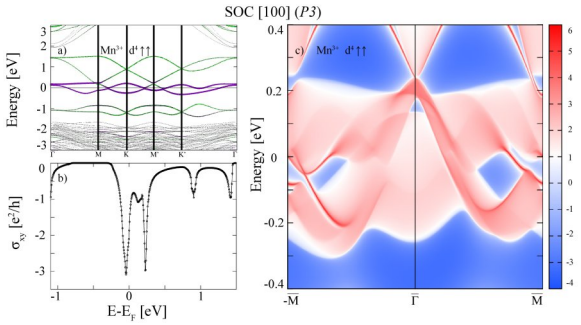

Figure 4: (a) GGA+*U*+SOC band structure for *X* = Mn with *P3* symmetry and magnetization direction along [100]; b) anomalous Hall conductivity σ_{xy} as a function of the chemical potential. (c) Edge states along the (100) surfaces: bulk band structure denoted by pink-white colors, edge states by solid red lines. Modified from Ref. [5].

High Chern numbers in a perovskite-derived dice lattice $(LaXO_3)_3/(LaAlO_3)_3$(111) with X = Ti, Mn and Co

Within this project we investigated the electronic and topological properties of $(LaXO_3)_3/(LaAlO_3)_3$(111) which contain a stack of three triangular lattices, the so-called dice lattice. For X = Mn with P3 symmetry and ferromagnetic coupling, DFT+*U* simulations including SOC with magnetization direction along [100] indicate the emergence of interface-dominated bands with high Chern numbers ±3 close to the Fermi level (see anomalous Hall conductivity in Fig. 4b) [5].

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.

Within CRC TRR 247, the investigation of the effect of surface transformation on the OER is currently extended to spinels. Furthermore, within IMPRS SUSMET we explore alternative ways to reduce iron oxides using hydrogen, a promising route toward green steel production.

References and Links

[1] https://www.uni-due.de/physik/pentcheva/
[2] Z. Liu, H. M. A. Amin, Y. Peng, M. Corva, R. Pentcheva, K. Tschulik, Adv. Funct. Mater. 33, 2210845 (2023).
[3] A. Füngerlings, M. Wohlgemuth, D. Antipin, E. van der Minne, E. M. Kiens, J. Villalobos, M. Risch, F. Gunkel, R. Pentcheva, C. Baeumer, Nat. Commun. 14, 8284 (2023).
[4] M. Verma, R. Pentcheva, Phys. Rev. Research 6, 013189 (2024).
[5] O. Köksal, L. L. Li, R. Pentcheva, Sci. Rep. 13, 10615 (2023).

AI revolution of materials discovery

Research Institution
Ruhr University Bochum

Principal Investigator
Miguel Marques

Researcher
Jonathan Schmidt

Project Partner
–

SuperMUC Project ID: pn25co

Introduction

Over the last few decades, ab-initio methods such as density-functional theory have become sufficiently accurate to allow for the prediction of many properties of new crystal structures. However, these predictions come at a significant cost, and due to the vastness of the space of possible materials, theoretical material discovery remains one of the most challenging questions in materials science. Machine learning methods, trained on existing databases of ab-initio calculations, have the potential to massively accelerate the process of theoretical materials discovery. One of the most important properties targeted is thermodynamic stability, which is used as a proxy to estimate the probability that a given compound can be synthesized.

Recently, we developed a neural network architecture specialized for predicting new stable materials, known as crystal graph attention networks. We used this network to predict the stability of around 15 billion compounds. The most promising compounds were then investigated with ab-initio theory in SuperMUC-NG, and added to our materials database alexandria [1].

Alexandria has grown over the past years to become the largest freely available database of computed properties of inorganic solids. It now includes entries for more than 4.4 million compounds, making it the largest (academic) database in the world. For comparison, AFLOW has around 3.5 million entries, the open quantum materials database has 1 million, and the materials project 155 thousand. Furthermore, alexandria combines this with a database of more than 140 thousand two-dimensional materials, 13 thousand one-dimensional materials, and 430 thousand entries calculated with a higher-accuracy method. Finally, all this information is available with a Creative Commons Attribution 4.0 License, that allows the sharing and adaptation of the data, and does not restrict any kind of academic or commercial use.

The effectiveness of alexandria relies on the abundance, diversity, and richness of crystal information, emphasizing the importance of size, structural diversity, and inclusion of computed properties for training robust machine-learning models in the field of materials science.

Results and Methods

In our workflow, we use our crystal-graph network model to predict the stability of a compound. This is measured by the energy distance to the so-called convex hull of thermodynamic stability, a hypersurface in the 100-dimensional materials space whose vertices are the stable compounds in the thermodynamic sense. A large distance indicates that the compound is highly unstable and is therefore of no practical use. A small distance, or if the compound is on the hull (zero distance) indicates stability towards decomposition, and a higher probability that the compounds can be synthesized experimentally. Factoring in errors in the theoretical framework and other physical effects such as temperature, defects, etc. one commonly uses a threshold of around 0.05-0.10 eV to distinguish the region of interesting compounds.

In the first step of the workflow, we select a known crystal structure that defines a “prototype”. This prototype specifies the crystal symmetry, the crystal unit cell, and the positions of the atoms in the cell. We then use a combinatorial engine to transmute the atoms in this structure into all other possible chemical elements. For example, the rocksalt crystal structure of the common NaCl kitchen salt will be transmuted in NaF, or LiCl (also naturally occurring salts) but also in RbU, TcH, and all other combinations. The neural network then decides if this is a viable compound, and, if so, passes it to the following stage.

The following stage consists of an ab initio calculation, using density functional theory as implemented in the code vasp, to calculate the basic properties of the proposed compound — and validate the stability prediction of the neural network. For each material we retain around 10 (text) files that contain the most output of the run. These are then compressed and transferred to our university where they are stored.

A calculation with vasp is around 1 million times slower than with the neural network, and it is in this step that the computational resources provided by SuperMUC-NG are essential. The compound is then screened for interesting properties (electronic, magnetic, mechanical, etc.) and inserted into the alexandria database. In the last two

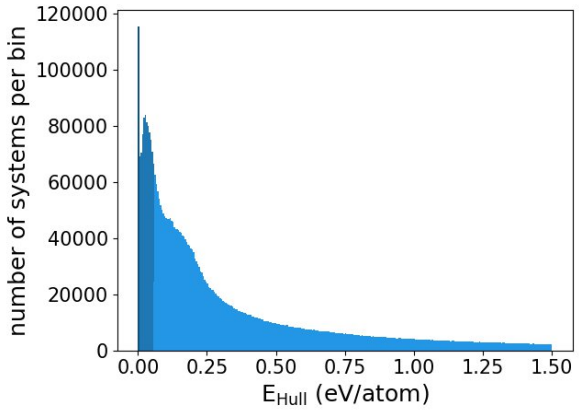

Figure 1: Distribution of the distance to the convex hull of thermodynamic stability of the compounds in alexandria. Only the compounds that are close to the hull (indicated by a darker shade of blue) have a high probability of being synthesized experimentally.

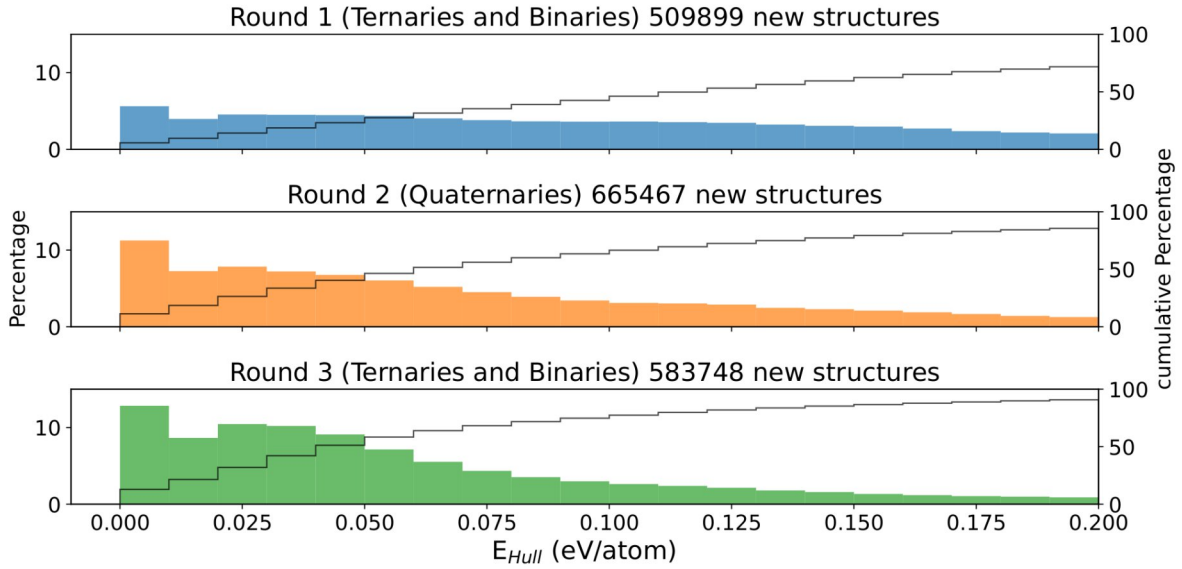

Figure 2: Distribution of the distance to the convex hull for the compounds calculated with vasp in three consecutive "rounds". Between each round, the neural network that proposes the materials to be calculated is retrained with all data from the previous rounds.

years we have used nearly 100 million core-hours in SuperMUC-NG, which allowed us to compute close to two million compounds. The process scales linearly with the number of compounds allowing us to make efficient use of the machine.

Of course, more data usually means better machine learning models, therefore after each "round" of vasp calculations we retrain the neural network to improve its accuracy. The improvement in the prediction of stable, or near stable compounds, can be seen in Figure 2. There we show the distribution of the distance to the convex hull of thermodynamic stability for three consecutive rounds. It is clear that in round 3 a much larger percentage of compounds are close to thermodynamic stability due to the improved training of the crystal graph network allowed by the extra data.

Besides allowing for better training of machine learning models, any of these newly discovered compounds may hold the key to important technological advances. For example, recently we scanned the compounds in alexandria for high-temperature conventional superconductors. The search yielded some surprising results, such as the compound $LiMoN_2$ that we predict to superconduct below 40 K, or the hydride Mg_2IrH_6 that, if synthesized, should be a superconductor at ambient pressure at temperatures surpassing the boiling point of nitrogen.

Ongoing Research / Outlook

At the moment, and in spite of the incredible advances of the past years, we only know a very small part of material space. As such, we are continuing this project in SuperMUC-NG to unveil more stable compounds with potential technological applications.

References and Links

[1] https://alexandria.icams.rub.de/
[2] J. Schmidt, N. Hoffmann, H.-C. Wang, P. Borlido, P.J.M.A. Carriço, T.F.T. Cerqueira, Si. Botti, and M.A.L. Marques, Adv. Mater. 35, 2210788 (2023).
[3] J. Schmidt, H. Wang, G. Schmidt, and M. Marques, npj Comput. Mater. 9, 63 (2023).
[4] T.F.T. Cerqueira, A. Sanna, M.A.L. Marques, Adv. Mater. 36, 2307085 (2024).
[5] A. Sanna, T.F.T. Cerqueira, Y.-W. Fang, I. Errea, A. Ludwig, M.A.L. Marques, npj Comput Mater. 10, 44 (2024).

Numerical simulations of topological and correlated quantum matter

Research Institution
[1]Institute for Advanced Study, Princeton

Principal Investigator
Giorgio Sangiovanni[1], Ewelina Hankiewicz[1]

Researcher
Fakher Assaad[1], Armando Consiglio[1], Matte Crispino[1], Lorenzo Crippa[1], Philipp Eck[1], Stefan Enzner[1], Max Fischer[1], Anika Götz[1], Florian Goth[1], Andreas Hausoel[2], Alexander Kowalski[1], Tobias Müller[1], Gabriel Rein[1], Manish Verma[1], Niklas Wagner[1]

Project Partner
[1]University of Würzburg
[2]IFW Dresden

Funding
SFB 1170, EXC2147

SuperMUC Project ID: pr94vu

Introduction

This report covers phenomena in the solid state central to Würzburg's SFB1170 on Topological and Correlated Electronics at Surfaces and Interfaces. The research makes avail of various established numerical techniques, such as the renowned Density Functional Theory (DFT) and Dynamical Mean-Field Theory (DMFT). DFT as it is capable of modeling ab initio the properties of realistic materials and heterostructures, such as moiré lattices, while DMFT can describe with an all-order accuracy local quantum fluctuations, which play an important role in systems with narrow orbitals or emerging flat dispersion (such as the famous Magic-Angle Twisted Bilayer Graphene), due to the strong Coulomb repulsion felt among the electrons that occupy those energy levels. Hence, the combined use of the two approaches is extremely powerful in describing novel features in both theoretical models and realistic compounds: among those we count exotic and topologicalphenomena related to the properties of the single-particle Green's function when the Coulomb repulsion is strong with respect to the noninteracting dispersion bandwidth, correlations effect and their interplay with ordering, phase transitions and topology in 2D systems, Hund physics in metallic systemsand the interesting physics of the very popular topic of kagome materials.

Results and Methods

Dynamical Mean-Field Theory and related approaches can be effectively employed to study the complex landscapes of phase transitions in the family of Hubbard models, and of two-particle response function in various channels (spin, charge, etc…). Using DMFT with the well-established CT-HYB impurity solver w2dynamics1 we computed stable and unstable solutions near the critical point of the transition between a “good” metal phase and the strongly correlated “Hund's metal” phase in a two-orbital model with Hund's coupling [1], partially including two-particle “generalized susceptibilities”. As seen in Figure 1, this allowed us to extend an analysis of the connection between the lowest eigenvalue λl of the latter and the charge compressibility beyond the threshold indicating the pole of κ (near “A”). We found that it also separates stable from unstable solutions (unfilled example point) and that all thermodynamic stability conditions can be expressed in the eigenbasis of the generalized susceptibility to show the link between stability and eigenvalues. In combination with DFT, DMFT is capable of describing the complex phase diagrams and thermodynamic observables of realistic structures, where an abundance of experimental data is already available and up for interpretation. In this direction, we progressed in the study of correlation effects in twisted bilayer graphene, making use of the topological heavy fermion model developed by Song and Bernevig in 2022, capable of replicating the band dispersion and symmetries of the continuum Bistritzer-MacDonald model while making use of localized Wannier functions. We have performed a CT-QMC study of local spin susceptibiliy at varying temperature [2], and a deeper analysis of the symmetric and symmetry-broken phases with focus on the origin and ordering of local moments, a cascade-like behavior of the spectral weight due to charge redistribution, entailing Lifshitz transitions at the Fermi level, and sawtooth patterns in charge compressibility. The full combination of DFT+DMFT was employed in the description of the physics of localized and itinerant electronic excitations in another bilayer system, an heterostructure of the Tantalum Dichalcogenide material TaS_2. This is obtained by juxtaposing a two dimensional Mott insulating 1T (tetragonal layer and a conducting 1H (hexagonal) layer [3] of the same compound. By our numerical simulations, we determined that the narrow zero-bias peak in spectral weight observed in various recent experiments is controlled by a doped-Mott behavior of the 1T layer driven by charge transfer, rather than by Kondo physics as generally believed. We confirmed this conclusion by the analysis of local spin susceptibility, local spectral functions and Fermi surface spectral weight distribution at varying temperature. We utilized Density Functional Theory (DFT) also to investigate kagome structures, which have gathered significant attention in the scientific community for their potential to host unique correlated and topologically nontrivial electronic states. Because of its peculiar geometry, the electronic structure of the kagome lattice exhibits features consisting a Dirac

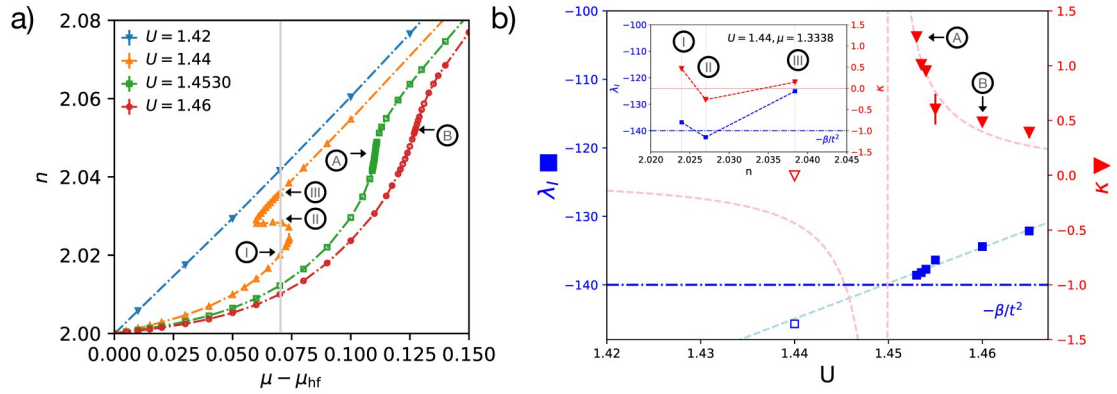

Figure 1: Densities *n* for chemical potential *μ* and interaction U around the phase transition. (b) At the threshold (blue dash-dotted line) for the lowest generalized susceptibility eigenvalue λ_I the compressibility κ diverges and only the unstable solutions (unfilled) are found below it.

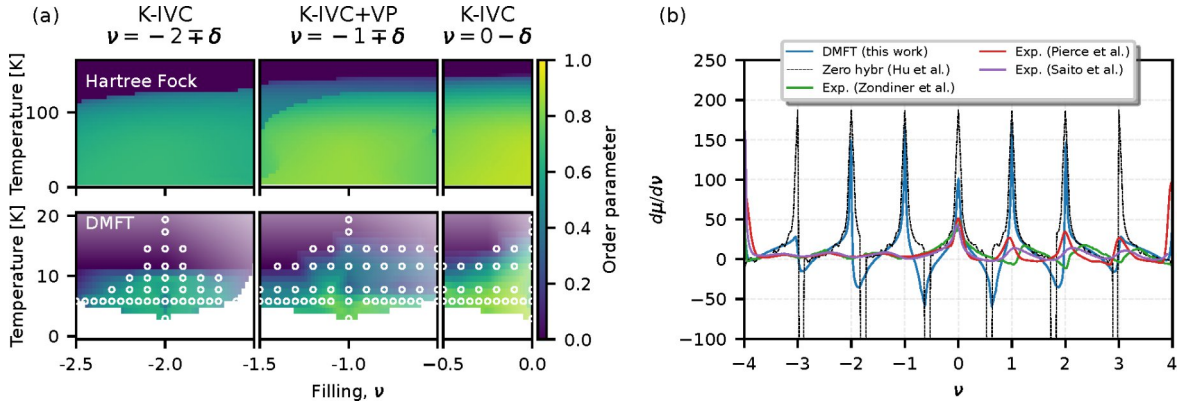

Figure 2: Magnetic phase diagram of the Topological Heavy Fermion model for twisted bilayer graphene (TBG), and inverse charge compressibility plot compared with recent experiments.

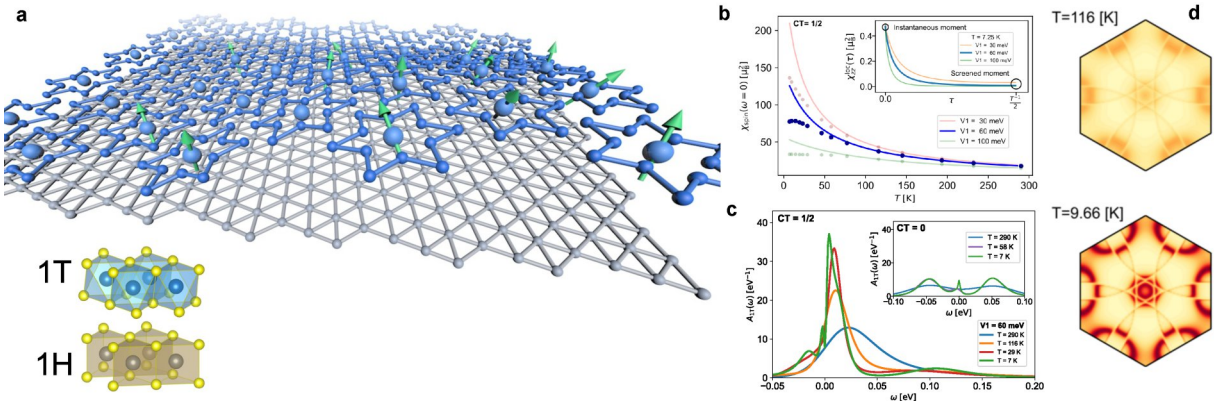

Figure 3: Stacking of the bilayer TaS_2 heterostructure and DMFT spectral function at different values of inter-layer charge transfer and temperature.

point, two van Hove singularities and a flat band across the whole Brillouin zone (BZ). Due to its versatility and accuracy, DFT stand as an indispensable tool for discovering the intricate physics of kagome materials involving superconductivity, charge density waves, spin bond orders, competing phononic and electronic degrees of freedom, just to name a few. Reflecting this growing interest, our DFT-based projects have increasingly directed their focus towards these novel kagome systems. Through successful collaborations with experimental colleagues, experts in angle resolved photoemission spectroscopy and time-resolved reflectivity, we have made significant progresses in understanding and unraveling the underlying physics [4,5].

Ongoing Research / Outlook

For our future researches, we plan to highlight the possibility to induce and enhance exotic physical properties in kagome materials through control of strain, doping and hydrostatic pressure. Contestually, we will benchmark as well the limits within which these compounds can be controlled. We will continue these calculations on SuperMUC-NG and we are hence grateful for the support provided over the years.

References and Links

[1] Chatzieleftheriou et al., PRL 130, 066401, Feb. 2023.
[2] Hu et al., PRL 131, 166501, Oct. 2023.
[3] Crippa et al., Nat. Comms., 15 no. 1, 1357, Feb 2024.
[4] Di Sante et al., Nature Physics, 19 no. 8, 1135-1142, Aug 2023.
[5] Mazzola et al., Nano Letters, 23 no. 17, 8035-8042, 2023.

Multiloop functional renormalization group for interacting fermions

Research Institution
Arnold Sommerfeld Center for Theoretical Physics, Ludwig-Maximilians-Universität München

Principal Investigator
Jan von Delft

Researcher
Anxiang Ge, Marcel Gievers, Gün Günal, Marc Ritter, Nepomuk Ritz, Elias Walter

Project Partner
–

Funding
DFG EXC-2111 (Project No. 390814868), Munich Quantum Valley, International Max Planck Research School for Quantum Science and Technology (IMPRS-QST), Studienstiftung des deutschen Volkes

SuperMUC Project ID: pn73mo

Introduction

The renormalization group (RG) is a powerful conceptual framework for an effective treatment of physical systems with several differing energy scales. Among the many implementations of RG ideas, the functional renormalization group (fRG) is a particularly viable tool for a systematic study of interacting many-particle systems. For a given physical system, it is constructed by introducing a scale parameter Λ, that smoothly connects the model of interest to one whose solution is known.

The fRG method is based on an exact hierarchy of coupled flow equations for the n-point vertices. In practice, most works truncate this infinite hierarchy in the one-loop fRG scheme, i.e., by neglecting six-point and higher vertices. The quantitative validity of this truncation is difficult to assess. However, a few years back, a scheme called multiloop fRG (mfRG) was developed, which enables systematic improvements by including all contributions of the six-point vertex that can be computed with numerical costs proportional to the one-loop flow. After a series of Master's and PhD projects, in which proof-of-principle results for the viability of using mfRG were obtained, in this project we scaled up the resources in an effort to obtain truly high-quality results in two different settings: (I) **pseudo-fermion fRG for interacting spin systems** in an imaginary-frequency Matsubara formulation and (II) **fRG for low-dimensional interacting fermionic systems** in a real-frequency Keldysh formulation.

Results and Methods

(I) Pseudo-fermion fRG for interacting spin systems
This subproject investigated interacting spin systems designed to model the magnetic behavior of materials. In some materials, all magnetic phenomena are due to local magnetic moments that usually stem from electrons that are tightly bound to each atom. These can be modeled in terms of just their magnetic moment interacting with other electrons' magnetic moments. The resulting so-called Heisenberg model seems simple, yet shows rich and unique behavior such as quantum spin liquid (QSL) phases on some lattice structures. QSL are states where no classical ordering is present, yet quantum mechanical correlations and entanglement are dominant enough to dictate even macroscopic behavior. The search for a material with unambiguous QSL behavior is still ongoing. Even in theoretical models it is challenging to determine whether the model does indeed lead to QSL behavior, since methods that are capable of simulating this class of models across all parameter choices are scarce. The pffRG is one of very few methods capable of doing so, and in the past we have used it to investigate Heisenberg models on the kagome and pyrochlore lattices to model two material classes generally seen as most promising candidates for QSL behaviour in the field.

In his Master's thesis [1], Gün Günal extended our simulation to models with long-ranged dipolar interactions between magnetic moments on the triangular and square lattice using the resources of SuperMUC-NG. There, we verified spin-liquid behavior using the momentum-resolved spin susceptibility, which is particularly valuable data, as it can be observed in experiment. This data was obtained using pffRG, and required integrating ca. 10^7 coupled differential equations, where each right-hand side is given by a strongly peaked frequency integral. These integrals are evaluated in massively parallelized fashion, using an adaptive Gauss-Kronrod scheme for each integral and then forwarding results to a master node via MPI. The ODE timestep is then performed on the master node with a step-size controlled 5^{th} order Runge-Kutta scheme.

We have also used SuperMUC-NG to benchmark the effect of some technical advances over previous implementations, most notably more efficient parametrization schemes and the consistent application of adaptive schemes throughout the code, which led to qualitative improvement of results in some cases [2].

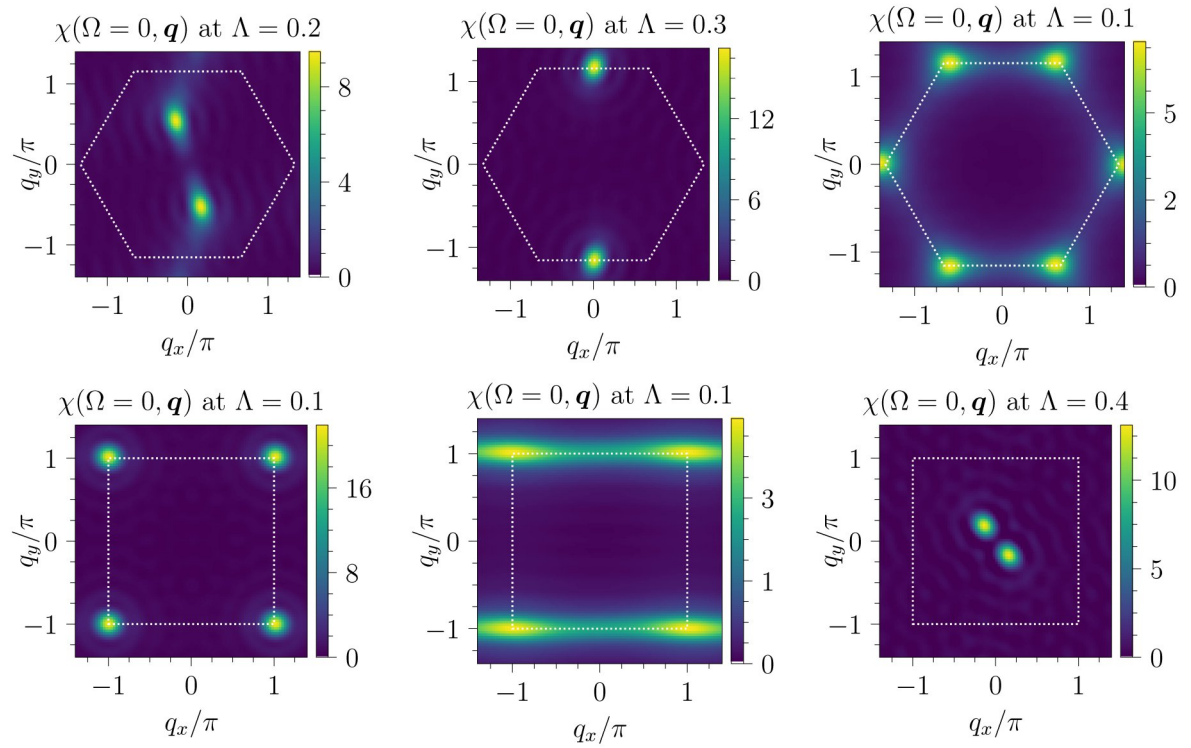

Figure 1: Momentum-resolved spin susceptibility of the Heisenberg model with dipolar interactions on a triangular (top row) and square (bottom row) lattice for various interaction parameters. Sharp features show presence of classical ordered phases, and broad features correspond to more quantum mechanical behavior. Adapted from [1].

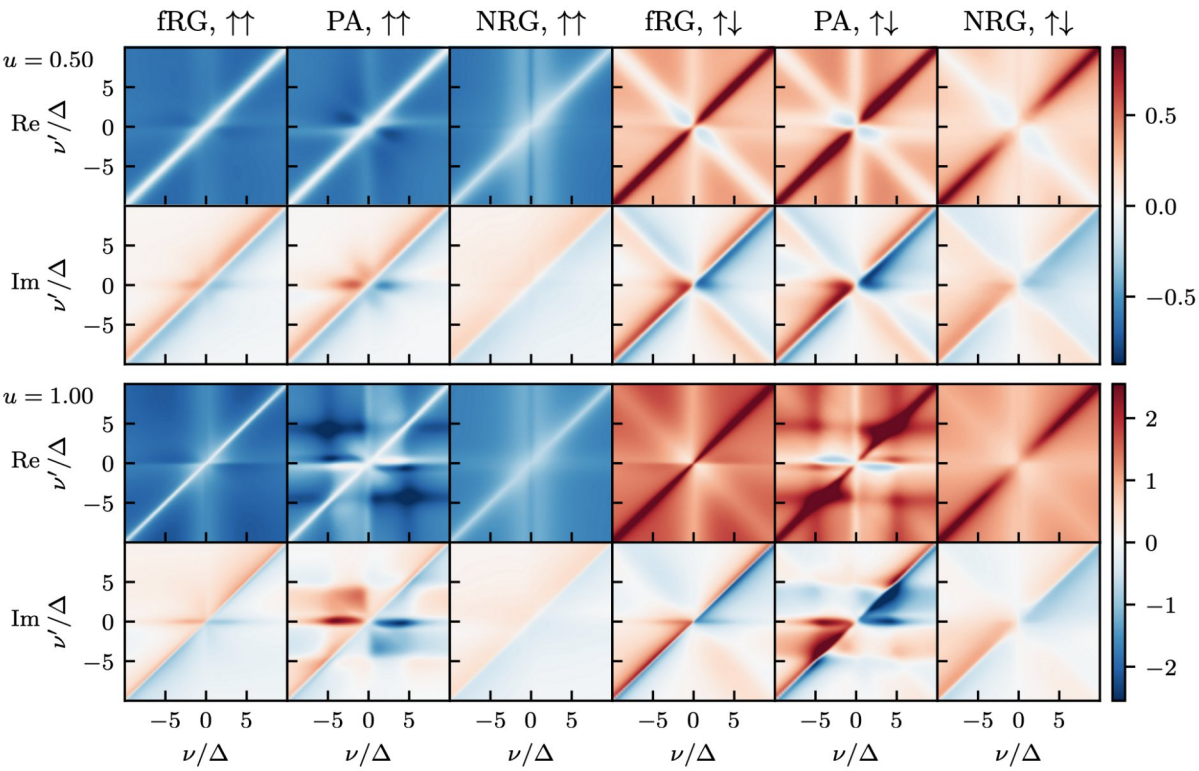

Figure 2: Two-dimensional cuts through the frequency dependence of both spin components of a fully retarded Keldysh component of the full vertex of the SIAM, computed using one-loop fRG, a direct solution of the parquet equations in the parquet approximation (PA) and benchmark data from NRG for two different interaction strengths. We observe good qualitative agreement throughout, as all methods capture all nontrivial features. Adapted from [3].

(II) Keldysh fRG for the single-impurity Anderson model

Though most of the fRG codes developed so far employ the imaginary-frequency Matsubara formalism, the real-frequency Keldysh formalism is often preferable. For example, when computing dynamical response functions, or the finite-temperature conductance through a quantum dot, one needs the self-energy and vertex as functions of real, not imaginary, frequencies. The Keldysh formalism yields these directly, without the need for the (mathematically ill-posed) analytical continuation of imaginary frequency data. Moreover, it can also deal with non-equilibrium situations. These benefits come at the cost of introducing an additional Keldysh index, taking two values, so that the self-energy and vertex contain 4 and 16 Keldysh components, respectively. Though this causes increased coding complications and numerical costs, these can be reduced by exploiting the fact that not all Keldysh components are independent.

In a recent study [3] we performed such calculations for the single-impurity Anderson model (SIAM) with the one-loop fRG scheme and, independently, solving the self-consistent parquet equations, whose solution corresponds to a fully converged mfRG flow. We chose to solve the parquet equations directly, as during the course of the project it turned out to be more feasible than converging an mfRG flow in the number of loops, due to technical intricacies of the real-frequency formalism. The SIAM was chosen because of its structural simplicity (no spatial indices), challenging dynamical properties (emergence of a dynamically generated scale, the Kondo temperature), and the possibility for benchmark comparisons against highly reliable numerical results from the numerical renormalization group (NRG). We hence compared our results to benchmark data obtained with our own NRG implementation and found that capturing the full three-dimensional frequency dependence of the four-point vertex significantly improves the fRG results compared to previous implementations, and that solving the parquet equations in the parquet approximation (PA) yields the best agreement with the NRG benchmark data, but is only feasible up to moderate interaction strengths.

Fig. 2 shows two-dimensional cuts through the frequency dependence of both spin components of a fully retarded Keldysh component of the full vertex of the SIAM, computed using the three different methods outlined above for two different interaction strengths. Physically, this corresponds to the effective interaction of two electrons on the impurity with equal or opposite spins, respectively, and energies ν, ν' without energy transfer. We observe good qualitative agreement throughout, as all methods capture all nontrivial features.

Ongoing Research / Outlook

For future projects we would like to investigate lattice models in intermediate to strongly correlated regimes. A promising method is DMF2RG which starts the flow with the solution of an impurity model and turns on the coupling tunneling between lattice sites along the flow.

However, calculations involving four-point vertices with momentum-dependence or more orbital indices require large amounts of resources. In a recent publication [4] some of us investigate a parsimonious representation which raise hope that we can strongly cut back both consumption of memory and computation time. Particularly patching schemes are useful for parallelization in said representation. While prototypes of such algorithms already exist, it might take us months until we can push for models which require large-scale parallalization.

Until these algorithms are ready we are planning to continue the project by investigating different cutoff-functions and variations of the model, such as the SIAM out of equilibrium and a two-impurity model with exchange interaction.

References and Links

[1] G. Günal, Multiloop Pseudofermion Functional Renormalization Group Study of the Heisenberg Model with Dipolar Interactions on the Triangular and Square Lattice, Ludwig-Maximilians-Universität, 2023.
[2] M. K. Ritter, D. Kiese et al., Eur. Phys. J. B 95, 7 (2022).
[3] A. Ge, N. Ritz et al., Phys. Rev. B 109, 115128 (2024)
[4] M. K. Ritter, Y. N. Fernández et al., Phys. Rev. Lett. 132, 056501 (2024).

Computational Fluid Dynamics and Engineering

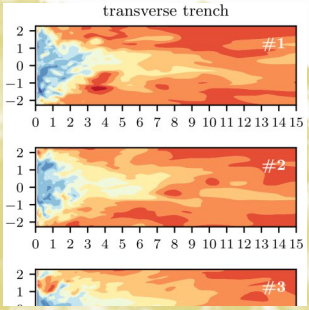

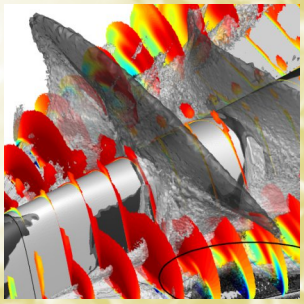

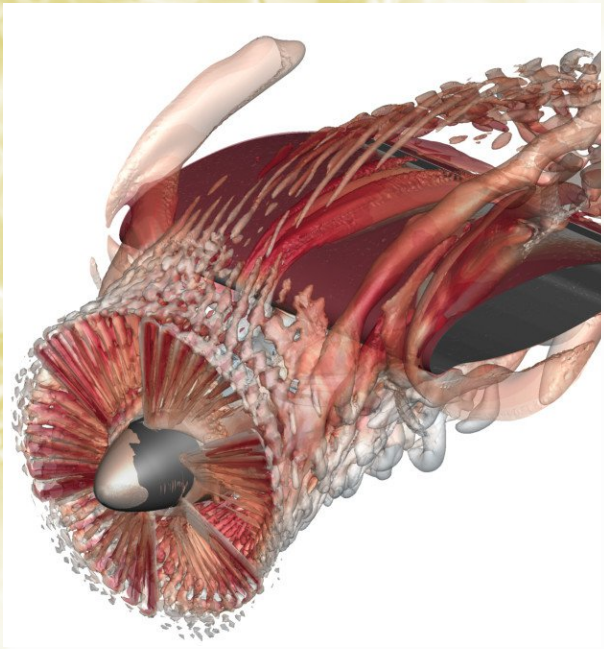

Turbulent convection at very small Prandtl numbers and very large Hartmann numbers

Research Institution
[1]Technische Universität Ilmenau

Principal Investigator
Jörg Schumacher[1]

Researcher
Dmitry Krasnov[1], Philipp P. Vieweg[1], Ambrish Pandey[2]

Project Partner
[2]Indian Institute of Technology Roorkee, India

SuperMUC Project ID: pn68ni

Introduction

Turbulent convection is one of the fundamental flows which are ubiquitous in nature and technology. One central parameter of the flow, the dimensionless Prandtl number *Pr* is often very small, $Pr << 1$, which implies that temperature diffusion proceeds much faster than momentum diffusion. Turbulent convection is often connected with other physical processes, such as with strong (external) magnetic fields. The relevance of the latter process is quantified by a further dimensionless parameter, the Hartmann number *Ha*. The Hartmann number values are very large, $Ha >> 1$, when the external magnetic field is strong and dominates the dynamics. Our fundamental understanding of such complex turbulent flow processes is still incomplete. High-resolution direct numerical simulations of the equations of Rayleigh-Bénard convection (RBC) in the Oberbeck-Boussinesq approximation can thus reveal new insights into the basic mechanisms of turbulent transport of heat and momentum and its connections to characteristic structures and patterns. The same holds for magnetoconvection and magnetohydrodynamic (MHD) flows.

In our present supercomputing project, we continue to study two extreme cases of RBC that cannot anymore be obtained in a controlled laboratory experiment. We use massively parallel direct numerical simulations that resolve all vortices and thermal plumes to their smallest extension and do not require parametrizations of unresolved turbulence. On the one hand, we pushed the Prandtl number to values as small as $Pr = 0.001$, which are not anymore obtainable in controlled laboratory experiments [1,2]. On the other hand, we studied convective turbulence in the presence of very strong external magnetic fields [1].

Results and Methods

We numerically solve the three-dimensional equations of thermal convection which couple the velocity and temperature fields. This is done either by spectral element method or finite difference schemes. The external magnetic field enters the Navier-Stokes equations as an additional Lorentz force term in the quasistatic model [3]. The electrical currents *j* have to be closed inside the fluid volume, i.e., $div\ j = 0$ has to be satisfied. Our MHD model exploits the quasi-static limit [3].

The simulation domains are cuboid cells or ducts with no-slip boundary conditions at all walls. The sidewalls are thermally insulated. The biggest direct numerical simulations of a non-magnetic RBC case at $Pr = 10^{-3}$ and a Rayleigh number of in $Ra = 10^7$ a domain with an aspect ratio 25:25:1 requires 144,000 SuperMUC-NG cores for a grid with 20,480×20,480×1,280 points [2]. The data analysis of such simulations is very demanding. One snapshot of the turbulent convection flow including the three velocity components and the temperature requires for example 17 TBytes storage.

Figure 1 shows a supergranule convection cell obtained from a long-term direct numerical simulation of RBC at a Prandtl number $Pr = 0.01$ in a very strongly extended horizontal domain. A contour plot of the temperature field close to the top plate is shown in both panels (cold = blue and hot = red). The dominant large-scale structure arises only when the thermal boundary conditions at the top and bottom walls are changed from a prescribed temperature to a prescribed heat flux. This study, which is published in [4], confirmed the omnipresence of a very slow aggregation process to a supergranule for the entire Prandtl number range of investigated fluids ranging from 0.01 all the way up to 100. The bottom panel of Figure 1 underlines this point by a plot of the temperature distribution at $Pr = 100$, i.e., 4 orders of magnitude larger. Note that the pair of supergranules is superposed by a smaller-scale network of granules, i.e., smaller convection cells. Moreover, we analysed the Prandtl number effects on the global heat and momentum transport. The simulations could clarify the role of a potential stable stratification in the bulk of the fluid layer which showed up for the higher Prandtl numbers. The ubiquity of the investigated mechanism of flow self-organisation underlines its relevance for pattern formation in geo- and astrophysical convection flows, the latter of which are often driven by

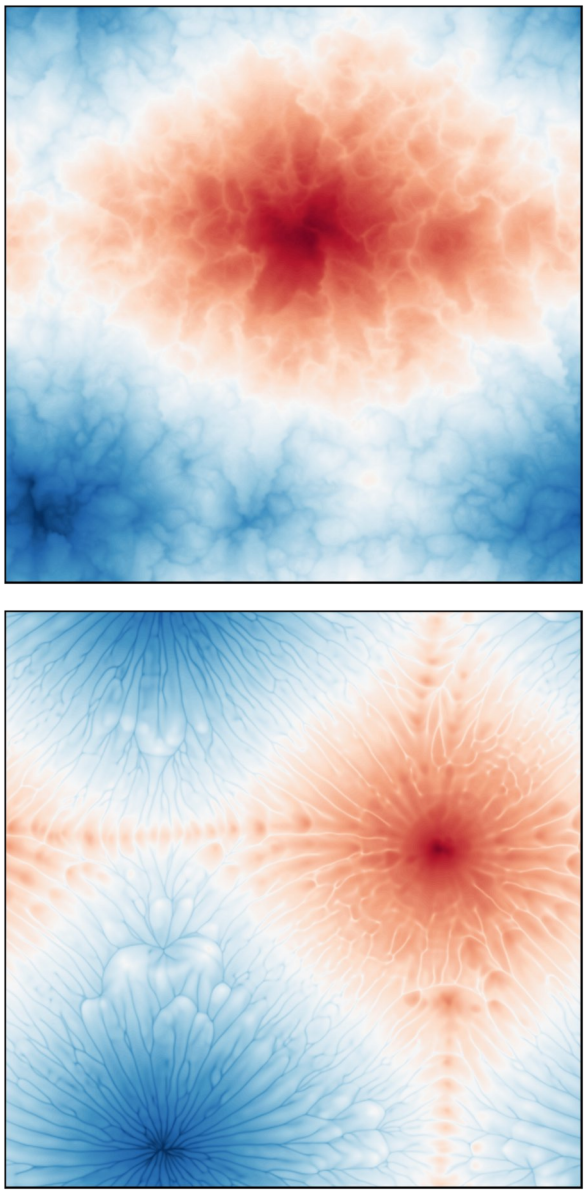

Figure 1: Supergranular cell formation in an extended convection layer of aspect ratio 60:60:1 at a low Prandtl number *Pr = 0.01* (top) and *Pr = 100* (bottom). Figures are adapted from [4].

prescribed heat fluxes (as in the present cases). In these simulations, it was also seen that a characteristic scale of the order of the mesoscale seems to saturate at a wavelength of $\Lambda = 3H$ with the height H of the convection layer for the lowest Prandtl numbers.

Figure 2 shows a recent result of our magnetohydrodynamic simulations in extended isothermal ducts with thin electrically conducting walls in presence of very strong magnetic field at a Hartmann number $Ha = 2{,}000$. Here, the electrical conductivities at the front and back sidewall differ from those at the top and bottom ones. An extension of the existing in-house finite difference code allowed us to implement this enhanced complexity that exists in all realistic technological applications. The figure demonstrates how the new Tensor-product-Thomas solver is applied to solve the elliptic problems for pressure, electrical potential, temperature, and velocity. New physical effects, such as the formation of these wall jets arise for example when the side walls are conducting.

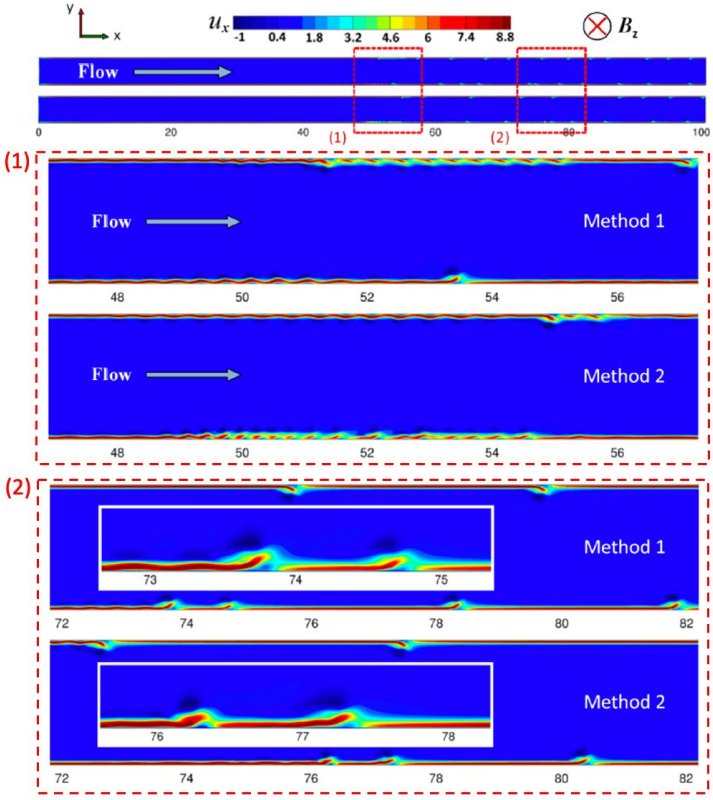

Figure 2: Isothermal duct flow with electrically conducting walls in the presence of a magnetic field. Method 1 and 2 stand for two arrangements of the Tensor-product-Thomas procedure. The top panel illustrates the whole duct (with 100 length units) and indicates the mean flow direction as well as the direction of the external magnetic field. The subsequent 4 panels zoom into the two annotated regions (dashed line boxes). They demonstrate the slow evolution of wall jets which will affect the turbulent mixing properties and thus the turbulent heat transfer in such a configuration. This flow is also termed a Hunt flow.

Ongoing Research / Outlook

The existing numerical database on low-Prandtl-number RBC simulations provides a unique possibility to determine turbulent eddy viscosities and diffusivities which are still the backbone for turbulence parametrizations in numerous real-life applications of thermal convection. Our series of high-resolution simulations that vary the Prandtl number Pr for several fixed Rayleigh numbers Ra enable a calculation of the thermal and kinetic energy dissipation rates which are required for such a task. The analysis of the data will give us unique insights into the small-scale turbulence statistics and the connected intermittency in such flows and can eventually improve the modeling of such flows.

In view to the isothermal magnetohydrodynamic or the magnetoconvection flows, control strategies can be developed to even further enhance the turbulent heat transfer. These studies have been initiated recently and will be continued in the coming year by further numerical simulations.

References and Links

[1] https://www.tu-ilmenau.de/tsm
[2] A. Pandey et al., J. Fluid Mech. 948, A23 (2022).
[3] D. Krasnov et al., J. Comp. Phys. 474, 111784 (2023).
[4] P. P. Vieweg, J. Fluid Mech., 980, A46 (2024).

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

Research Institution
Numerical methods in Aerospace Engineering, Bundeswehr University Munich

Principal Investigator
Markus Klein[1]

Researcher
Marco Herbert[1], Michael Pfitzner[1], Nilanjan Chakraborty[2], Isaac Boxx[3]

Project Partner
[2]Newcastle University, Newcastle upon Tyne, UK, [3]RWTH Aachen

SuperMUC Project ID: pn69ga

Introduction

Combustion in most engineering applications, such as spark ignition engines and gas turbines, often involves elevated pressure conditions and non-unity Lewis number *Le* fuel-blends. However under these extreme conditions, the flame morphology becomes increasingly complex and turbulent, persistent with convoluted structures arising due to the presence and interactions of inherent flame instabilities [1]. In this project, direct numerical simulation (DNS) analysis is performed to evaluate and complement existing modelling approaches to account for realistic operating conditions for combustion applications. In this regard, a wide parameter database of Bunsen flames is generated using the higher-order compressible DNS code SENGA [1-5]. The database is representative of wrinkled/corrugated flamelets and thin reaction zones regimes, comprising variations of operating pressure, turbulence intensity and Lewis number conditions.

Results and Methods

DNS, where the flame front is adequately resolved, are used to generate a comprehensive numerical database to analyze and quantify the flame structure and physics under varying operating conditions.Compressible Navier-Stokes equations coupled with a scalar transport equation for the reaction progress variable are numerically solved using the finite difference-based DNS code SENGA, which implements 10^{th}-order finite differences for space and a third-order Runge-Kutta method for time discretization. The code is fully parallelized using domain decomposition. A generic single-step Arrhenius-type irreversible chemistry was shown to be accurate enough for the analysis in this project. The numerical configuration comprises a burner of diameter n_d in a cubic domain of sides $2d_n$ and characteristic mesh sizes of about 750^3 cells, with all boundaries modelled as partially non-reflecting outlets (except the burner inlet). At the inlet, a hyperbolic-tangent like mean velocity distribution with superimposed pseudo turbulence is imposed.

The database has been established in [1] and analyzed [2] and extended later on with focus on modelling turbulent premixed combustion [3-5]. The present report illustrates a more recent comparison between DNS of a lab scale burner (see Fig. 1) with experiments conducted at RWTH Aachen for three hydrogen diluted atmospheric methane flames. The instantaneous snapshots in Fig. 2 show an excellent qualitative agreement between experiment and DNS confirming the validity of the chosen approach. The development of flame instabilities with increasing hydrogen content which results in an increased turbulent flame speed and a shortening of the flame brush, can be clearly seen. A more quantitative analysis in terms of turbulent flame wrinkling, turbulent flame speed and validity Damköhler's hypothesis [2] is shown in Fig. 3. The effects of hydrogen dilution can be well captured with the present methodology and illustrate a very good agreement between experiment and DNS at laboratory scale.

References and Links

[1] M. Klein, H. Nachtigal, M. Hansinger, M. Pfitzner, N. Chakraborty. FLOW TURBUL COMBUST, 101(4):1173–1187, Dec 2018.
[2] N. Chakraborty, D. Alwazzan, M. Klein, R.S. Cant. Proc. Combust. Inst., 37(2):2231–2239, 2019.
[3] R. Rasool, N. Chakraborty, M. Klein. Combust. Flame, 231:111500, 2021.
[4] R. Rasool, M. Klein, N. Chakraborty. Combust. Flame, 239, 111766, 2022.
[5] V. Mohan, M. Herbert, M. Klein, N. Chakraborty. Energies, 16:2590, 2023.

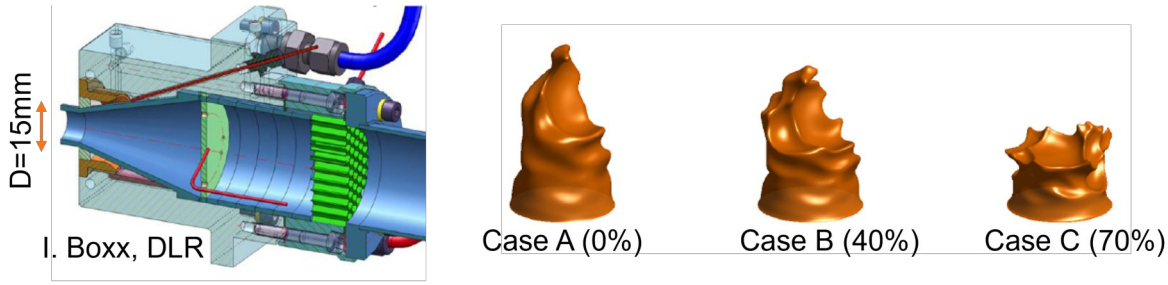

Case	ϕ	$H_2[\%]$	$U_0\ [m/s]$	$S_L[m/s]$	$T_{ad}\ [K]$	$\delta_{th}[mm]$
A	1.0	0	2.13	0.3654	2229	0.4063
B	0.8	40	2.18	0.3665	2026	0.4112
C	0.65	70	2.12	0.3456	1818	0.4474

Figure 1: Experimental setup and 3D isosurfaces from DNS for Cases A-C.

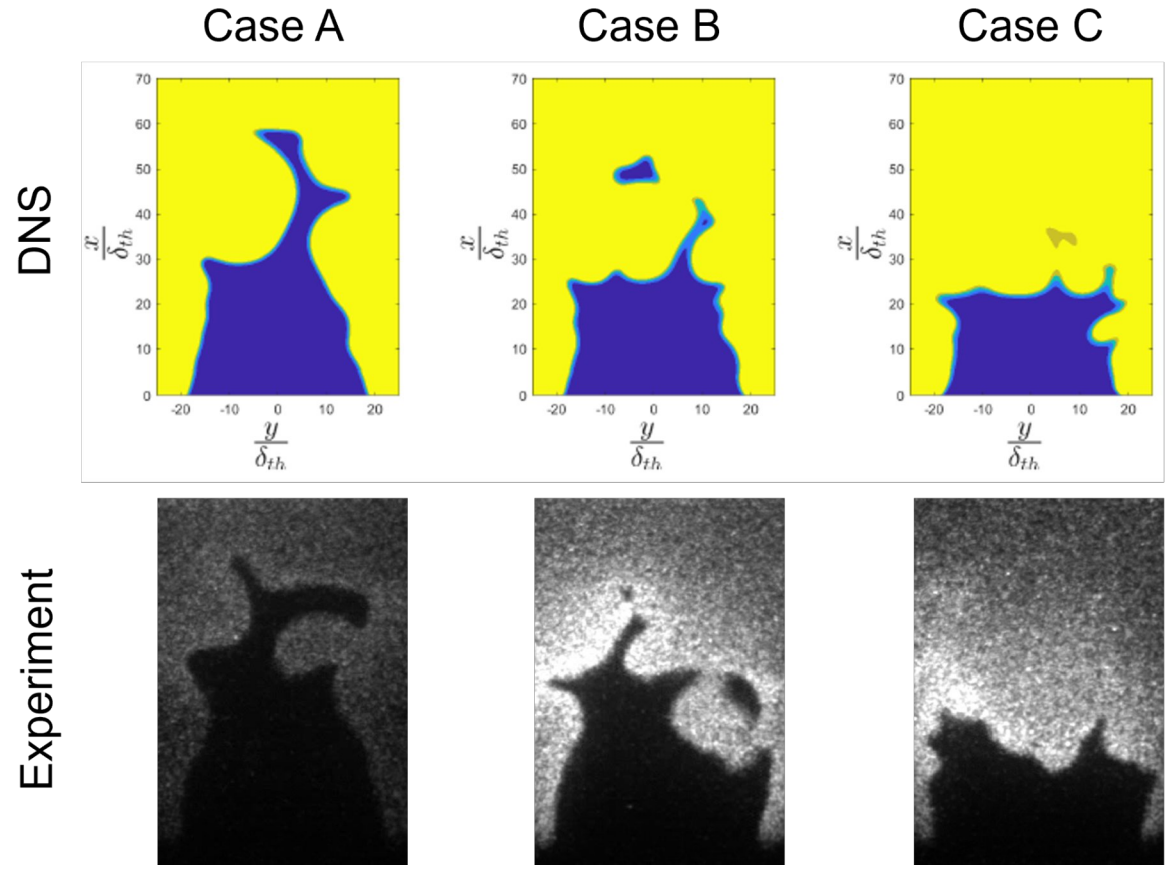

Figure 2: Instantaneous snapshots of flame contours: DNS versus Experiment.

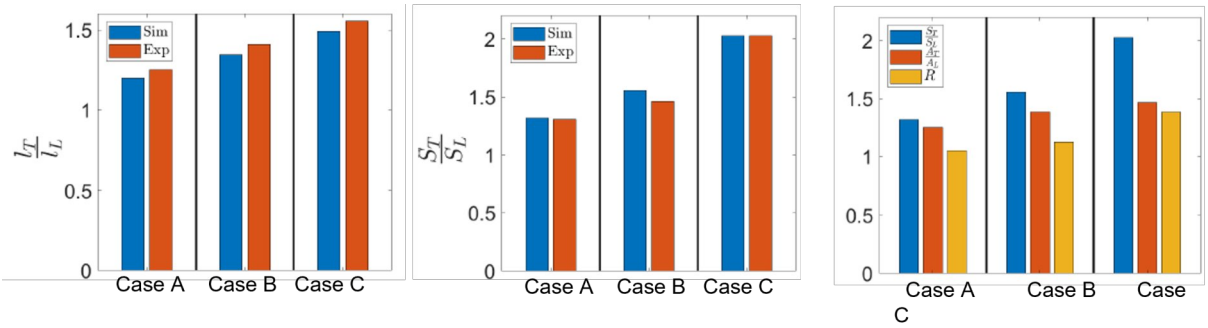

Figure 3: 2D turbulent flame wrinkling, turbulent flame speed; *R* = $(S_T/S_L)/(A_T/A_L)$, Damköhler's hypothesis: *R* = 1.

Film cooling of walls at gas turbine-like conditions

Research Institution
Institute of Thermodynamics, University of the Bundeswehr Munich

Principal Investigator
Lukas Fischer

Researcher
Michael Pfitzner

Project Partner
—

Funding
DFG project number PF 443/7-1

SuperMUC Project ID: pn73ji

Introduction

The air stream in a gas turbine is firstly compressed and delivered to the combustion chamber, where fuel is mixing in and burnt, releasing a tremendous amount of heat. The hot turbulent burnt gases expand through the turbine placed downstream and the exhaust nozzle. Over the last decades, the turbine inlet temperature has increased because this leads to a higher efficiency of the gas turbine. The temperature of the hot gas of the combustion chamber (2,200 °C) and turbine section (1,700 °C) surpasses the material's maximum temperature limit (900 °C). In order to safeguard the metal walls from damage, they are covered by a ceramic thermal barrier coating (TBC) but this is not sufficient to protect the metal components from overheating. Film cooling techniques are employed to generate a protective coolant film, shielding the component surface from direct contact to the hot gas. A common film cooling method introduces the coolant compressor air through discrete cylindrical holes in the walls of the cooled surfaces (effusion cooling) at a temperature of around 600 °C. The use of air to cool the turbine section decreases the efficiency of the gas turbine. Hence, an objective to enhance the gas turbine's efficiency is to reduce the coolant mass flow of the compressor. Twenty years ago, Bunker et al. [1] suggested removing some of the thermal barrier coating near the cooling holes' outlets to form a transverse trench, as shown in Figure 1. These trenched film cooling designs have shown, in several investigations, to outperform the state-of-the-art design. Improved designs, such as a segmented trench [2] as shown in Figure [2], were proposed in the literature. Most investigations were performed under conditions in which the hot gas flow was featuring low, wind tunnel-like levels of turbulence, which is much different from the turbulent flow field in a gas turbine. To investigate trenched discrete film cooling designs under conditions at low turbulence and at more realistic flow conditions, high-performance computing resources from the LRZ were necessary. The investigation aims to enhance the understanding of film cooling flows, ultimately contributing to the reduction of fuel consumption of gas turbines in the future.

Results and Methods

Unsteady Computational Fluid Dynamics (CFD) simulations were performed using the finite volume method. The film cooling domain was discretized with up to 14 million hexahedral elements. Based on this, the conservation equations for mass, momentum, and energy, known as the Navier-Stokes Equations (NSE), was solved using the open-source CFD software OpenFOAM. Large Eddy Simulations (LES) were conducted, which resolved over 95% of the turbulent eddies within the film cooling flow domain. This allowed for a high accuracy in predicting film cooling performance. The simulations for each design under laminar and gas turbine-like conditions were run on approximately 1,000 cores of SuperMUC-NG per run for 2-8 weeks. Throughout this investigation, we were able to utilize up to 5,000 CPU cores simultaneously to simulate different film cooling designs, enabling us to produce the results [3] within a reasonable timeframe. First, we discuss the time-averaged results for discrete film cooling holes and for the two investigated trenched designs. In Figure 3, a side view of the temperature field in the center plane of a film cooling hole configuration is presented. The color blue indicates low temperature, while the hot gas temperature is represented by red color. The solid line (-) corresponds to gas turbine-like high turbulence level conditions, whereas the results of the laminar hot gas flow simulations are depicted by the dotted lines (- -). The objective of film cooling is to maintain a low temperature at the wall. The state-of-the-art, simple effusion hole design features a film cooling jet emerging from the cooling hole that lifts off the wall, which is typically encountered in effusion-cooled combustor walls. With normal effusion cooling (top), an increased turbulence level reduces the jet lift-off slightly, resulting in a positive cooling effect as the coolant is transported in closer proximity to the wall. By employing a trench, the jet makes direct contact with the trench wall. Subsequently, the coolant flow can more effectively adhere to the downstream wall, creating a better cooling film. However, this efficiency is compromised slightly by elevated levels of free-stream turbulence for both designs. The next set

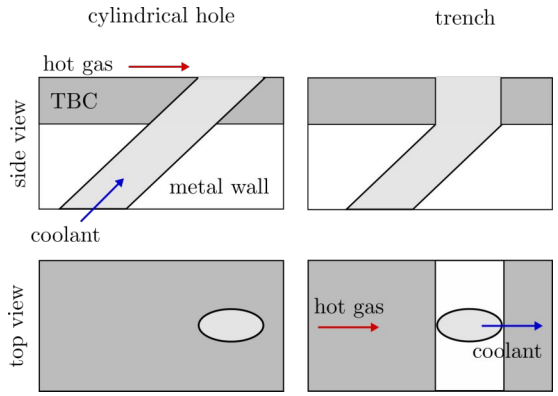

Figure 1: Schematics of commonly used cylindrical hole and of a transverse trenched film cooling design [3].

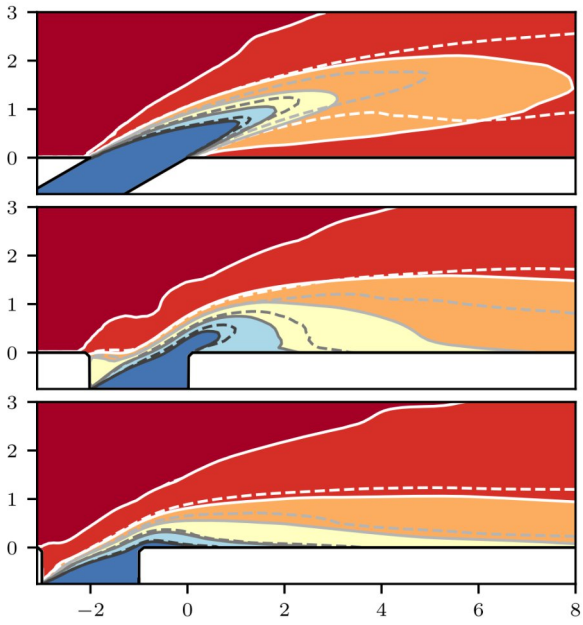

Figure 3: Side view of time-averaged temperature (blue: low temperature; red: high temperature) depending on low turbulence (- -), high turbulence (-) for standard effusion cooling (top), transverse trench (middle) and segmented trench (bottom) [3].

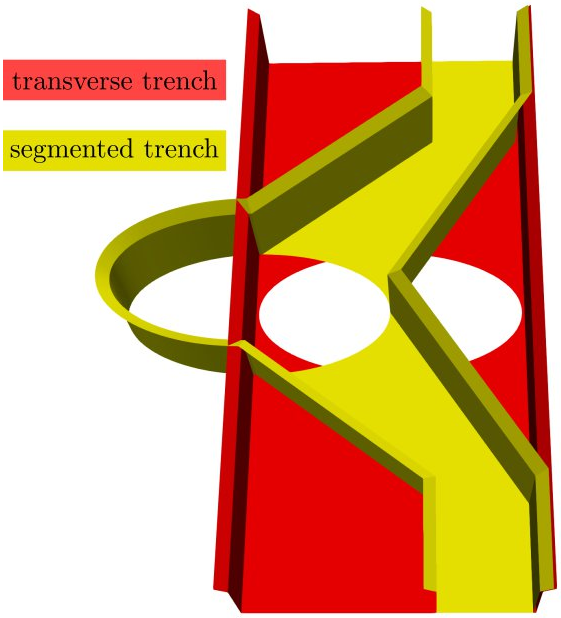

Figure 2: Comparison of transverse and segmented trench designs [3].

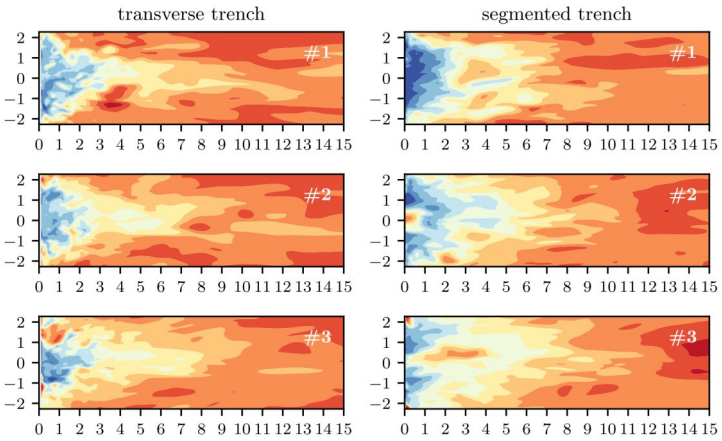

Figure 4: Top view of the instantaneous film coolant distribution (#1–#3) at the wall downstream of the trenches (blue: low temperature; red: high temperature) [3].

of results focuses on the film cooling coverage immediately downstream of the trenches. Instantaneous images of the film coolant distribution near the wall for both investigated trench designs are shown in Figure 4. In snapshot #1, the axial and lateral distribution of the coolant at the wall showed a low temperature downstream of the trenches for both designs. In a subsequent instant (#2), the axial extent of the film cooling efficiency was reduced for both configurations. Additionally, a hot spot emerged just downstream of the segmented trench. The transverse trench exhibited less sensitivity in the central region of the wall due to the higher trajectory of the coolant, as shown in Figure 3. However, it displayed significantly lower coolant coverage at the sides compared to the segmented trench. In instant #3, the hot spot downstream of the segmented trench shifted even further downstream. Moreover, brief hot spots emerged at both lateral sides. These results demonstrate how Large-Eddy-Simulations allow a detailed evaluation of the unsteady features of such highly turbulent flow fields and their effect on the wall heat fluxes.

Ongoing Research / Outlook

Subsequent investigations may explore the influence of potential manufacturing tolerances, including aspects such as the contour of trench edges. Additionally, there is scope for development and comparison of innovative trench designs [4]. In addition, alternative cooling configurations combining trenches with fan-shaped cooling holes may be investigated.

References and Links

[1] R.S. Bunker, Film cooling effectiveness due to discrete holes within a transverse surface slot, Proceedings of ASME Turbo Expo 2002.
[2] P. Schreivogel, B. Kröss, M. Pfitzner, Study of an optimized trench film cooling configuration using scale adaptive simulation and infrared thermography, in: Proceedings of ASME Turbo Expo 2014.
[3] L. Fischer, M. Pfitzner, Int. J. Heat Mass Transfer, 2023.
[4] L. Fischer, D. James, S. Jeyaseelan, M. Pfitzner, Int. J. Heat Mass Transfer, 2023.

Numerical investigations of turbulent emulsions

Research Institution
Numerical methods in Aerospace Engineering, University of the Bundeswehr Munich

Principal Investigator
Theresa Trummler, Markus Klein

Researcher
Alexander Begemann, Oscar Krzeczek, Elias Trautner

Project Partner
–

SuperMUC Project ID: pn73zu

Introduction

Emulsions are suspensions of immiscible liquids, such as oil and water, where the dispersed liquid is present in form of poly-disperse droplets in the carrier liquid. Emulsions play an essential role in various industrial processes, such as food processing, oil production or pharmaceutical processes and recent research is exploring the application of fuel-water emulsions for more efficient and environmentally friendly power generation. An example is gasoline-water direct injection (GWDI) for future gasoline engines, which can be realized by injection of gasoline-water emulsions. Other examples are fuel-water emulsions for small gas turbines or diesel engines. For these applications, a better understanding of emulsions is important. Additionally, a detailed knowledge of the stability of such emulsions without any further energy input and of the time scale of the segregation process is crucial for the realization of GWDI in order to know after which time at rest the emulsion can no longer be injected. One way to gain a deeper insight into the underlying mechanisms is by conducting numerical flow simulations using computational fluid dynamics (CFD).

This project investigates turbulent emulsions in homogeneous isotropic turbulence (HIT) using direct numerical simulations (DNS) in combination with the volume of fluid method (VOF). First a detailed study of emulsions and their droplet size distribution has been performed, focusing on the analysis of the effects of different density ratios and volume fractions [1-2]. Secondly, the segregation process has been analyzed in detail [2-3].

Results and Methods

The simulations are conducted as visualized in Figure 1. First, single phase simulations have been performed to obtain a fully developed single-phase HIT. Then, we initialize the dispersed phase as spherical droplets. This approach results in a significant reduction of emulsification time and thus computational cost. The droplets break up due to the energy input from the forcing and an emulsion is generated. To study the segregation in decaying turbulence under gravitational acceleration, we switch off the forcing, turn on gravitational acceleration, and insert no-slip walls in the direction of the gravitational force as visualized in Fig. 1 (d). The simulations have been conducted on Cartesian equidistant meshes with up to $1{,}024^3$ computational cells. The finite volume code PARIS [4] is second order accurate in time and space and parallelized using domain decomposition.

The distribution of bubble sizes in waves follows certain laws, with their validity also being discussed for droplets. The distribution of bubble sizes in the sub Hinze region is different to the distribution of bubbles larger than the Hinze scale because breakup is more probable. Fig. 2 shows the droplet size distributions for four different cases. The scaling laws $d^{-3/2}$ for $d < d_h$ and $d^{-10/3}$ for $d > d_h$ can be clearly observed for two cases only. The dependence of the critical Weber number on the density ratio is discussed in [3].

The segregation process takes place through coalescence and gravitational forces and is illustrated in Fig. 3. A time scale for segregation has been derived in [2].

Ongoing Research / Outlook

The Lundgren forcing used to generate the HIT in this work results in a turbulent length scale which is a fixed fraction of the domain size. A new method has been developed to overcome this limitation using an efficient replacement of the filtered forcing known from literature. The new method will be used to study emulsification and segregation for varying turbulent length scales.

References and Links

[1] A. Begemann, T. Trummler, E. Trautner, J. Hasslberger, M. Klein. Can. J. Chem. Eng. , 100:3548-3561, 2022.
[2] T. Trummler, A. Begemann, E. Trautner, M. Klein. Phys. Fluids , 34:113324, 2022.
[3] O. Krzeczek, T. Trummler, E. Trautner, and M. Klein. . Energies, 16:3160, 2023.
[4] W. Aniszewski, T. Arrufat, M. Crialesi-Esposito, S. Dabiri, D. Fuster, Y. Ling, J. Lu, L. Malan, S. Pal, R. Scardovelli, et al., Comput. Phys. Commun. , 263, 107849, 2021.

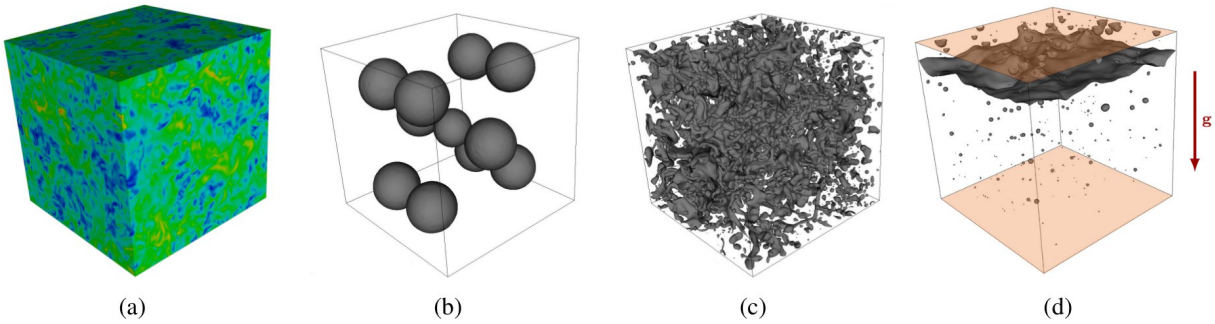

Figure 1: Simulation setup. (a) Single phase forcing, flow field visualized by the velocity magnitude, (b) initialization droplets with diameter $d_0 \approx 0.14L$, (c) turbulent emulsion at statistically stationary state, (d) segregation under gravitational acceleration.

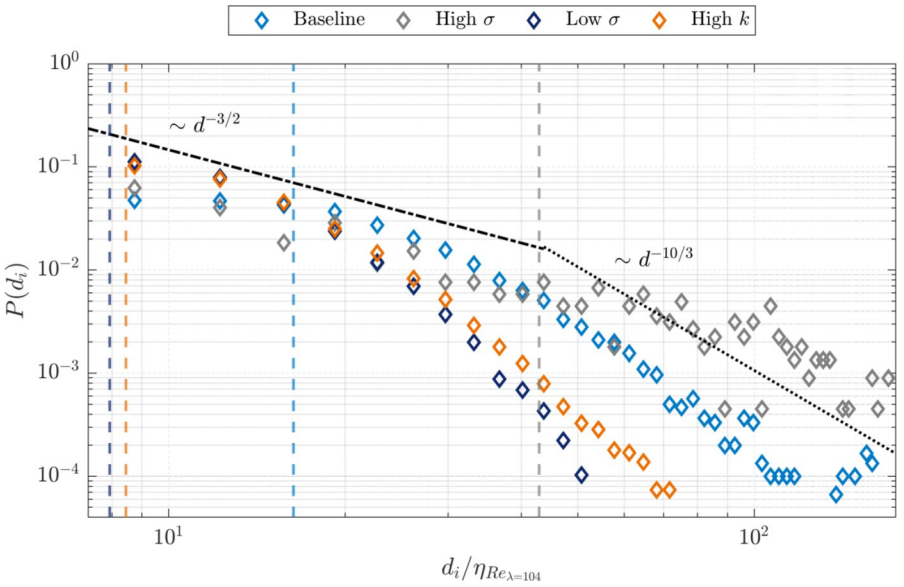

Figure 2: Droplet size distribution for four different cases. The data has been sampled once per eddy turn-over time and is averaged for 6 time instants each after reaching the statistically stationary state. The dashed vertical lines indicates the Hinze scale and the color corresponds to the respective case.

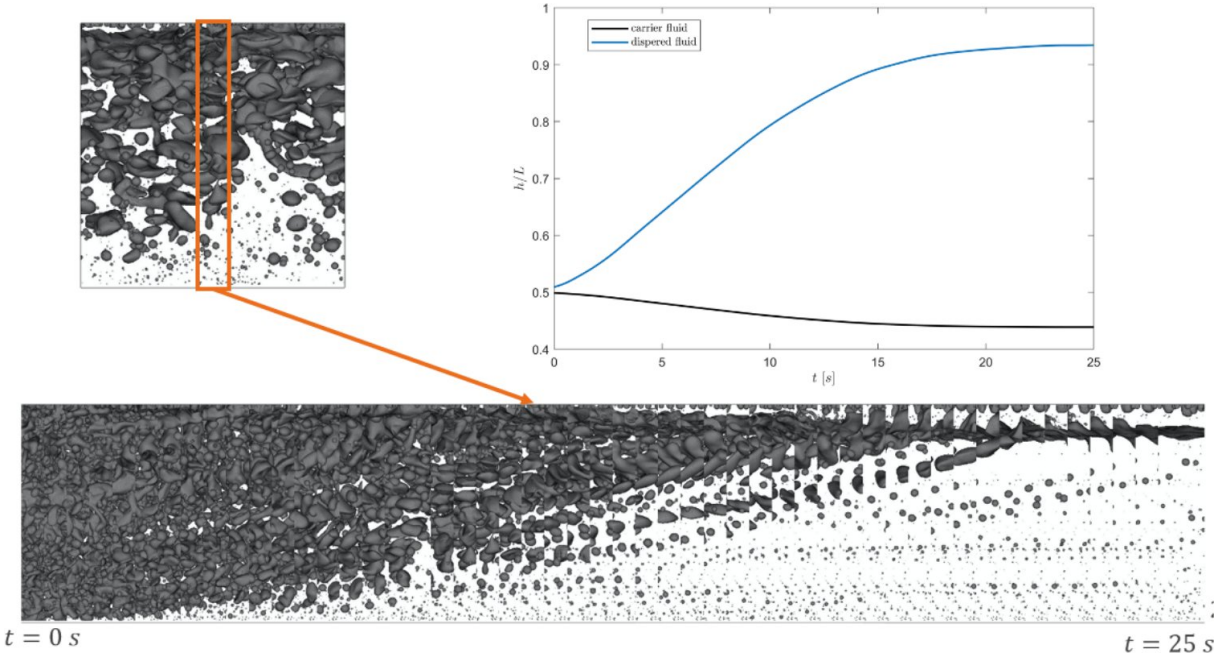

Figure 3: Visualization of the segregation process. Section of the simulation domain shown with a frame rate of 2 frames/s. In the top right panel: recorded evolution of the center of mass of the dispersed phase.

Roughness-induced laminar-turbulent transition for hypersonic vehicles

Research Institution
Technische Universität München

Principal Investigator
Christian Stemmer

Researcher
Friedrich Ulrich, Giuseppe Chiapparino

Project Partner
–

Funding
TUM IGSSE 13.09

SuperMUC Project ID: pn36bi

Introduction

For the design of sustainable hypersonic flight at Mach numbers above 5 or re-entry vehicles like space capsules or the Space Shuttle, heat management is one of the crucial aspects of engineering. Besides the application of insulating materials and their development to be used in the heat shield, the generation of the heat through fluid dynamics processes in the boundary layer along the vehicle has to be known. As experiments for these flight conditions are either extremely expensive or subject to large errors, numerical simulations can bridge the gap between the lab environment and the full-scale application.Laminar-turbulent transition at hypersonic speeds is one of the key feature to understand the physics of the flow at high speeds and high temperatures as the heat load on the vehicle can change more than 10-fold if the flow becomes turbulent compared to a laminar flow. In this project, we investigate through high-precision and high-fidelity numerical simulations including chemical reactions (dissociation) and non-equilibrium the influence of surface imperfections on the laminar-turbulent transition. The surface roughness either comes from holes for screws, gaps between heat-shield tiles or from the ablating material subject to burning at these extreme temperatures.

Results and Methods

Two flow configurations have been studied, namely the compression-corner (or ramp) geometry, representative of control surfaces or engine intakes of a hypersonic vehicle, and an Apollo-like capsule geometry. In both cases, surface roughness was modeled at specific locations (upstream of the compression corner and on the shoulder of the capsule, respectively) to study its influence on the downstream evolution of the high-speed boundary layer. The Direct Numerical Simulations are performed with the code NSMB, which allows to make full use of the parallel-computing capabilities of SuperMUC-NG. In the first configuration, the roughness shows a dramatic impact on the ramp-induced Shock-Wave/Boundary-Layer Interaction (SWBLI). The roughness induces the formation of streamwise vortices in its wake, which are convected downstream in the boundary layer. Once the flow separates due to the corner-induced shock and then reattaches on the ramp, the vortices are strongly amplified and generate a mushroom-shaped structure characterized by large velocity gradients (Figure 1, taken from [1]). Underneath this vortical structure, hot streaks appear at the wall, where the surface temperature reaches values much higher than those achieved without roughness. As known from the literature [2,3,4], strong roughness-induced velocity gradients can lead to flow instabilities thatcan initiate the transition process. Hence, a similar flow evolution can be expected in the reattached flow on the ramp, once freestream disturbances are injected in the baseflow. For the space capsule configuration, our research focused on a restricted domain specifically targeted at the shoulder region of the Apollo-like spacecraft. This region exhibits a high concentration of observed turbulent spots, making it a critical area for analysis. An example of such a restricted domain can be found in Fig. 2 from [5].It displays a random roughness patch and cross-flow-like vortex formed in the wake region of the roughness-patch. To explore the impact of surface roughness on flow transition, we generated a comprehensive database which consists of over 10,000 Direct Numerical Simulations (DNS) results. Each simulation featured unique, individually randomized roughness patches, crafted using a superposition of three distinct wave-numbers. This approach aimed at modeling surface imperfections, akin to those encountered during the ablation process of an actual re-entry flight. Our analysis delved into the crossflow-like vortices generated downstream of each roughness patch [6]. These vortices play an important role in flow transition, and understanding their formation process was paramount. Furthermore, we investigated the potential of machine learning (ML) in quantifying the strength of these crossflow-like vortices. We evaluated various ML algorithms, ultimately achieving the most promising results with a multi-layer perceptrons and convolutional neural networks. This integration of ML offers exciting possibilities for future research and real-time monitoring

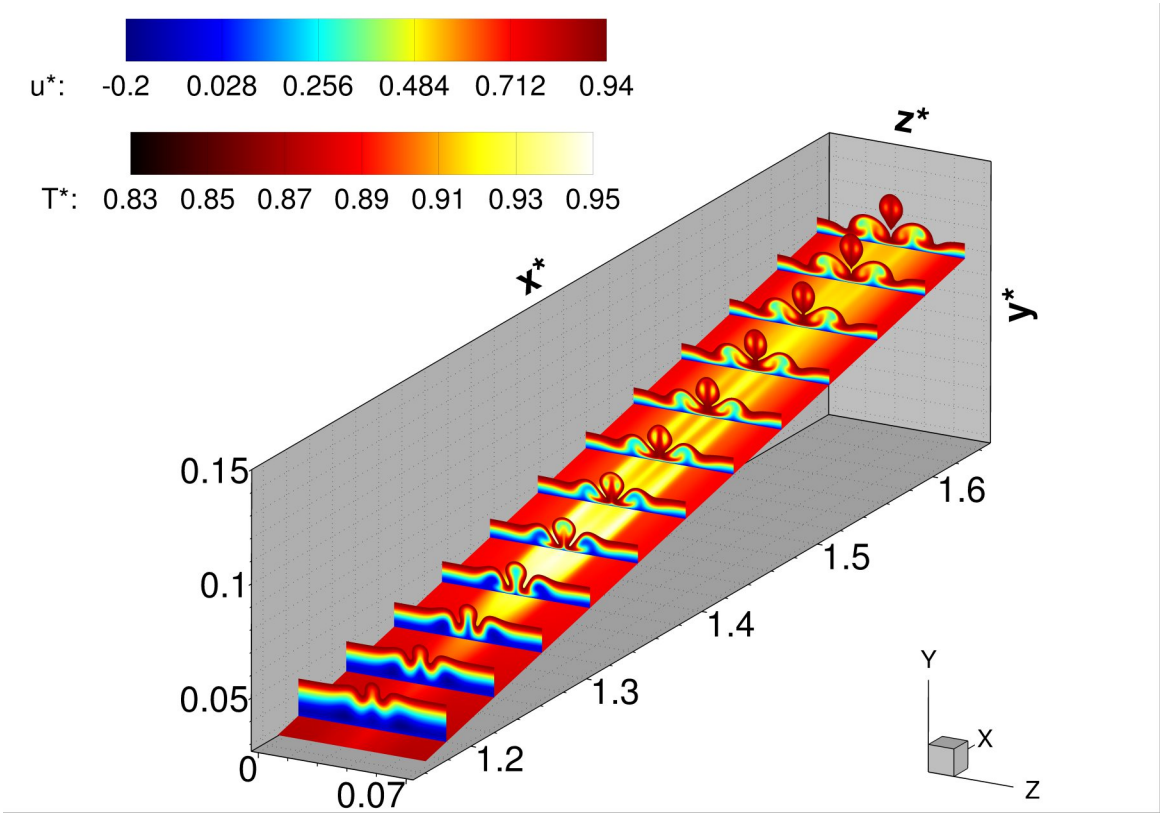

Figure1: Mushroom-shaped structure forming in the wake of the roughness as the flow reattaches on the ramp.

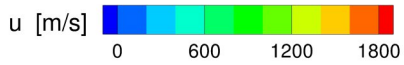

of re-entry conditions. The complete DNS database as well as our ML models can be downloaded and used from our GIT repository [7].

Ongoing Research / Outlook

On the ramp-induced SWBLI side, the current efforts focus on the unsteady response of the roughness-modified baseflow to different types of freestream disturbances prescribed upstream of the flow interaction with the corner shock. The study aims at characterizing the mechanisms, specifically associated with the presence of surface roughness, that lead to laminar-turbulent transition downstream of reattachment. For the capsule, our future work expands upon this investigation, exploring the influence of higher-complexity roughness patches on re-entry flow. We will incorporate a broader range of wavelengths, amplitudes, and phases within the random roughness model, striving for enhanced realism in representing ablated heatshields. High-resolution CFD simulations will delve deeper into the flow dynamics, offering refined insights into turbulence initiation and heat transfer. Additionally, we made our improved DNS database public [7], facilitating further research and model development. We want to find more partners in order to expand our database. Finally, we aim to advance our machine learning models using sophisticated algorithms to improve the accuracy and efficiency of crossflow-like vortex characterization. This comprehensive approach will contribute to a deeper understanding of the blunt body paradox, ultimately promoting safer and more predictable re-entry experiences.

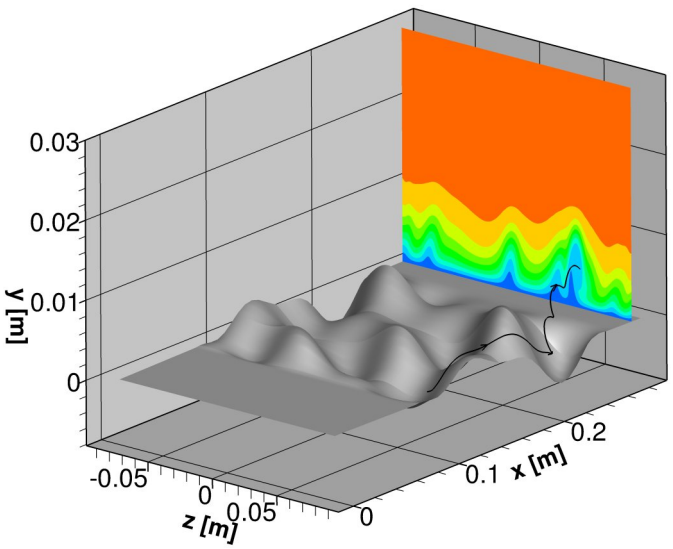

Figure 2: Restricted domain with a random distributed roughness patch and a cross-flow-like vortex in the wake.

References and Links

[1] G. Chiapparino, C. Stemmer, Int. J. Heat & Fluid Flow, 103:109193, 2023.
[2] P. Wassermann, M. Kloker. J. Fluid Mech. 456, 49-84, 2002.
[3] G. Groskopf et al. In: Proceedings of the Summer Program 2008 – Center for Turbulence Research, pp. 55-72.
[4] G. Groskopf, M. Kloker, O. Marxen. In: 7th IUTAM Symposium on Laminar-Turbulent Transition, pp. 171-176.
[5] F. Ulrich and C. Stemmer, CEAS Space Journal, 2023, 971988 (2023), doi: 10.1007/s12567-023-00510-2.
[6] A. Di Giovanni et al., J. Fluid Mech. (Vol. 856), 2018.
[7] https://gitlab.lrz.de/fulrich/dns-roughness-database-for-ai

Investigating the Safety of Constructions During Wind Events

Research Institution
[1]Technical University of Munich

Principal Investigator
Máté Péntek[1]

Researcher
Guillermo Martínez-López[1,2], Anoop Kodakkal[1], Philipp Bucher[1]

Project Partner
[2]Universitat Politècnica de València

Funding
"La Caixa" Foundation fellowship (ID 100010434, code LCF/BQ/EU20/11810069); AiF-ZIM ZF4066103MC8

SuperMUC Project ID: pn56ba

Introduction

Wind events represent a critical load state on various civil engineering structures, such as buildings, bridges and stadia. Their all-time safety is of general interest, may it be considerations related to the integrity of constructions or occupational comfort. High-performance computing (HPC) has enabled a new era of tools and methods, to assist and drive corresponding activities. Computational structural mechanics (CSM) represents the knowledge base anchored in numerics focusing primarily on structures. It helps deliver insights into the internal stress state and deformation, alongside many other details. Meanwhile, computational fluid dynamics (CFD) is the area of interest providing the means to simulate wind flow. Constructions often exhibit a strong interaction with the forcing term, in particular in the case of wind events. HPC sets the appropriate premise for investigating this intricate connection, technically termed as fluid-structure interaction (FSI). The scientific maturity of the proposed methods rapidly evolved during the last twenty years, with the trajectory being critically influenced by the development of appropriate hardware infrastructures. While the theoretical bases and early implementations have been available for more than half a century, detailed investigations applying new developments and extensive validation campaigns are only possible by endeavors such as the current project. It is key for these advances in applied science to pose the right questions, ensure due diligence during investigations and formalize a framework for HPC systems for upcoming activities to rely on. The computational resources available enable in-depth investigations of wind-related events, which were previously not possible or viable. Our work dedicated effort to improving methods used to investigate flutter, the phenomenon known as being the cause of collapse of the Tacoma Narrows bridge near Seattle. We also focused on wind-induced vibrations of tall buildings and their mitigation by added devices, such as tuned mass dampers, a solution also used in the case of 432 Park Avenue in New York. Here, the underlying physics is related to mechanical resonance, counteracted by an additional mass. The restoring force is often modeled assuming an analogy with a swinging pendulum and corresponding inertia. Lastly, the Olympic Stadium in Munich was investigated in detail during the project. This is a landmark construction, which has been influencing scientific and technical development since its preparatory phase more than fifty years ago. It still proves to be challenging, and dedicated computational methods were needed to ensure a realistic model from a physical point of view, which can be assessed robustly and reliably on this HPC infrastructure.

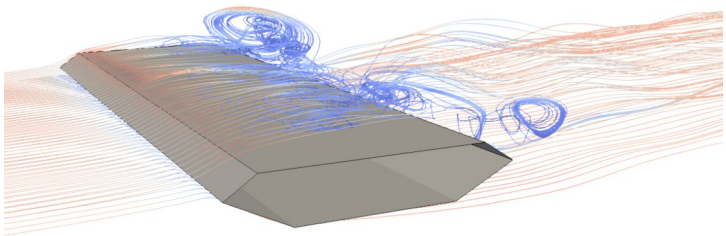

Figure 1: Streamlines around a bridge cross-section.

Results and Methods

Numerical developments were contributed to the community-driven open-source project Kratos Multiphysics [1], which provides a framework for building parallel numerical simulations. Modularity and extensibility for HPC systems are the stated objectives of the community. Licensing terms favor scientific and industrial collaborations alike. We were also involved in driving technical developments, such as performant input-output operations and formats, job profiling and balancing, bundling of investigations and improved scalability. Our activities resulted in substantially reducing the numerical effort for a certain class of investigations related to bridge flutter [2]. We also outlined a workflow [3] and provided use-cases for assessing and mitigating wind-induced resonance for buildings [4]. Further results are related to computational method development and analysis of stadia, motivated by the Olympic Stadium in Munich [5]. Figures 1, 2, and 3 provide some snapshots on the flavor of investigations for all these topics.

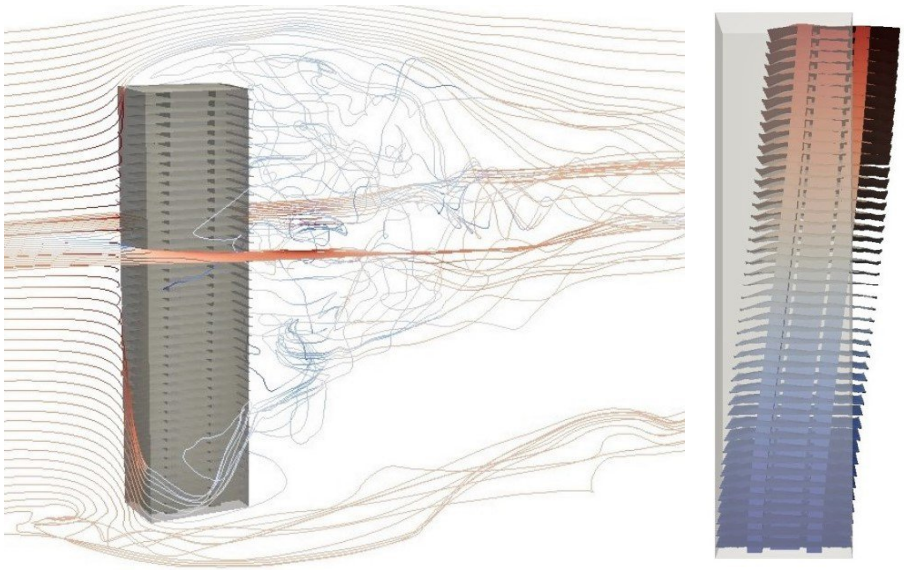

Figure 2: Wind flow patterns around a tall building using CFD (left) and structural deformations using CSM (right).

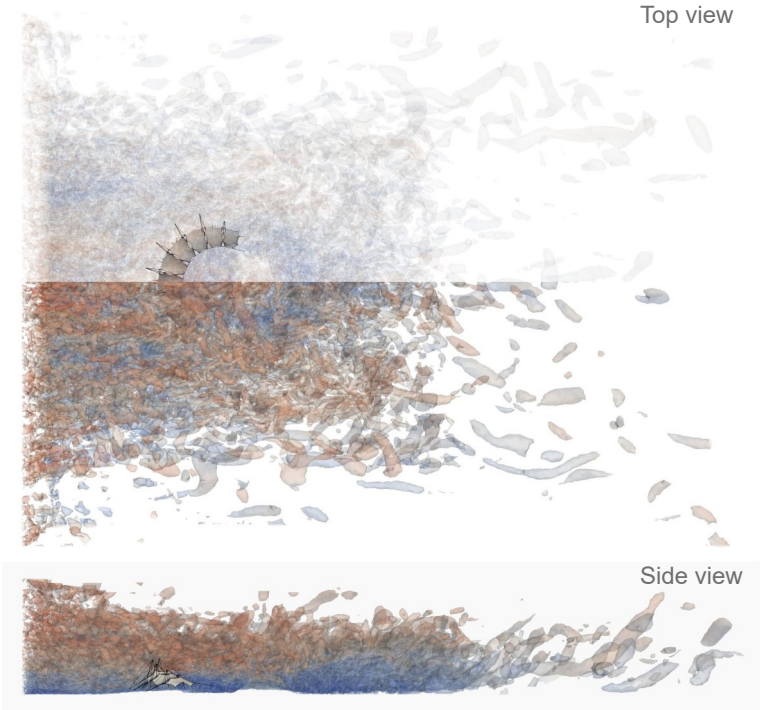

Figure 3: Top and side view of the vorticity contours around the Olympic Stadium in Munich, during a strong wind event (direction: left to right, corresponding to West/ Southwest).

Ongoing Research / Outlook

A subsequent project on SuperMUC-NG is already supporting further developments. These focus on additional measures against wind-induced structural motions. Adaptive bridge cross sections are investigated, as a promising alternative to delay the onset of flutter. Other innovative solutions may rely on the active change of shape for constructions. Considered wind events could include special types, such as tornadoes and down-bursts. These efforts contribute to the safety and resilience of structures in the built environment – as well as the wellbeing of their occupants – in case of natural hazards. The vast numerical capabilities will be combined with experimental possibilities provided by two ERIES projects [6]. Subsequent activities will be carried out in collaboration with the WindEEE Dome at the University of Western Ontario (Canada) and the GS-WinDyn at the University of Genoa (Italy).

References and Links

[1] https://github.com/KratosMultiphysics/Kratos
[2] G. Martínez-López et al., Reducing the computational cost of forced-motion simulations to obtain flutter derivatives, 16th International Conference on Wind Engineering, 2024, Florence, Italy.
[3] M. Péntek et al., A multiply-partitioned methodology for fully-coupled computational wind-structure interaction simulation considering the inclusion of arbitrary added mass dampers, J. Wind Eng. Ind. Aerod. 177 (2018) 117135, DOI: https://doi.org/10.1016/j.jweia.2018.03.010
[4] M. Péntek et al., Investigating the Vibration Mitigation Efficiency of Tuned Sloshing Dampers Using a Two-Fluid CFD Approach, Appl. Sci. 2022, 12(14), 7033 , DOI: https://doi.org/10.3390/app12147033
[5] M. Péntek et al., Numerical modeling and simulation of lightweight structures using the example of the Olympic Stadium in Munich on its 50th anniversary, Proceedings of 33. Forum Bauinformatik, 2022. DOI: https://doi.org/10.14459/2022md1686600
[6] https://eries.eu/

Turbulence Inside the Tesla Turbine

Research Institution
Universität der Bundeswehr München

Principal Investigator
Stefan Klingl, Stefan Lecheler, Michael Pfitzner

Researcher
–

Project Partner
–

SuperMUC Project ID: pn29ki

Introduction

Invented and patented by Nikola Tesla in 1913 [1], the Tesla turbine is a unique type of turbomachine that features a rotor made from stacked disks with small gaps in between, as shown in Figure 1. It is driven by a fluid that enters the disk gaps at the outer circumference and travels on a spiral path radially inward towards central exit openings. Figure 2 visualizes the path of the fluid. Rotor torque is generated purely from friction between the working fluid and the rotor disk surfaces.

To make Tesla turbines more accessible to engineers, researchers in the past developed and validated methods for predicting turbine performance mathematically. An important aspect of this is beforehand knowledge whether the flow inside the disk gap is going to be laminar or turbulent, since modelling approaches usually target only one of those. So far, a stability boundary that separates laminar and turbulent flow regimes has been mapped out through experiments [2,3] and mathematically through linear stability analysis [4,5]. The latter searches frequencies of the flow that are amplified in time and space and possibly trigger transition to turbulence. Figure 3 shows the stability boundary from linear theory as function of Reynolds number Re* and Taylor number Ta* [5]. These can be interpreted as mass flow and rotational speed parameters respectively. Linear theory and experimental results so far do not fully agree with each other, because transition in experiments was observed earlier than predicted by theory, slightly into the laminar region of the chart. In order to verify results from linear theory and clear up the connection between stability theory and experiment, the present study describes direct numerical simulation of the turbine rotor flow.

Results and Methods

The computations are conducted without any turbulence modelling and thus accurately simulate oscillations and other unsteady features of the flow. In the turbulent case, this is computationally very expensive because the mesh has to be fine enough to resolve even the tiniest features of the flow field. In addition to the high resolution in time and space, the simulation has to run for a long period of simulation time, so that the turbulent flow field can be analyzed statistically. To cut down on computational cost, simulations are restricted to a periodic sector of the ring-shaped space between two disks. In the inlet region of the domain, the flow is artificially perturbed by a volume force that causes velocity and pressure oscillations in the flow field downstream. In the laminar region of the stability chart, it is expected that all oscillations are dampened and the flow returns to a steady, laminar state. In the transitional and turbulent regions, it is expected for certain perturbation wavelengths to be amplified further and further until the oscillations break down into a chaotic, turbulent flow field. For the turbulent operating conditions, this process can be observed in Figures 4 and 5.

Indeed, during the simulations with laminar operating conditions, perturbations were dampened across all wavelengths. For some wavelengths however, the dampening rate turned out to be so low, that unsteady oscillations remained in the flow field throughout the whole domain. Based on this observation, it is plausible, that the unsteady, seemingly transitional flow that was observed at these conditions in experiments is a consequence of remaining flow perturbations that originated before or during entry of the fluid into the rotor gap. For the transitional case, a small amplification rate was found for certain perturbation wavelengths in accordance with predictions from linear stability analysis. For a more detailed look into the results, please see reference [6].

Ongoing Research / Outlook

Future research could clear this up further and target other points in the stability plane. In particular, it would be interesting to investigate the role of rotation in the transition mechanism. Other than that, the simulation domain could be extended with the rotor inlet edge and possibly a nozzle to produce a more realistic flow field without the need for an artificial perturbation force.

References and Links

[1] Nikola Tesla, Turbine, Patent No. US1061206, New York, 1913.
[2] L. L. Pater et al., J. Fluids Eng., 96, no. 1, 29–34, 03 1974.
[3] G. Dibelius, D. Nendl, Reibungsturbomaschinen. Forschungsberichte des Landes Nordrhein-Westfalen, VS Verlag für Sozialwissenschaften, Wiesbaden, 1973.
[4] S. Klingl et al., Eur. J. Mech. B Fluids, 84, 455–469, 2020.
[5] S. Klingl et al., Eur. J. Mech. B Fluids, 91, 226–232, 2022.
[6] S. Klingl et al., Eur. J. Mech. B Fluids, 105, 119-137, 2024.

Figure 1: Sample Tesla turbine rotor with multiple disk gaps and outlet openings in the center.

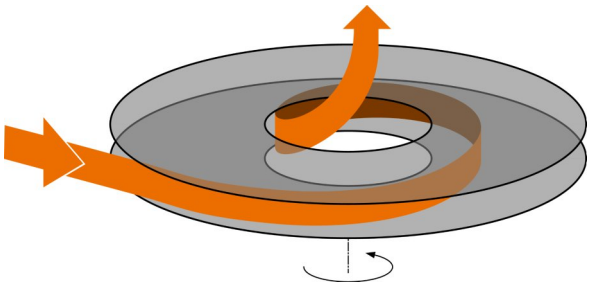

Figure 2: Simplified model of a single disk gap.

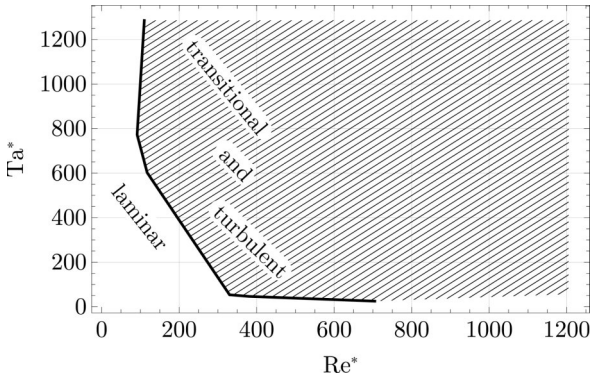

Figure 3: Stability regions from linear stability analysis.

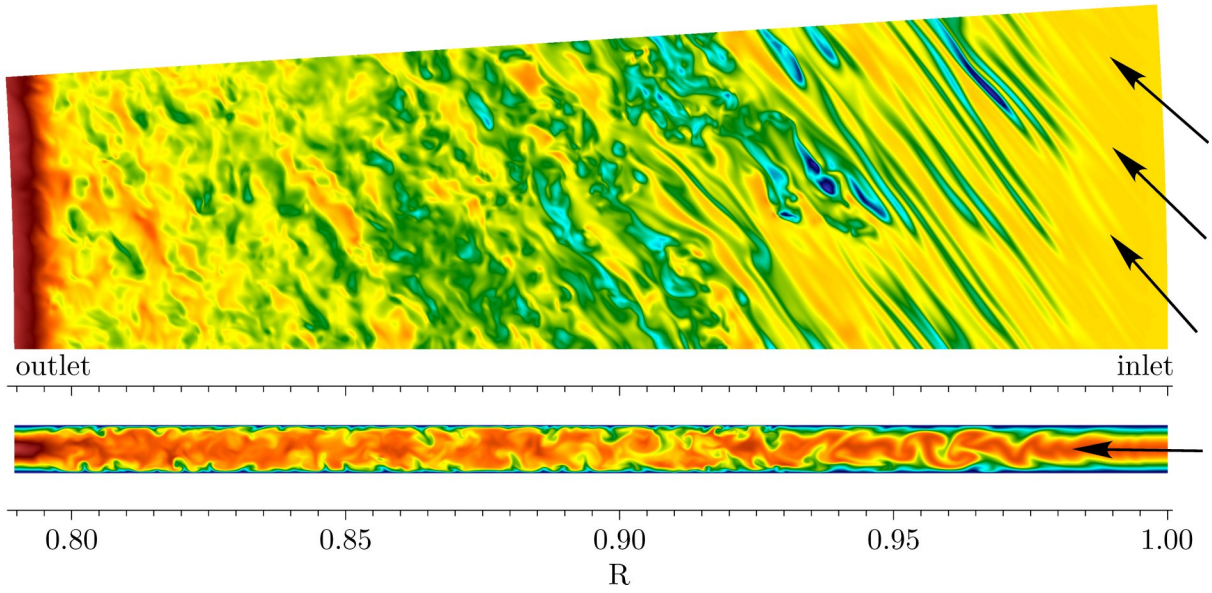

Figure 4: Velocity magnitude parallel (top) and normal (bottom) to the disk surfaces for the turbulent conditions.

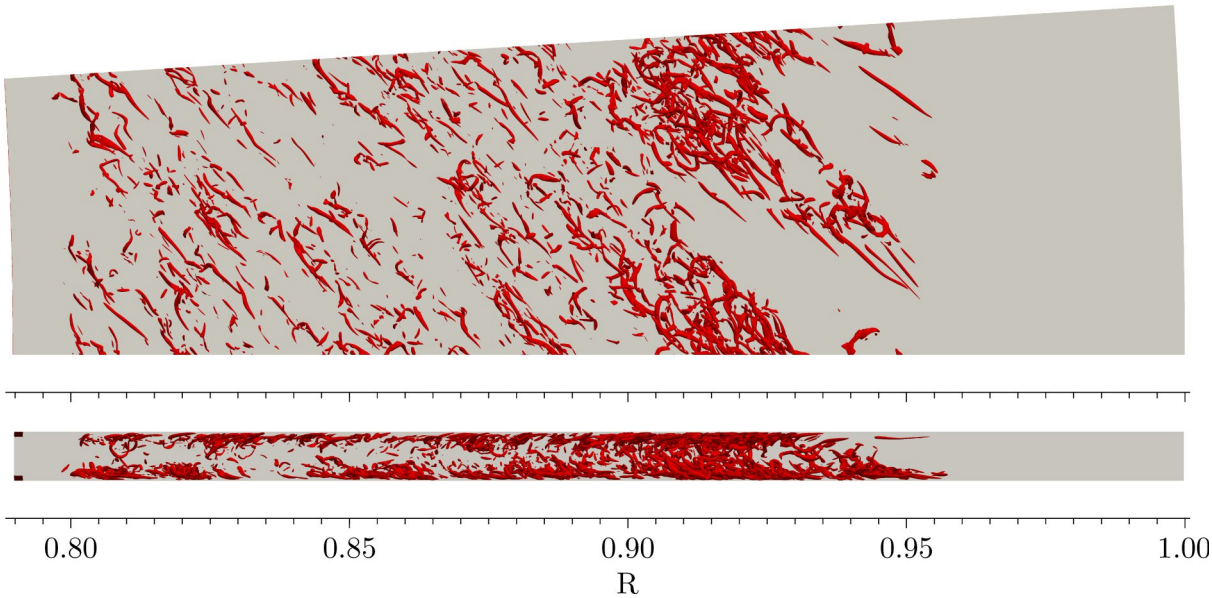

Figure 5: Visualization of turbulent structures in the velocity field from Figure 4, using the λ_2-criterion.

Particle Behavior in Oscillatory Fluid Flow

Research Institution
[1]Institute of Urban and Industrial Water Management, Technical University Dresden
[2]Leichtweiß-Institute for Hydraulic Engineering and Water Resources, Technical University Braunschweig

Principal Investigator
Fabian Kleischmann[1,2]

Researcher
Paolo Luzzatto-Fegiz[3], Eckart Meiburg[3] and Bernhard Vowinckel[1,2]

Project Partner
[3]University of California, Santa Barbara

Funding
DFG grant VO2413/2-1

SuperMUC Project ID: pn25wa

Introduction

Oscillating flows, which involve the rhythmic back-and-forth movement of fluid, hold immense potential across various industries and medical fields. For instance, they are used in water treatment plants to clear out particles and in diesel engines to enhance combustion efficiency. Moreover, oscillating flows are employed in medicine for drug delivery through non-invasive methods. Understanding how particles interact with these oscillating fluids is crucial for predicting their behavior accurately and developing precise applications. Researchers have extensively studied the behavior of individual particles freely moving in oscillating fluids, as well as the resulting flow patterns, through analytical, experimental, and numerical investigations. Building on this existing knowledge, our research delves deeper by examining how two mobile particles interact in such environments and how their behaviors are intertwined with the accompanying flow patterns. Several factors come into play in these interactions. The properties of both the particles and the fluid are key, including particle size, distance between the particles, fluid viscosity, and their respective densities. Additionally, the characteristics of the oscillation, such as its amplitude and frequency, play crucial roles in shaping the behavior of particles and the resulting flow dynamics. By exploring these factors, we aim to enhance our understanding of fluid-particle interactions in oscillating environments, paving the way for more precise and efficient applications across various domains.

Results and Methods

In order to conduct these investigations, we performed a set of particle-resolved direct numerical simulations. This approach involves computations of both particle and fluid motion with high fidelity by geometrically resolving the flow around the particles and considering their interaction by coupling both phases. The two particles were placed in a cubic domain with a size of 10 times the particle diameter. These particles were allowed to move freely in response to a horizontal oscillation applied to the domain. The particles were not affected by the walls of the domain, meaning they could move around without hindrance. Gravity was neglected in the considered setups to focus solely on the fluid-particle interactions. We varied the frequencies of the oscillations from 1 to 30 Hz, and the amplitude of the oscillation was set at 10% of the diameter of the particles. The particles were initially arranged horizontally, aligning with the direction of the oscillation. The systematic study involves different initial distances between the particles, ranging from 25 to 300% of the particle diameter, and three different particle densities were examined. To perform these simulations, we used specialized numerical methods and algorithms implemented in an in-house C-based computational code. This code, which is freely editable and modifiable, runs on standard CPUs and is parallelized using MPI functions. It relies on external libraries such as FFTW and HDF5 [1,2]. Running one simulation required a significant amount of computational power, utilizing 768 cores and lasting approximately 5 hours. In total, we conducted 420 simulations to thoroughly explore the parameter range mentioned earlier. The findings from our numerical simulations reveal some interesting patterns. When the oscillation frequencies are low, the particles tend to move closer to each other, which would lead to an aggregation of the particles. However, as the frequencies increase, they start pushing away from each other, resulting in a segregation. Another intriguing observation is that when the particles are initially placed closer together, they tend to attract each other, while if they are further apart, they tend to repel. Moreover, we discovered that the density of the particles affects the intensity of the interaction but not the general outcome in terms of attraction or repulsion. The behaviors of these particles are also reflected in the flow patterns they create, as depicted in Figures 1 and 2. These flow patterns show how the fluid moves around the particles over the course of one oscillation period. In the figures, the flow directions are indicated by the black arrows, while the strength of the rotations, known as vorticity, is represented by the blue and red coloring. Blue indicates clockwise rotations, while red signifies anticlockwise rotations. In Figure 1, which corresponds to a scenario with a lower oscillation frequency (5 Hz), we observe flow patterns for which attractions between the

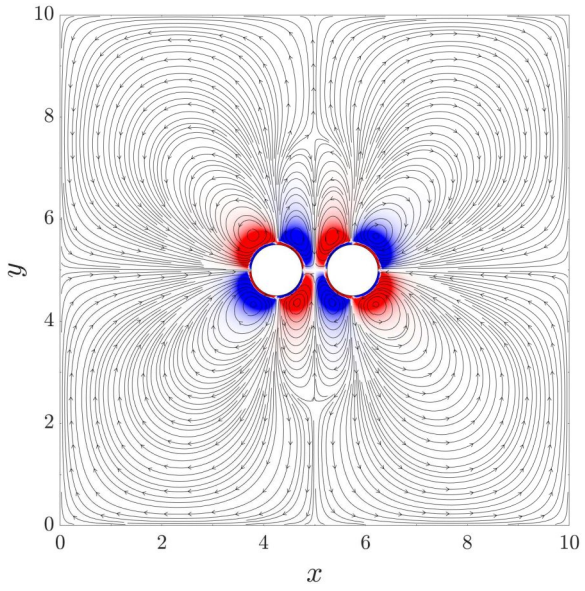

Figure 2: Flow pattern and vorticity for a setup with a frequency of 15 Hz.

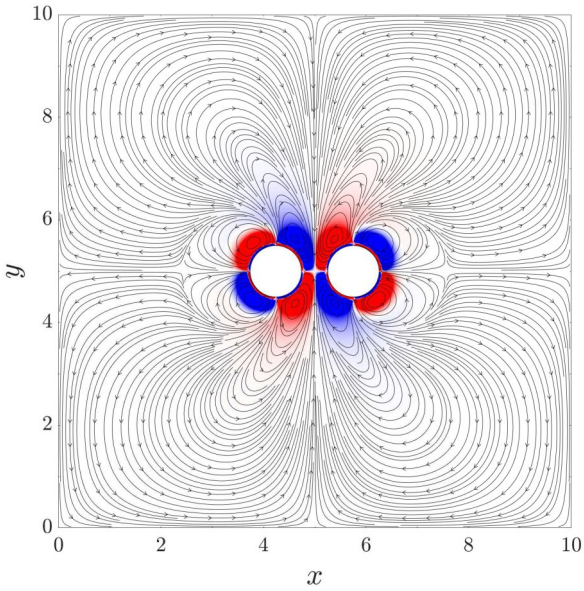

Figure 1: Flow pattern and vorticity for a setup with a frequency of 5 Hz.

particles occurs. Conversely, Figure 2 illustrates a situation with a higher frequency of oscillation (15 Hz), resulting in repulsion between the particles. In both cases, each particle is surrounded by four vortexes, represented by the structures of the black arrows, varying in size. In the case of attraction (Figure 1), the outer vortex structures dominate, indicating that the fluid tends to push the particles together. However, in the repulsive setup (Figure 2), the inner vortexes become the dominating flow feature. These vortices push fluid into the gap that separates the particles, which exerts forces that let particles drift away from each other. These flow patterns provide visual insights into how the fluid interacts with the particles, influencing their behavior in the process. The findings of our study highlight the relationship between different oscillation properties and the resulting behaviors of particles, particularly in terms of whether they are drawn together or pushed apart. By altering factors like the frequency of the oscillation and the initial distance of the particles, we observed distinct changes in how particles interacted with each other. This comprehensive investigation has shed new light on the complexities of particle dynamics in oscillating environments. The study also provides a means to manipulate oscillatory flows to guide the particle interaction. The entire study was recently published in the Journal of Fluid Mechanics [3].

Ongoing Research / Outlook

Expanding upon the insights gained from our presented investigation involving two particles, we started to examine the dynamics of systems containing a significantly larger number of particles. Our current studies involve simulations encompassing particle populations ranging from 600 to 2,500 individual entities. By scaling up the complexity of our simulations in this manner, we aim to explore how the intricate interplay between oscillations and particle interactions influences the collective behavior of these systems. Through these ongoing investigations, we seek to uncover how the oscillatory motion of the fluid environment impacts the aggregation behavior of a large number of particles. By analyzing the trajectories and interactions of each individual particle within these densely populated systems, we strive to elucidate emergent patterns and phenomena that may arise as a result of the oscillatory forces at play. This expanded scope of our research endeavors represents a significant step forward in our quest to unravel the fundamental principles governing the behavior of particle-laden fluids under oscillatory conditions. We anticipate that the insights gleaned from these endeavors will contribute to a deeper understanding of complex fluid.

References and Links

[1] E. Biegert et al., J. Comp. Phys. 340 (2017) 105-127.
[2] B. Vowinckel et al., J. Fluid Mech. 858 (2019) 208-223.
[3] F. Kleischmann et al., J. Fluid Mech. 984 (2024) A57.

Study of turbulent flame properties for ammonia blended with DME by Direct Numerical Simulation

Research Institution
Laboratory of Fluid Dynamics and Technical Flows, Otto-von-Guericke-University Magdeburg

Principal Investigator
Dominique Thévenin

Researcher
Wei Guan, Abouelmagd Abdelsamie, Cheng Chi

Project Partner
—

Funding
SPP2419 (grant number 523880888), DFG (grant number 512043261)

SuperMUC Project ID: pn36gi

Introduction

Ammonia (NH_3) as an alternative fuel has recently attracted growing interest to reduce green-house emissions in combustion applications, being carbon-free and a high-density hydrogen energy carrier. However, the reactivity of NH_3 is comparatively poor, yielding low burning velocity and high auto-ignition energy. Co-firing with dimethyl ether (DME) in ammonia combustion is an outstanding route to shorten ignition delay time and increase flame speeds significantly. Leveraging the computational power of modern supercomputers, our project employs Direct Numerical Simulations (DNS) to delve into the turbulent flame characteristics and turbulence-chemistry interactions of NH_3/DME flames, utilizing our in-house low-Mach number combustion solver, DINO. In this research, a reduced kinetic mechanism for NH_3/DME blend fuels made up of 30 species, 15 species with QSS assumption and 294 reactions was first developed, improving the computational speed while ensuring calculation accuracy of DNS simulations. Additionally, a modified bandwidth-filtered forcing method in physical space is proposed to cope with turbulence decay and maintain the prescribed turbulence intensity in the fresh gases ahead of the flame. The aim of this project is to improve our understanding of flame propagation and structure of turbulent premixed NH_3/DME/air combustion. The reduced mechanism has been validated extensively against experimental data and numerical predictions obtained using other full kinetic mechanisms. Turbulent flame speed statistics are investigated in terms of global fuel consumption rates and local flame displacement speed, followed by an analysis of the turbulent flame characteristics. Our findings reveal intriguing insights, particularly concerning the effects on different turbulent combustion regimes for NH_3/DME/air flames propagating into forced turbulent flows. The main turbulent parameter influencing the turbulent flame speed is the size of the most energy-containing eddies rather than the Karlovitz number. This work significantly contributes to understanding turbulence-chemistry interactions for the utilization of NH_3/DME fuels in future engines.

Results and Methods

A DNS of a NH_3/DME/air premixed flame is carried out in a three-dimensional box using the in-house low-Mach number combustion solver DINO [2]. A sketch of the DNS configuration is depicted in Figure 1. Inlet and outlet boundaries are specified in the direction of mean flame propagation (x) while all other boundaries are set to be periodic. An isotropic turbulence velocity field is generated by a Passot-Pouquet analytical spectrum and then superimposed on the initial velocity field. A modified bandwidth-filtered forcing method in physical space is applied in the fresh gases ahead of the flame, in the region where the temperature-based progress variable is less than 0.01. To qualitatively inspect the flame structure, Instantaneous fields of heat release rate, temperature, and mass fractions of NH_3, DME, NO, NO_2, CO_2, and H_2O_2 through the central vertical plane of the DNS domain are shown in Figure 2. The mass fraction of NH_3 decreases from reactant to product and approaches 0 at the right boundary (i.e., very small amount of unburnt NH_3). The mass fraction of DME is also found negligible on the product side, which shows an even earlier oxidation than NH_3 [3]. The pollutants NO and CO_2 dominate in this NH_3/DME flame and show peak mass fractions near the flame front. Comparison of A and B indicates that the flame brush of cases B has far more wrinkles due to larger and more intensive small-scale eddies. A comparison of cases A, C, and D, which have different fluctuation velocities, shows a more blurred and thickened flame front in case D, which is mainly owing to an increased corrugation in the upstream region of the flame. In summary, the integral length scale is crucial in determining the level of wrinkling, and high Ka number leads to a wider and more disturbed preheated zone.

Ongoing Research / Outlook

In the coming phases of our research, we aspire to elevate our numerical procedure to simulate an even broader turbulent levels for NH_3/DME flames with increased accuracy and faster calculation speed.A

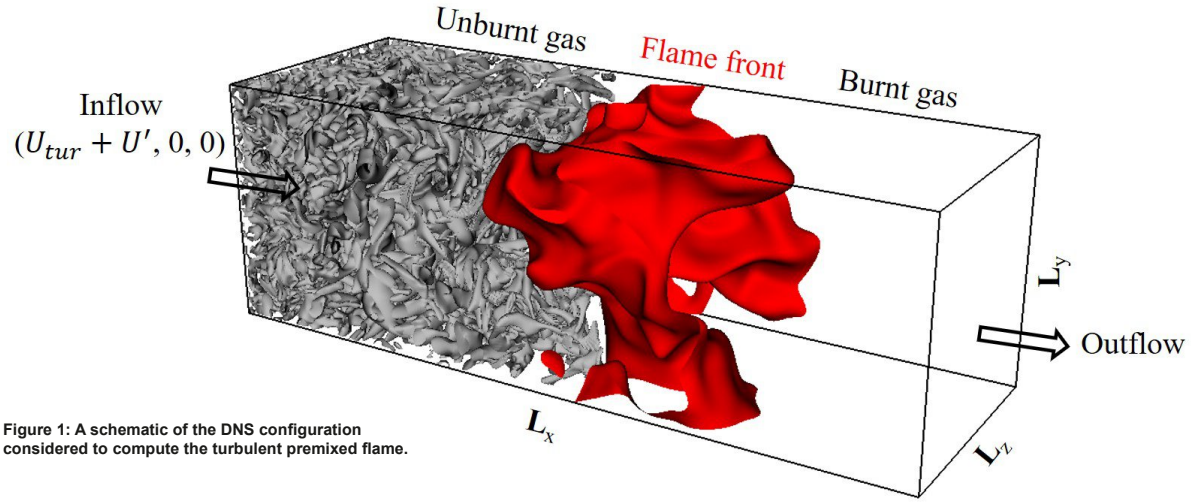

Figure 1: A schematic of the DNS configuration considered to compute the turbulent premixed flame.

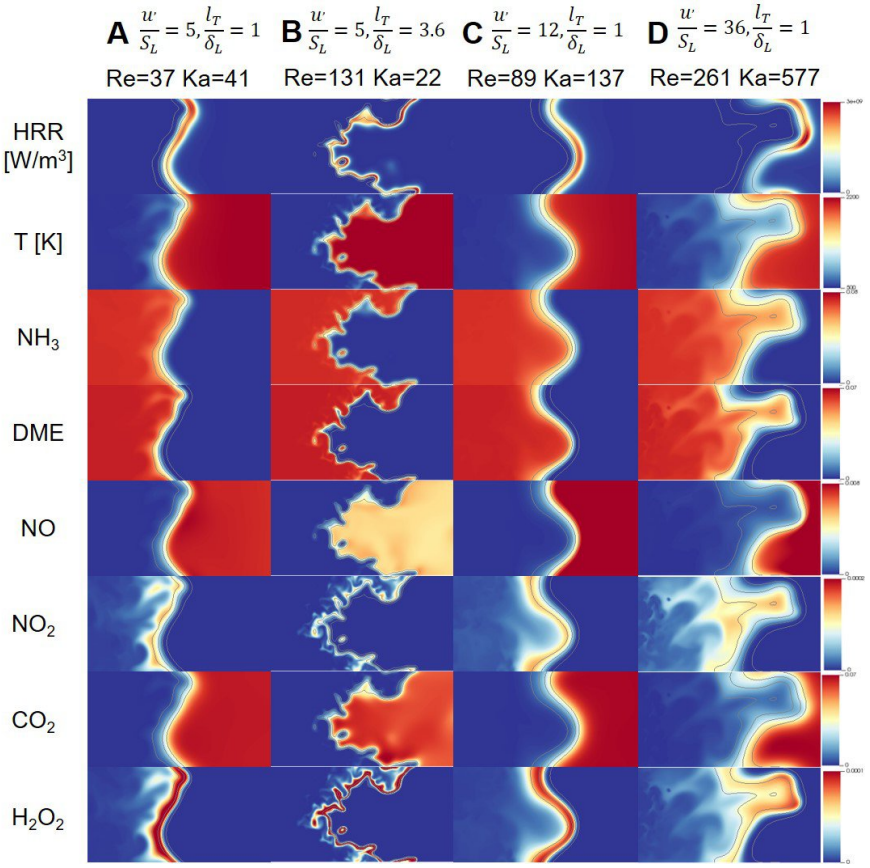

Figure 2: Two-dimensional slices of heat release rate, temperature, and mass fractions of NH_3, DME, NO, NO_2, CO_2, and H_2O_2 (from top to bottom) for different Reynolds and Karlovitz numbers (from left to right).

flamelet-based neural network will be used to store the tables for a reduced manifold modeling of the thermo-chemical system, reducing considerably CPU cost and memory storage [4]. Furthermore, our future study will extend to encompass numerical simulations involving real burners for NH_3/DME, such as bluff-body burner. This expansion is motivated by the need to evaluate the effectiveness of existing and emerging NH_3/DME blend fuel to reduce fossil fuel dependence for alleviating energy crisis. By exploring pollutant emissions and turbulent-chemistry interactions, we aspire to assess device stability, curb emissions, and explore opportunities for optimization. In summary, our future research directions aim to not only enhance the accuracy and scope of our simulations but also to explore pollutant emissions and flame characteristics for other process conditions.

References and Links

[1] https://www.lss.ovgu.de/lss/en/
[2] A. Abdelsamie et al., Computers & Fluids 131:123-141 (2016).
[3] W. Guan et al., Combustion and Flame 257:113002 (2023).
[4] C. Chi et al., Combustion and Flame 245:112325 (2022).

Uncertainty Quantification of Buoyancy-Induced Mixing Processes

Research Institution
[1]Numerical Methods in Aerospace Engineering, University of the Bundeswehr Munich

Principal Investigator
Philipp J. Wenig[1]

Researcher
Ruiyun Ji[1,2], Stephan Kelm[2] and Markus Klein[1]

Project Partner
[2]Institute of Energy and Climate Research (IEK-6), Forschungszentrum Jülich GmbH

Funding
BMUV (project no. 1501595), DAAD

SuperMUC Project ID: pn29ce

Introduction

The primary scientific and technical objective of the present project is to establish efficient methodologies for quantifying the propagation of uncertainties in Computational Fluid Dynamics (CFD) simulations and the further development of existing methods to optimize the computational effort. Imprecise knowledge of physical properties, model parameters along with initial and boundary conditions leads to uncertainties in the results of numerical simulations. In some applications, where safety issues are of great importance for the engineering design such as in plant and reactor design, the consideration of uncertainties is essential. Within this project, Uncertainty Quantification (UQ) of an application-oriented CFD validation case was conducted, which investigates buoyancy-induced mixing processes between two miscible fluids within a vessel of height $h \approx 10m$ and volume $V \approx 60m^3$. Various uncertainties in initial and boundary conditions, as well as in the thermo-physical properties affect the CFD model results. The impact of these uncertainties on responses needs to be quantified. Therefore, by using a generic test case, different methods were initially developed and qualified as suitable for the application to engineering applications [1,2,3]. Stochastic spectral methods, such as Polynomial Chaos Expansion (PCE) and Karhunen-Loève Expansion (KLE), demonstrated great potential as viable techniques for UQ of technical-scale CFD computations.

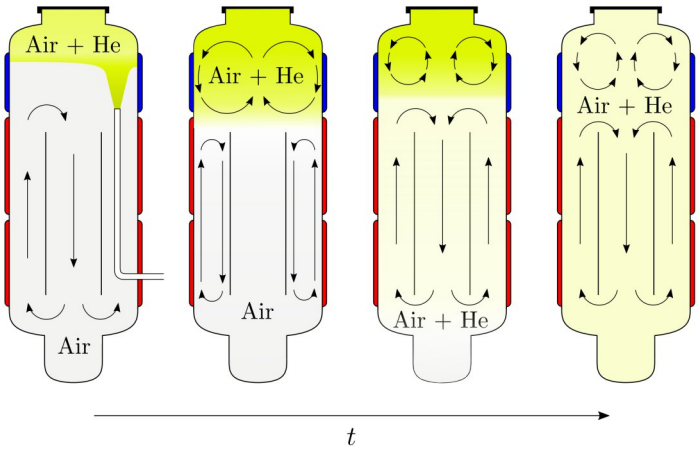

Figure 1: Consecutive phases of the THAI TH32 experiment.

Results and Methods

Investigation of flow and transport phenomena within a unique, technical-scale experimental facility was conducted using Unsteady Reynolds-Averaged Navier-Stokes Simulations (URANS) and Large Eddy Simulations (LES). The experimental facility refers to the THAI (Thermal-hydraulics, Hydrogen, Aerosols and Iodine) test containment [4]. The experiment comprises multiple consecutive phases, which are shown in Fig. 1 and mimic the release and dispersion of hydrogen in a nuclear reactor containment. Initially, air circulates in the vessel as a result of the integrated cooling and heating system. Subsequently, helium is injected and the circulation is suppressed. Following the completion of the injection process, a stable stratification forms and the erosion phase of the helium layer begins. Once the circulation establishes again, the stratification is eroded rapidly. The CFD simulation of the THAI TH32 experiment requires the definition of many input parameters. However, the definition of initial and boundary conditions is subject to uncertainties in the present setup, since measurement errors in the experiment occur or a lack of required input variables exists. Therefore, the UQ analyses comprised uncertainties in initial structural temperatures, heat transfer coefficients at the outer walls of the vessel and the helium mass flow rate during injection. The set of equations, which govern the buoyancy-driven mixing process within the test vessel, is solved using the containmentFOAM solver [5], a tailored solver and model library based on OpenFOAM. Utilizing the Dakota software framework, global sensitivity analyses and UQ analyses were conducted based on Stroud-cubature, linear regression and sparse grid approaches. The UQ studies required a large number of simulation runs of the resource-intensive URANS and LES of the presented

technical-scale application. The parallelization capability and high computing performance of the SuperMUC-NG were crucial for the success of this research endeavor.

In the following results from UQ studies based on URANS are presented. The underlying stochastic models were derived through linear regression. Figure 2 shows a comprehensive stochastic representation of selected response quantities, which include a local concentration in the test vessel, a local temperature and a local gas velocity. These quantities provide information about the progress of the mixing process, the heat transport and the convection mechanisms, respectively. The plots show the probability density function (PDF) and discrete levels of the cumulative distribution function (CDF) at different points in time. The CDF levels are visualized with colors according to a colormap.

From Fig. 2a one can see that the local helium concentration increases at the beginning, which corresponds to the injection of helium. After the injection phase, the concentration gradually decreases with a narrow uncertainty band. In this phase the mixing is mainly driven by diffusive mass transport. At a certain point in time, the erosion phase occurs, and the helium cloud is eroded. This leads to a steeper decrease and ultimately results in the accumulation of probability at the equilibrium concentration. However, the initiation of the erosion is subject to uncertainty and therefore takes place at different points in time, which leads to a wide uncertainty band. The UQ results for a vessel temperature are shown in Fig. 2b. The initial uncertainty up to $t = 0s$ corresponds to the predefined uncertainty in the structural temperature. Afterwards the vessel temperature decreases through the injected colder helium gas and the suppression of the circulation flow. During the mixing, the gas heats up, consequently raising the temperature of the vessel. The results regarding a local velocity are shown in Fig. 2c. The initial phase is characterized by uncertainty due to turbulent fluctuations. After the initial phase, the circulation flow is suppressed and the velocity decreases. As the mixing proceeds, the circulation flow is reestablished, and the velocity increases. The uncertainty of the velocity at this stage is mainly caused through the uncertainty of the onset of erosion and by turbulent fluctuations. In summary, the analyses have shown that Stochastic Spectral Methods are highly efficient for approximating responses. Additionally, they enable a comprehensive reliability and risk assessment of engineering applications at an early stage of the design process.

Ongoing Research / Outlook

A High-Dimensional Model Representation (HDMR) technique for conducting efficient UQ studies in an industrial context is currently being developed based on data collected from computations on the SuperMUC-NG.

References and Links

[1] P.J. Wenig et al., FTaC (2023). doi:10.1007/s10494-023-00467-6
[2] P.J. Wenig et al., NED (2023). doi:10.1016/j.nucengdes.2023.112317
[3] P.J. Wenig et al., Fluids (2021). doi:10.3390/fluids6040161
[4] M. Freitag et al., in: NURETH-19 (2022).
[5] S. Kelm et al., Fluids (2021). doi:10.3390/fluids6030100

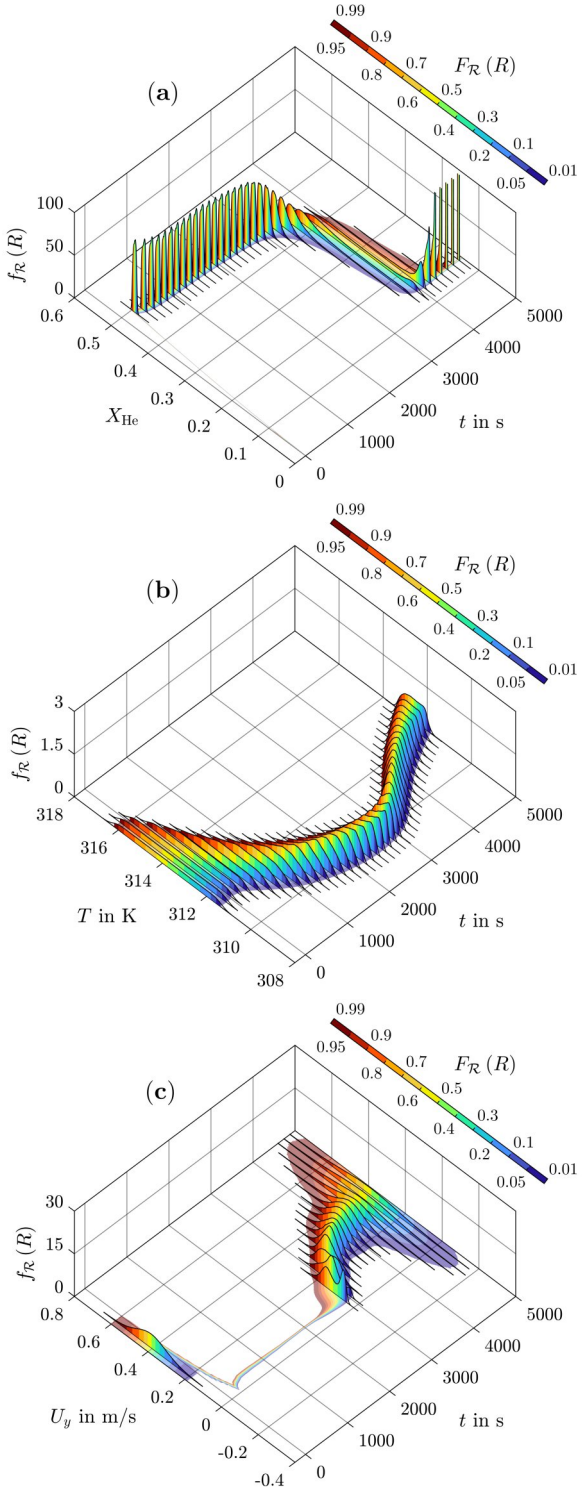

Figure 2: Representation of stochastic processes using PDFs $f_R(R)$ and discrete levels of CDFs $F_R(R)$: (a) local concentration (b) local temperature and (c) local velocity.

Unsteady propeller-wing interactions and wake-tailplane interactions

Research Institution
University of Stuttgart, Institute of Aerodynamics and Gas Dynamics

Principal Investigator
Thorsten Lutz

Researcher
Mário Firnhaber Beckers, Johannes Kleinert

Project Partner
–

SuperMUC Project ID: pn98wo

Introduction

A (partial) electrification of the aircraft powertrain allows for greater flexibility regarding the position of the propulsive units. This is because electric motors are comparably simple and lightweight and can therefore be positioned more freely than conventional turbofans or turboprops. One resulting configurational concept found in ongoing research is “distributed electric propulsion” (DEP). One main advantage of DEP is the ability to augment a wing’s lift and thus to improve the high-lift behavior of an aircraft, i.e. for take-off and landing. The Clean Sky 2 project DISPROP investigated this high-lift aspect of DEP, with wind tunnel experiments (carried out by project partners) and extensive CFD work [1]. The high-lift configuration examined in the wind tunnel experiments consists of a wing segment with 3 propellers and a Fowler flap. On the CFD side, the wind tunnel environment was reproduced in simulations for validation purposes. Also, multiple CFD investigations were carried out in a simplified 2.5D environment, where a setup with a single propeller was simulated, while the influence of the neighboring propellers was modelled via periodic boundary conditions. One such 2.5D investigation will be discussed in this report, where the unsteady propeller-wing interaction effects were analyzed [1].

The flight envelope of transonic aircraft is limited in terms of Mach number and angle of attack. At high Mach numbers typical for cruise, flow separation already begins at moderate angles of attack due to shocks that form above the wing, which is regarded as high-speed stall. This can lead to a coupled oscillation of the shock and the corresponding boundary layer separation, denoted as transonic buffet, which in turn may induce structural vibrations that can be critical in terms of fatigue or control. This is avoided by incorporating corresponding margins into aircraft design. A better understanding of the phenomena would allow for a more accurate forecast of the latter in order to reduce these margins, decreasing aircraft weight and increasing aerodynamic efficiency, resulting in a lower fuel consumption. The DFG research unit FOR 2895 “Unsteady flow and interaction phenomena at High Speed Stall conditions” thus aims to further the understanding of transonic buffet and the corresponding phenomena. This goal is pursued by means of experimental and numerical investigations performed on the XRF-1 research configuration designed by Airbus and a generic tandem wing configuration consisting of two untapered and unswept wing sections. The results discussed here are part of subproject 4 of the research group, which focuses on the interaction of the separated wake with the tail plane in high-speed stall.

Results and Methods

All the simulations that were conducted on SuperMUC-NG used the DLR TAU finite-volume code. The simulations for the DISPROP investigation considered a simplified 2.5D setup. The Actuator Line method, implemented into TAU by Schollenberger [2] was employed to model the unsteady propeller blade forces in the URANS simulations. In total, the simulations solve a time of 12 propeller revolutions, which is required to obtain a periodic state.

Best practice for unsteady propeller-wing simulations was determined [1]: among other parameters, the time step (denoted with respect to the blades' angular movement per time step) was varied. Figure 1 showcases the influence of the time step on the formation of blade tip vortices in the propeller slipstream: with a time step equivalent to 8° (left image), no blade tip vortices can be observed, while the vortical structures are very clear for a time step equivalent to 0.5° (right image). On the other hand, smaller time steps also require significantly more computational resources: A simulation with a time step of 0.5° requires about 47,000 core hours, which is more than twice the requirement of a simulation with a time step of 8°.

In the following, the observed unsteady interaction effects will be briefly discussed. The simulations were run with a time step of 0.5°. Figure 2 again displays the vorticity and the blade tip vortices can be clearly identified. In the areas marked as “central vortex sections”, the movement of blade tip vortices is influenced by the presence of the wing: they are slowed down, wrap around the wing’s

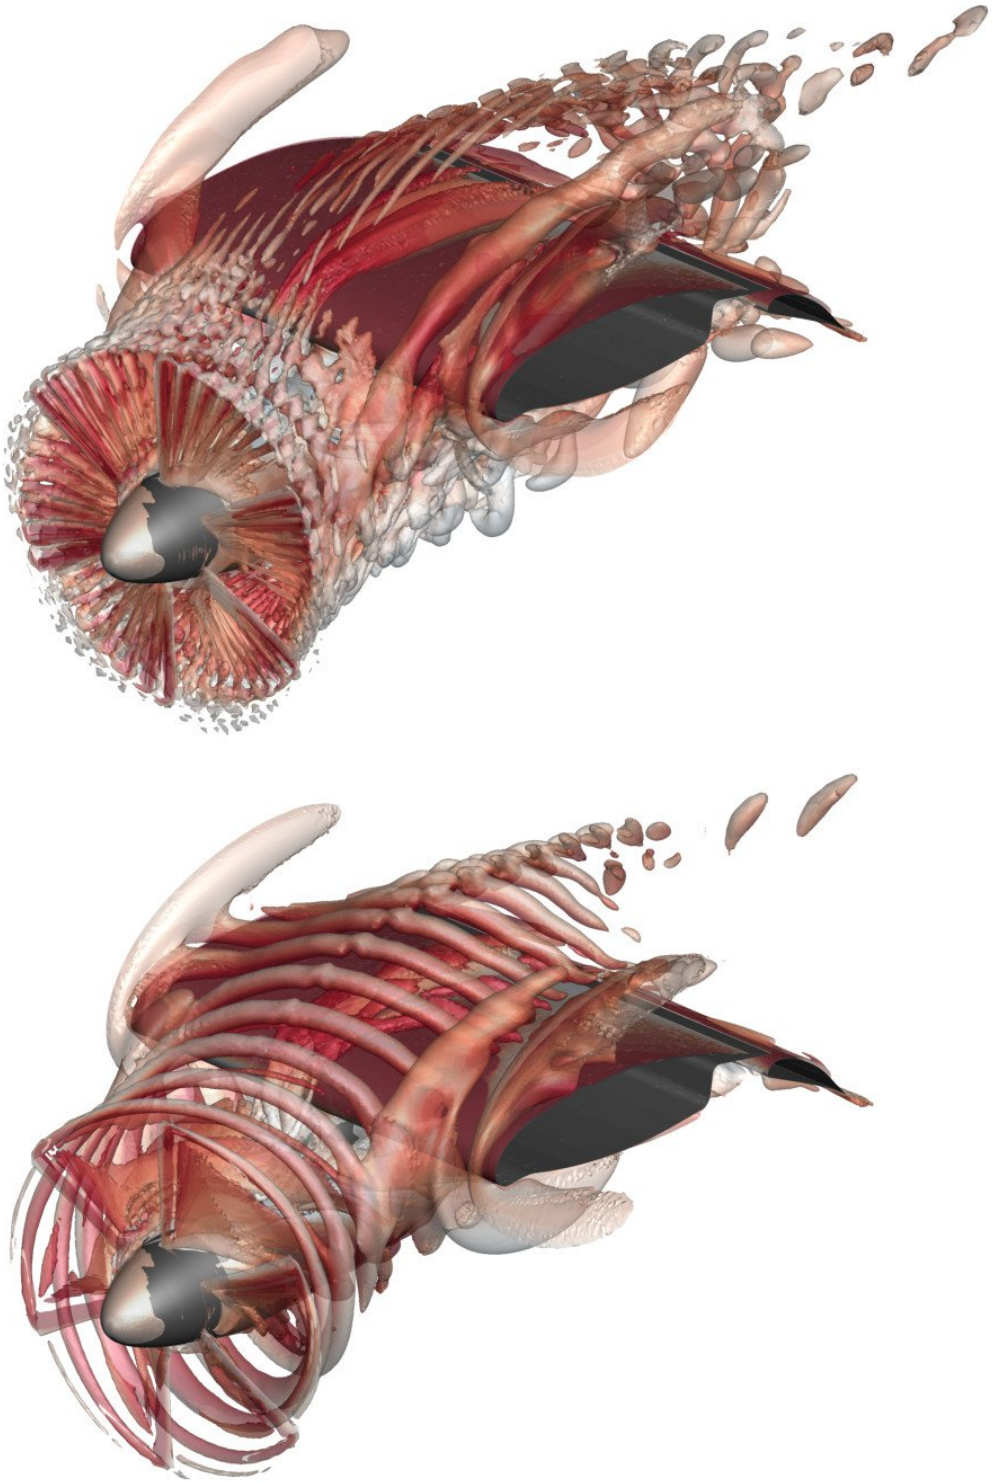

Figure 1: Vortex visualization for different time steps: 8° (left) and 0.5° (right) of blade rotation per step, from [1].

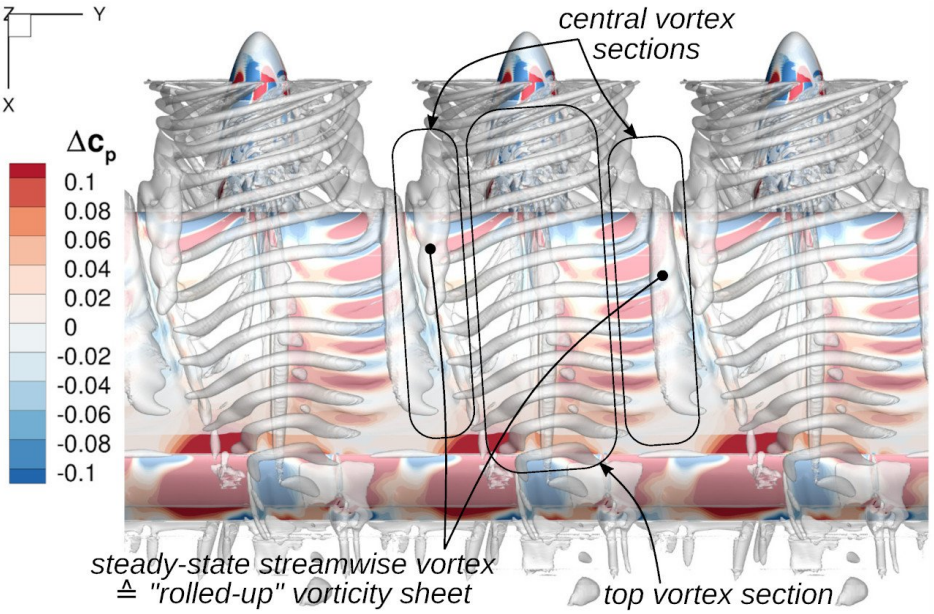

Figure 2: Visualization of unsteady propeller-wing interaction effects (vortices and surface pressure deviations), from [1].

leading edge, and then move “outwards” (i.e. in spanwise direction, left or right). Finally, all the blade tip vortices in this area merge into a large steady-state streamwise vortex, which can also be observed in steady-state CFD simulations.

The simulations for the DFG FOR 2895 were conducted as hybrid RANS/LES, and the zonal hybrid method AZDES (Automated Zonal DES) developed in-house was applied. For a typical simulation run, about 600,000 core hours are needed for the tandem and 1.8 million core hours for the XRF-1 configuration, respectively, to compute a sufficient number of convective time scales to achieve convergence of the turbulence statistics. As first results have already been shown in the last results book, several additional findings are summarized in the following.

The tandem wing configuration [3] allows to study the impingement of the separated wake of the front wing, which experiences pronounced buffet for the investigated inflow conditions, onto the rear wing (or model tail plane), as depicted in Figure 3. This impingement causes strong fluctuations of the loading of the latter. Here, a low-frequent oscillation with a period equal to the buffet period of the front wing can be distinguished from fluctuations of higher frequency. The high-frequent oscillations are due to the impingement of turbulent structures, which evolve from the shear layers at the boundaries of the wake. The footprint of the turbulent structures is notable especially in the spectrum of the fluctuation of the vertical velocity in the wake shown in Figure 4 (pink curve), indicated by the peak at a reduced frequency of around two to three, which can be attributed to a vortex street that emerges from the separated flow. The load spectrum of the tail plane, as well as the spectrum of the pressure fluctuations on its surface, also shown in Figure 4, exhibit a dominant peak at the same frequency, emphasizing the correlation between the wake fluctuations and the loads created by the impact of the latter. The interaction of the separated wake from the wing with the tail plane is more involved for the XRF-1 configuration [4] due to a strong three-dimensionality of the flow. For the high off-design Mach number considered, shock-induced separation begins at moderate angles of attack in the mid-wing section. As the angle of attack is increased, the separated area grows towards both tip and root. The wake originating from the inboard section then forms a vortex-like structure that impinges onto the tip of the tail plane, with a notable swirl, that is introduced by the gradient in lift between the mid and outer wing sections with strong and the root area with less separation. The wake impingement causes pressure fluctuations on the tail plane’s surface leading to unsteady loads, shown in Figure 5 (left), that are concentrated in the leading edge area of the tip where the vortices impinge upon. In addition, the footprint of the movement of the shock on the tail plane’s lower surface induced by the variation of the inflow is visible. Compared to the tandem configuration, the spectrum of wake fluctuations is more broadband, as depicted in Figure 5 (blue curve). A dominant peak at a reduced frequency of around 1.5 is also prominent in the spectrum of the pressure fluctuation on the tail plane’s tip (black curve) caused by the vortex impingement.

Ongoing Research / Outlook

The discussed unsteady propeller methods will be applied in the wind tunnel environment for validation purposes, using also fully resolved propeller blades. For the XRF-1 configuration, the impact of an UHBR nacelle will be investigated in more detail. In addition, the comparison with wind tunnel test data of the experimental campaigns will be furthered.

References and Links

[1] Firnhaber Beckers et. al., "Numerical Simulation of the Unsteady Slipstream of Lift Augmenting Distributed Propellers" (2023).
[2] Dieter Schwamborn, Thomas Gerhold, and Ralf Heinrich. THE DLR TAU-CODE: RECENT APPLICATIONS IN RESEARCH AND INDUSTRY, ECCOMAS CFD 2006 CONFERENCE. 2006.
[2] Schollenberger et al., “Boundary Condition Based Actuator Line Model to Simulate the Aerodynamic Interactions at Wingtip Mounted Propellers”, 10.1007/978-3-030-25253-3_58 (2020).
[3] Kleinert et al., CEAS Aeronautical Journal 15, 79-103 (2024).
[4] Lutz et al., Journal of Aircraft 60(3), 623-636 (2023).

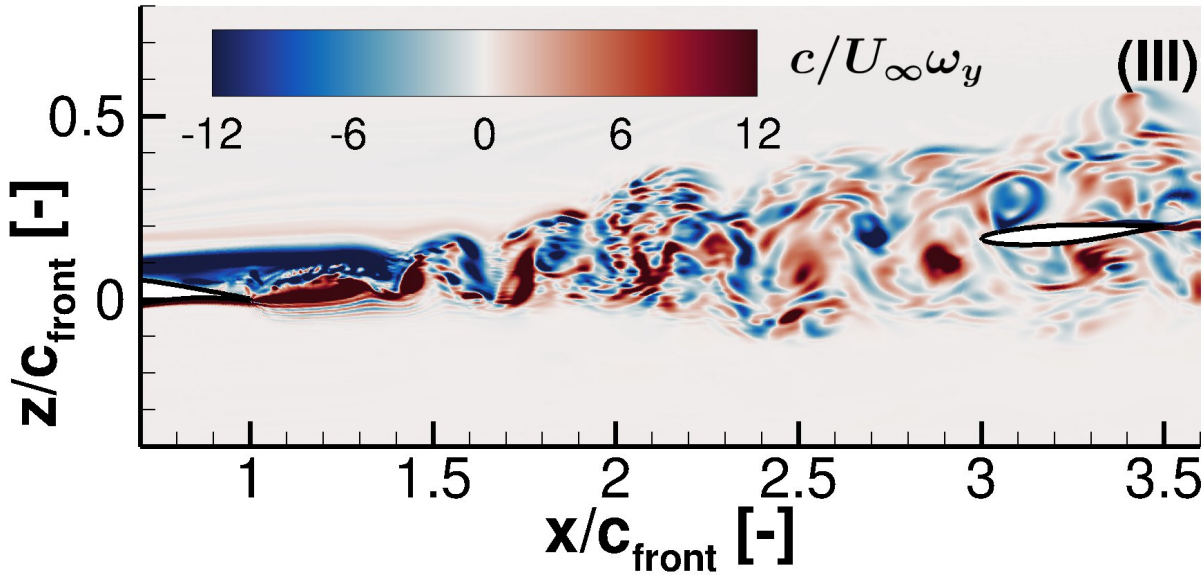

Figure 3: Snapshot of the flow field around the tandem configuration (spanwise vorticity shown).

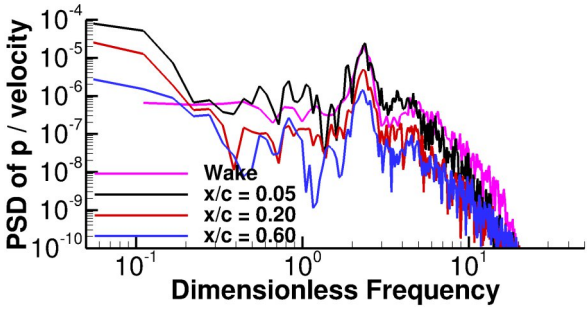

Figure 4: Spectrum of the vertical velocity fluctuation in the wake (pink) and of the surface pressure on the rear wing.

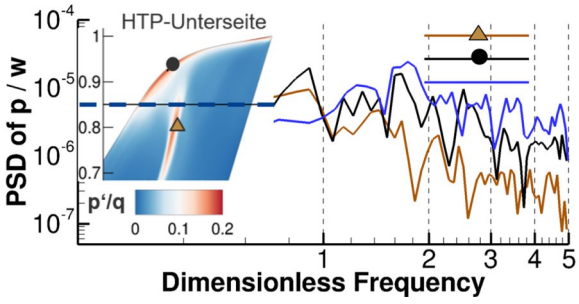

Figure 5: Spectra of velocity fluctuations in the wake (blue) and pressures on the tail plane surface (rms shown left).

Detailed Large-Eddy Simulation of multiregime combustion and Cyclical-Variation in IC-engines

Research Institution
[1]Fluid Dynamics, Energy and Material Process Institute (EMPI, former Institute of Combustion and Gasdynamics), University of Duisburg-Essen, Duisburg

Principal Investigator
Andreas M. Kempf[1]

Researcher
Linus Engelmann[1], Judith Laichter[2], Patrick Wollny[1], Markus Klein[3], Sebastian A. Kaiser[2]

Project Partner
[2]Reactive Fluids, University of Duisburg-Essen, Duisburg
[3]Department of Aerospace Engineering, University of the Armed Forces, Neubiberg

Funding
DFG, FOR 2687

SuperMUC Project ID: pn25vi

Introduction

This project examines cyclic variations in flame propagation within spark-ignited engines, with a particular emphasis on leveraging Large Eddy Simulations (LES) alongside advanced imaging techniques. These methodologies enable an in-depth investigation into the interactions between turbulence and combustion processes, which are pivotal for optimizing engine efficiency and minimizing harmful emissions. The comparison of experimental data with the outcomes from simulations uncovers critical insights into the nuances of combustion dynamics, shedding light on the complex mechanisms that influence engine performance and environmental sustainability. This research underscores the importance of sophisticated simulation tools in advancing our understanding of combustion phenomena, offering valuable perspectives on improving propulsion systems for a greener future. Through this detailed analysis, the study contributes significantly to the body of knowledge in automotive engineering, highlighting the potential for innovative approaches to enhance the design and operation of modern engines.

Results and Methods

Utilizing the SuperMUC-NG, our project embarked on a detailed simulation campaign, employing Large Eddy Simulation (LES) to capture the complex phenomena of turbulent combustion in spark-ignited engines. These simulations were meticulously designed to mirror the physical conditions inside engine cylinders, accounting for the intricate interplay of fluid flow, chemical reactions, and heat transfer. By simulating a range of operating conditions, we were able to isolate and study the effects of various parameters on the efficiency and stability of the combustion process. The technical backbone of our research hinged on the development of sophisticated computational models that accurately reflect the physics of the turbulent flow and combustion. This involved integrating state-of-the-art turbulence models, advanced combustion kinetics, and high-fidelity geometry representations. The computational demand of these simulations was substantial, requiring efficient use of SuperMUC-NG's parallel processing capabilities and innovative algorithmic solutions to optimize simulation runtimes without compromising accuracy. Our results highlight the significant influence of turbulence on flame propagation and combustion efficiency in spark-ignited engines. Key findings demonstrate that by optimizing turbulence characteristics, combustion can be made more efficient, leading to reduced emissions and improved engine performance. The research also uncovers the importance of cyclic variations in combustion, attributing these to changes in turbulence intensity and flow patterns within the engine cylinder. A close agreement between Large Eddy Simulation (LES) results and experimental observations validates the effectiveness of LES in capturing critical aspects of flame propagation and turbulence interaction, marking a significant advancement in simulation techniques for engine combustion.

Ongoing Research / Outlook

The support of SuperMUC-NG was instrumental in achieving our research objectives, providing the computational resources necessary to conduct high-fidelity simulations that push the boundaries of our understanding of combustion processes. Looking forward, we recognize the need for further exploration into more complex chemical kinetics and the effects of alternative fuels on combustion efficiency and emissions. Despite the groundbreaking insights gained, our work also encountered limitations, particularly in simulating even finer scale phenomena within the combustion process, which are critical for capturing the complete spectrum of turbulence-chemistry interactions. Future projects are already in the planning stages, aiming to leverage next-generation HPC resources to address these challenges, explore new combustion models, and further refine our understanding of engine performance.

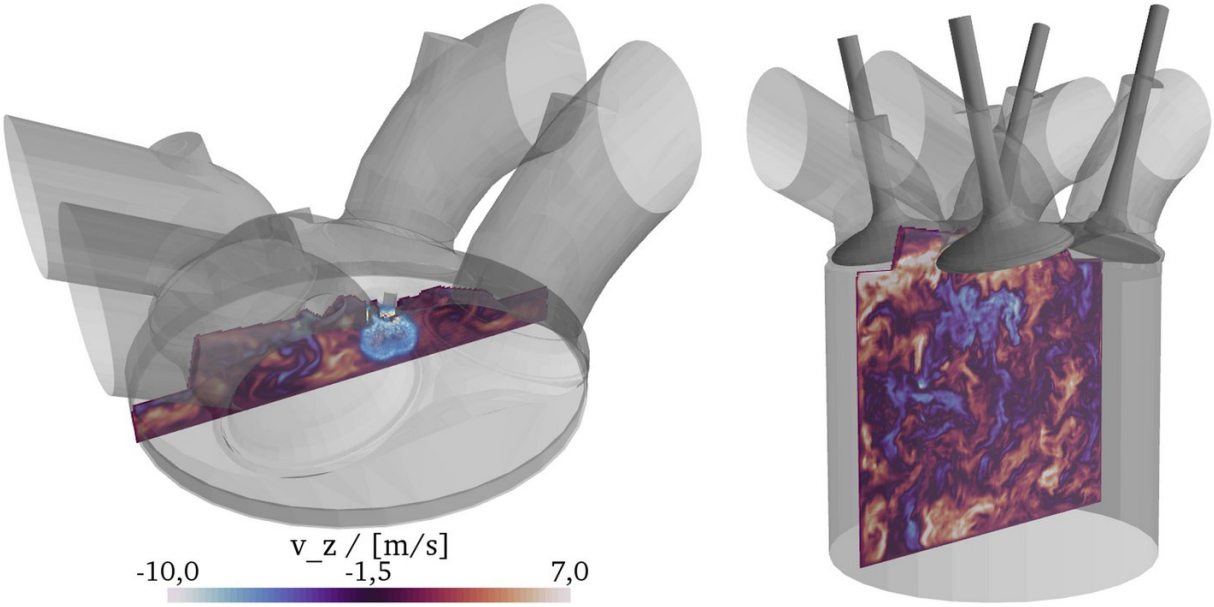

Figure 1: Image-normal velocity component during intake stroke (−120°CA).

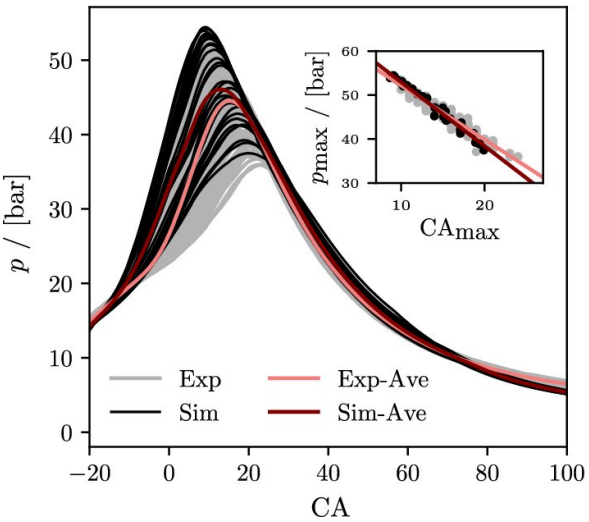

Figure 2:Experimentally and numerically obtained pressure traces from all cycles. Averages are given by the colored lines. The insert shows the correlation between the peak pressure.

References and Links

[1] L. Engelmann, C. Welch, M. Schmidt, D. Meller, P. Wollny, B. Böhm, A. Dreizler, A. Kempf, Applied Energy 342 (2023).
[2] L. Engelmann, J. Hasslberger, S.J. Baik, M. Klein, A. Kempf, Scientific Reports 13 (1), 11202.
[3] L. Engelmann, J. Laichter, P. Wollny, M. Klein, S.A. Kaiser, A.M. Kempf, Flow, Turbulence and Combustion 110 (1), 91-104.
[4] L. Engelmann, J. Hasslberger, E. Inanc, M. Klein, A. Kempf, Computers & Fluids 240, 105441.
[5] M. Klein, S. Ketterl, L. Engelmann, A. Kempf, H. Kobayashi, International Journal of Heat and Fluid Flow 81, 108496.

Exchange of mass and momentum across the interface between a turbulent flow and a porous medium

Research Institution
Technische Universität München

Principal Investigator
Michael Manhart

Researcher
Simon v. Wenczowski

Project Partner
–

Funding
DFG grant MA2062/15-1, KONWIHR F.3-M7635.1/20/8

SuperMUC Project ID: pn68vo

Introduction

In riverine ecosystems, the hyporheic zone comprises the bio-geo-chemically active topmost layers of a river bed, where the water in the pore space is in permanent interaction with the overlying turbulent flow. Various organisms in the hyporheic zone depend on the supply of specific nutrients as well as the removal of waste products of their metabolism. Therefore, the health of the aquatic ecosystem is strongly influenced by the exchange of mass and momentum across the interface between a turbulent flow and the porous sediment bed. Using the hyporheic zone as an example, our research aims to advance the mechanistic understanding of such exchange processes.

Results and Methods

For the numerical simulations, the sediment bed is represented by a mono-disperse random sphere pack. To obtain the flow field, the incompressible Navier-Stokes equations are solved. The flow computation is coupled with the advection-diffusion equation, to predict the transport of two passive scalars with different boundary conditions. In a single-domain pore-resolved Direct Numerical Simulation (DNS), all spatial and temporal scales of the flow are resolved. We use our in-house code MGLET [1], which employs an energy-conserving central second-order Finite Volume Method. The variables are defined at staggered positions within a Cartesian grid on different refinement levels, which allows a local grid refinement. The time integration is accomplished by an explicit third-order low-storage Runge–Kutta (RK) method. MGLET represents the complex geometry of the porous medium by means of a ghost-cell immersed boundary method, which reaches second-order spatial accuracy, while ensuring mass conservation. The performance of MGLET on SuperMUC-NG is enhanced by SIMD-optimized core routines [2] and efficient HDF5-based parallel I/O [3]. During runtime, VTK-based visualization pipelines output flow visualization at high frequencies. Figure 1 shows the weak scaling behavior of MGLET up to 135,168 cores spread across 2,816 nodes of SuperMUC-NG. Most production runs of this project were carried out on one island of the SuperMUC-NG, as shown in Table 1.

case	# cells	# nodes	file size [GB]
S-150	7.474×10^9	704	6,309.9
S-300	7.474×10^{10}	704	6,309.9
S-500	7.474×10^9	704	6,309.9
M-150	8.437×10^9	176	1,577.5
M-500	8.437×10^9	786	7,122.6
L-180	0.708×10^9	225	7,122.6
L-300	3.185×10^9	375	2,688.8

Table 1: Computational parameters of simulation cases. One node hosted 48 processes.

In comparison to flow over smooth walls, roughness and permeability change the mean flow and the turbulence structure in direct proximity of the sediment bed, as shown by the vortex visualization in Figure 2. With increasing distance from the rough and permeable interface, an outer-layer similarity to flow over smooth walls can be recovered, if the wall-distance is defined in a dynamically consistent way. The specification of a wall-distance, however, requires the definition of an interface between the porous medium and the turbulent free flow. The evaluation of our high-resolution simulation data indicates that the centroid and the width of the drag distribution on the porous medium provide a physically meaningful interface definition. This drag-based definition of the interface between porous medium and turbulent free flow also reveals a similarity of flow profiles in direct proximity of the sediment bed surface [4]. Besides hyporheic momentum exchange, the conducted simulation shed light on the hyporheic transport of scalar quantities like dissolved substances.

Ongoing Research / Outlook

Figure 3 provides an impression of instantaneous scalar concentration fields for the different boundary conditions

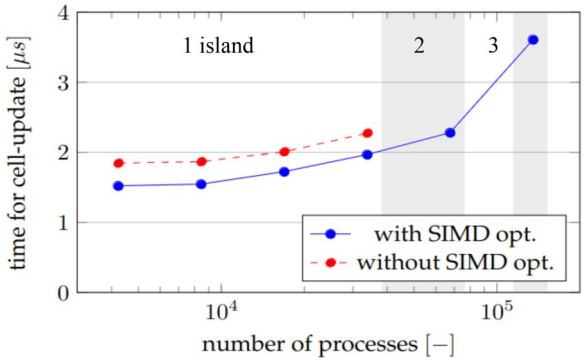

Figure 1: Weak scaling behavior of MGLET on SuperMUC-NG. Changes in the background color mark island boundaries.

on the surface of the sediment grains. The statistical simulation results allow us to distinguish turbulent scalar transport, dispersive scalar transport and transport due to molecular diffusion. Whereas turbulent transport is dominant in the free flow region, dispersive transport plays a major role in the sediment bed [5].

The gathered statistics of the flow and scalar field provide the basis for the investigation of various aspects. We are currently about to evaluate the budgets of different quantities. These budget evaluations are expected to aid the development or improvement of models for hyporheic exchange, which serve as tools for practitioners.

References and Links

[1] www.github.com/tum-hydromechanics/tum-mglet-base
[2] www.konwihr.de/konwihr-projects/simd-optimisation-of-the-cfd-software-package-mglet-for-supermuc-ng/
[3] Sakai, Y. et al., 2019. Performance Optimisation of the Parallel CFD Code MGLET across Different HPC Plat-forms. In Proceedings of the Platform for Advanced Scientific Computing Conference (PASC '19). Association for Computing Machinery, New York, Article 6, 1–13.
[4] Wenczowski, S., Manhart M., 2024. Turbulent flow over random sphere packs - an investigation by pore-resolved Direct Numerical Simulation, Submitted to Journal to Fluid Mechanics
[5] Wenczowski, S., Manhart, M., 2024. Direct Numerical Simulation of Scalar Transport Across the Interface Between a Porous Medium and a Turbulent Flow. In: Marchioli, C., Salvetti, M.V., Garcia-Villalba, M., Schlatter, P. (eds) Direct and Large Eddy Simulation XIII. DLES 2023. ERCOFTAC Series, vol 31

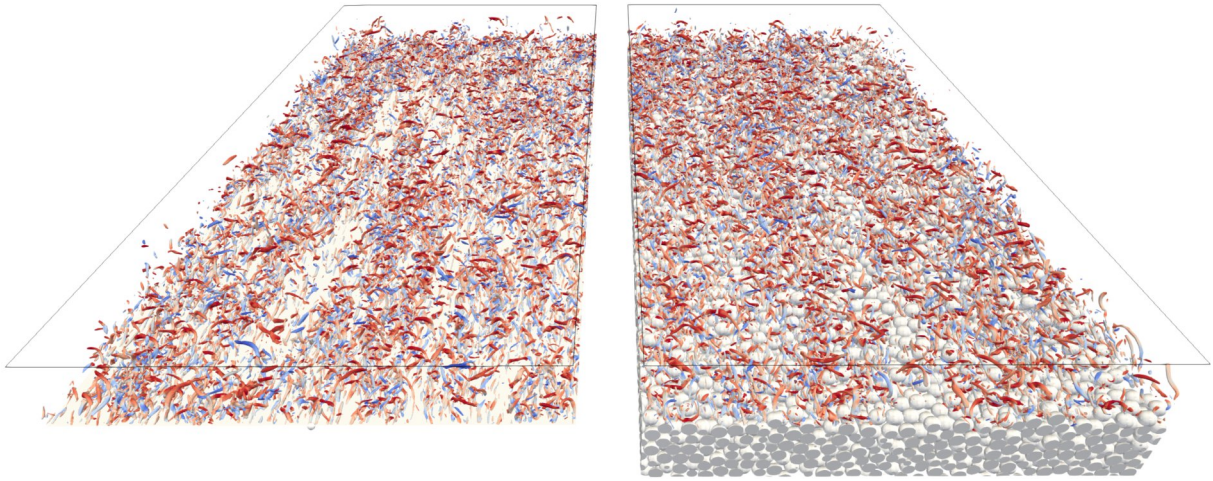

Figure 2: isualization of vortex structures in turbulent flow (Re_τ = 500) over smooth and impermeable bed (left) and over rough and permeable bed (right, case M-500) in side-by-side comparison.

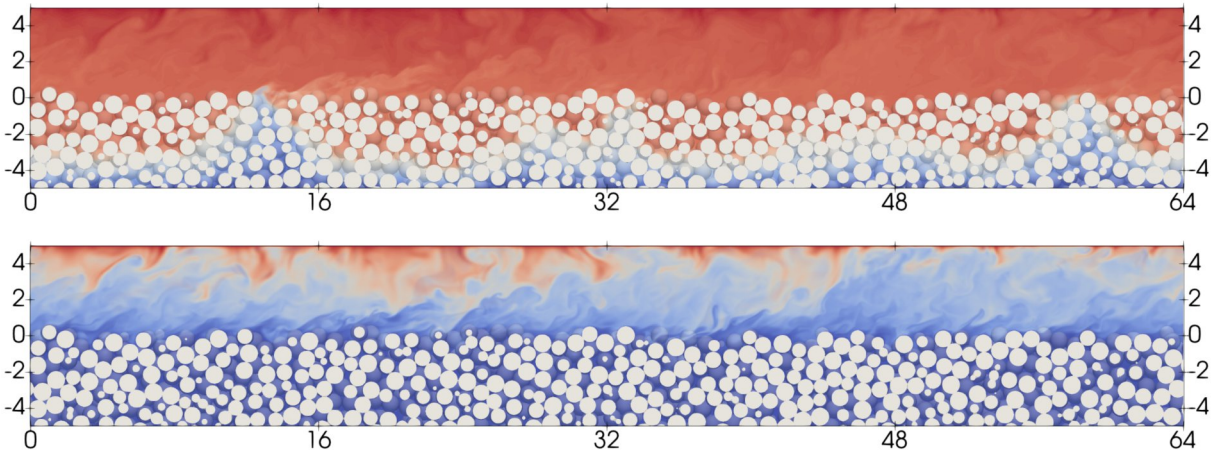

Figure 3: Instantaneous scalar concentration fields within an arbitrarily chosen slice. As passive scalars are considered, different boundary conditions for different scalar can be treated within the same flow simulation. Top: Zero-flux Neumann boundary conditions on the sphere surfaces. Bottom: Fixed-value Dirichlet boundary conditions on the sphere surfaces.

Turbulence model conditioning and flow control for aircraft vortical flow

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

Principal Investigator
Christian Breitsamter

Researcher
Dominik Sedlacek, Julius Stegmüller

Project Partner
–

Funding
DFG

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. The controllability and stability investigations on delta wing configurations and the fluid-structure coupling have 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 operation at extreme flight conditions. The configuration is based on wings with medium to high sweep angles and leading edges with small radii. Even at low angles of attack, the flow separates at the leading edge or multiple swept leading edges, forming a vortex or vortex system, respectively. Extensive investigations on hybrid delta wing configurations were carried out. The project can be divided into two subprojects.

In the first subproject, “Multiple Swept Wings”, funded by the Deutsche Forschungsgemeinschaft (DFG), the vortex-dominated flow on hybrid-delta-wing configurations is investigated. The focus is the analysis of the influence of the leading-edge radius (LER) on the vortex systems.

The second subproject, “Aeroelastic Wind Tunnel Model – Full Span Configuration (AWTM-F)”, also funded by the German Research Foundation (DFG), deals mainly with the analysis of vortex-induced buffeting effects on a generic strake trapezoidal wing configuration with horizontal and vertical tail planes. The focus of the simulations is on the numerical prediction of fluid-structure interactions (FSI) with weak coupling and strong coupling simulation methods.

Using the High-Performance Computing system has played a fundamental role in properly developing the simulations. Numerical simulations are routinely employed in parallel to experimental data in the design and investigation of new aerodynamic solutions. Considering a complex geometry and a progressively demanding level of accuracy, the computational grids must be refined and the numerical methodology improved, 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.

Results and Methods

In the subproject “Multiple Swept Wings”, one focus is to investigate the influence of the LER on the formation, interaction, and breakdown of vortex systems. For example, the pressure loss isosurfaces at $\alpha = 24°$ and $\beta = 5°$ for the double-delta-wing configuration with a leading-edge sweep of $\varphi_1 = 75°$ and $\varphi_2 = 52°$ are shown in Figure 1. Regarding the sharp-edged variant, see Figure 1a, at the windward side (left side), vortex breakdown is observable for the strake vortex and downstream, as well as for the main wing vortex. With the vortex breakdown of the main wing vortex, the cross-section of the isosurface increases significantly. At the leeward side (right side), the two vortices are stable over the whole wing and interact with each other. For the variant with a leading edge radius of $r_N/c_r = 0.25\%$, the vortex breakdown at the windward strake vortex breakdown occurs more downstream. At the leeward side, the strake vortex burst in the rear part of the wing. Therefore, a change in LER can have stabilizing as well as destabilizing effects. With an increase in LER, see Figure 1c, the windward vortex breakdown occurs more

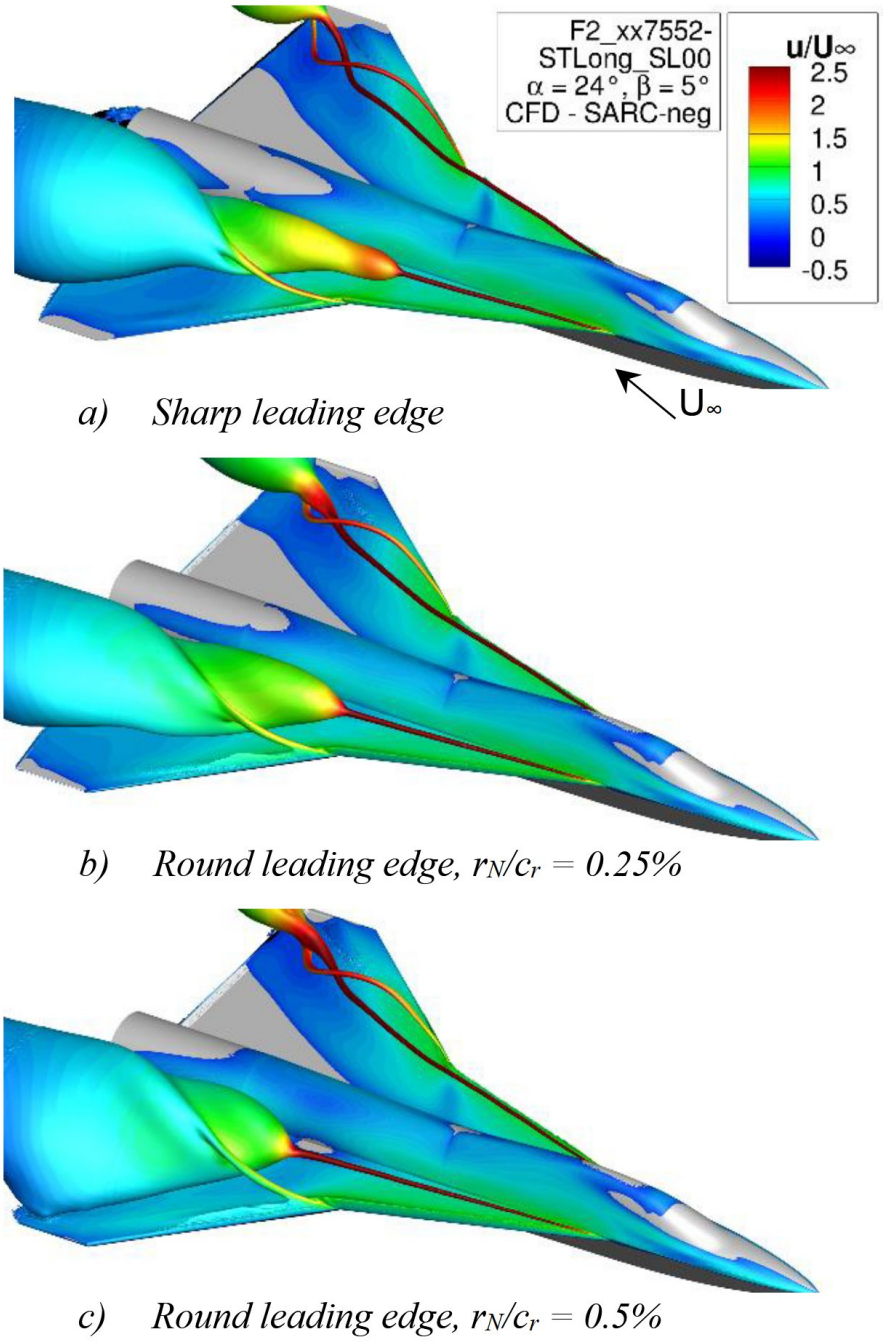

Figure 1: Isosurfaces of the pressure loss of 2% colored with the non-dimensional axial velocity u/U_∞ for the double-delta-wing configuration at α = 24°, β = 5°, and Re = $3.0 \cdot 10^6$.

downstream compared to the other configurations and the leeward side features a stable vortex system.

For a more detailed analysis, the non-dimensional axial velocity in the vortex core as function of the downstream position x/c_r is shown in Figure 2 for the leeward and windward sides. The sharp-edged variant features the highest level of axial velocity. For the leeward side, the three variants behave quite similarly; only the largest LER exhibits high core velocities in the strake vortex close to the trailing edge. At the windward side, the vortex breakdown shifts more downstream with increasing LER, and the gradient in axial velocity decreases in the area of vortex breakdown. Therefore, the LER greatly impacts the vortex systems at hybrid-delta-wing configurations and further investigations will be conducted.

In the project “Aeroelastic Wind Tunnel Model – Full Span Configuration”, numerical simulations for different angles of attack and for different deflection angles of the HTP were performed to analyze the flow field and the unsteady pressure fluctuations caused by vortex breakdown on the generic strake trapezoidal wing configuration with horizontal and vertical tailplanes. The DLR-TAU

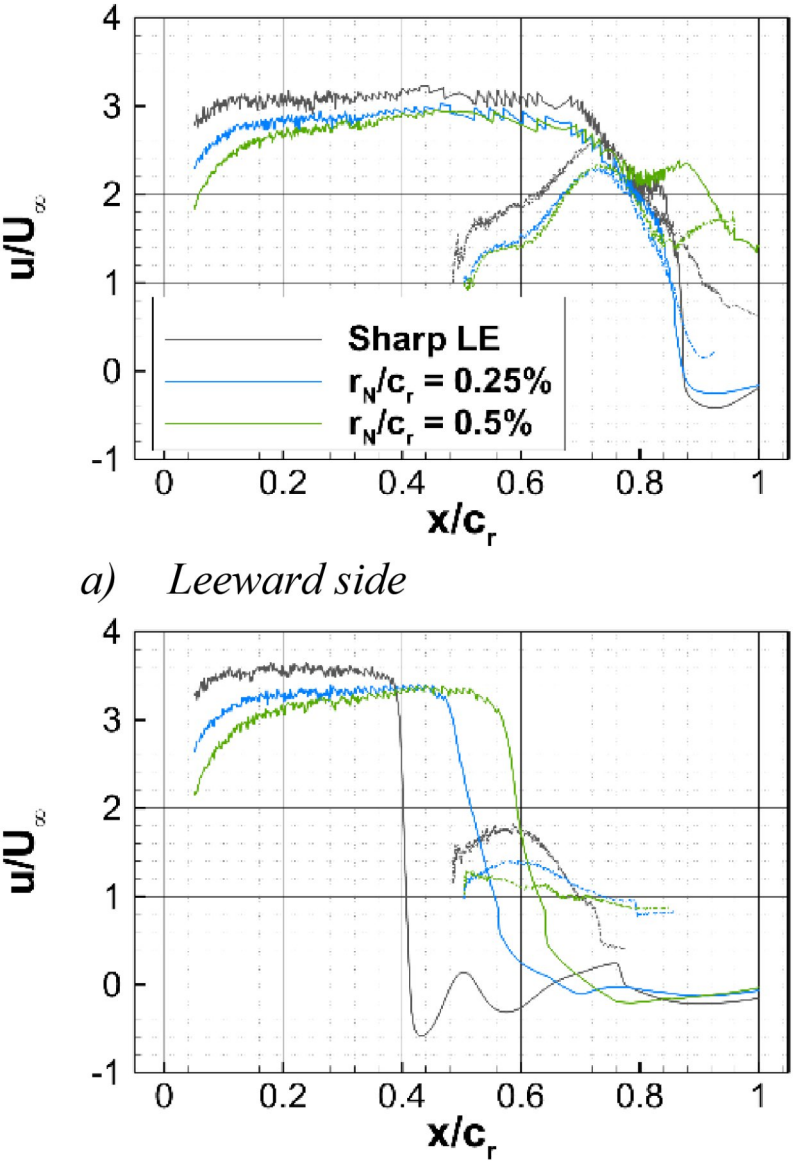

Figure 2: Distribution of the non-dimensional axial velocity u/U_∞ in the vortex core for the double-delta-wing configuration at α = 24°, β = 5°, and Re = 3.0·10⁶.

Code is used as flow solver for URANS and DDES simulations. Figure 3 shows exemplarily the results of a DDES simulation using the Q-criterion of the time-averaged (left) and the instantaneous flow field (right) for an angle of attack of α = 35° and a tail deflection angle of δ = 0° at subsonic speed. It can be seen that the unsteady flow field downstream of vortex breakdown hits the vertical tailplane. The resulting pressure fluctuations on the configuration surface lead to a structural excitation of the aircraft components.

As a next step, several FSI simulations for different deflection angles of the HTP and angles of attack of α = 25°, α = 30° and α = 35° are performed to analyze the fluid-structure interactions and to trim and evolve the coupled methods by comparing them with experimental results. For the numerical simulations, existing weak and strong coupling methods of partitioned formulation are used and further developed [2, 3]. MSC Nastran is used as a structural solver for modal and displacement analysis. For analyzing the influence of flexible structures on the flow field, several configurations are used in the FSI simulations by using either rigid lifting surfaces or flexible lifting surfaces or by mixing rigid and flexible components.

Ongoing Research / Outlook

The use of SuperMUC-NG is essential to this project. Major investigations have been performed in the past, and important and interesting results have been 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.

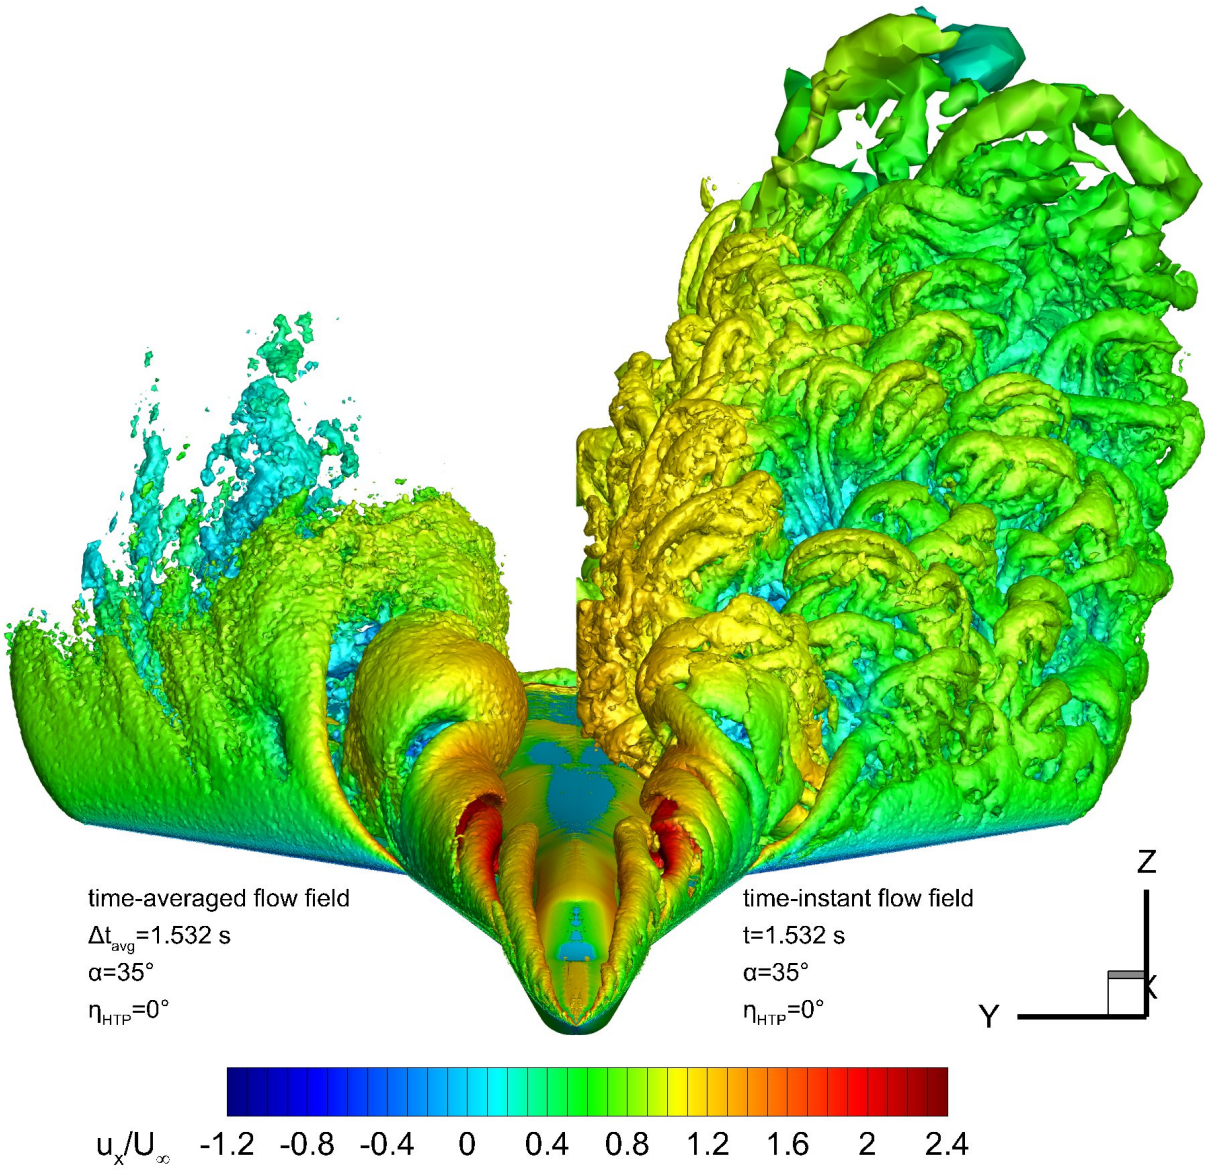

Figure 3: *Q* - criterion flow visualization of time-averaged (left)and instantaneous flow field (right) for the AWTM-F configuration at α = 35°, M = 0.15 and Re = $3.6 \cdot 10^6$.

References and Links

[1] D. Sedlacek et al., Influence of the Leading-Edge Radius on Vortex Development at Hybrid-Delta-Wing Configurations, Aerospace Europe Conference, Lausanne, 2023. DOI: https://doi.org/10.13009/EUCASS2023-418

[2] Katzenmeier, L., Vidy, C., Benassi, L., Breitsamter, C.: Prediction of horizontal tail buffeting loads using URANS and DES approaches. In: International Forum on Aeroelasticity and Structural Dynamics (IFASD) 2019, Savannah, Georgia (2019).

[3] Katzenmeier, L., Vidy, C., Breitsamter, C., J. Fluids Struct. 99 (2020). DOI: https://doi.org/10.1016/j.jfluidstructs.2020.103178

Aerodynamics of Multiple Swept Delta Wings

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

Principal Investigator
Christian Breitsamter

Researcher
Konstantin Bantscheff, Moritz Zieher

Project Partner
–

Funding
DIGIfly LuFo V-I FKZ 20X1909I

SuperMUC Project ID: pn36xo

Introduction

A focus area for the aircraft aerodynamics group at the Chair of Aerodynamics and Fluid Mechanics (AER) is the prediction and understanding of complex vortex flow physics of mid to low aspect-ratio aircraft configurations. This includes the computation and assessment of aerodynamic characteristics, turbulence modeling with regard to separated flows, and the investigations of aeroelastic effects and aeroservoelastics induced by fluid-structural couplings. Prediction and understanding of the complex aerodynamics of high-agility aircraft play a major role in the development of future aircraft systems pushing the boundaries of current flight envelopes.

Results and Methods

Sub-project 1: Adaptive Turbulence Model
Eddy-viscosity-based turbulence models provide the most used modeling approach for computational fluid dynamics simulations in the aerospace industry. These models are very accurate at relatively low costs for many cases but lack accuracy in the case of leading-edge vortex flows for mid to low aspect-ratio wings. An enhanced adaptive turbulence model based on the one-equation Spalart-Allmaras turbulence model is fundamental to this work. This model employs several additional coefficients and source terms, specifically targeting vortex-dominated flow regions, where these coefficients can be calibrated by an optimization procedure based on experimental or high-fidelity numerical data.

The enhanced turbulence model is implemented in the DLR TAU-solver, which is a widely adopted computational fluid dynamics solver for aerospace applications. In Figure 2, different coefficients and their influence regions are shown for reference. From the influence regions, it can be seen that the different coefficients are active in different flow regions.

When comparing the original Spalart-Allmaras turbulence model to an optimized version of the enhanced turbulence model, there are various improvements

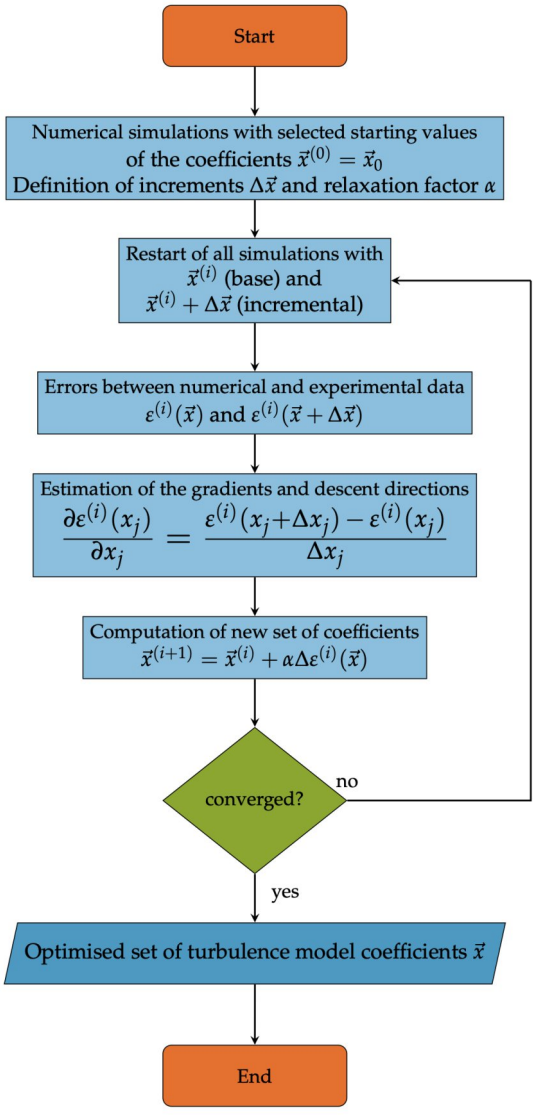

Figure 1: Flowchart of the coefficient optimization procedure.

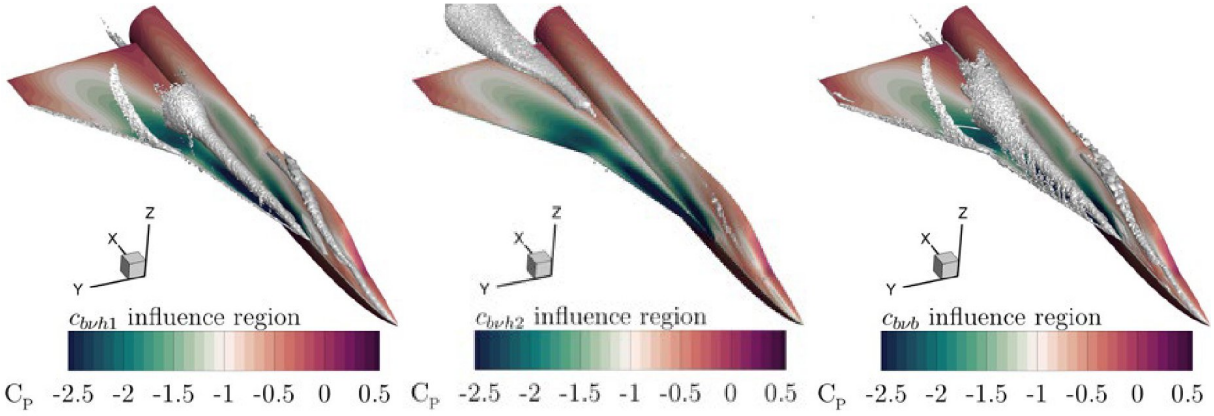

Figure 2: Double delta wing configuration; Re = 3·10^6, M_∞ = 0.15, α = 32°. Additional turbulence model coeffcient influence regions.

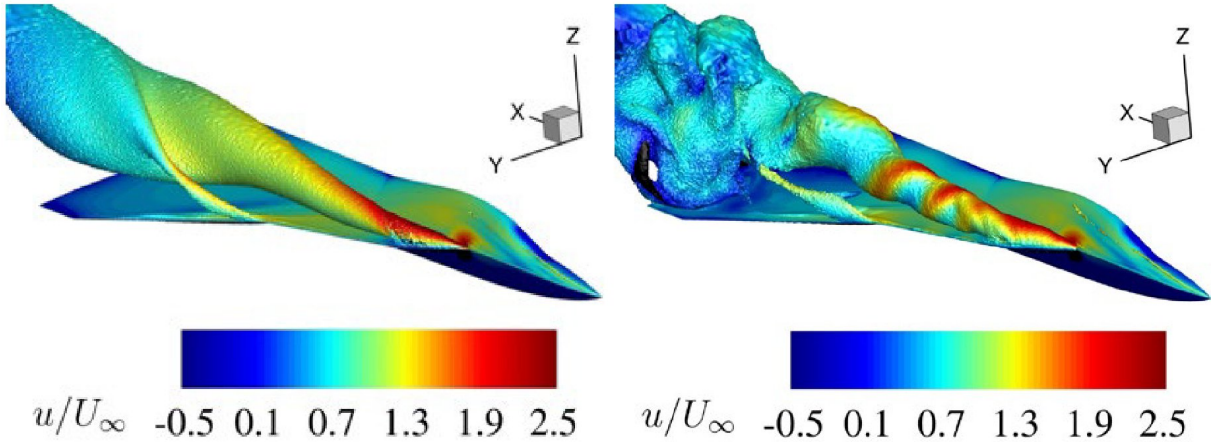

Figure 3: Triple delta wing configuration; Re = 3·10^6, M_∞ = 0.15, α = 32°. Vortical flow field structures by the means of total pressure loss.

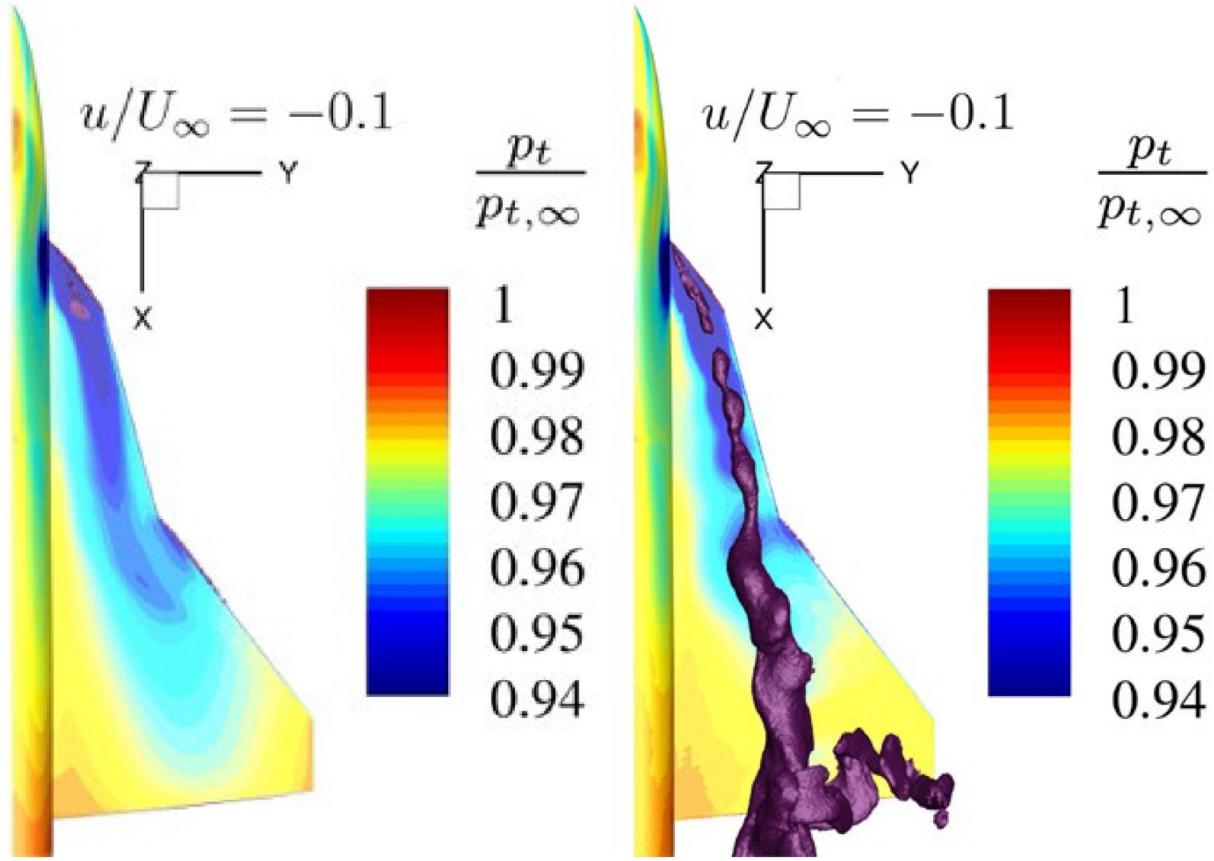

Figure 4: Double delta wing configuration; Re = 3·10^6, M_∞ = 0.15, α = 32°. Comparison of reverse flow.

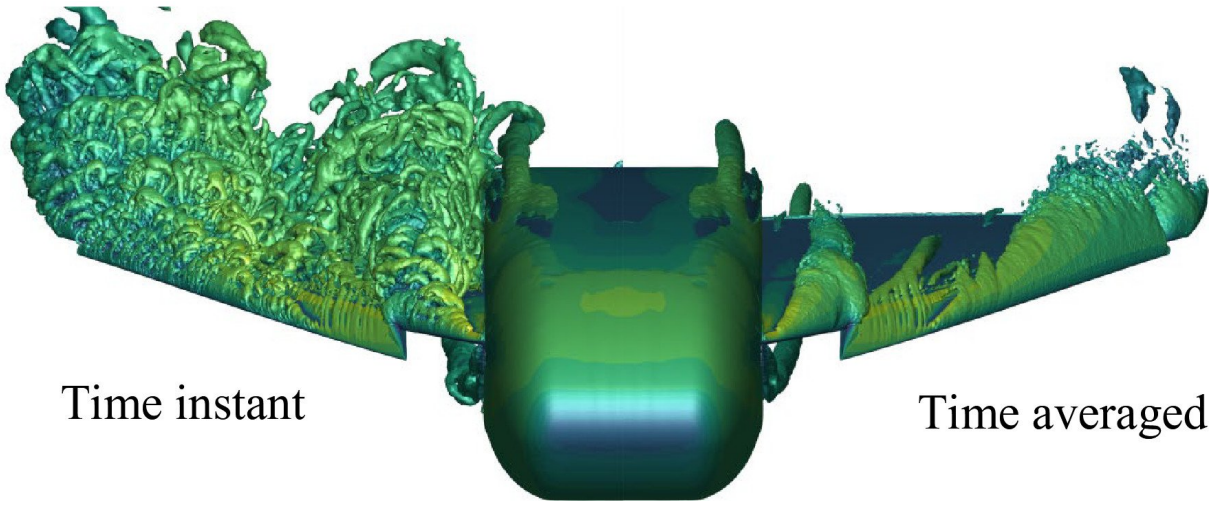

Figure 5: Vortical flow field structures for the investigated delta wing configuration. Time instant vortical structure on the left and time averaged vortex structure on the right.

visible. Figure 3 shows the normalized total pressure loss for the original turbulence model (left), and the adaptive version (right) at an angle of attack of α = 32° and a Mach number M_∞ = 0.15. A comparison of the two results shows that the adaptive model represents in more detail the vortical structures emanating on the multiple-swept delta wing. In contrast, the original model does not show a vortex-like behavior of the flow field.

Figure 4 highlights another important improvement of the adaptive turbulence model. Reverse flow regions by means of normalized axial velocity $u/U_\infty \leq -0.1$ are shown. The original Spalart-Allmaras turbulence model exhibits nearly no reverse flow, whereas the adaptive turbulence model shows some cone-like reverse flow region over the whole wing area. The proper resolution of these regions is important since vortex breakdown and aerodynamic behavior are directly linked to these phenomena.

Sub-project 2: Aeroelastic Simulation
The influence of fluctuating aerodynamic forces on a wing caused by the fluid flow around it leads to dynamic structural deformations of the wing. In particular, for high angles of attack and large aerodynamic loads, the deformation of aerodynamic surfaces reach a significant value. For those cases, the deformation of the wing needs to be accounted for as the structural deformation again influences the aerodynamics. As a result, fluid-structure interactions (FSI) need to be performed to predict the aircrafts aerodynamics correctly. In the project at hand, a delta wing model is investigated regarding the dynamic deformations of the wing and trailing edge flaps in the high angle of attack regime.

Therefore, the simulation software ANSYS is utilized. The overall simulation procedure includes three major processes. First, the structural solver is addressed. Therefore, the mechanical components of the wing are modeled. The wing with its entire inner structure, including ribs and spars, is meshed using the finite element meshing tool, which is integrated into ANSYS Mechanical. The second task involves the classical CFD simulation. Here, the two equations k-ω SST turbulence model is applied to solve the unsteady Reynolds Averaged Navier-Stokes equations. The System Coupling tool is applied for the third process, which links the mechanical part and the aerodynamic simulation. This tool enables a data transfer between the mechanical and the aerodynamic models. As a result, aerodynamic forces can be transferred to the structural model. Vice versa deformations from the finite element solver can be used to adapt the mesh and thus the wing geometry of the CFD simulation.

As transient two-way FSI simulations are performed, multiple data exchanges between the two models are conducted for each time step. Due to the resource intensiveness of the FSI simulation, classical CFD simulations for an undeformed wing are used as a starting point for the aerodynamic side of the FSI process.

Following, the results of the simulations are presented at a Mach number of M_∞ = 0.15 and an angle of attack of α = 20°. A visualization of the emerging vortical structure above the investigated undeformed delta wing is shown in Figure 5. Figure 6 shows the result of the FSI simulation. Here, the influence of the aerodynamic loading on the structural layout is visible. In particular in the wing tip region, a significant influence on the twist and bending is noticeable.

Ongoing Research / Outlook

Sub-project 1: Adaptive Turbulence Model
To extend the usability of the adaptive turbulence model, the goal of this research is to implement a machine learning framework that can predict the additional turbulence model coefficients based on a small number of inputs rather than rely on the optimization procedure. For the proper training of such a machine learning framework, an extended training dataset is needed, which will involve numerous CFD simulations of different aircraft configurations.

Sub-project 2: Aeroelastic Simulation
Further research will focus of the comparison between the aerodynamic results for the undeformed and de-

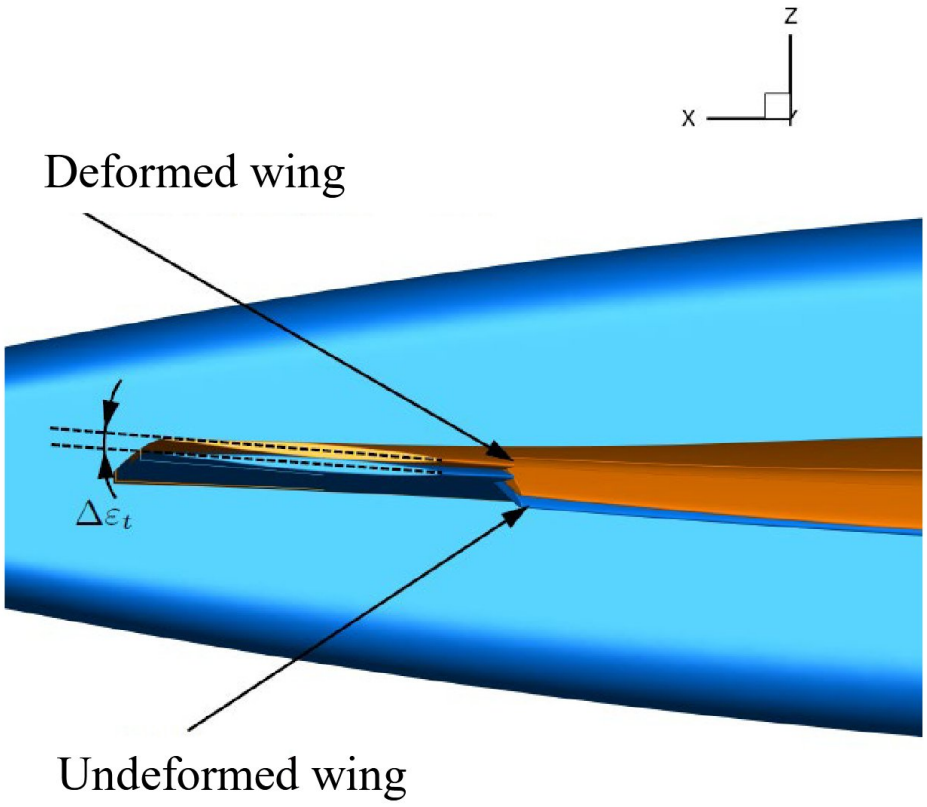

Figure 6: Result of the FSI simulation for the investigated delta wing configuration.

formed structure as well as a comparison with conducted wind tunnel measurements. In that regard the focus lies on the interpretation of unsteady characteristics of the flow and the structure of the wing. The influence of different flap angles on the structural behavior of the wing under aerodynamic loading will be investigated in multiple following CFD simulations.

References and Links

[1] Moritz Zieher, Christian Breitsamter, “Adaptive Turbulence Model for Leading Edge Vortex Flows Preconditioned by a Hybrid Neural Network”, Aerospace 2024. (submission ID: aerospace-2866419).

[2] Moritz Zieher, Christian Breitsamter, “Prediction of Aerodynamic Coefficients for Multi-Swept Delta Wings via a Neural Network”, Aerospace Europe Conference 2023 – 10TH EUCASS – 9TH CEAS, Lausanne, July 2023. DOI: 10.13009/EUCASS2023-38.

[3] Bantscheff K. and Breitsamter C. Dynamic structural scaling concept for a delta wing wind tunnel configuration using additive manufacturing. Aerospace, 10(7):581, June 2023. DOI: 10.3390/aerospace10070581.

[4] Bantscheff Konstantin, Christian Breitsamter, “Design and Investigation of an Aeroelastically Scaled Delta Wing Wind Tunnel Model for Dynamic Load Analysis”, 72. Deutscher Luft- und Raumfahrtkongress, Stuttgart 2023.

Aerodynamics of Helicopter and Rotor Configurations

Research Institution
Chair of Aerodynamic and Fluid Mechanics, Technical University of Munich

Principal Investigator
Christian Breitsamter

Researcher
Ulrich Hartmann, Andreas Kümmel, Catharina Martinez Moreira

Project Partner
–

Funding
LuFo VI-2 (NANNY, ETHAN)

SuperMUC Project ID: pn69ni

Introduction

A major research topic at the Chair of Aerodynamics and Fluid Mechanics (TUM-AER) is the investigation of helicopter aerodynamics as well as propeller and rotor aerodynamics. The main goals of these investigations are the increase in lift and the decrease in drag. In addition, the reduction of noise emissions is another major issue. High-quality flow simulation results help to understand and analyze specific flow phenomena, leading to innovative and environmentally friendly aircraft that will help to develop green and safe air transport.

One focus of this project lies in the aerodynamic investigation of a helicopter wind tunnel model. The aerodynamic complexity of the rotor and the influence of the rotor downwash with other parts of the helicopter require high-quality flow simulations to capture the flow phenomena correctly. Especially in forward flight, the rotor wake interacts with the horizontal and vertical stabilizer, affecting the pitching moment characteristics of the whole helicopter. Therefore, correctly predicting the turbulent flow around the rotor head and its wake is essential in helicopter development.

Aerodynamic investigations of propeller configurations build another research field of TUM-AER. The upcoming demand for electrically powered Vertical Take-Off and Landing (eVTOL) vehicles requires efficient propulsion to ensure long air time and cruising range. Electrically driven propellers allow the distribution of the propulsion system to be different from the classical main-tail rotor design for helicopters. These setups lead to new aerodynamic interaction phenomena, for example, propeller-propeller, propeller-wing, and propeller-main body interactions. With high-quality flow simulations, the interactional aerodynamics between VTOL parts can systematically be investigated in detail, leading to a further understanding of the flow physics. With this knowledge, new and innovative aircraft configurations can be designed which helps to accomplish the increasing demand for Urban Air Mobility (UAM).

Results and Methods

High-quality CFD simulations were used to analyze the complex aerodynamics around the 1:7 scaled helicopter wind tunnel model [1,2]. For scaled rotor head configurations, the dominant Mach number and Reynolds number effects cannot be represented along with the similar advance ratio. Furthermore, most scaled rotor blades are too big for common wind tunnel facilities. Therefore, the rotor blades are cut in spanwise direction. By truncating the blades, the blade tip vortices arise closer to the main body compare to a rotor blade with full length. Therefore, the wake of the blade stubs interacts with the rear part of the helicopter, especially the fins, differently than in full-scale. This leads to unintended aerodynamic behavior on the vertical and horizontal stabilizers as well as the inflow of the Fenestron.

ANSYS Fluent was used for the meshing and solution process. The mesh is an unstructured poly-hexcore mesh containing approximately 33.7 Mio cells. The numerical simulations are based on the unsteady Reynolds Averaged Navier-Stokes equations. The turbulence is modeled by the Scale-Adaptive Simulation (SAS) approach. The rotation of the rotor head is modeled using the sliding mesh technique. For the helicopter simulation, 1,056 cores were used on 22 nodes. Overall, ten rotor revolutions were simulated.

Figure 1 shows the instantaneous vortical structures occurring around the rotor head. The wake patterns are visualized by using the Q-criterion at the time instant at which one blade is positioned in the forward direction. Figure 1 depicts the resulting iso-surface colored with the dimensionless axial velocity. One can detect the blade tip vortex on the advancing blades. In addition, highly turbulent structures arise around the hubcap.

A detailed flow field understanding can be achieved when averaging the flow quantities for every blade passing. The resulting iso-surface is shown in Fig. 2. Similar to the instantaneous velocity field, the tip vortex is detected as a dominant flow structure. It originates at the advanc-

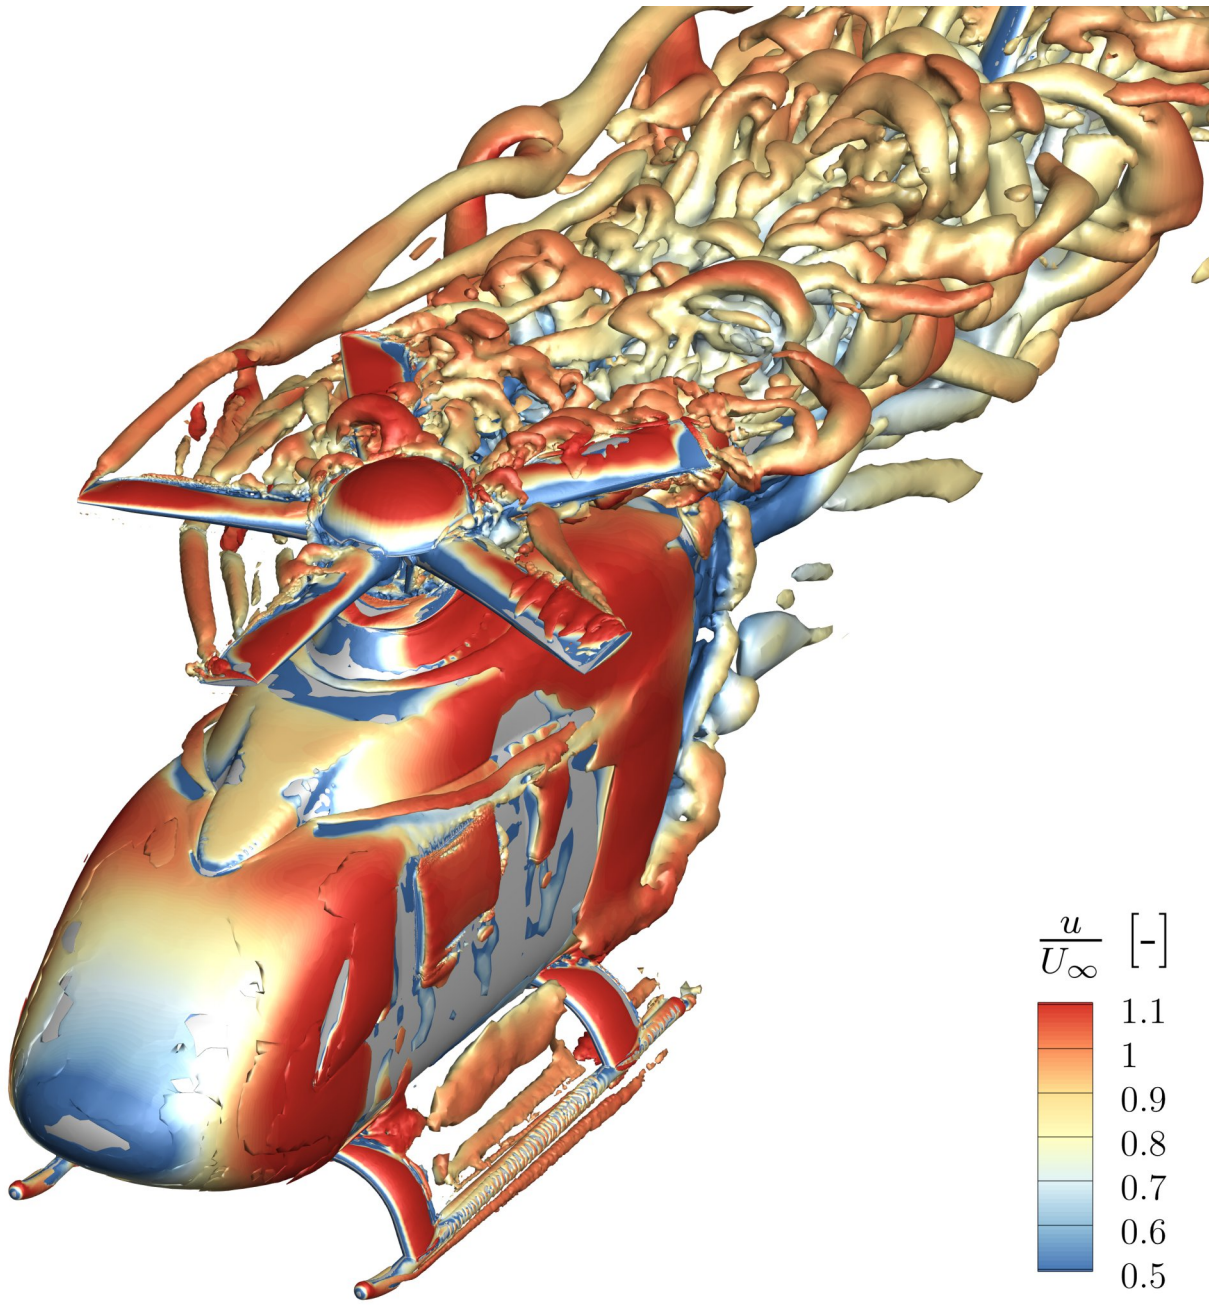

Figure 1: Light utility helicopter at cruise flight condition; flow scenario at corresponding 1:7 scaled wind tunnel model: Iso-surface of the Q-criterion colored by the axial velocity component.

ing blade and propagates downstream until it passes the horizontal tail. On the retreating blade side, distinct vortical structures are detected, which extend in the lateral direction. Further downstream, these structures interact directly with the vertical and horizontal tail. This interaction leads to high vibrations and can affect the pitching moment characteristics of the wind tunnel model.

Propeller simulations are another field of interest within this project. To overcome the challenges that arise with UAM vehicles, the interactional aerodynamics, especially between propeller-propeller and propeller-wing, are investigated in the second subproject.

An example of propeller-wing interaction is shown in Fig. 3. The three-bladed propeller operates in front of the wing and is connected to it via a nacelle. The wake of the propeller is strongly interacting with the wing.

The mesh was performed using ANSYS Fluent Meshing, consisting of approximately 22.1 Mio Cells. The propeller counts approximately 9.9 Mio Cells, while 12.2 Mio cells resolve the outer domain including the wing. For the propeller simulation, 864 cores were used on 18 nodes. Overall, 17 propeller revolutions were simulated.

The instantaneous flow field around the propeller-wing configuration is demonstrated by the non-dimensional

Q-criterion. The iso-surface is colored by the axial velocity component. In addition, the surface of the wing and propeller is colored by the static pressure.

The tip vortices of the three-bladed propeller are identified as the dominant structures. In addition, small-scaled vortical structures arise on the blade's root position. On the wing's leading edge, the helical structure of the tip vortices is subdivided and propagates on the upper and lower part of the wing. In addition, the wake of the spinner further influences the pressure distribution on the wing's surface.

Ongoing Research / Outlook

SuperMUC-NG enables the complex and CPU-intensive flow simulation performed within this project. At the current stage, detailed analyses of different flight conditions are conducted for the helicopter configuration as well as for propeller configurations shown in this article. More high-quality simulations need to be performed to understand the arising flow phenomena further and, therefore, contribute to developing green and safe air transport.

References and Links

[1] Hartmann, Ulrich; Breitsamter, Christian: Helicopter Rotor Wake Investigations on a Wind Tunnel Model with Varying Blade Stub Lengths. FORUM 2023 - Vertical Flight Society 79th Annual Forum and Technology Display, Vertical Flight Society, 2023, https://www.scopus.com/inward/record.uri?eid=2-s2.0-85167699511&partnerID=40&md5=c7546b1bc754cb20e856704f71d166aa

[2] Boldrin, Manuel: Numerical Investigation of Helicopter Rotor Wake, 2023, Master Thesis.

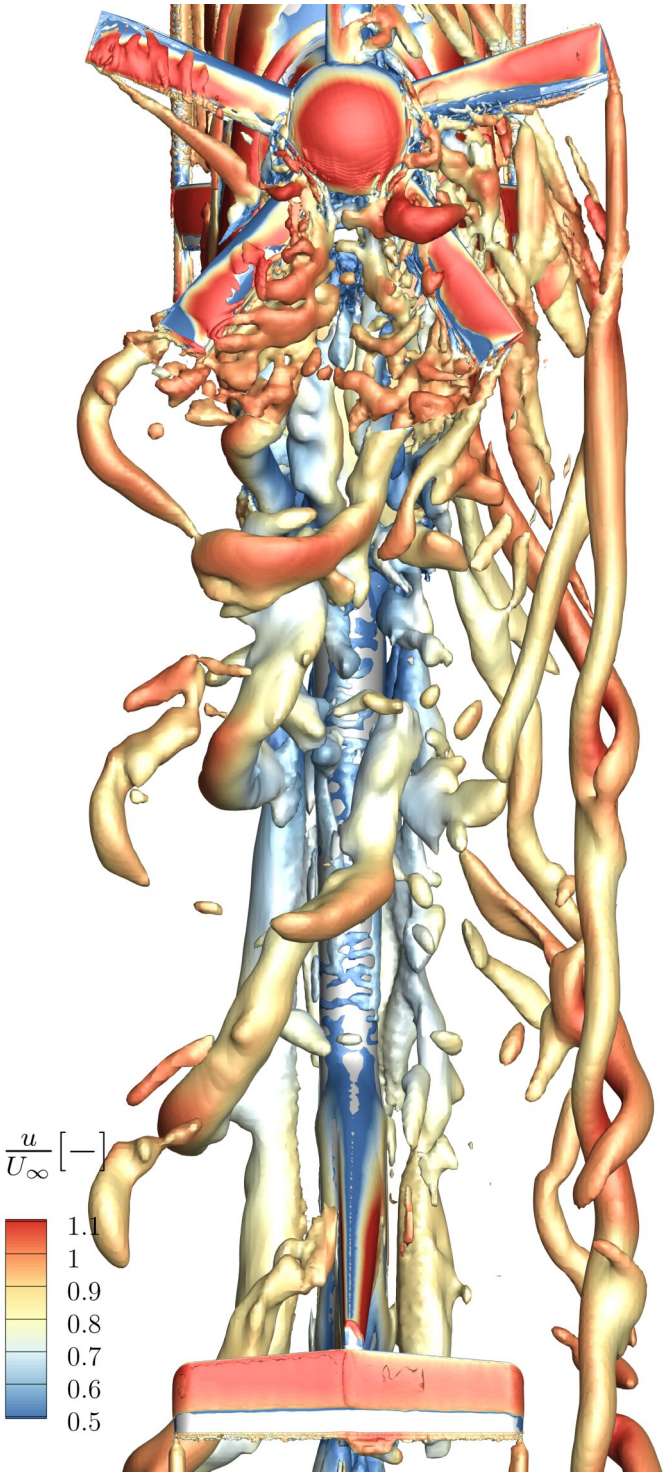

Figure 2: Rotor head wake flow of light utility helicopter at cruise flight condition: Iso-surface of the mean Q-criterion colored by the axial velocity component.

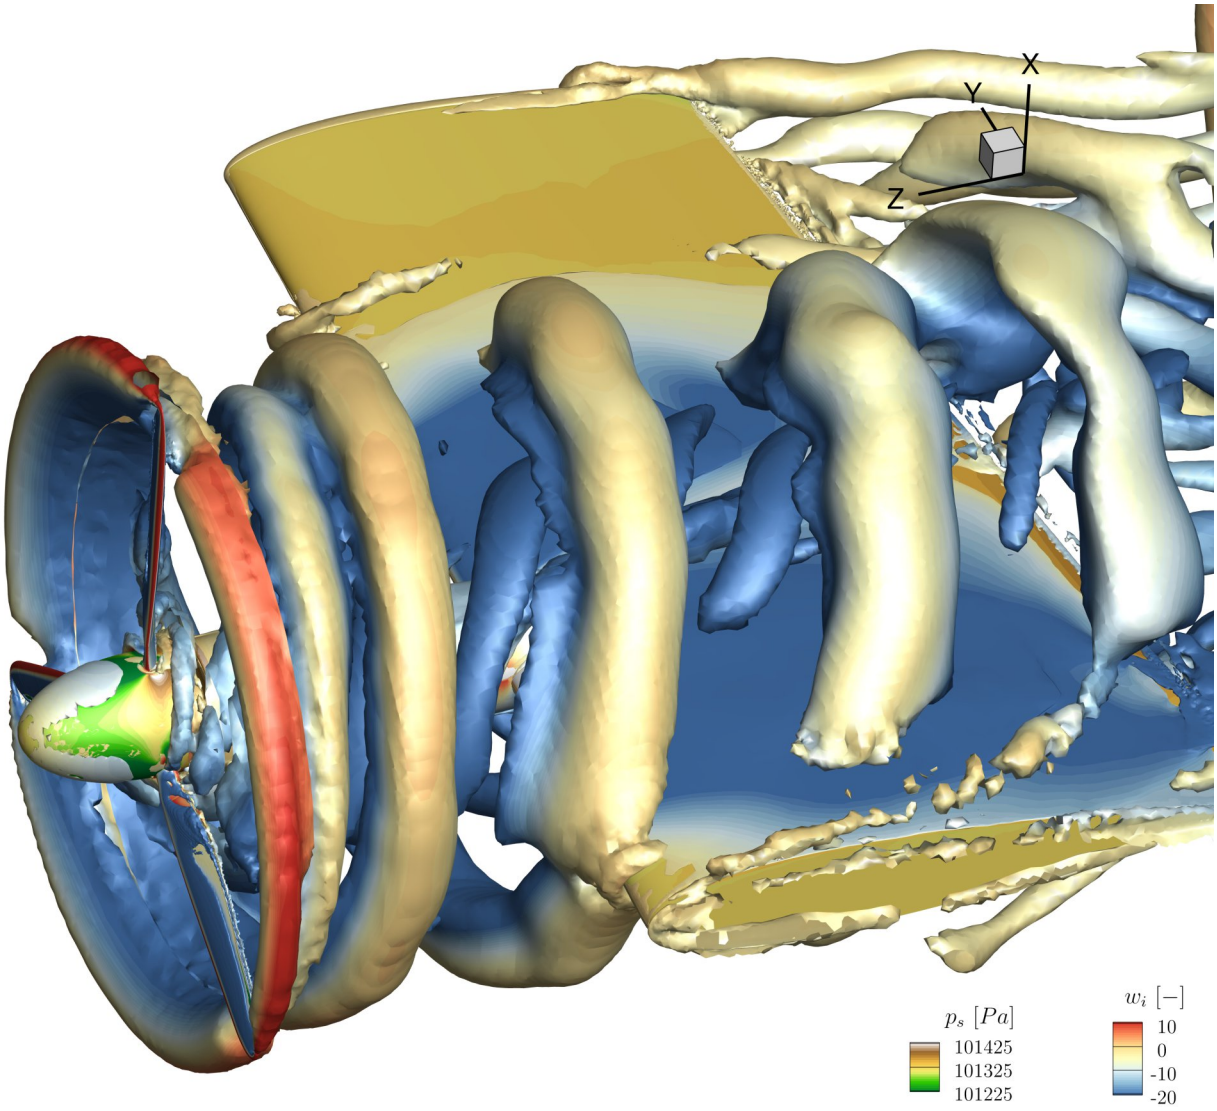

Figure 3: Propeller-wing interaction: Iso-surface of the Q-criterion colored by the axial velocity component.

Generation of a High-Fidelity Database for Delta Wing Flow

Research Institution
Numerical Methods in Aerospace Engineering, University of the Bundeswehr Munich

Principal Investigator
Tony Di Fabbio

Researcher
Karthick Rajkumar, Eike Tangermann, Markus Klein

Project Partner
–

SuperMUC Project ID: pn29xa

Introduction

Modern combat aircraft development integrates cutting-edge technology like stealth capabilities, advanced avionics and enhanced flow controls systems for improving maneuverability at high speed. The FCAS project represents a collaborative effort among European nations, bridging industry and academia, to develop a new generation of fighter aircraft. Agile delta wing aircraft configurations operate at high angles of attack, with complex flow fields dominated by vortex systems. Aircraft configurations, particularly those featuring agile delta wings operating at high angles of attack, exhibit complex flow fields dominated by vortex systems.

The precise understanding of the generation and evolution of leading-edge vortices, shock waves, and their interactions under transonic conditions are of critical importance for advanced aircraft configurations. High-fidelity simulations using the SuperMUC-NG computational resources have been instrumental in enhancing the accuracy of predictions concerning flow physics phenomena. Employing the DLR-TAU flow solver, a computational tool developed by the German Aerospace Center, these simulations have provided comprehensive insights into the aerodynamics of complex configurations, including open cavities and multiple delta wing designs.

Results and Methods

CFD is vital for forecasting complex flow phenomena but remains costly. The complex turbulence fluctuations in the flow field are captured by the underlying turbulence models. Classical RANS models, while computationally efficient, often fall short in accurately capturing the flow dynamics, especially in predicting leading-edge vortices. To address these challenges and improve the numerical predictions, hybrid RANS/LES computations have been employed. Various turbulence models (detailed in refs. [2-5]), including RANS and hybrid RANS/LES, such as SA, k-ω SST, SAneg-IDDES, and SAS, have been compared to analyze the flow physics of these complex configurations.

Cavity flows
In exploring cavity flows, the study yielded pivotal insights. Firstly, determining the signal sample length is crucial for experimental validation, ensuring accurate results. Instead of modeling the entire turbulent spectrum, it is possible to resolve parts of the spectrum by means of a scale resolving simulation. SA-IDDES provides precise spectral predictions despite accompanying noise. Implementing a wall function can enhance computational efficiency by 50% while boosting robustness, underscoring the benefits of optimizing strategies. Investigation into SAS method variants broadened understanding of their applications. The wall-resolved approach showed 90% computational efficiency in predicting oscillation and cavity physics, highlighting its potential in simulating complex flow dynamics precisely. Careful setup of shear layer forcing was identified as effective, emphasizing meticulous highlights setup importance. Presence of spanwise waves, especially under sideslip conditions, was observed, strengthening coherence structures and mode number correlation understanding. Despite challenges like higher time step sizes and coarse resolutions, frequency accuracy remained intact, with noted moderate overprediction suggesting areas for refinement. These findings significantly contribute to fluid dynamics, particularly cavity flow study, offering a comprehensive understanding of factors influencing experimental and computational outcomes.

Unsteady Effects
In the exploration of leading-edge vortex flows, a study revealed pivotal observations contributing to aerodynamic understanding. Both the kwSST and SAS models accurately predict shocks and vortex breakdown, emphasizing turbulence model selection importance for capturing complex flow features. A critical finding is the need for mesh refinement near vortex breakdown, impacting aerodynamic coefficient oscillation and simulation accuracy. In double-delta configurations, shocks and vortex breakdown oscillate around half the chord length, requiring longer flow-through times for accurate capture. Fundamental aerodynamic differences between double and triple delta wings were identified. double-delta investigations didn't conclusively determine shocks' role

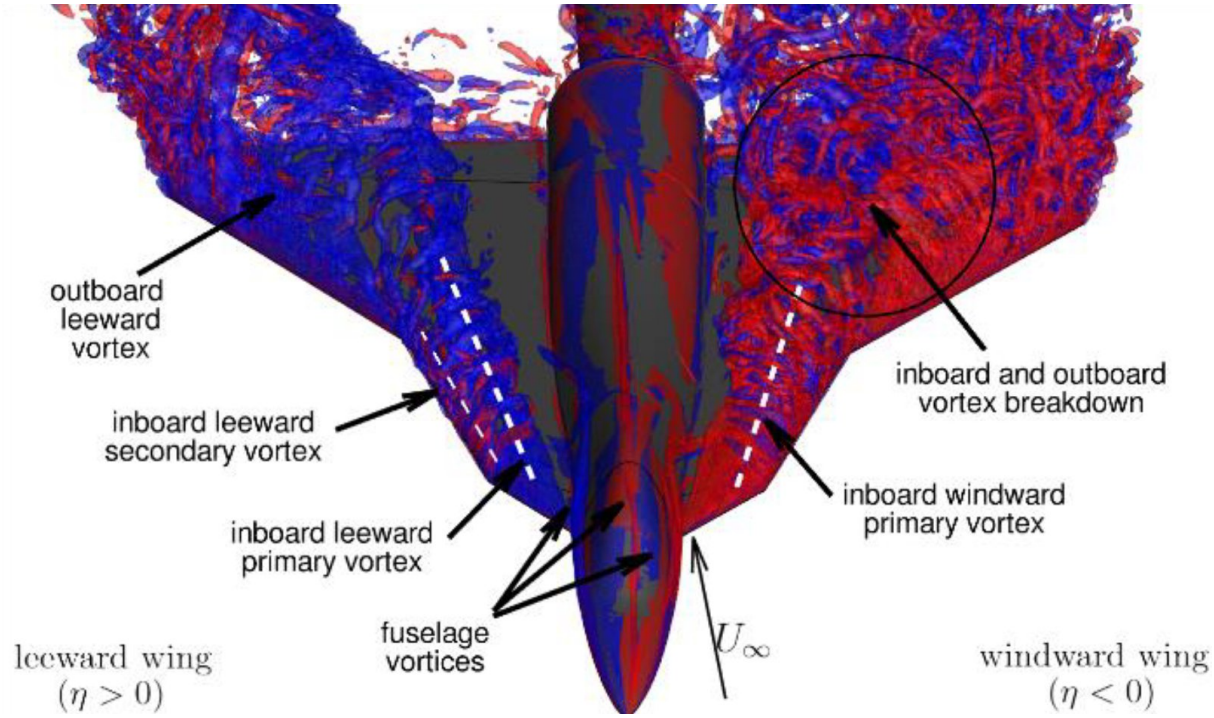

Figure 1: Multiple delta-wing aerodynamics [5].

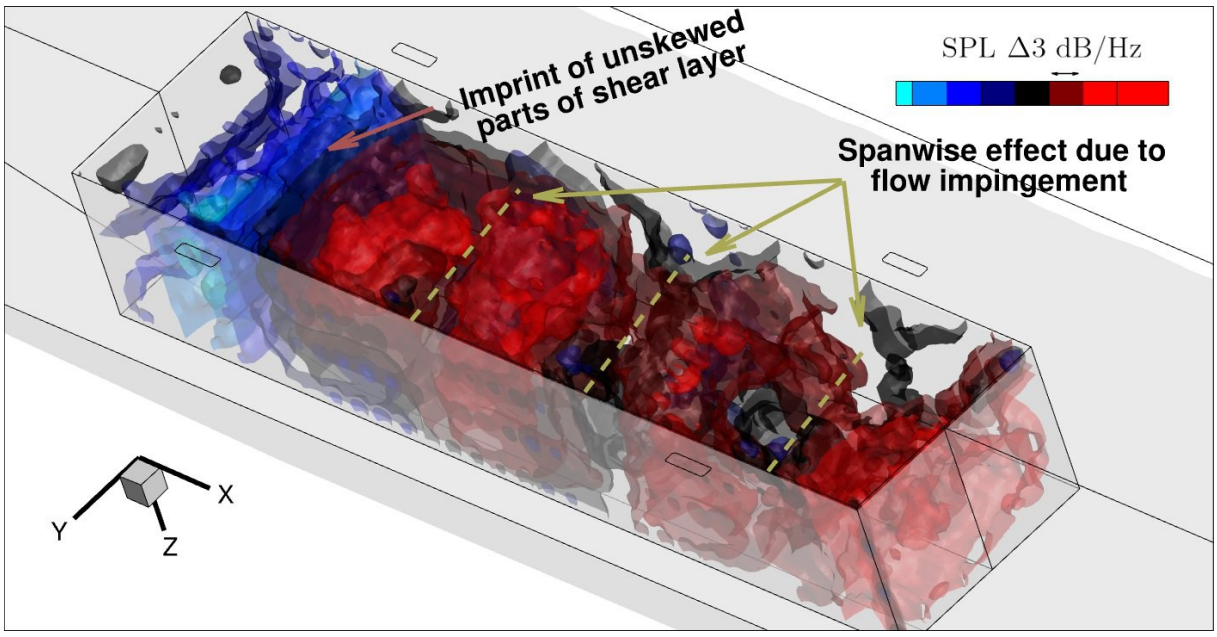

Figure 2: Spectral shapes in the cavity [3].

in vortex breakdown, while triple-delta configurations pinpointed kink-induced shocks as primary factors. Adding a strake in triple-delta mitigated shock effects on the inner vortex boundary, highlighting the potential of aerodynamic modifications in flow stability improvement. Post-processing scripts/templates have been developed for extracting vortex core properties, visualizing shock systems, and analyzing vorticity and enstrophy transport equations, aiding in parametric studies involving circulation and shock strength. These advancements enhance understanding of leading-edge vortex flows, providing insights into shock and vortex interactions' mechanisms.

Turbulence Modeling

In the context of generic (multiple-) delta-wing configurations [4,5], scale-resolving simulations have shown a pronounced sensitivity to both spatial and temporal resolutions. These simulations, particularly around the VFE-2 delta wing, have been instrumental in

depicting the pattern of leading-edge vortices and their subsequent interactions with shocks under transonic conditions. Through the analysis of several flow physics phenomena, including vortex-shock interaction, cross-flow shock waves, stream-wise separation, and the primary feeding of the vortex, the research has illuminated the complex mechanisms governing these flows. Notably, the turbulent shear-layer and stream-wise boundary layer separation have been identified as significant factors contributing to the vortex breakdown phenomenon, providing fresh insights into its underlying causes [4,5]. Consequently, a new theory and explanation for the phenomenon of vortex breakdown, a topic that has intrigued scientists for decades, has been discussed.

Further analysis has addressed the limitations in the predictive accuracy of RANS simulations, with a particular focus on the SA one-equation model. Despite its widespread adoption in various aerodynamic studies, some of the SA model's limitations in accurately capturing the intricate details of flow phenomena have been identified, primarily attributed to its treatment of turbulence. This includes the model's handling of eddy viscosity and the Reynolds stress tensor. The examination reveals that the simplifications inherent in the SA model's approach to turbulence modeling contribute to its inability to predict complex flow dynamics with high fidelity. A straightforward modification to the SA model is proposed to enhance its performance and provide deeper physical insight into flow phenomena. This modification aims to improve the model's capacity to accurately capture stream-wise boundary layer separation and vortex breakdown, thereby enhancing the overall prediction of aerodynamic coefficients [5].

Ongoing Research / Outlook

In the rapidly advancing field of aerodynamics research and turbulence modeling, the integration of machine learning techniques with CFD represents a notable progression. A novel framework using machine learning within a CFD-driven context to refine turbulence modeling in aerodynamic flows, aiming to foster a comprehensive understanding of the application of machine learning in fluid dynamics, will be introduced. The exploration extends to highlight the expansive application and significant impact of machine learning methodologies in aerodynamics research. Finally, this study will aim to enhance the framework's generalizability across various test cases, thereby improving its adaptability and prediction accuracy across a wide range of aerodynamic scenarios.

References and Links

[1] https://www.unibw.de/numerik/forschung/forschung
[2] K. Rajkumar et al. “Time-efficient simulations of weapon bay in fighter aircraft”, CEAS Aeronautical Journal (2023).
[3] K. Rajkumar et al. “Efficient Scale-Resolving Simulations of Open Cavity Flows for Straight and Sideslip Conditions”, Fluids (2023).
[4] T. Di Fabbio et al. “Analysis of the vortex-dominated flow field over a delta wing at transonic speed”, The Aeronautical Journal (2023).
[5] T. Di Fabbio et al. “Towards the understanding of vortex breakdown for improved RANS turbulence modeling”, Aerospace Science and Technology (2024).

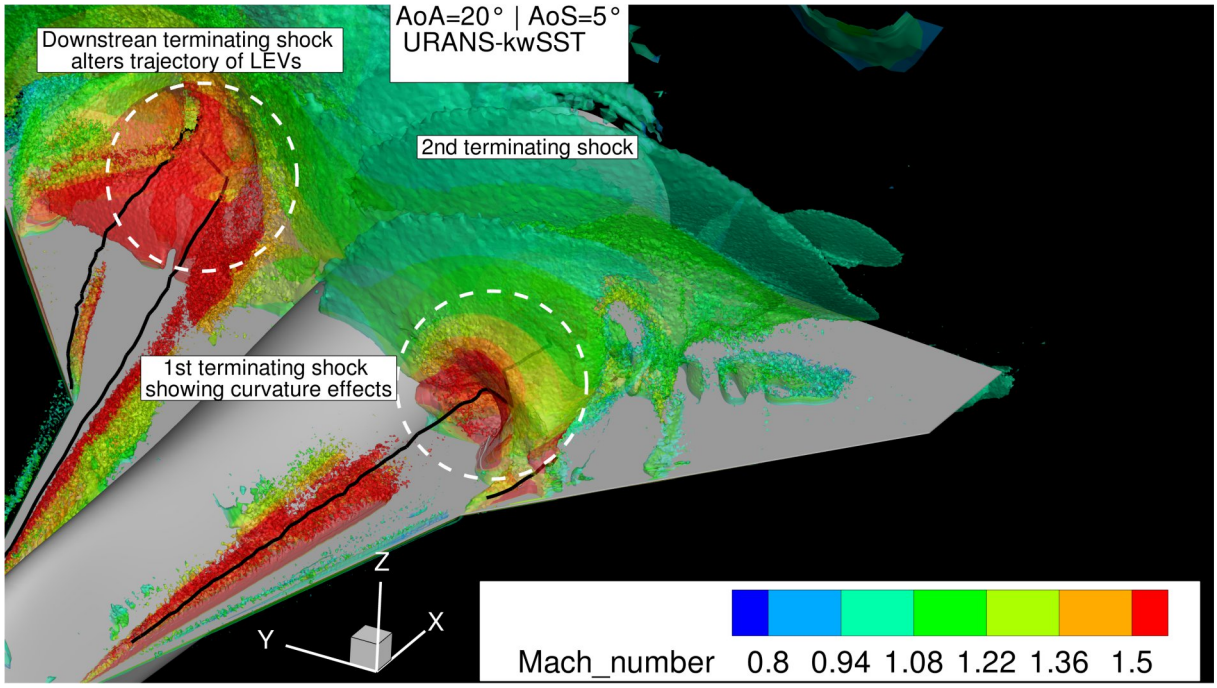

Figure 3: Shock-vortex interaction in a delta-wing planform.

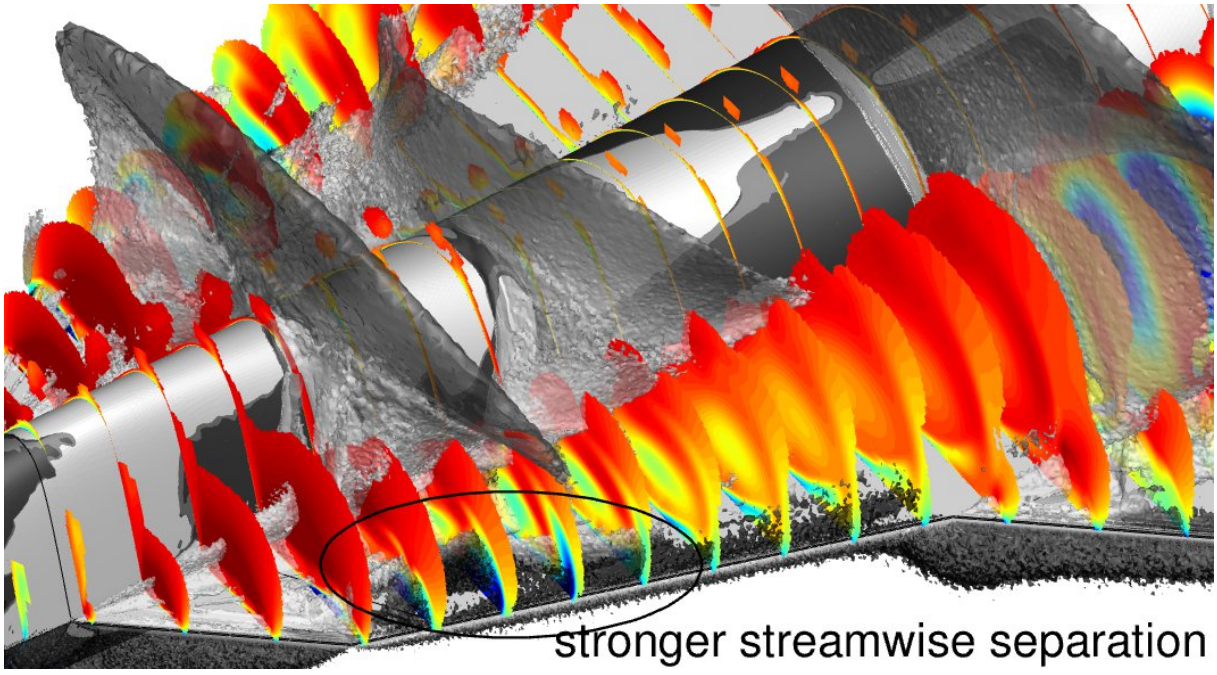

Figure 4: Stream-wise boundary layer separation and vortex breakdown prediction [5].

Particle-resolved simulation of antidunes in free-surface flows

Research Institution
[1]Chair for Computer Science 10 - System simulation, Friedrich-Alexander-Universität Erlangen-Nürnberg

Principal Investigator
Harald Köstler[1]

Researcher
Christoph Schwarzmeier[1], Christoph Rettinger[1], Samuel Kemmler[1], Jonas Plewinski[1], Francisco Núñez-González[2], Ulrich Rüde[1,3], Bernhard Vowinckel[4]

Project Partner
[2]Universitat Politécnica de Catalunya – BarcelonaTech (UPC)
[3]CERFACS, Toulouse
[4]Technische Universität Braunschweig

Funding
DFG: 408062554, 434946896, 433735254 and 428445330; SCALABLE project; European High-Performance Computing Joint Undertaking (J.U.) Grant Agreement No. 956000

SuperMUC Project ID: pn29po

Introduction

This project presents the first numerical attempt to simulate upstream-migrating antidunes with geometrically resolved particles and a liquid–gas interface. We base this section on the article of Schwarzmeier et al. [2]. To manage the simulations' computational costs and physical complexity, we employed the cumulant lattice Boltzmann method in conjunction with a discrete element method for particle interactions, as well as a volume of fluid scheme for tracking the deformable free surface of the fluid. The simulations carried out for this sub-project aimed to reproduce the laboratory experiments of Pascal et al. [3] who performed their experiments with water and a sediment bed of natural gravel.

Results and Methods

Figure 2 illustrates a three-dimensional view of the total simulation domain. The undulations of the free surface are in phase with those of the sediment bed, conforming the appearance of antidunes.

The alternating blue and yellow (dark and light) diagonal strips denote the troughs and crests of the bedforms, respectively. The experimental and numerical diagrams with the evolution of the bottom elevation in the spatial–temporal domain are in good qualitative agreement. A similar upstream migration trend of the bedforms is clearly visible due to the negative slope of the strips. In some regions the strips bend, indicating acceleration/deceleration, or even stationarity of the bedforms. Overall, the simulated bedforms appear stable and compare well to their experimental counterparts. The vast amount of data generated by our simulations in a non-intrusive manner can be of great use to supplement experimental measurements under challenging supercritical flow conditions.

The detailed data of particle and fluid motion available through our simulation approach opens a multitude of possibilities. The present study, therefore, encourages further simulation campaigns to understand antidune mechanics and the physical controls on bedform initiation and morphodynamics in supercritical flows.

Ongoing Research / Outlook

Based on our validated proof-of-concept simulation model, we aim to perform further large- scale simulation studies in future work. These simulation campaigns will help understanding antidune formation and propagation. We aim to use the available compute time of this LRZ project to work on this research direction. Accordingly, we extended our project.

References and Links

[1] https://www.cs10.tf.fau.de/
[2] Schwarzmeier, C. et al., B. (2023). Particle-resolved simulation of antidunes in free-surface flows. Journal of Fluid Mechanics, 961, R1.
[3] Pascal, I., Ancey, C., & Bohorquez, P. (2021). The variability of antidune morphodynamics on steep slopes. Earth Surface Processes and Landforms, 46(9), 1750-1765.

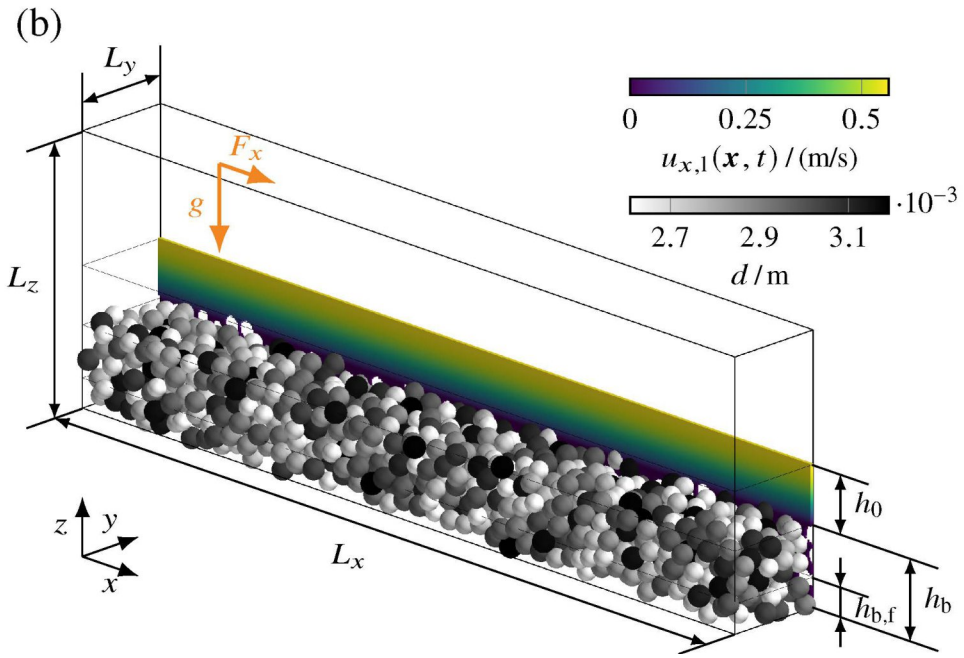

Figure 1: Simulation setup in its initial condition as a zoom into the computational domain covering only 16 % of its streamwise extent.

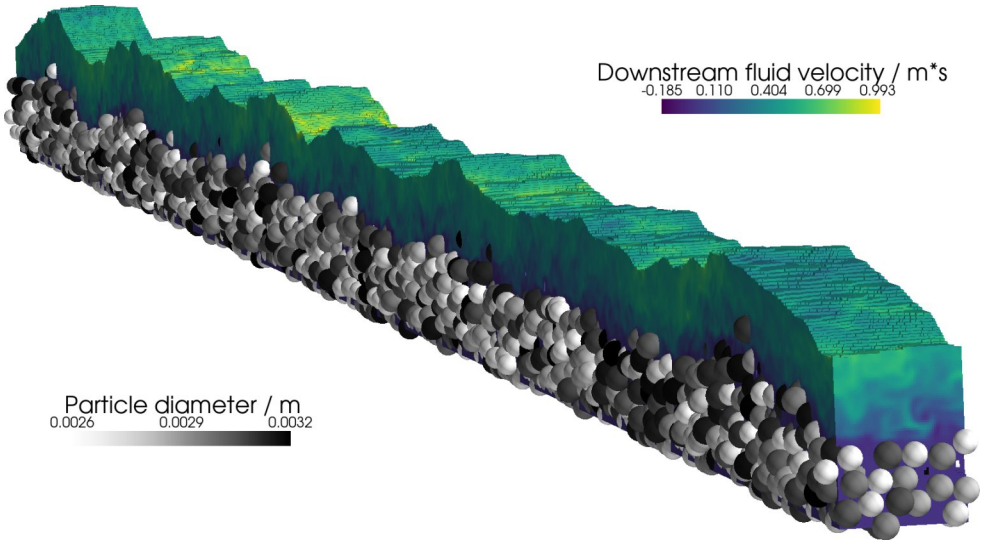

Figure 2: Visualization of the antidunes simulation with the simulated velocity in streamwise direction and the diameters of the spherical particles. The undulation of the sediment bed and free surface of the liquid are in phase, conforming to the definition of antidunes.

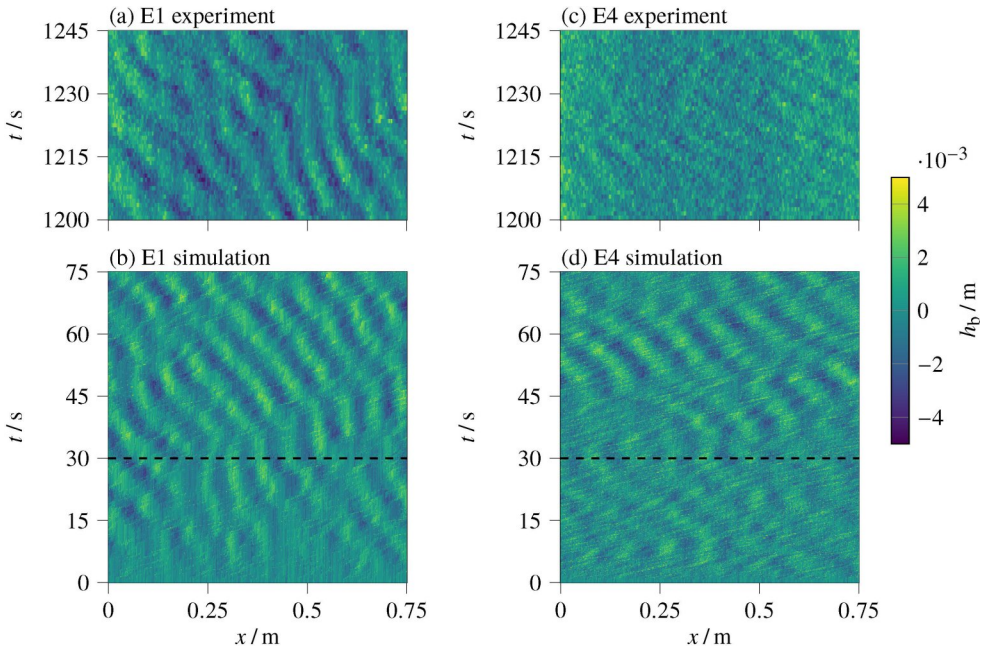

Figure 3: Sediment bed elevation for experiment E1 (left) and experiment E4 (right): a and c data from Pascal et al. [2], and b and d data from the numerical simulations. We consider the system fully developed after t = 30 s, as illustrated by the dashed black line.

Magnetic-field-gradient and sidewall effects in magnetoconvection

Research Institution
[1]Technische Universität Ilmenau

Principal Investigator
Thomas Boeck[1]

Researcher
Shashwat Bhattacharya[1,2], Dmitry Krasnov[1]

Project Partner
[2]Amity University, Noida, India

Funding
Alexander von Humboldt Foundation

SuperMUC Project ID: pn49ma

Introduction

Flows that are driven by buoyancy forces are a common occurence in nature as well as in technological applications. The driving mechanism is the temperature dependence of the fluid density, which leads to density variations when heat is transported through the fluid. If the flowing fluid is electrically conducting and is under the influence of magnetic fields, electric currents are induced in the fluid due to Faraday's law, which, in turn, induce magnetic fields by the virtue of Ampere's law. These electric currents interact with the applied and induced magnetic fields to generate a Lorentz force distribution that acts on the fluid. Such flows, called magnetoconvection, are typically encountered in liquid-metal batteries, cooling liquid-metal blankets in fusion reactors, and magnetic stirring and braking of liquid metal melts in metallurgy. Magnetoconvection is also encountered in convection in the Sun, stars, and in planetary dynamos.

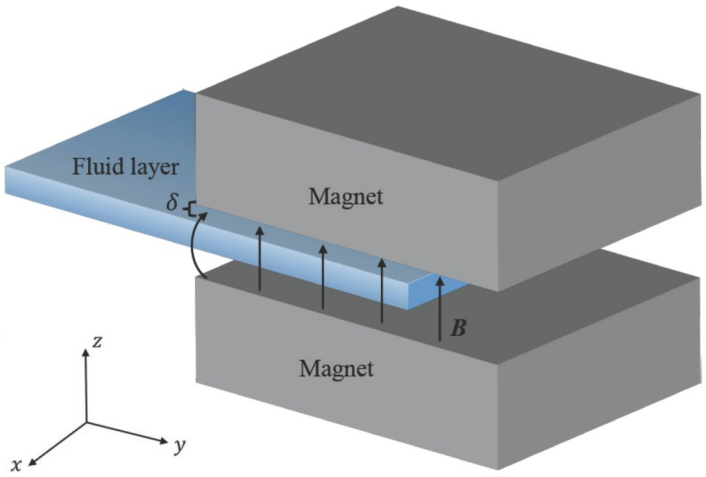

Figure 1: A schematic of the setup of RBC under spatially varying magnetic fields. The magnetic poles are semi-infinite in *y* and *z*-directions, and infinite along *x*-direction.

Our work is concerned with buoyancy-driven flows of a liquid metal (mercury) in a geometrically simple planar horizontal layer, where the bottom and top walls are heated and cooled, respectively. Such flows are called Rayleigh-Bénard convection (RBC). They emerge from a linear instability of the purely conductive state. An imposition of horizontal magnetic fields in an RBC setup results in the convection rolls aligned along the direction of the magnetic field. On the other hand, a vertical magnetic field suppresses convection, and above a certain threshold of magnetic field strength, the convection ceases completely in the bulk and gets confined near the sidewalls (if present) [1]. Our project focuses on the effect of magnetic field gradients on large-scale features of turbulent convection, particularly on their response to different strengths of forcing. Moreover, we explore how additional horizontal magnetic field components and finite wall conductivity change the flow structure of wall-attached convection for very strong vertical magnetic fields.

Results and Methods

We performed direct numerical simulations (DNS) of turbulent convection in a horizontally extended domain of dimensions $16 \times 32 \times 1$ under the influence of fringing magnetic fields. The simulations were conducted using a second-order finite-difference solver developed by Krasnov et al [2]. The magnetic field strength varied along the horizontal *y*-direction and was generated in the gap between semi-infinite poles of permanent magnets (see Figure 1). The convection domain was located near the edge of the gap.The gap δ between the magnetic poles and the thermal plates was nine times the height H of the cell. The DNS were conducted on a grid of $4,800 \times 9,600 \times 300$points, employing 9,600 cores in SuperMUC-NG and consuming a total of 5 million core-hours.

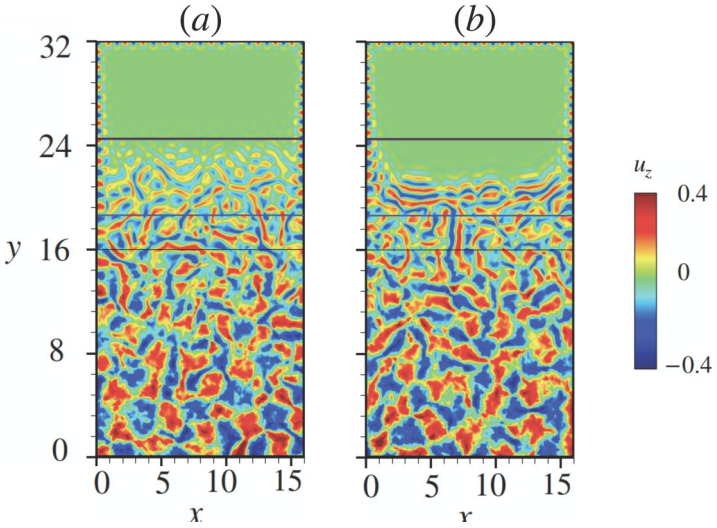

Figure 2: Contour plots of time-averaged vertical velocity u_z for $\delta = 9$ in the horizontal midplane for two different initial conditions (a) and (b). The magnetic poles extend from $y = 16$ to $y = \infty$. The large-scale patterns are qualitatively similar but differ in the alignment in the regions of large gradients in magnetic field ($y \approx 16$).

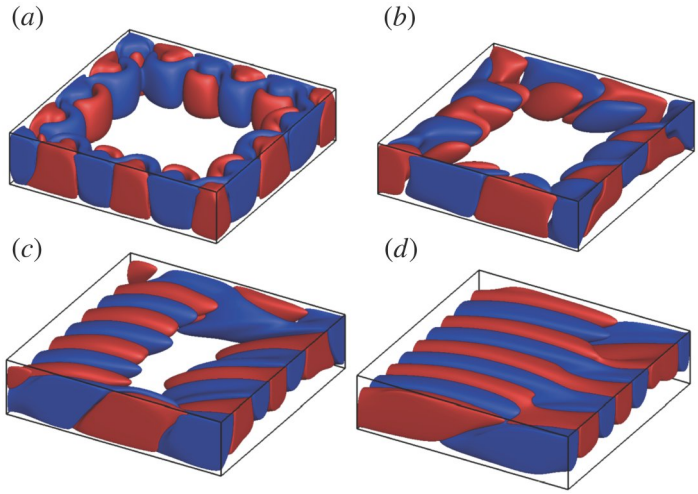

Figure 3: Results of DNS of wall-attached convection in a homogeneous magnetic field. Isosurfaces of the vertical velocity $u_z = 0.01$ (red) and $u_z = -0.01$ (blue) when (a) there is no horizontal magnetic field, (b) the horizontal and vertical magnetic fields are of equal strength, and the horizontal magnetic field is (c) twice and (d) thrice the vertical magnetic field. The wall modes get elongated along the direction of the resultant magnetic field.

We studied the evolution of large-scale structures and the resultant heat transport in the fringing magnetic field for two different initial conditions. We observed qualitative similarities between the large-scale patterns in the weak magnetic flux region for both initial conditions (see Figure 2). However, there was a difference in the alignment of the convection patterns in the regions of large gradients in the magnetic field. Despite this difference, the other results such as the global Nusselt and Reynolds numbers, and the variations of anisotropy and the local heat transport along the horizontal direction remained largely unchanged with initial conditions [3].

We further studied the effects of strong inclined magnetic fields on wall-attached magnetoconvection using a combination of linear stability analysis and DNS. The stability analysis assumed periodicity in the spanwise direction perpendicular to the plane of the homogeneous magnetic field. The study revealed that for a fixed vertical magnetic field, the imposition of horizontal magnetic fields results in an increase of the critical buoyant force required to sustain wall-attached convection. Further, a decrease in the wavelength of the wall modes was also observed. The wall modes become tilted along the direction of the resulting magnetic fields and therefore extend further into the bulk as the horizontal magnetic field is increased. However, once the modes localized on the opposite walls interact, the critical buoyant force decreases again and eventually drops below the value for onset with a purely vertical field.We find that for sufficiently strong horizontal magnetic fields, the steady wall modes occupy the entire bulk, i.e. convection is no longer restricted to the sidewalls.

The stability results are confirmed by DNS of the nonlinear evolution of magnetoconvection (see Figure 3). These DNS were conducted with the same code [2]. A total of 12 simulations were performed, with 6 simulations employing a grid-resolution of 1,200×1,200×300 points and utilizing 1,200 cores each, and the remaining 6 employing a grid resolution of 1,800×1,800×400 points and utilizing 3,600 cores each. The mesh was non-uniform with stronger clustering of the grid points near the boundaries. A total of 9 million core-hours were used in our simulations. The DNS results also reveal that, at least for large values of horizontal magnetic field, the wall-mode structures and the resulting heat transfer are dependent on the initial conditions [4].

Ongoing Research / Outlook

Our work highlights the effects of magnetic field's gradients on the evolution of large-scale convection patterns and the resultant heat and momentum transport. Further, we also show the characteristics of wall-attached convection under the influence of strong horizontal magnetic fields. Such flows are expected to be relevant for confined geometries, such as cooling blankets in fusion reactors and liquid metal batteries.

The remaining core hours will be required for analyzing the influence of different strengths of forcing on turbulent convection with spatially varying magnetic fields. The aspect ratio will be reduced along the horizontal x-direction perpendicular to the gradient of the magnetic field in order to reduce the computational costs per simulation.

References and Links

[1] W. Liu, D. Krasnov and J. Schumacher, J. Fluid Mech 849 (2018) R2.
[2] D. Krasnov et al, J. Comput. Phys. 474 (2023) 111784.
[3] S. Bhattacharya et al, J. Fluid Mech 964 (2023) A31.
[4] S. Bhattacharya et al, J. Fluid Mech 979 (2024) A53.

Numerical investigation on flashback mechanisms in premixed H_2/air swirl combustion

Research Institution
Institute for Simulation of reactive Thermo-Fluid Systems, TU Darmstadt

Principal Investigator
Christian Hasse

Researcher
Wang Han, Arne Scholtissek

Project Partner
–

SuperMUC Project ID: pn29so (Gauss Large-Scale project)

Introduction

Hydrogen (H_2) has the potential of emerging as the leading energy carrier for next-generation, zero-carbon power generation, and hence has received considerable attention. H_2 can offer significant benefits over hydrocarbon fuels, such as wide flammability range, low ignition energy, and high diffusivity. However, the use of H_2 in gas turbines poses considerable challenges, e.g., the risk of flashback due to its high flame speed, which adversely affects the performance of H_2 combustion [1].

Flashback, a problem that occurs in premixed combustors, is the upstream propagation of the flame from the combustor into the premixing tube due to the change in mass flow rate, which could change the combustion process and pollutant emissions as well as cause considerable damage on the combustor [1]. Therefore, a deep understanding of the flashback mechanisms is required before applying pure H_2 or H_2 enriched fuels in combustion systems.

Flashback can occur because of four main mechanisms [1]: (1) boundary layer (BL) flashback in non-swirling or low swirling flow, (2) combustion induced vortex breakdown (CIVB) flashback in a swirling flow without a bluff-body, (3) combustion instability induced flashback, and (4) combined BL/CIVB flashback in a bluff-body swirl burner. The focus of this work is the latter one, because bluff-body swirling flows are typically employed to enhance mixing and flame stabilization.The systematic investigation of BL flashback started with the work by Lewis and von Elbe [2], who proposed a classical critical gradient model to evaluate the flashback in channel flows. In the confined BL flashback configurations where the flame is already inside the premix duct before flashback, the formation of reversed flow pockets ahead of laminar/turbulent flames is one of the main characteristics prior to the onset of flashback, which is attributed to a pressure rise induced by the formation of flame bulges and has been reported both experimentally and numerically. On the other hand, flashback in a swirl combustor without a central bluff body has been studied extensively. It is concluded that the flashback is related to vortex breakdown (i.e., CIVB flashback). This is because the baroclinic torque can produce a negative azimuthal vorticity and then induce a negative axial velocity along the vortex axis, leading to the vortex breakdown and facilitating flashback.

Despite its high relevance for various practical combustor designs, there are only a few studies on flashback in a bluff-body swirl burner. The Clemens group [3] experimentally identified two modes of flame propagation during flashback, i.e., (1) small-scale bulges propagating in the negative streamwise direction; and (2) large-scale flame tongues swirling with the bulk flow while leading flashback. More recently, a new flashback mode was discovered by Ebi et al. [4], and they found that the upstream propagation of flame can be led by flame bulges. Due to the inherent three-dimensional and transient nature of flame flashback, 3D parameter measurements and high-speed detecting devices are required, which leads to a significant challenge for experimental investigations of flashback.

It is noted that the above investigations of flashback in bluff-body swirl burners only provided 2D or local 3D flow field (without flame) measurements due to the limitation of laser-based techniques, which cannot comprehensively capture flow-flame interaction during the swirling flashback. The overarching objective of this project is to provide for the first time a detailed understanding of complex flow-flame interaction during flashback in a bluff-body swirl burner using Direct Numerical Simulation (DNS), which would help to reveal the mechanism(s) allowing the flame to propagate upstream during flashback and to improve practical swirl burner designs.

Results and Methods

In this project all simulations and analyses are based on direct numerical simulations (DNS) which are conducted with the DNS code DINO. The solver is designed for the

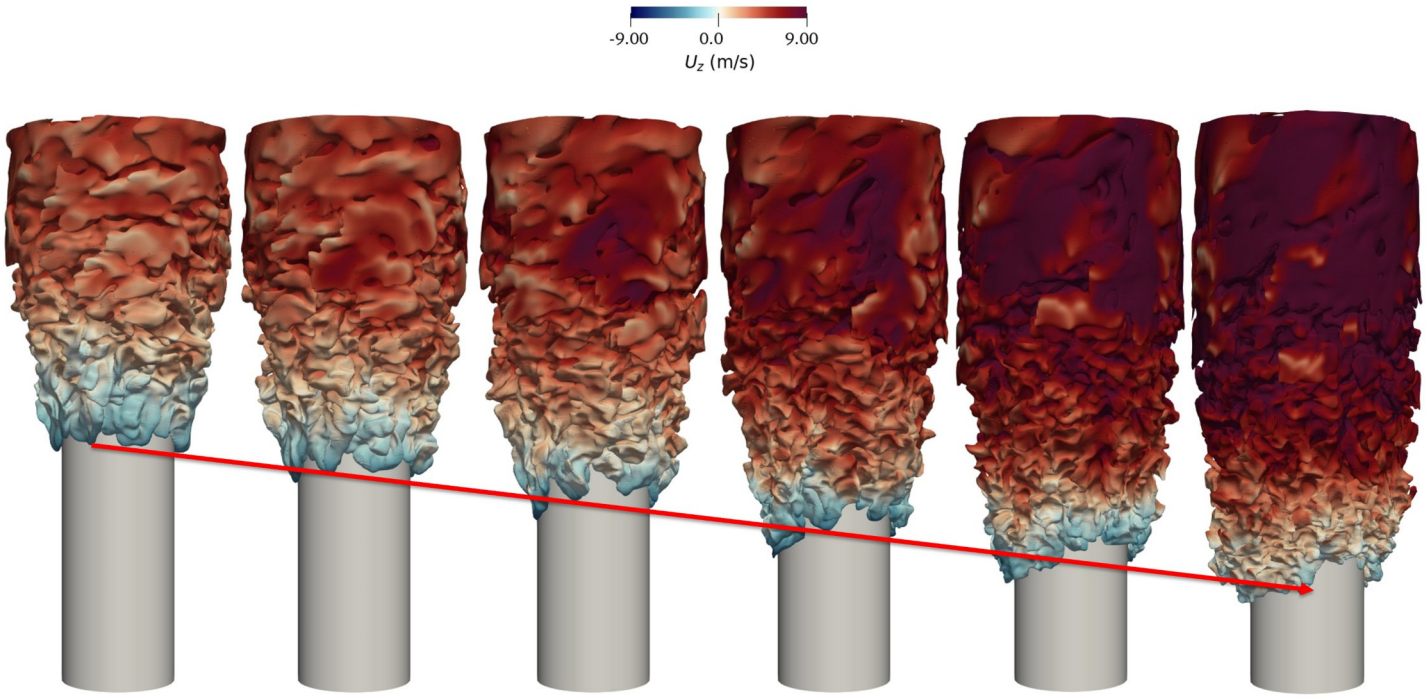

Figure 1: Evolution of flame front colored by the axial velocity in the flashback case.

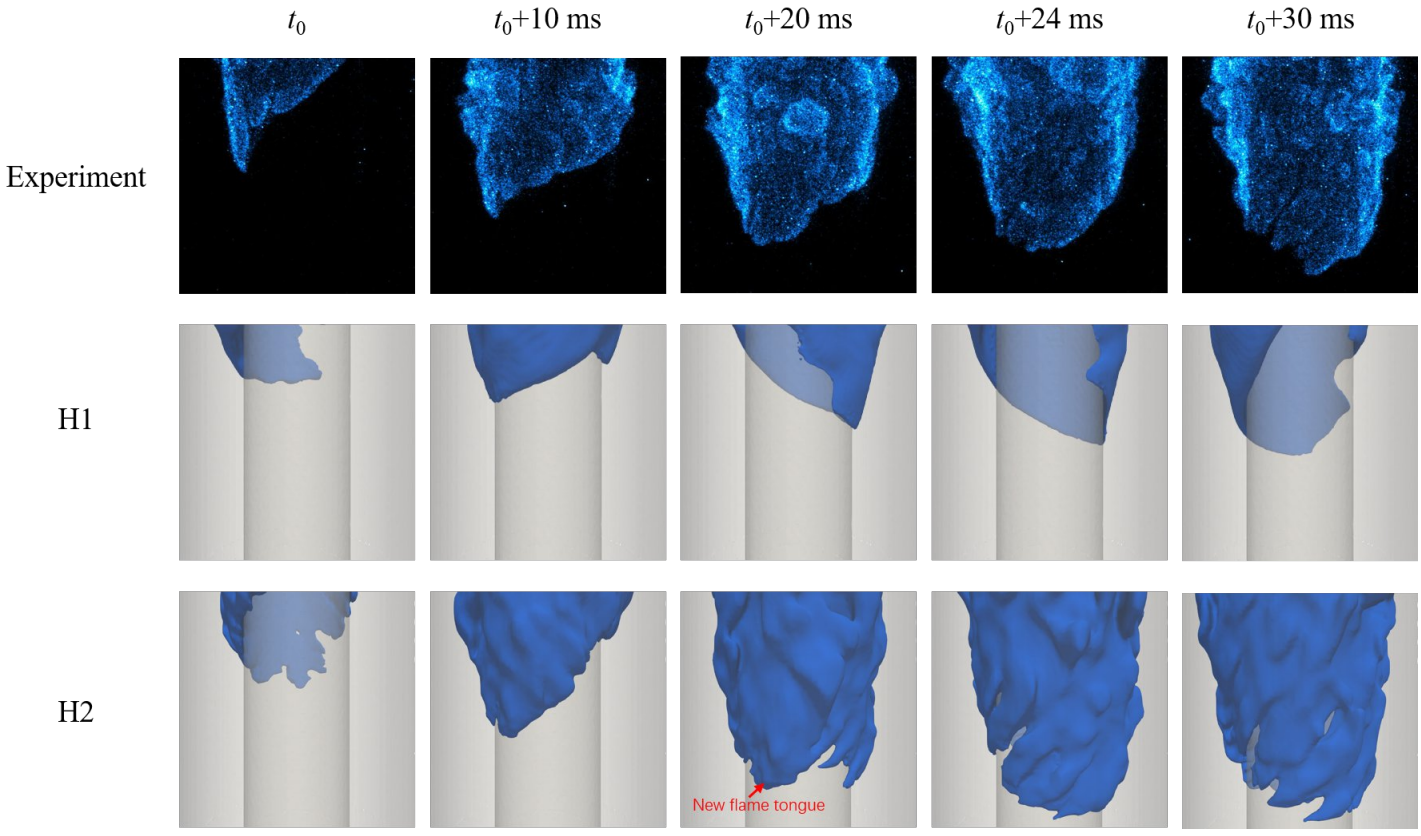

Figure 2: Comparison between instantaneous flame front over time of case H1 (without considering differential diffusion) [5], case H2 (considering differential diffusion) [5] and experimental chemiluminescence images [3].

simulation of low-Mach number reactive flows, where spatial derivatives in the governing equations are discretized with 6th order finite differences. The temporal integration is done by a 3rd order semi-implicit Runge-Kutta scheme. Based on the distributed memory architecture of SuperMUC parallelization of the solver is achieved by the message passing interface (MPI), where an excellent scalability up to 65,536 cores is achieved.

The DINO 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. An implicit time integration of the stiff chemical source terms has been implemented, relying on a semi-implicit Runge-Kutta 3rd-order. In the DINO code, 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.

It is noted that a cylindrical DINO version has been developed with the help of the project to achieve DNS calculations of flashback efficiently.

DNSs of burner stabilized flame and flashback
In this project, three DNSs are performed: one DNS of a stable, lean premixed swirling hydrogen/air flame in the bluff-body swirl burner; and two DNSs of lean premixed swirling hydrogen/air flame flashback in the bluff-body swirl burner with/without considering differential diffusion effects, respectively. Flashback is triggered by disturbing the stable burning flame through impulsively increasing the equivalence ratio uniformly at the inlet plane from 0.5 to 0.8. The computational domains comprise a region of 291×72×72 or 350×88×88 flame thicknesses for the first DNS or second DNS, respectively. They require 3,200×900×900 and 3,500×1,000×1,000 grid points for the first DNS and second DNS, respectively. Concerning the overall simulation time, the DNS is estimated to run for about 150 ms and 200 ms for stable flame and flashback DNS calculations, respectively. The stable burning case is found to have reached a statistically steady state after 150 ms and the flashback process lasts over a time span of 200 ms which was assessed qualitatively by a visual inspection of the flame front evolution during flashback.

It is seen from Fig. 1 that for the flashback case the flame can propagate upstream. The flame front is swirling around the center body. It is noted that owing to a higher axial velocity near the outside wall, the flame mainly propagates near the bluff-body wall.

Differential diffusion effects
The transient flame front of case H1 (without considering differential diffusion) and case H2 (considering differential diffusion) as well as the instantaneous chemiluminescence images of the same flame extracted from experiments are shown in Fig. 2 [5]. When unity Lewis number is considered, the flashback speed is far from the experimental value (0.94 m/s). Besides, compared with experimental chemiluminescence results, the flame front of case H1 is much less wrinkled. Different from CH_4/air flames, the preferential diffusion plays an important role in the flashback of hydrogen-enriched flames. From the bottom row of Fig. 2, it is obvious that the flame front is much more wrinkled after considering the preferential diffusion. This is considered because of the preferential diffusion of hydrogen.

Besides, the flashback speed of case H2 is about 1.02 m/s, which is much larger than case H1 and closer to the flashback speed calculated from the experimental results (0.94 m/s). This indicates that the flamelet-based methods should account for heat loss and preferential diffusion effects for capturing the flashback process of the H_2/air flame in the bluff-body swirl flame accurately.

Ongoing Research / Outlook

The preliminary results shown in this report imply that differential diffusion has a significant impact on flashback characteristics of hydrogen flames. To understand the underlying mechanisms, we are investigating the coupling relationship between flame stretch and differential diffusion in details.

References and Links

[1] A. Kalantari, V. McDonell, Prog. Energy Combust. Sci. 61 (2017) 249–292.
[2] B. Lewis, G. von Elbe, J. Chem. Phys. 11 (1943) 75-97.
[3] D. Ebi, N.T. Clemens, Combust. Flame 168 (2016)39-52.
[4] D. Ebi, R. Bombach, P. Jansohn, Proc. Combust. Inst. 38 (2021) 6345–6353.
[5] H. Xia, W. Han, W. Zhang et al., Combust. Flame (2024), *in revision*.

Transition and turbulence in magnetohydrodynamic duct and channel flow

Research Institution
[1]Technische Universität Ilmenau

Principal Investigator
Thomas Boeck[1]

Researcher
Mattias Brynjell-Rahkola[1], Yohann Duguet[2]

Project Partner
[2]Laboratoire Interdisciplinaire des Sciences du Numérique, Université Paris-Saclay

Funding
DFG project no. 470628784

SuperMUC Project ID: pn49su (Gauss Large-Scale project)

Introduction

Magnetohydrodynamics (MHD) is concerned with flows of electrically conducting fluids under the influence of magnetic fields, which appear in a wide variety of situations ranging from microfluidics to astrophysics. The transition between laminar and turbulent wall-bounded shear flows is of considerable fundamental and practical interest and has received much attention in the fluid mechanics literature. However, the understanding of the corresponding phenomena in MHD notably lags behind. Reasons for this are the high costs and considerable difficulties of liquid-metal experiments, which cannot make use of optical flow measurement. To this end, direct numerical simulations (DNS) and high-performance computing play a central role in advancing our understanding of the topic.

In a rectangular duct, the presence of a magnetic field yields a Lorentz force that changes the characteristics of the boundary layers next to walls orthogonal to the magnetic field relative to the ones that develop along walls parallel to the magnetic field. These are commonly referred to as Hartmann and Shercliff layers, respectively. While laminar-turbulent transition in MHD ducts has traditionally been characterized by a parameter based on the thickness of the Hartmann layers, DNS frequently report that the Shercliff layers are most prone to instabilities [1]. In the present project, these two apparently contradictory findings are investigated using techniques from the field of dynamical systems. In particular, the concept of edge states, i.e. unstable solutions between laminar and turbulent flow [2], is used in the investigation of the MHD duct flow.

Results and Methods

All simulations have been performed using the open-source spectral element fluid solver Nek5000. This code offers high accuracy, flexibility with respect to geometry and boundary conditions, and performs well on modern large-scale computer architectures such as SuperMUC-NG. In order to simulate MHD flows, the solver capabilities have been extended to compute the additional eddy current and the Lorentz force density fields in the quasi-static approximation [3], which increases the computational cost relative to ordinary fluid flows by at least 40%. As can be seen in Figure 1, the scaling of the code is ideal. The size of the computational grids used in the project varies between approximately 6 million and 421 million degrees of freedom. The largest simulations have been performed using over 16,000 cores.

In order to compute the unstable edge states by means of DNS, the bisection method is used. For the parameters investigated, the edge state may either be fully spatially and temporally chaotic, or cycle through recurrent quiescent and bursting phases. It is, however, always found to reside in the Shercliff layer as shown in Figure 2. By perturbing the laminar Shercliff and Hartmann layers independently in DNS of transition to turbulence, it is observed that the evolution from laminar to turbulent flow always involves the Shercliff layer, and no evidence for a second transition route involving the Hartmann layer has been found [4]. These investigations were performed in a single quadrant of the duct, i.e. the flow has mirror symmetry with respect to the two midplanes.

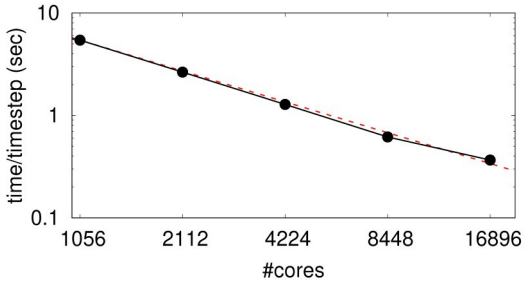

Figure 1: Strong scaling of Nek5000 on the SuperMUC-NG for DNS of the MHD duct flow. The dashed line corresponds to ideal scaling.

To further investigate the role of the induced Lorentz force density on the edge states and on the turbulent states, an idealized setup consisting of a doubly periodic channel without Hartmann walls has been studied as well [5]. In this simplified configuration, the edge state becomes increasingly chaotic and energetic upon strengthening the magnetic field, while the opposite effect is observed for the turbulent state. Eventually, the edge and the turbulent solution branches coalesce and disappear in a similar way as for a saddle-node bifurcation of equilibrium solutions in simple dynamical systems. The simultaneous computation of the edge state and the turbulent state for increasing magnetic field strength therefore turns out to be a successful means of determining the tipping point, i.e. the strongest magnetic field beyond which only laminar flow exists.

Ongoing Research / Outlook

While the present research on the MHD duct has mainly focused on the dynamics in symmetry-reduced small domains, the validity of these results in 'full' ducts awaits confirmation. This would, however, require a four-fold increase in the required computational resources. Similarly, increases in domain length that would capture a spatial localisation of edge states in the streamwise direction are expected to increase the computational requirements of the project significantly.

The effect of a transversal magnetic field is mainly to inhibit the streamwise streaks and rolls, both central to the common self-sustaining processes of wall turbulence. A future research direction may therefore be to consider a streamwise magnetic field instead. In such a case, the streaks would be completely unaffected by the Lorentz force. Hence, the physics encountered may differ considerably from the findings of the present project.

References and Links

[1] O. Zikanov et al., Appl. Mech. Rev. 66 (2014) 030802.
[2] T. Schneider et al., Phys. Rev. Lett. 99 (2007) 034502.
[3] M. Brynjell-Rahkola, Int. J. Numer. Meth. Fluids, *accepted for publication*, doi: 10.1002/fld.5321.
[4] M. Brynjell-Rahkola et al., Bull. Am. Phys. Soc. 67, vol.19 (2022).
[5] M. Brynjell-Rahkola et al., Phys. Rev. Res., *accepted for publication* (2024).

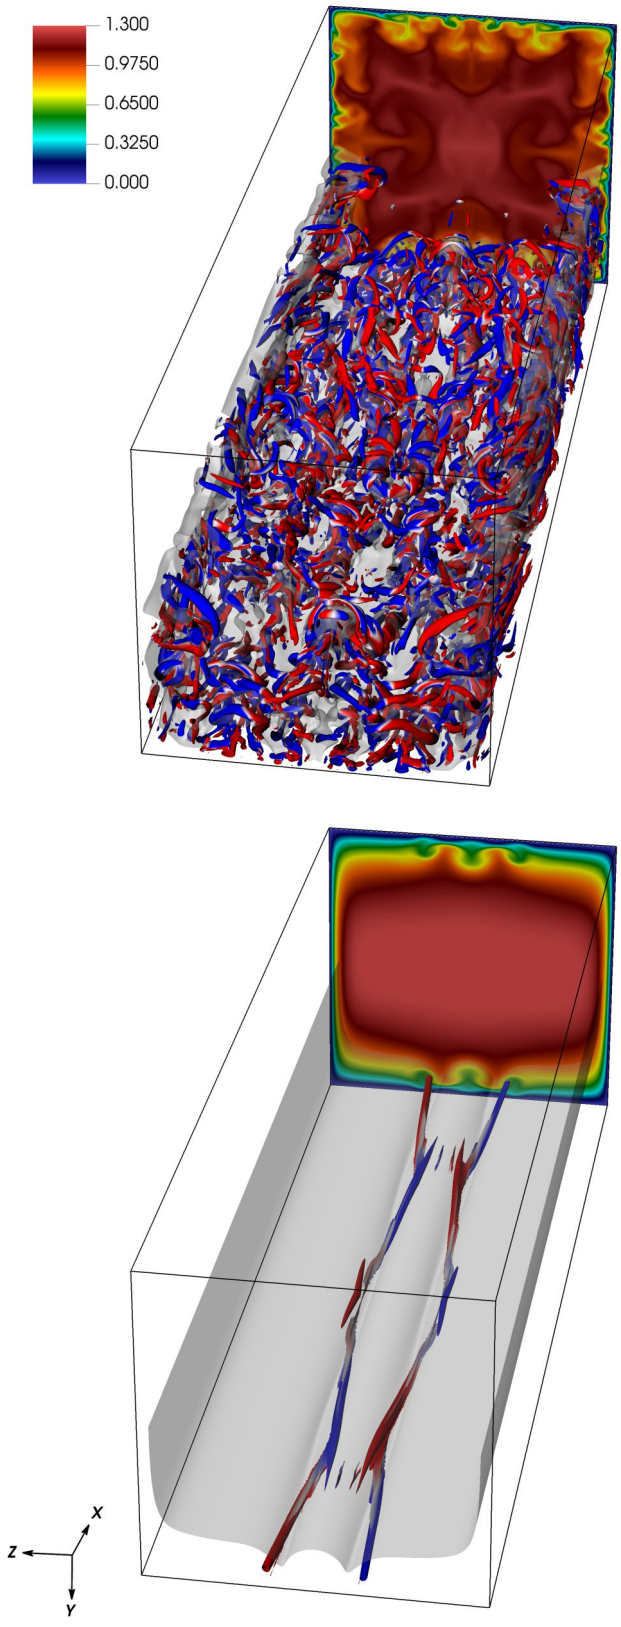

Figure 2: Comparison of the turbulent state (upper panel) to the edge state (lower panel) for the same flow parameters. A two-dimensional cross-section of the streamwise velocity is shown together with three-dimensional isosurfaces that correspond to 95% of the bulk velocity (grey) and vortical structures colored by their sense of rotation around the *x*-axis (red - positive, blue - negative). Two streamwise periods are shown, and although only a quarter of the domain is simulated, the flow structures are replicated for visualisation purposes. The magnetic field is along the *z*-direction.

Direct numerical simulation of turbulent round jet flows at high Reynolds numbers

Research Institution
Technical University of Darmstadt

Principal Investigator
Martin Oberlack

Researcher
Cat Tuong Nguyen

Project Partner
—

SuperMUC Project ID: pn73fu (Gauss Large-Scale project)

Introduction

Turbulence, characterized by its chaotic and complex fluid motion, has been an important point of research in fluid dynamics for decades. Despite significant efforts, understanding turbulence remains a formidable challenge due to the non-linear nature of the underlying Navier-Stokes equations (NSEs), which lack general analytical solutions.

Direct numerical simulations (DNS) offer detailed insights into turbulence by solving the NSEs without turbulence models. However, their high computational costs limit practical use to academic settings, whereas simplified turbulence models in industrial applications sacrifice accuracy for efficiency.

Turbulent round jets, being canonical flows and a subject of extensive research, find widespread application in engineering contexts such as jet engines, industrial mixing processes, and environmental dispersion. They are often investigated using DNS to understand their flow dynamics. Comprehending the mixing behavior of turbulent round jets is crucial for optimizing combustion efficiency in jet engines and controlling pollutant dispersion in the atmosphere. Moreover, DNS of turbulent round jets are important for the design and performance of industrial mixing devices, including chemical reactors and wastewater treatment systems. Simulating turbulent round jets play a vital role in advancing our comprehension of turbulence and its practical applications across various engineering disciplines.

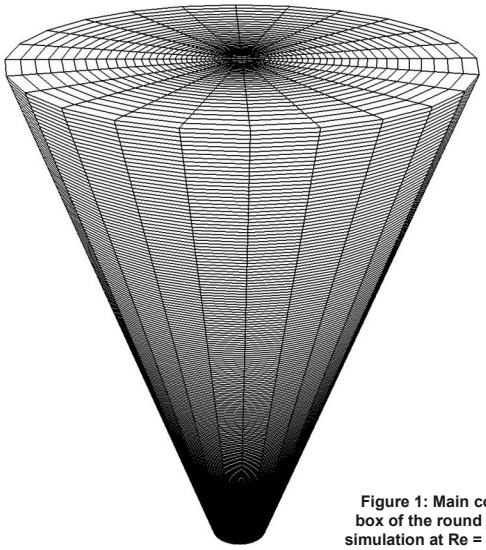

Figure 1: Main computational box of the round jet flow simulation at Re = 3,500.

In parallel to these efforts, researchers are exploring advanced mathematical techniques, such as Lie-symmetry theory, to develop more robust turbulence models. By leveraging the inherent symmetries within the NSEs, this approach aims to improve the accuracy of predictions. This interdisciplinary effort holds promise for advancing our understanding of turbulence.

In recent works [2,3], we conducted large-scale DNS of a plane channel flow at a Reynolds number of $Re = 10{,}000$ on SuperMUC-NG to validate new turbulent scaling laws for the log and core region derived with Lie-symmetry analysis. We generated symmetry-invariant solutions from the infinite-dimensional hierarchy of moment equations derived from the NSEs for moments of arbitrary order. Our key finding showed a power-law scaling of the streamwise velocity moments in the channel center, with an exponent dependent only on the first and second-order moments. We also expanded on this with symmetry-based turbulent scaling laws for streamwise velocity and temperature moments of arbitrary order for the core region of a turbulent channel [4]. This knowledge shall be extended to turbulent round jet flows in this project and validated against DNS data.

Results and Methods

The CFD code Nek5000 has been used for the DNS of the turbulent round jet flow. This code is based on the spectral element method (SEM) and solves the NSEs on a hexahedral mesh. The code is well established in the scientific turbulence community and is highly scalable

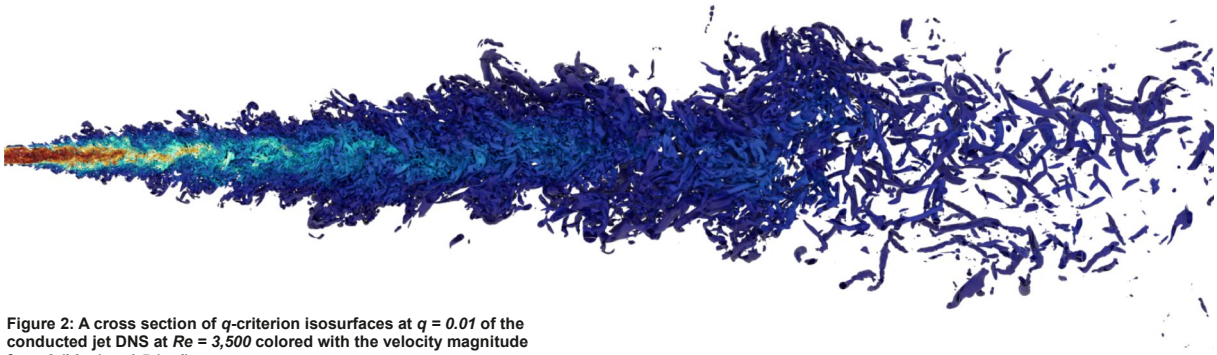

Figure 2: A cross section of *q*-criterion isosurfaces at *q = 0.01* of the conducted jet DNS at *Re = 3,500* colored with the velocity magnitude from 0 (blue) to 1.5 (red).

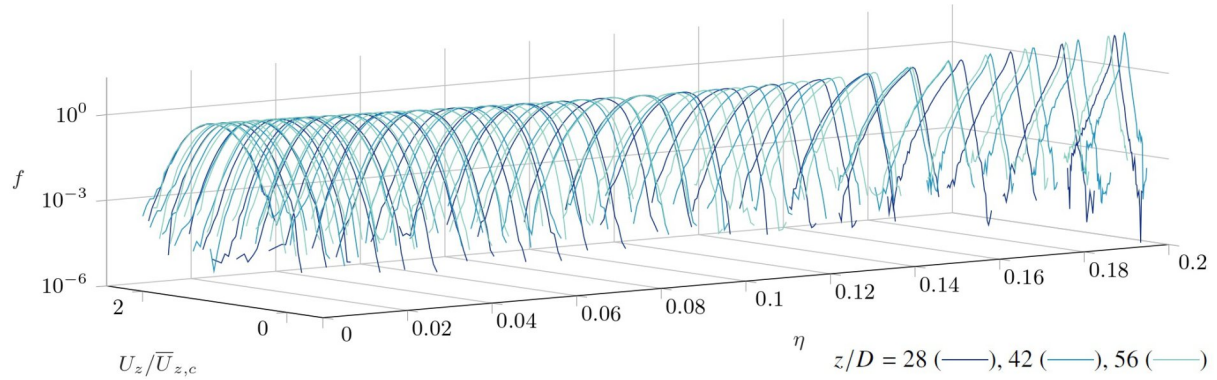

Figure 3: Probability density functions of $U_z/\overline{U}_{z,c}$ for three axial distances from the orifice *z/D*.

over a million ranks using Message Passing Interface (MPI) for parallelization. As a timestepping method the implicit second order backward differentiation formula (BDF2) scheme is used. The timestep is dynamically controlled and so the computational costs are further optimized. SEM features an exponentially growing accuracy with increase of the order of the polynomial basis functions used.

We have run two DNS cases on SuperMUC-NG for a turbulent round jet flow which are probably the largest round jet flow DNS to date. One is at *Re = 3,500* [4] and the second is at *Re = 7,000*. Additionally, a scalar transport equation is solved at a Prandtl number of *Pr = 0.71* [5]. The simulations are carried out on two separate domains. On the first domain, a turbulent pipe flow is generated. Then, the velocity field of the turbulent pipe flow is interpolated onto the main computational box at each time step to generate a turbulent round jet. The main computation is carried out on a truncated cone box with a length of 75 jet nozzle diameters *D* and a radius of 4 nozzle diameters at the inlet and 32 jet nozzle diameters at the outlet (see Fig. 1). The size of the computational box ensures the full capture of the spreading of the jet and minimizes the influence of the far-field boundary conditions. Altogether, for the *Re = 3,500* case the box consists of around 245 million degrees of freedom (DOFs), while for the *Re = 7,000* case the box has 1.31 billion DOFs.

Fig. 2 depicts isosurfaces of the *q*-criterion, commonly used for vortex visualization, at *q = 0.01* for the *Re = 3,500* case. This figure reveals that jets spread linearly with the data showing an inverse decay to that. Therefore, the radial coordinate is often rescaled by the axial distance from the orifice showing self-similarity of the velocity profiles. This is beautifully depicted in Fig. 3 by the probability function *f* of the axial velocity U_z scaled by the axial centerline velocity $U_{z,c}$, where the rescaled radial coordinate is η.

Ongoing Research / Outlook

From Lie-symmetry theory, we discovered that the scaling laws of round jet flows contain a free parameter allowing for a variation of the inlet conditions [6]. However, this parameter appears to be unaffected by the Reynolds number. Therefore, we intend to conduct further DNS to explore alternative inlet conditions, such as fractal grids or chevrons.

References and Links

[1] https://www.fdy.tu-darmstadt.de/
[2] Hoyas, S., Oberlack, M., Alcántara-Ávila, F., Kraheberger, S. V., Laux, J. Phys. Rev. Fluids 7, 014602 (2022).
[3] Oberlack, M., Hoyas, S., Kraheberger, S. V., Alcántara-Ávila, F., Laux, J., Phys. Rev. Lett. 128, 024502 (2022).
[4] Nguyen, C. T. & Oberlack, M., Similarity analysis of a turbulent round jet based on new direct numerical simulation data at large box and high Reynolds number. Phys. Rev. Fluids (2024), *in press*.
[5] Nguyen, C. T. & Oberlack, M., Passive scalar statistics in a turbulent round jet: symmetry theory and direct numerical simulation. J. Fluid Mech. (2024), *in review*.
[6] Nguyen, C. T. & Oberlack, M., Hidden intermittency in turbulent jet flows. Phys. Rev. Lett. (2024), *in review*.

Heat Transfer Phenomena in Cooling Channels

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

Principal Investigator
Steffen Schmidt

Researcher
A. Doehring, T. Kaller, N. A. Adams

Project Partner
–

Funding
SFB TRR 40, DFG

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 channel 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 correct prediction of the heat transfer within rocket engines cooling channel is crucial for safety reasons and in order to improve the engine's performance. The following results have already been obtained within this SuperMuc project:

- High fidelity simulations of a transcritical channel flow analyzing the influence of strong property variations on the turbulent flow and heat transfer [2]
- High fidelity simulations of a high aspect ratio duct studying the influence of secondary flow on the heat transfer [3].

The development of additive manufactured (AM) technologies as for instance Selective Laser Melting (SLM) offers new possibilities designing regenerative cooling channels. As a consequence, an increased surface roughness is observed in AM cooling channels, leading to an enhancement of the heat transfer but also to an increased pressure drop.
For this reason, we investigate the influence of a varying wall shear stress stemming from rough walls on the secondary flow in cooling ducts [4]. Although, the motivation for our numerical investigation emanate from rough duct flows in context of additive manufacturing, this study intend to show and assess the effect of a varying wall shear stress, which arises from any source. For this study, we performednumerically expensive well-resolved LES developed with our in-house LES solver within the SuperMUC-NG project pr48me.

Results and Methods

Our LES are carried out solving the three-dimensional compressible continuity, momentum and total energy equations. A finite-volume method is applied in order to spatially discretize the governing equations on a block structured, curvilinear grid. An explicit second-order low-storage four-stage Runge-Kutta method with enhanced stability region is applied for time advancement. A compact four cell stencil approach is used to compute the convectivefluxes. A discontinuity detecting sensor functional is used to switch the flux calculation between a linear fourth- order reconstruction for high accuracy and a more stable upwind-biased scheme. A physically consistent subgrid-scale turbulence model based on the Adaptive Local Deconvolution Method (ALDM) is included in the convective flux calculation. Viscous fluxes are determined by a linear second-order centered scheme. The ideal gas law is used to close the system and the viscosity is modeled by the Sutherland law. A generic duct configuration with periodic boundary conditions in the streamwise direction is used for the analysis in order to focus on the influence of an asymmetric wall shear stress modulation on secondary flow. Isothermal no slip smooth boundary conditions are applied at all four walls. This case has a bulk Reynolds number of $Re_b = 17{,}800$ and a friction Reynolds number $Re_\tau = 519$. Periodic boundary conditions are used in streamwise direction. The bulk pressure is set to $p_b = 1.01325$ bar. The duct geometry is $10h \times 2h \times 2h$ in streamwise, wall-normal and spanwise direction with the duct half height h, respectively. A sketch of the duct's cross section is shown in Fig. 1.

The wall shear stress of the lower wall in y direction (Y1) is artificially changed. The modifications is achieved by multiplying the obtained shear stress within the simulation by a constant factor $\tau_{w,mod}\,(Y1) = C(z) \times \tau_{w,orig}$. This modification is not a stress boundary conditions, but a no-slip boundary condition is applied at each wall. This is achieved by adjusting the ghosts cells at the walls after every Runge–Kutta step based on the updated solution. The targeted wall shear stress profiles normalized with the perimeter averaged wall shear stress of the reference case are shown in Fig. 2. The grid consist of approxi-

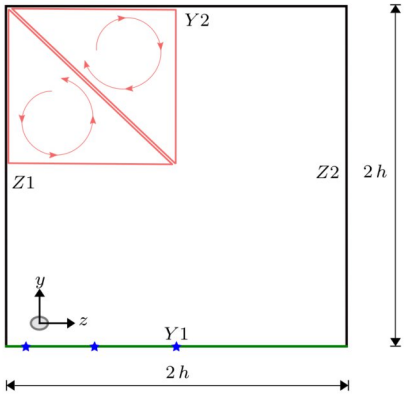

Figure 1: Cross section of the computational domain. The modified lower *y*-wall (Y1) is indicated green. In each duct corner two counter rotating vortices of Prandtl's second kind are formed, which are indicated in red.

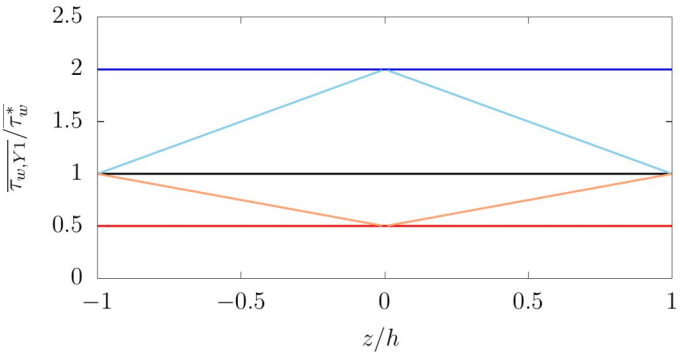

Figure 2: Intended mean wall shear distribution for different cases. The reference case is in black.

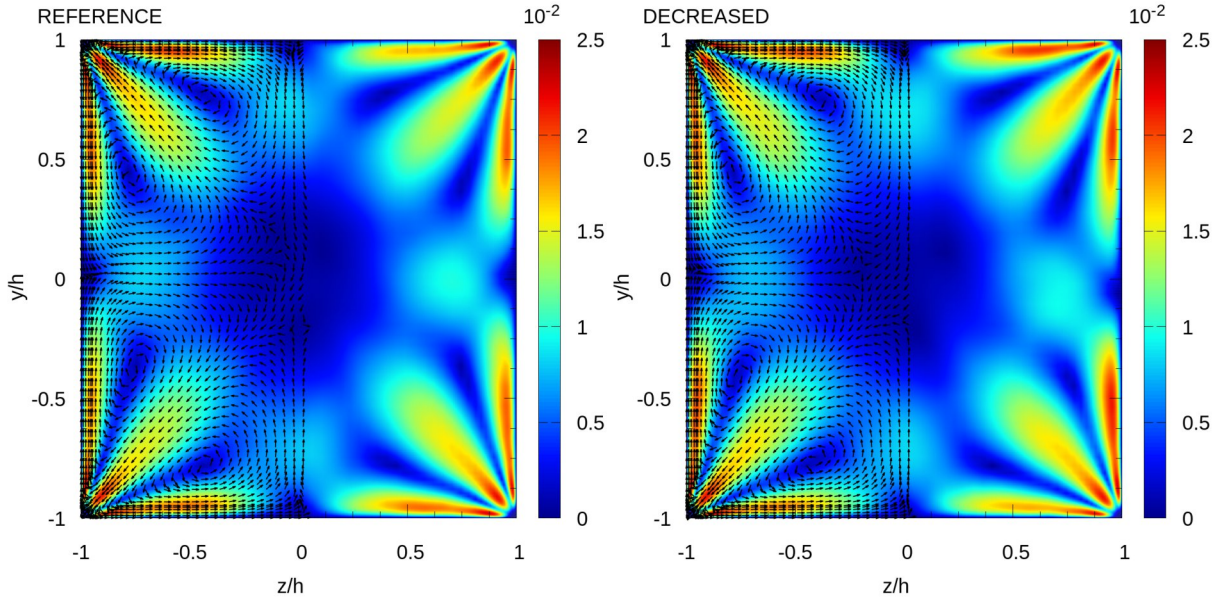

Figure 3: Mean cross-sectional velocity distribution. Reference case (left) and case with decreased wall shear stress (right). The counter-rotating corner vortex system is visualized by velocity vectors.

mately 0.9 million cells and the LES is using 480 cores on Supermuc-NG. In total seven simulations have been performed. The time needed to obtain statistical data is 12 runs of 48h. Figure 3 shows the mean cross-sectional velocity distribution $\overline{u}_{cf} = \sqrt{\overline{v}^2 + \overline{w}^2}$ for the reference case without modification and with a decreased wall shear stress at the lower wall. A weakening of the secondary flow can be observed at the lower wall compared to the reference. On the other hand, the secondary flow is increasing at the sidewalls. This change intensifies the streamwise vorticity and the resulting corner vortices. Aside from this, a further study [5] has been conducted regarding transcritical channel flows analyzing the influence of strong property variations on the turbulent flow and heat transfer. This analysis was founded on priorconducted investigations [2]. For this purpose, LES with 16 million cells have been performed on 1,024 cores on Supermuc-NG. In total seven simulations have been performed. The time needed to obtain statistical data is 40 runs of 48h.

References and Links

[1] http://www.sfbtr40.de/
[2] Doehring, A., Kaller, T., Schmidt, S.J., Adams, N.A. (2021). Large-eddy simulation of turbulent channel flow at transcritical states. Journal of Heat and Fluid Flow 89, 108781. doi:10.1016/ j.ijheatfluidflow.2021.108781.
3] Kaller, T., Pasquariello, V., Hickel, S., Adams, N.A. (2019). Turbulent flow through a high aspect ratio cooling duct with asymmetric wall heating. Journal of Fluid Mechanics, 860, 258-299. doi:10.1017/jfm.2018.836
[4] Doehring, A., Kaller, T., Schmidt, S. J., & Adams, N. A. (2024). Influence of wall shear stress on the secondary flow in square ducts. International Journal of Heat and Fluid Flow, 105, 109240.
[5] Doehring, A., Schmidt, S. J., & Adams, N. A. (2023). Momentum boundary layers in transcritical channel flows. International Journal of Heat and Fluid Flow, 103, 109201.

LES of the coaxial injection of liquid nitrogen and gaseous hydrogen under supercritical pressures

Research Institution
Technische Universität München

Principal Investigator
Christian Stemmer

Researcher
Raffaele Olmeda, Alexander Döhring

Project Partner
–

Funding
DFG SFB TRR40

SuperMUC Project ID: pr92lo

Introduction

Rocket combustion chamber are machines with the largest energy density man has ever built. This requires enormous efforts in terms of the supply with propellants (e.g. 326 kg/s for the Ariane VULCAIN engine) as well as heat management to maintain structural integrity. The engines are tested at DLR Lampoldshausen but a proper design ahead of the extensive and expensive validation procedure reduced risk of failure during flight and increases reliability of the entire launcher system. Large-Eddy simulations are currently capable of simulating the reacting flow from injection to the exhaust under supercritical conditions including the injection, drop breakup, evaporation and combustion. These numerical challenges require state-of-the-art supercomputers in order to get the proper spatial and temporal resolution for these time-resolved simulations. In order to ease on the computational effort, i.e. reduced combustion modeling is employed as well as models for the treatment of the flow close to the wall which is critical for accuracy but requires the largest resolution. To overcome these limitations, this project tackled the improvement of flamelet models for the harsh environments as well as the wall-modelling, especially in presence of wall roughness, for these environments.

Results and 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 species 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

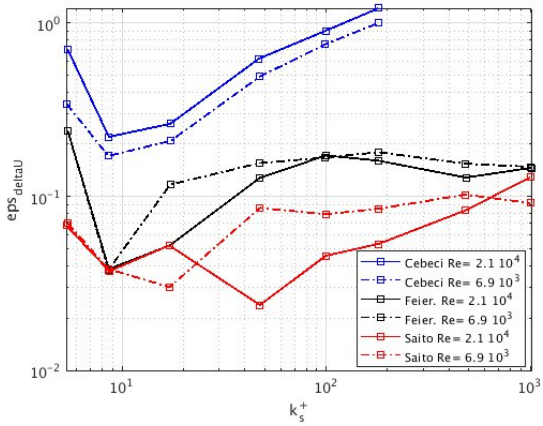

Figure 1: Performance of the three tested roughness methods with respect to different values of Reynolds number and sand grain roughness.

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. A correct estimation

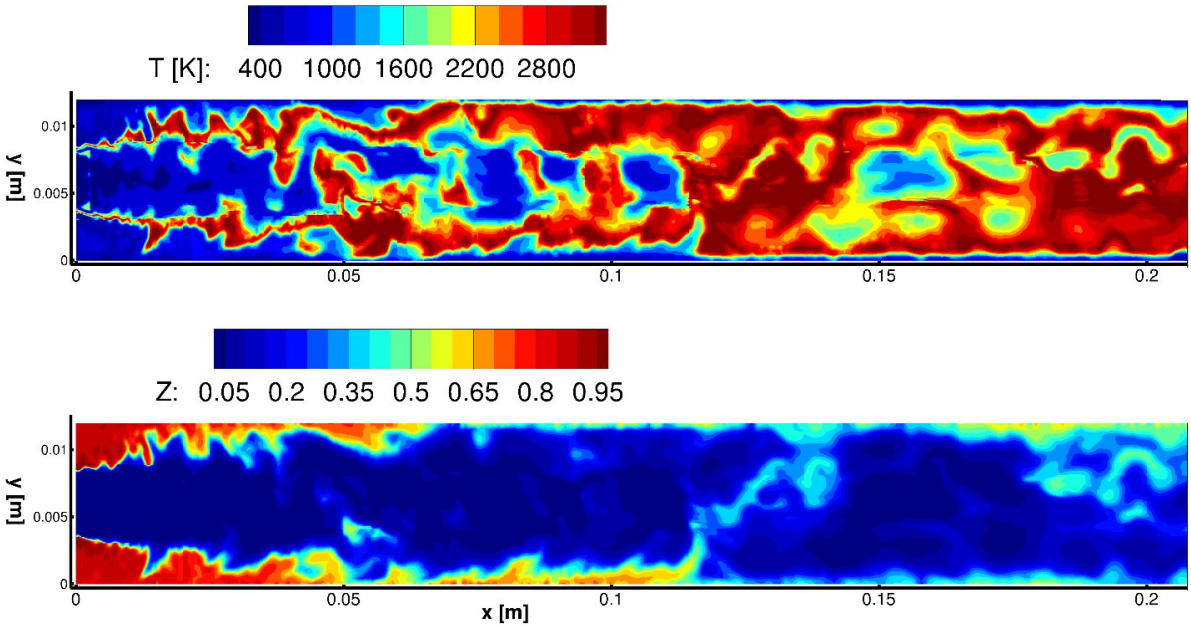

Figure 2: Snapshot of the temperature field (upper) and the mixture fraction (lower) in the methane-based combustion chamber.

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. 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 Mio core hours were used, divided in meshes of different refinement. For the roughness simulations, 2.9 Mio core hours 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, "Calculation of incompressible rough-wall boundary-layer flows," AIAA Journal, pp. 730–735, 1978.
[3] W. J. Feiereisen and M. Acharya, "Modeling of transition and surface roughness effects in boundary-layer flows," AIAA Journal, pp. 1642–1649, 1986.
[4] N. Saito, D. I. Pullin and M. Inoue, "Large eddy simulation of smooth-wall, transitional and fully rough-wall channel flow," 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.

Towards simulations of ultimate convection

Research Institution
[1]Max Planck Institute for Dynamics and Self-Organization, Göttingen
[2]Max Planck Center Twente for Complex Fluid Dynamics and Physics of Fluids Group, University of Twente, The Netherland

Principal Investigator
Detlef Lohse[1,2]

Researcher
Richard Stevens[2], Robert Hartmann[1], Roberto Verzicco[2]

Project Partner
–

Funding
ERC starting UltimateRB, ERC advanced MultiMelt

SuperMUC Project ID: pr74sa (Gauss Large-Scale project)

Introduction

Rayleigh-Bénard convection, a fluid layer heated from below and cooled from above, is a key model in fluid dynamics. The system is widely used for exploring concepts like instabilities, chaos, and turbulence. It is relevant to understanding various natural and industrial processes, building ventilation, and convection in the atmosphere, oceans, and planets. Therefore, this system is relevant for multiple scientific fields such as geology, oceanography, and astrophysics. Rayleigh-Bénard convection is particularly suitable for gaining fundamental insights into the interaction between boundary layer and bulk interactions in wall-bounded turbulent flows.

The Grossmann-Lohse theory, a central concept for understanding thermal convection, models the Nusselt number (heat transfer) and the Reynolds number (flow strength) as a function of the model's control parameters. These control parameters are the Rayleigh number (Ra), which indicates the temperature difference, and the Prandtl number, which describes the fluid's properties. The existence of an ultimate regime with enhanced heat transport properties and turbulent boundary layers had been predicted by Kraichnan (1962) and later by Spiegel (1972) and Grossmann & Lohse (2000). In this ultimate regime, turbulence occurs not only in the bulk fluid but also within the boundary layers. This transition occurs when the shear-induced Reynolds number is strong enough to make the boundary layers turbulent and results in significantly enhanced heat transfer, see Figure 1.

Results and Methods

Both experiments and simulations have found it enormously challenging to achieve the ultimate regime. In this ultimate regime, the scaling exponent γ, present in the relationship $Nu \sim Ra^{\gamma}$, increases, see Figure 1. The critical Ra number (Ra^*) for this transition is $Ra^* \approx 2 \times 10^{13}$ in the Göttingen experiments conducted under the leadership of Bodenschatz and Ahlers. The highest Ra number achieved in corresponding direct numerical simulations is $Ra = 2 \times 10^{12}$ [4]. For further background, we direct the reader to the invited general audience article in the Physik journal by Detlef Lohse titled "Turbulence in the Ultimate Regime" [2], which was written upon receiving the Max Planck Medal, and the Physics Today article [3].

We performed innovative direct numerical simulations of Rayleigh-Bénard convection up to $Ra = 10^{13}$ to allow one-to-one comparison with experiments performed in Göttingen. Snapshots of the temperature field from these simulations are provided in Figures 2 and 3. To accurately capture all flow features, we use grids with up to 100 billion grid points. This simulation was performed on tens of thousands of cores, requiring over 50 million core hours per simulation. Capturing a complete snapshot of the flow field requires up to 3 TB. The entire project database, comprising various large cases, exceeds one petabyte. This extensive database facilitates comprehensive flow field analysis, revealing intricate flow dynamics. Excellent agreement between simulations and the Göttingen experiments is observed for the heat transfer, as illustrated in Figure 1, and local flow characteristics, specifically the mean and temperature variance close to the sidewall.

The complexity of these simulations stems from the requirement to execute them within a completely enclosed domain. This constraint significantly increases the computational effort and complexity compared to simulations conducted in periodic domains, which benefit from optimized libraries. Regrettably, these optimization methods are not readily applicable to our specific challenge. Thus, to conduct these groundbreaking simulations, we developed a custom second-order finite-difference flow solver tailored for cylindrical geometries. The code is programmed in Fortran 90, leveraging extensive parallelization. We use MPI for distributing the computational workload across the broader computational loops and OpenMP for two-dimensional parallelization within these loops. This approach is strategically chosen

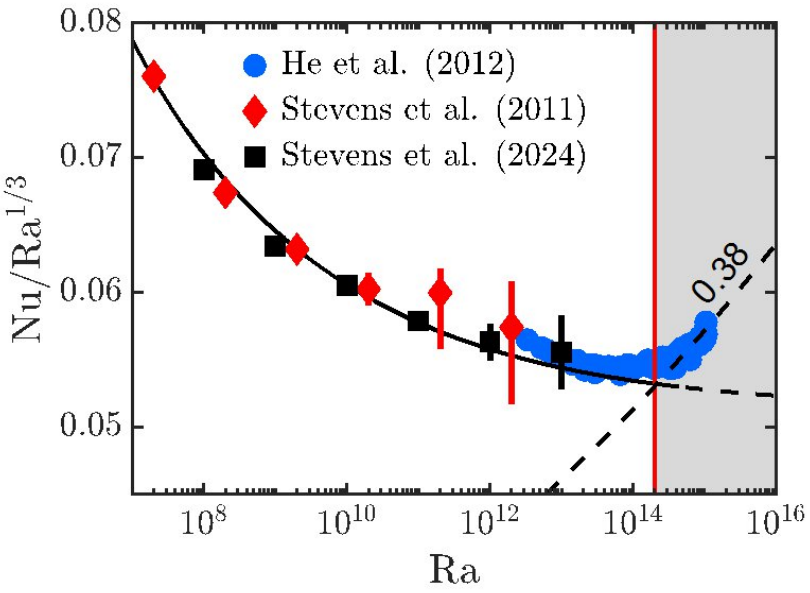

Figure 1: Comparison of Nu versus Ra obtained from simulations to the Göttingen measurements (He et al. (2012). Note the enhanced heat transfer in the ultimate regime [5].

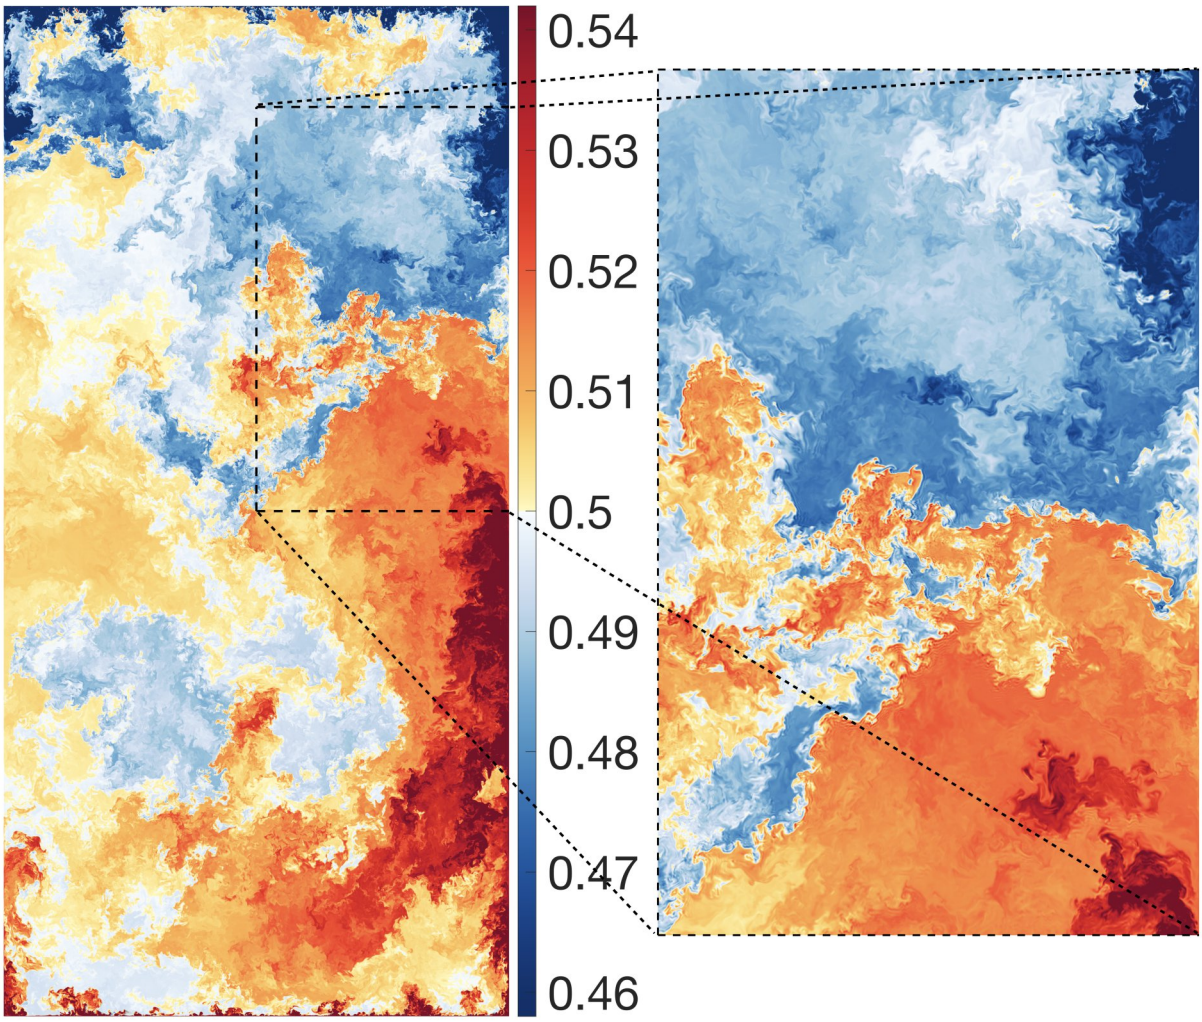

Figure 2: Visualization for Ra = 10^{13}. The right panel provides a zoom in of the indicated area. Figure 3 shows corresponding snapshots of the horizontal temperature field at thermal BL height.

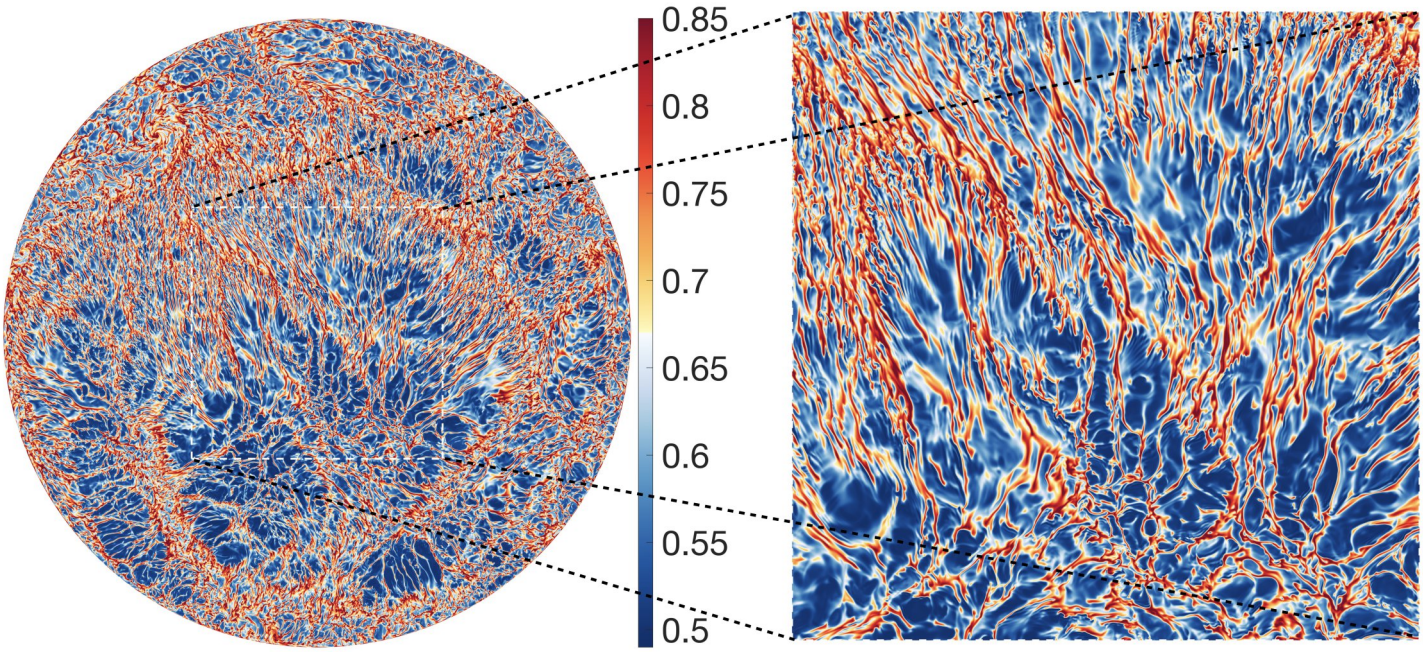

Figure 3: Snapshots of the temperature field at the boundary layer height just above the hot bottom plate for the case shown in Figure 2. One can see tiny plume structures, which require massive computational grids.

to minimize network communication between computational nodes by concentrating computations within nodes in a shared memory framework. Efficiency is further ensured by meticulously vectorizing the most computationally demanding segments of the code. We take additional steps to minimize memory movement through manual coding and inlining of computational operations and data transfers. The code also efficiently uses various cache levels and replaces standard library functions with custom routines to achieve specific performance enhancements.

Ongoing Research / Outlook

SuperMUC-NG enabled us to undertake unparalleled simulations, which we believe to be the most extensive turbulence simulations ever attempted in a fully enclosed domain. The computational requirements of simulating such turbulent flows stem from a large range of length and time scales that must be resolved. Therefore, these pioneering simulations can only be performed on the most powerful supercomputers, such as SuperMUC-NG. A key factor in making these simulations feasible was the development of algorithms that minimize communication between different processors, significantly enhancing the computational and parallel efficiency of our code. To ensure long-term data storage and accessibility, we employed the open-source HDF5 data format.

In future work, we plan to extend our simulation approach to study melting and dissolution in liquid systems. Melting objects create differences in temperature and concentration in liquids. These act back on the melting or dissolving objects, implying that such systems are extremely complex, namely multicomponent, multiphase, and phase transitions. Additionally, we plan to study flow dynamics in wind farms. Extending simulation approaches to these real-world problems will necessitate novel supercomputers.

References and Links

[1] Group website: pof.tnw.utwente.nl
Project details: pof.tnw.utwente.nl
[2] D. Lohse, Physik Journal 18, Nr 8/9, 39-45 (2019).
[3] D. Lohse and O. Shishkina, Phys. Today 76, 26–32 (2023).
[4] R.J.A.M. Stevens, D. Lohse, R. Verzicco, J. Fluid Mech. 688, 31-43 (2011).
[5] R.J.A.M. Stevens, D. Lohse, R. Verzicco, Direct and Large Eddy Simulation 12, ERCOFTAC Series 2019, volume 27, 215-224 (2020).

Massive-Parallel Simulations of Turbulent Premixed Hydrogen Flames with Thermodiffusive Instabilities

Research Institution
Institute for Combustion Technology, RWTH Aachen University

Principal Investigator
Heinz Pitsch

Researcher
Terence Lehmann, Thomas Howarth, Lukas Berger, Michael Gauding

Project Partner
–

Funding
ERC Advanced Grant (HYDROGENATE)

SuperMUC Project ID: pn36ga (Gauss Large-Scale project)

Introduction

The transformation of energy systems towards hydrogen creates unique opportunities to decarbonize thermochemical energy conversion processes. However, it also poses new challenges that arise from the specific molecular transport and combustion properties of hydrogen, which differ significantly from conventional fossil fuels.

In technical applications, utilizing lean premixed flames for hydrogen combustion is favored to mitigate the formation of nitrogen oxides (NO_x). However, a significant challenge in this approach is the sensitivity of the flame to intrinsic flame instabilities (IFI). Under lean premixed conditions, hydrogen flames are affected by two different instabilities: firstly, the hydrodynamic instability, which originates from the abrupt density change across the flame front, and secondly, thermodiffusive instabilities that arise from the disparity between heat and mass fluxes at the reaction zone, inducing differential diffusion effects across the flame front. In lean hydrogen flames, thermodiffusive instabilities amplify initial corrugations of the flame front, leading to a quickly growing flame surface area and, hence, faster flame propagation speeds. This behavior results from the low Lewis number of hydrogen, which quantifies the ratio of thermal diffusivity to mass diffusivity. Flame instabilities significantly affect flame dynamics and heat release rates, creating about ten times higher burning velocities than in settings where flame instabilities are suppressed.

Figure 1: Snapshot of the temperature field showing small-scale cellular structures and large-scale finger-like structures along the flame front.

This situation is even more complex for turbulent hydrogen flames. Turbulent motion interacts with the flame surface by stretching and wrinkling it over a wide range of different length scales. Berger et al. [1] demonstrated that synergistic interactions between thermodiffusive instabilities and turbulence further enhance the turbulent flame speed and increase the heat release per unit flame area by a factor of 2.5 compared to laminar unstable flames. A better understanding of the physical mechanism behind these synergistic effects in lean premixed hydrogen flames will enable the development of combustion models for thermodiffusive instabilities. This is highly relevant for various technical applications, ranging from gas turbines to industrial burners. The main objective is to leverage high-fidelity simulation data to develop a quantitative understanding of the relevant physical principles governing hydrogen flames, with the aim of systematically constructing advanced combustion models.

Results and Methods

Recent advances in high-performance computing and numerical methods have enhanced our ability to perform high-fidelity simulations of turbulent flames under realistic conditions. Direct numerical simulations (DNS) directly solve the governing equations for all relevant scales without relying on models for the turbulent motion. In reactive flows, this also includes the direct solution of chemistry, which is usually represented using a reduced reaction mechanism. This approach allows for an almost exact simulation of turbulent flames and is therefore also referred to as a numerical experiment. The data obtained provide unique insights into the physical mechanisms of reactive flows and are essential for the development of

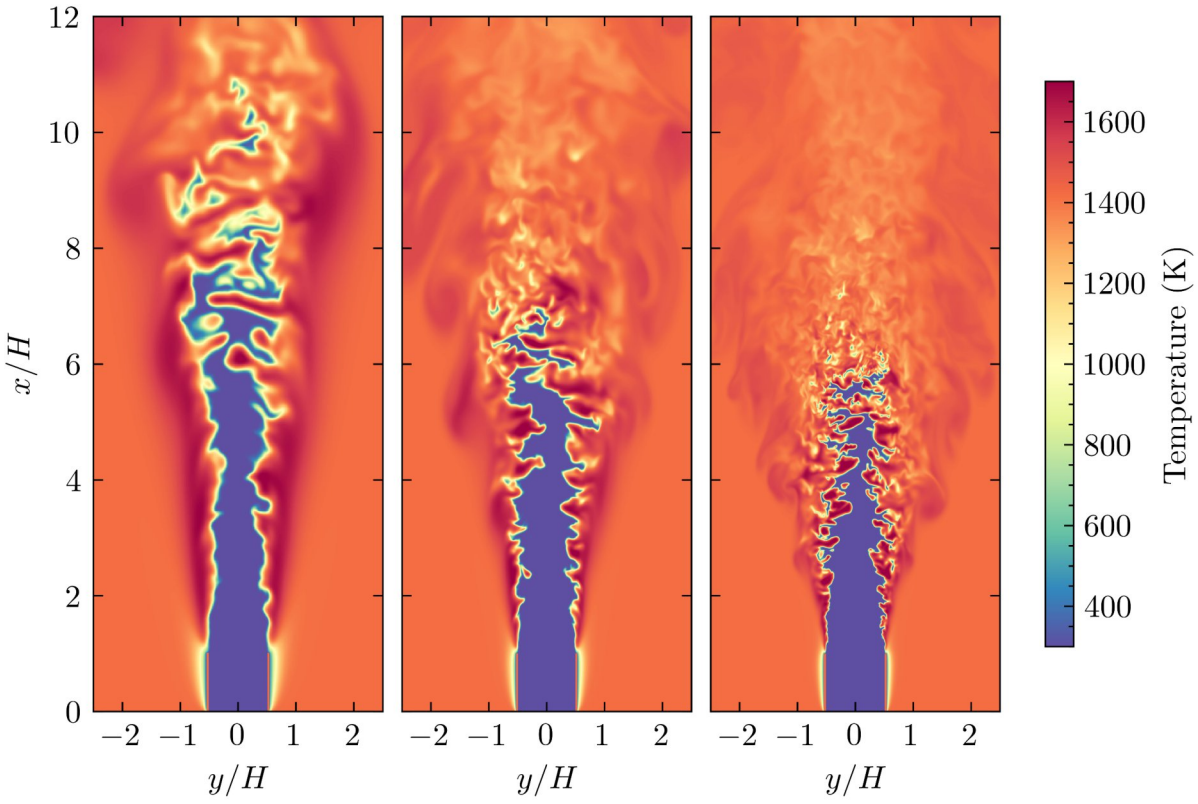

Figure 2: Instantaneous slice of temperature for jet Reynolds numbers of 5,500, 11,000, and 22,000 (from left to right).

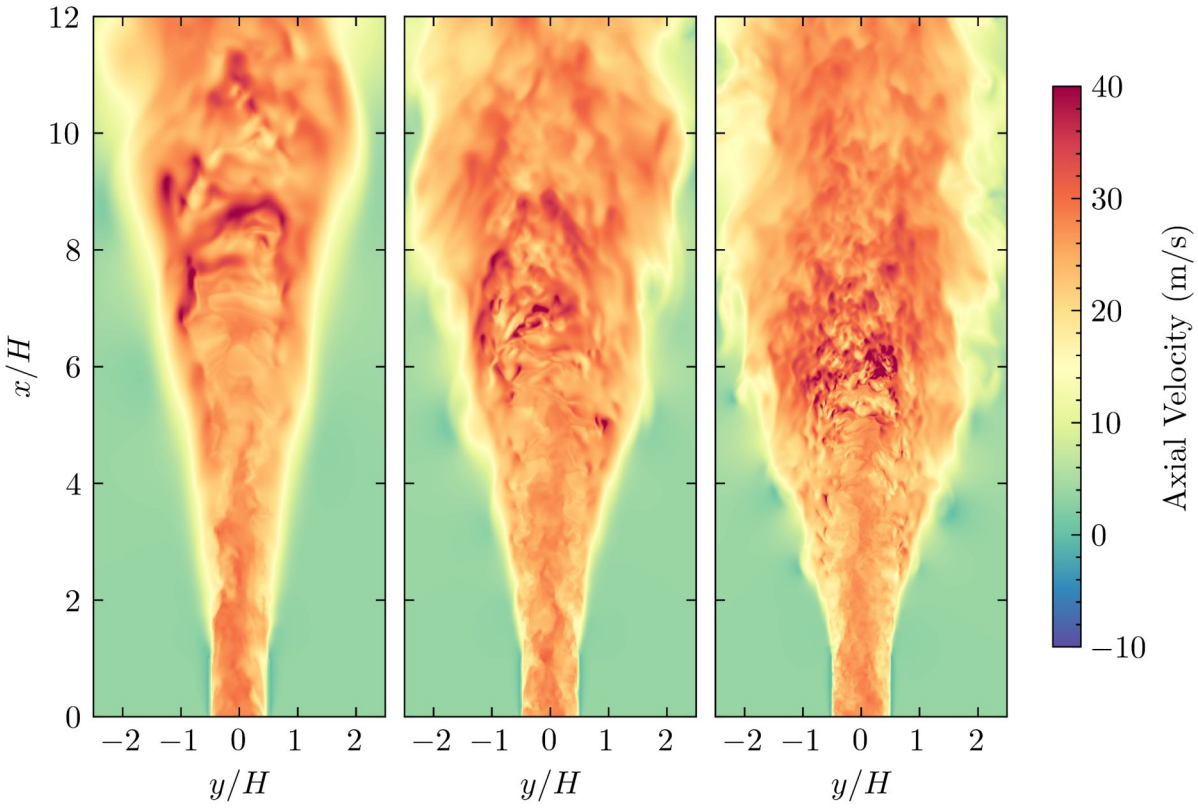

Figure 3: Instantaneous slice of axial velocity field for jet Reynolds numbers of 5,500, 1,100, and 22,000 (from left to right).

novel combustion models, especially for the simulation of hydrogen flames. However, conducting direct numerical simulations requires significant computational resources and is possible only on the most powerful supercomputers using efficient parallelization strategies. At the same time, precise and robust numerical solution methods are required to accurately resolve the smallest chemical and turbulent structures.

For this purpose, the massive-parallel high-fidelity code “CIAO” has been developed at the Institute for Combustion Technology. CIAO solves the reactive Navier-Stokes equations in the low-Mach number limit using arbitrary-order finite differences on a structured grid. Species and temperature equations are discretized by the WENO5 scheme for bounded solutions. Spatial and temporal staggering of flow variables enhance numerical accuracy. Temperature and species equations employ the symmetric operator split of Strang. CIAO is highly parallelized using MPI and scales well up to more than 3,072 compute nodes on SuperMUC-NG. All I/O operations are performed by parallel HDF5 to efficiently read and write terabyte-sized files.

To better understand thermodiffusive instabilities and to enable their modeling, we have conducted direct numerical simulations of turbulent premixed hydrogen jet flames in a slot burner configuration at different jet Reynolds numbers between 5,500 and 22,000. The Reynolds number has been increased solely by increasing the slot width while keeping all other quantities unchanged. The unburned mixture is set to an equivalence ratio of 0.4, an unburned Temperature of 298 K, and a pressure of 1 bar. Under these conditions, massive thermodiffusive instabilities are present. This is illustrated in Fig. 1 for a laminar flame under the same conditions, where the formation of a strongly corrugated and wrinkled flame front due to thermodiffusive instabilities is seen [2]. The associated increase of the flame surface leads to a significant increase in flame speed.

The simulations have been conducted on SuperMUC-NG, utilizing up to 3,072 compute nodes. The numerical mesh resolves the flame thickness with 10 grid points to accurately capture the molecular transport processes within the flame. The smallest turbulence scales are resolved with two grid points. For the largest case, this results in a computational mesh with more than 10 billion grid points.

A visualization of the turbulent hydrogen flames is shown in Fig. 2 for jet Reynolds numbers of 5,500, 11,000, and 22,000. The DNS reveals a remarkable increase of flame surface area with increasing Reynolds number. Berger et al. [1] explained this finding by synergistic interactions among the turbulent motion and thermodiffusive instabilities. Figure 2 displays the axial velocity field to illustrate the multi-scale nature of the turbulent motion and the acceleration of the jet flow at the flame tip. The DNS data provide highly resolved fields that enhance our understanding of the underlying physical principles and enable the development of predictive combustion models for thermodiffusively unstable flames. Such models are not yet available but necessary to advance the transition of energy systems towards hydrogen-based fuels.

Ongoing Research / Outlook

The ongoing research focuses on high-fidelity direct numerical simulations of blends between hydrogen and ammonia. Blending hydrogen with ammonia leads to effects that are entirely different from pure fuels due to the very different chemical scales of the mixture. The challenges of ammonia combustion include nitrogen oxide formation, potential toxicity of the burned gas, and combustion-specific characteristics such as low flame speeds and very low reactivity. Nevertheless, these properties can be controlled when combined with hydrogen to improve the blend's combustion characteristics.

References and Links

[1] Berger, L., Attili, A. and Pitsch, H., 2022, Combustion and Flame, 244, p.112254.
[2] Berger, L., Attili, A. and Pitsch, H., 2022, Combustion and Flame, 240, p.111936.

Development of Disruptive Technologies for Aero Engine Compressors

Research Institution
Institute of Turbomachinery and Flight Propulsion

Principal Investigator
Samuele Giannini

Researcher
Volker Gümmer, Jannik Eckel, Philipp von Jeinsen, Mattia Straccia

Project Partner
–

Funding
LuFo - Luftfahrtforschungsprogramm VI DLR

SuperMUC Project ID: pn73ri

Introduction

Air travel, while instrumental in shaping contemporary life, is under increasing scrutiny due to its environmental impact. The expanding sector has contributed to a steady increase in emissions despite a decrease in specific emissions due to improvements in efficiency. Continuous innovation becomes a necessity in the face of escalating political pressure and demand for more environmentally friendly air travel. One significant focus of this innovation is the reduction of fuel consumption and the downsizing of aircraft engines. It is estimated that saving 1kg in weight can lead to a reduction of around 2,900 liters of aviation fuel consumption annually [2]. The main drive for such innovation is twofold: firstly, the airlines ambition to cut operating costs linked to fuel consumption, and secondly, the overarching goal to decrease CO_2 emissions. Advancements in engine technology can yield smaller, lighter engines that maintain their power output and thermodynamic efficiency. A viable approach to meet decarbonization requirements involves further improving the performance of compressors, which constitute a significant share of the weight and axial length of power trains. An increase in total pressure ratio of the compressor leads to an increase in thermal efficiency of the whole engine. Accordingly, the use of tandem airfoils in compressor stages is identified as an emerging solution. Tandems have demonstrated the ability to significantly increase the achievable loading level thanks to an advanced lift distribution given by the front and rear airfoils. While tandem airfoils present certain benefits, they also introduce challenges in practical compressor applications. To address these obstacles, three-dimensional optimizations were carried out in this work with the aim of optimizing the airfoil aerodynamic performance and make use of innovative techniques like the non-axisymmetric contouring. Findings indicate that innovative tandem designs, which offer enhanced performance across various operational conditions, can be generated through the optimization processes. The optimization constraints employed, were found suitable for typical multi-stage compressor configurations. Thus, this research offers pivotal guidance for efficient aerodynamic designs of axial compressors. Through an understanding of the aerodynamic implications of tandem vanes, strides towards more efficient and environmentally-friendly engines can be made. The enhanced designs and fresh perspectives on performance determinants are anticipated to be foundational for the progression of future compressor technology and aeroengine evolution.

Results and Methods

The primary tool used for the aerodynamic optimizations in this work was an in-house optimization software, called ATOMIC, which uses AutoOpti, a meta-heuristic optimizer developed at the DLR's Institute of Propulsion Technology, with a focus on turbomachinery applications [3]. AutoOpti is an MPI-parallelized multi-objective optimizer that serves as the core for the research of optimum solutions. AutoOpti stands as a gradient-free optimization tool that utilizes evolutionary strategies and is enhanced with metamodels to expedite the optimization process. Every optimization in this study employed Latin Hypercube Sampling for initial sampling. A distinct advantage of AutoOpti is its asynchronous communication capability between the primary and secondary processes. To enhance and speed-up the optimization process, multiple approximative models were incorporated, called metamodels. A metamodel predicts fitness values of potential solutions, drawing on solutions evaluated previously in the search space. For every optimization in this research, the Kriging model was employed as the metamodel, balancing swift computation with consistent function accuracy. Once the members' quality is assessed through the Pareto rank, this new member's details are cataloged in the database and new members are generated accordingly. This loop persists until further significant improvements in the target functions are not achievable. The significant enhancement in the code for this research emerged from the introduction of the DLR's developed TRACE [4] as the CFD solver in place of the commercial software CFX. TRACE's integration enabled users to bypass the node limit constraint, aligning with

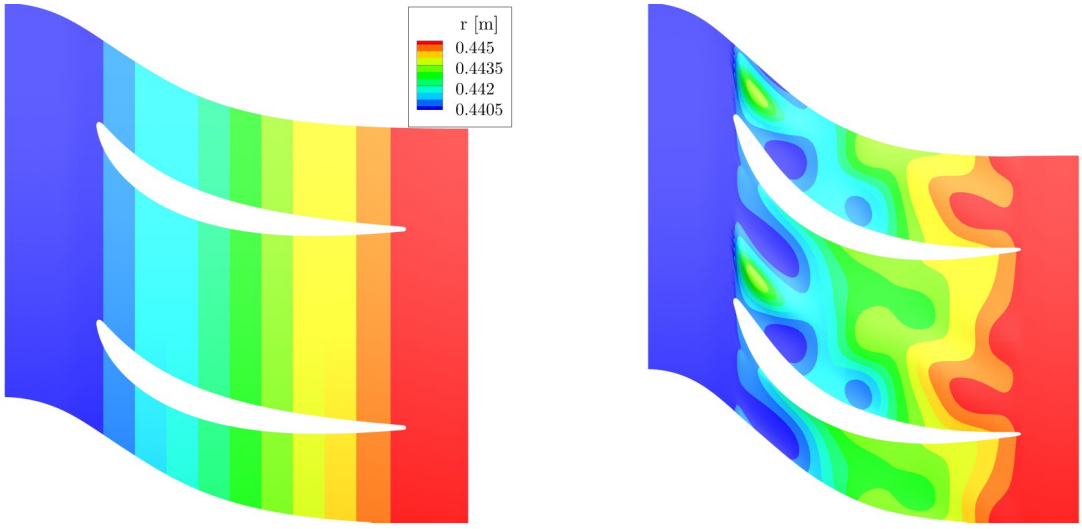

Figure 1: Radius Contours Comparison at the Hub Section of the Compressor Rotor. Left: Reference. Right: Optimized.

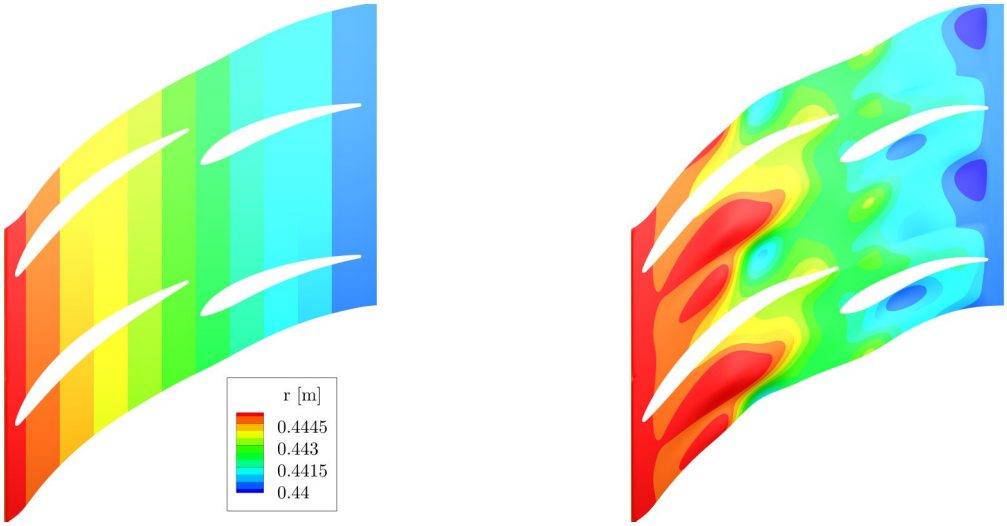

Figure 2: Radius Contours Comparison at the Hub Section of the Compressor Stator. Left: Reference. Right: Optimized.

the available licenses. As a result, the only restriction on the code's parallelization originated from the AutoOpti software, capped at a maximum of 100 nodes per optimization. This necessitated ensuring the process chain displayed apt scaling behavior. To validate this, various optimization routines were executed on the Thin Nodes of the SuperMUC-NG, beginning with 96 cores (equivalent to 2 Thin Nodes) and escalating to 3,072 cores. The scaling pattern was deemed satisfactory, endorsing the employment of the high-performance computing nodes of the SuperMUC-NG.

Starting from the reference, highly loaded low-speed axial compressor of the Institute of Turbomachinery and Flight Propulsion, a fully 3D robust optimization strategy of the isolated tandem stator is carried out. Two main optimized stator geometries have been obtained, both achieving an improvement in aerodynamic performance. It was shown that the optimized configurations manage to reduce losses by innovative three-dimensional geometries where the relative location of the front and rear airfoil is carefully tailored according to the specific aerodynamic needs, resulting in innovative tandem geometries. The study was then extended to include a combined optimization of the airfoil geometries of the whole machine. Building on these results, the potential benefits of the non-axisymmetric contouring technique was assessed. When the optimization extends to the 1.5 stage configuration, the optimized geometries exhibit enhanced stage efficiency and work increase, thereby augmenting efficiency further. The constraints of the optimization routines are demonstrated to be ideal for the multistage configuration of the machine, favoring their use in modern, compact engine configurations. Finally, the incorporation of the non-axisymmetric contouring technique emerges as a key factor to further augment the gains of the optimized geometries using a multi-stage configuration. The contoured hub of the stator emerges as the only means available to address the increasing losses of rear stages, predominantly evident under off-design conditions.

Figure 1 depicts the hub radius contour for both the reference and the optimized rotor. The refined hub

reveals specific areas of alteration. A noticeable increase in radius near the leading edge on the pressure side is evident. This is followed by radius reduction closer to the blade's pressure side. Intriguingly, these specific depressions on the pressure side align with findings from a comparable study on a high-speed rotor [5]. As the trailing edge is approached, there's an evident increase in the radius, aiming to locally augment the endwall blockage primarily near the pressure side—this area emerges as the main focus of the contouring modifications. Similarly, Figure 2 depicts the hub radius contour for both the reference and the optimized stator. Figure 3 illustrates the contours of the non-dimensional total pressure at the stage outlet of the reference and optimized geometry. The most significant achievement of the optimization process is the reduction of losses, indicated by low total pressure values, near the stator endwalls. Different axial velocity contours at different axial points for both the reference and optimized stators are also plotted in Figure 4. The main contributor to losses, originating from the low momentum fluid interaction with the rear airfoil, is significantly reduced by the optimized geometry. Notably, the optimized hub design seems to prevent the circumferential movement of this loss core (highlighted in black on the left of Figure 4, leading to the high loss region 'A' in the reference case) by using a larger gap size in the tandem profile at the hub.

Ongoing Research / Outlook

Many research projects are currently ongoing at the chair of Turbomachinery and Flight propulsion for the generation of new, highly efficient axial compressors. Some examples include the concept of boundary layer laminarization and passive boundary layer stabilization, designed to increase the lift and efficiency of compressor airfoils; the new concept of so-called "hybrid blades" is also being developed and numerically tested. In this perspective, the SuperMUC-NG computational resources can further help and support the execution of highly-resolved and time consuming optimization routines. However, some aspects need to be addressed. Some commercial software compatibility issues were faced on the SuperMUC infrastructure. Difficulties such as the parallelization challenges and the 48-hour time limit posed a significant challenge for the comprehensive optimization tasks.

References and Links

[1] https://www.asg.ed.tum.de/en/ltf/research/current-research-topics/driver/

[2] Dickens, A. (2008). Highly Loaded Axial Compressors. PhD thesis, University of Cambridge.

[3] Siller, U., Voß, C., and Nicke, E. (2009). Automated multidisciplinary optimization of a transonic axial compressor. In: 47th AIAA Aerospace Sciences Meeting Including The New Horizons Forum and Aerospace Exposition.

[4] Becker, K., Heitkamp, K., and Kügeler, E. (2010). Recent progress in a hybridgrid cfd solver for turbomachinery flows. In: Proceedings fifth European conference on computational fluid dynamics ECCOMAS CFD.

[5] Reising, S. and Schiffer, H.-P. (2009b). Non-axisymmetric end wall profiling in transonic compressors—part ii: Design study of a transonic compressor rotor using non-axisymmetric end walls—optimization strategies and performance. In Turbo expo: power for land, sea, and air.

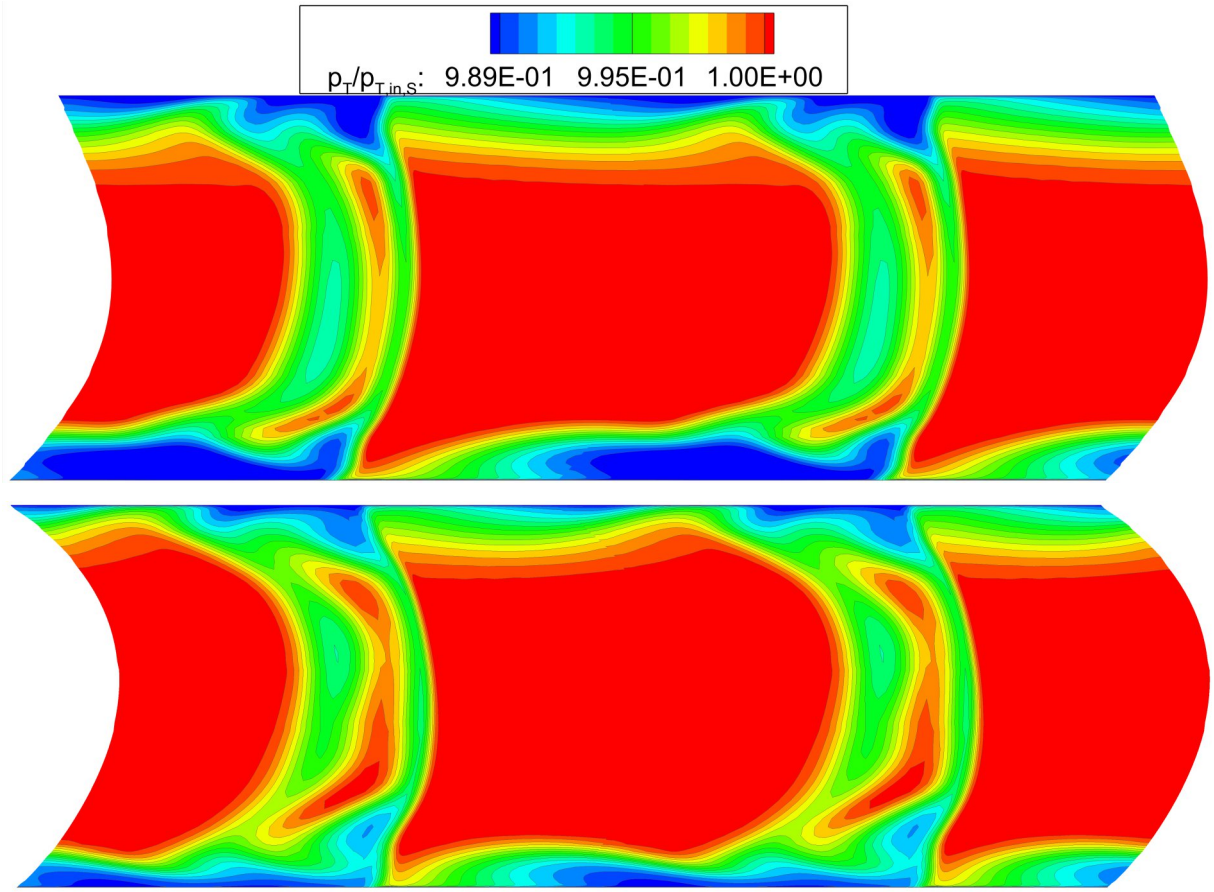

Figure 3: Contours of Non-dimensional Total Pressure at Stage Outlet for the Reference and Optimized Geometries.

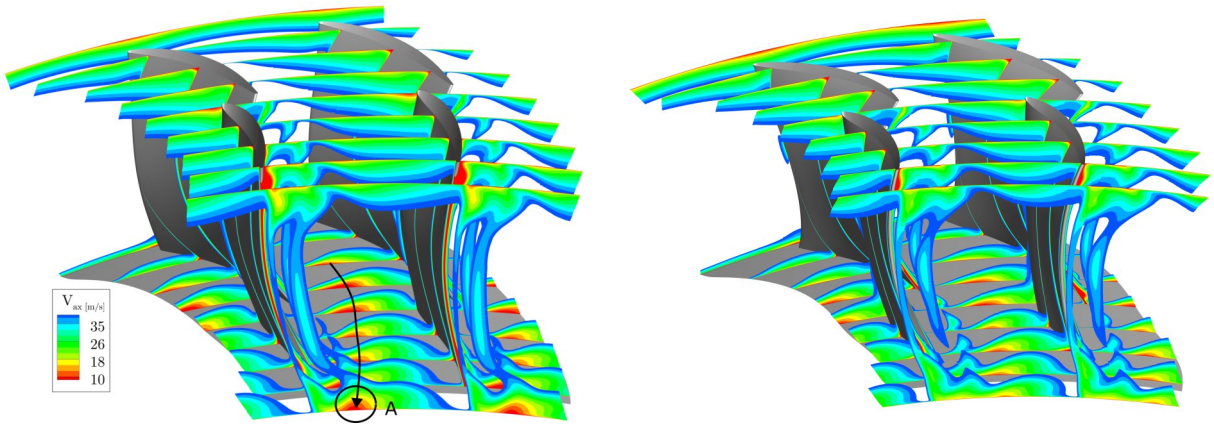

Figure 4: Contours of Axial Velocity for the Reference and Optimized Compressor Stators. Left: Reference. Right: Optimized.

Intrinsic instability and NO_x formation in 3D thermodiffusively unstable premixed hydrogen flames

Research Institution
Institute for Combustion Technology, RWTH Aachen University

Principal Investigator
Heinz Pitsch

Researcher
Xu Wen, Lukas Berger, Michael Gauding

Project Partner
–

Funding
ERC Advanced Grant (HYDROGENATE)

SuperMUC Project ID: pn49yo (Gauss Large-Scale project); pn36mi

Introduction

Hydrogen is drawing renewed and rapidly growing attention around the world. One of the most important advantages of hydrogen usage is that its thermochemical conversion does not lead to greenhouse gas (CO_2) emissions. For hydrogen combustion, the most challenging issues are the augmented effects of differential diffusion and the induced intrinsic instability, which can substantially change the flame dynamics and heat release rates, and lead to safety problems. The intrinsic instability of hydrogen combustion has two mechanisms: (i) the hydrodynamic instability due to the density jump across the flame front, and (ii) the thermodiffusive instability due to the disparity between the heat flux leaving the reaction zone and the mass flux entering the reaction zone, which is related to the effective Lewis number of the unburnt mixture. In a thermodiffusively unstable flame, the flame front is strongly corrugated over different length scales, especially at elevated pressures. The distributions of the thermo-chemical variables in such flames are closely related to the strain rate and curvature of the flame front.

Previous numerical and experimental studies reported that flames with various fuel-air mixtures become intrinsically unstable with increasing pressure. For highly diffusive fuel species such as hydrogen, the intrinsic flame instability would be further pronounced due to the introduced thermodiffusive instability of a lower Lewis number. Various direct numerical simulations (DNS) have been conducted for the thermodiffusively unstable premixed hydrogen flames focusing on different research aspects, e.g., evaluation of asymptotic theory, evaluation of confinement on flame dynamics, parametric study of pressure, equivalence ratio, and temperature, composition space modeling, etc. However, most of these studies adopted two dimensional (2D) configurations to approximate the combustion behaviors of the unstable premixed hydrogen flames. Although the characteristic patterns can be approximated to a certain extent, the three dimensional (3D) flame front corrugations cannot be approximated by a 2D configuration. To the authors' best knowledge, the 3D DNS has not been conducted for thermodiffusively unstable premixed hydrogen based on a sufficiently large computational domain without the dependence of the fuel consumption speed.

Although hydrogen is regarded as a "green" fuel as it does not emit greenhouse gases, nitrogen oxides (NO_x) can be formed when burning in an air atmosphere. As the distributions of the thermo-chemical quantities are strongly influenced by the local strain rate and curvature, the NO_x formation pathways can be significantly changed throughout the corrugated flame front. The NO formation mechanism in a 3D thermodiffusively unstable premixed hydrogen flame has not been investigated, particularly for elevated pressure conditions. The objective of this project is to conduct 3D DNS for thermodiffusively unstable premixed hydrogen flames in sufficiently large 3D computational domain for different pressure conditions, in which the NO_x formation is considered.

Results and Methods

For investigation of the 3D thermodiffusively unstable premixed hydrogen flame, the computational domain shown in Fig. 1 is adopted. It is a direct extension of the 2D computational domain of a planar configuration in [1], i.e., a box with periodic boundary conditions in the *x*- and *z*-directions, while it has an inflow and outflow boundary in the *y*-direction. As discussed by Berger et al. [1], the domain size affects the flame front wrinkling. Therefore, a domain length in the periodic direction of at least 100 thermal flame thicknesses is required to exclude spurious effects of domain size confinement on the flame front's evolution. The domain size is set to be 100 flame thermal flame thicknesses in the lateral *x*- and *z*-directions, with the length in the streamwise *y*-direction 150 thermal flame thicknesses. Specifically, 1,024 × 1,536 × 1,024 grid

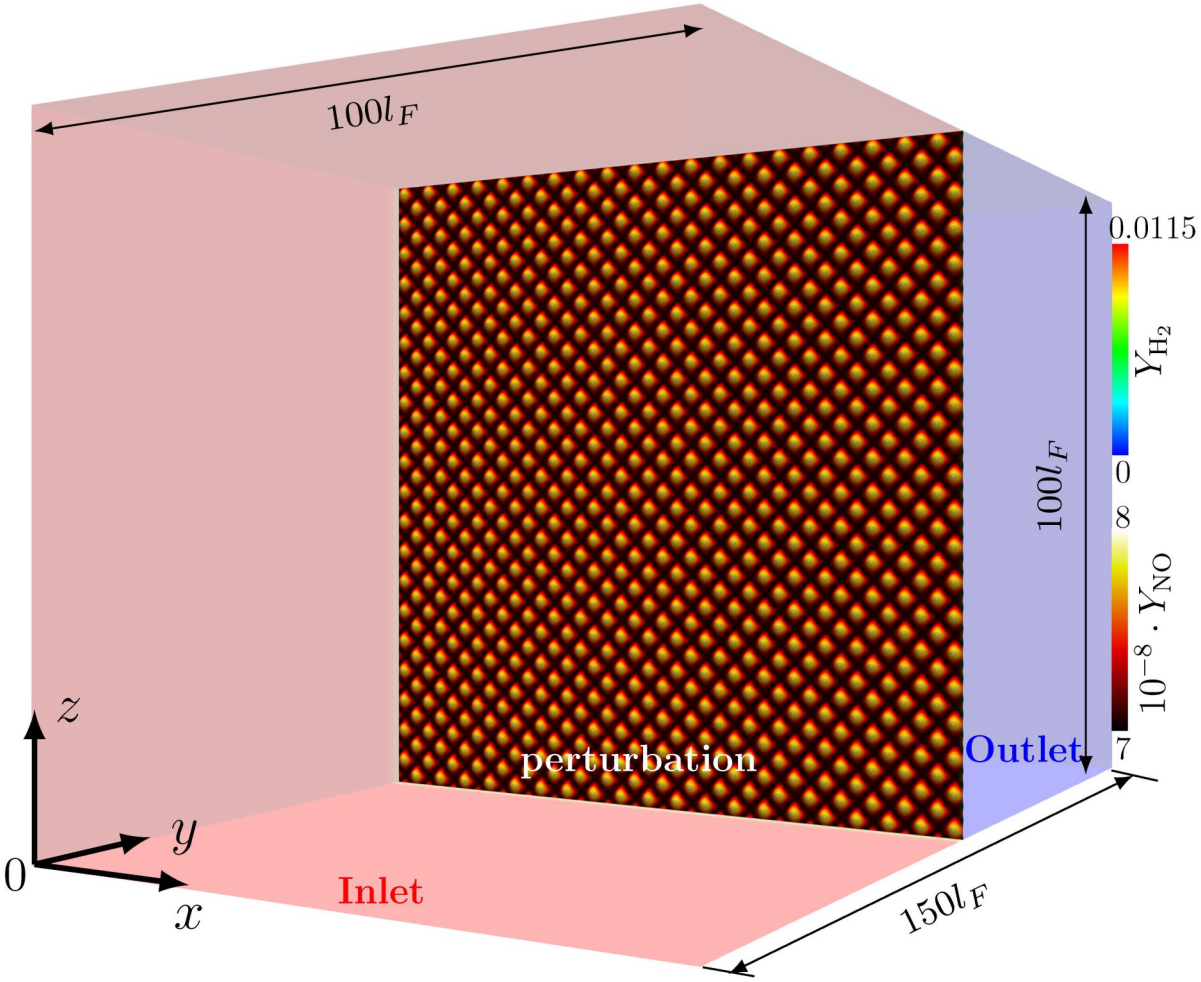

Figure 1: Computational setup of the 3D thermo-diffusively unstable premixed hydrogen flame.

points are uniformly set in the *x*-, *y*-, and *z*-directions, respectively, resulting in about 1.6 billion grid points.

The inflow is characterized by an unburnt lean premixed hydrogen/air mixture that is burnt within the simulation domain under the impact of thermodiffusive instabilities. Lean premixed hydrogen/air mixtures close to the flammability limit are characterized by small Lewis numbers, so the thermodiffusive instabilities are strongly enhanced. The inflow velocity was chosen such that the flame is stabilized within the simulation domain for a sufficiently long time interval, and the flame dynamics are not influenced by the inlet boundary conditions. The simulation was initialized by a planar flame that is obtained from a 1D unstretched flamelet solution. It was weakly perturbed in order to trigger the instabilities. Based on the findings of Berger et al. [1], the simulation has run for about 400 ms to study the non-linear regime, when the different flame front wrinkles interact with each other and yield a strongly corrugated flame front.

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 fourth order scheme. Species and temperature equations are discretized with a fifth order WENO (WENO5) 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. The code uses the message passing interface (MPI) standard for parallelization. The species diffusion velocity appearing in the species and temperature equations is modeled with the Curtiss-Hirschfelder approximation. The fluid is assumed to be an ideal gas, the Soret effect is included. Species diffusivities are determined by imposing spatially homogeneous, non-unity Lewis numbers. The Lewis numbers are taken in the burnt gas region of one-dimensional unstretched premixed flames, which was found to yield the best approximation for the unstretched laminar burning velocity.

The low-Mach number formulation of the Navier-Stokes equations removes acoustic effects from the governing

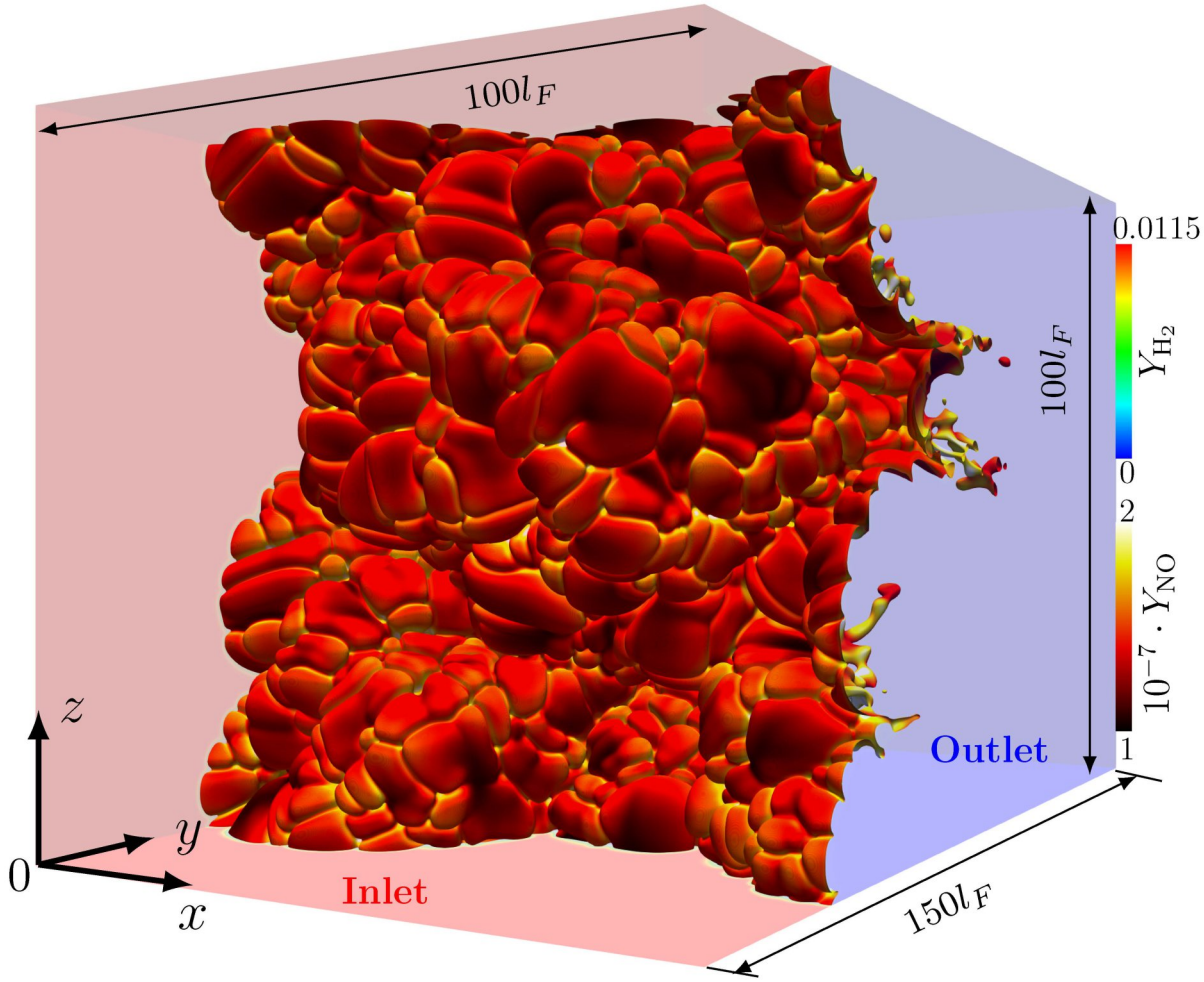

Figure 2: 3D cellular structure of the thermo-diffusively unstable premixed hydrogen flame. The color of the flame front indicates the local value of NO mass fraction [2].

equations and, therefore, avoids the limitations of the acoustic CFL number. As a result, the time step can be set to a significantly higher value compared to a fully compressible solver, so a lower number of total iterations is required to run the proposed DNS and study the long-term behavior of thermodiffusively unstable flames.

The 3D cellular flame structure is visualized in Fig. 2. In this project, the characteristic patterns of the 3D thermo-diffusively unstable premixed hydrogen flame have been analyzed in detail, and the effects of domain size and computational setup (2D vs. 3D) on the flame consumption speed, flame surface area, stretch factor, and the cell size have been quantified. In addition, a new flamelet model was proposed to predict the NO_x formation in thermo-diffusively unstable premixed flames, in which the wide range of curvature associated with the strongly corrugated flame fronts is considered by solving the premixed flamelet equations in composition space. The details of the results are referred to the published work in [2].

Ongoing Research / Outlook

The ongoing researches include (i) the a posteriori flamelet modeling of NO_x formation in thermodiffusivley unstable premixed hydrogen flames at atmospheric and elevated pressure conditions; (ii) further flame structure analysis and flamelet model development.

References and Links

[1] L. Berger et al., Proc. Combust. Inst. 37 (2019) 1879-1886. https://doi.org/10.1016/j.proci.2018.06.072
[2] X. Wen et al., Combust. Flame (2024) *in press*. https://doi.org/10.1016/j.combustflame.2023.113278

Towards understanding and prediction of strake vortices on aircraft

Research Institution
[1]Institute of Aerodynamics and Flow Technology, German Aerospace Center, Göttingen

Principal Investigator
Axel Probst[1]

Researcher
Philip Ströer[1,3], Elrawy Soliman[1], Tobias Knopp[1], Nils Rathje[2], Rolf Radespiel[2]

Project Partner
[2]TU Braunschweig
[3]University of Twente

Funding
DFG KN 888/22

SuperMUC Project ID: pn69mu

Introduction

On modern transport aircraft the front part of the high-lift system for take-off and landing, called slat, needs to have a gap where the jet engine is mounted. In this region, the flow on the main wing is more prone to flow separation. To counteract this effect, the engine nacelles are equipped with small vortex generators, called strakes. The longitudinal vortices generated by these strakes interact with the flow on the main wing upper surface and have a stabilizing effect on it. However, the physical mechanism of this vortex/boundary-layer interaction is not yet fully understood. Moreover, for the design of new aircraft the involved flow physics need to be accurately predicted by engineering simulation methods.

Thus, one goal of this project was to generate highly-accurate simulation data of the flow around a simplified strake - an inclined Delta wing - by means of an academic simulation method (wall-resolved large-eddy simulation, LES). With these data the physical mechanisms of vortex development and wandering (i.e. slow circular motion of the vortex core [1]) has been studied in detail. Besides, the data has been used to improve and validate a more efficient, but less accurate simulation method (hybrid RANS/LES, [2]) that is feasible for engineering purposes.

Results and Methods

The main part of the computational campaign on SuperMUC-NG involved a single wall-resolved LES of the flow around an inclined Delta wing geometry. In this approach, the turbulent motion in the flow, which occurs at extremely small scales at the typical high Reynolds numbers of aircraft flow, is resolved both in space and time throughout the entire flow field. This requires a very fine computational mesh with around 800 million cells and a correspondingly fine temporal resolution of the time-accurate simulation [3]. Moreover, the overall physical time of the flow simulation needs to be sufficiently long to obtain meaningful statistical data, including a correction of the turbulent field by the vortex wandering.

The simulations were conducted using DLR's flow solver TAU. Its unstructured data layout offers high flexibility in terms of the design of the computational mesh (comprising different cell types like tetrahedra, hexahedra or prisms) and the parallel partitioning of the problem using either geometric or graph-based domain decomposition. In this case, around 10,000 domains were created yielding roughly 80,000 grid points per CPU-core on SuperMUC-NG. Using classical MPI-based inter-process communication for parallel data exchange at the domain boundaries, this falls well in TAU's linear scaling range.

For the given problem size an amount of approximately 1.5 Terabyte of data was generated per two days of simulation. Therefore, much attention was put on efficient data extraction and reduction, e.g. storing and post-processing only sampled data as required for the respective evaluation. Data extraction and processing was realized using DLR's simulation framework FlowSimulator.

A snapshot of the resolved turbulent structures on and in the wake of Delta wing is shown in Figure 1.Based on these flow data, several important observations on the nature of vortex development and wandering could be made: Besides the small-scale turbulent fluctuations, the vortex in the wake of the Delta wing is indeed subject to a wandering motion. A vortex-core detection method could be employed to localize the instantaneous vortex centers (see Figure 2) and to characterize the wandering strength and speed. Moreover, by correcting the flow field according to the instantaneous vortex center, the flow velocity could be decomposed into a mean (see

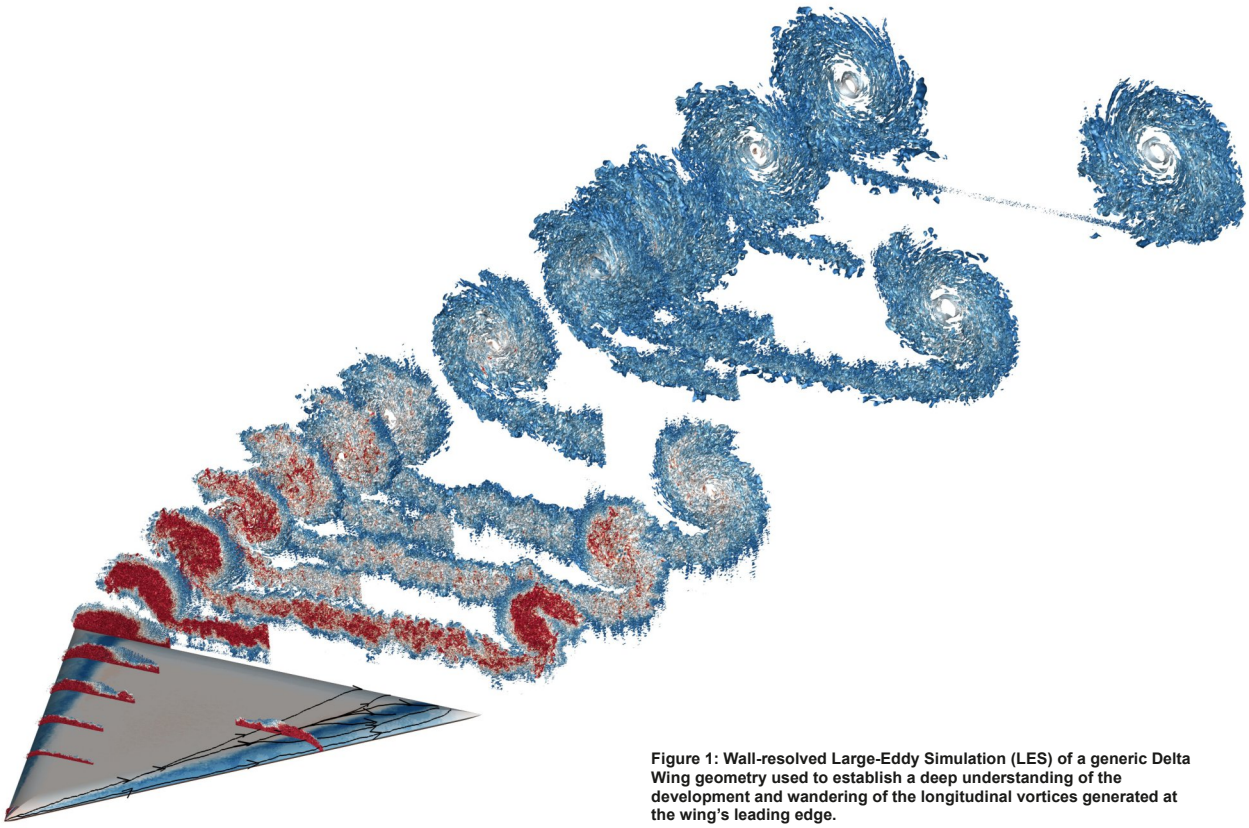

Figure 1: Wall-resolved Large-Eddy Simulation (LES) of a generic Delta Wing geometry used to establish a deep understanding of the development and wandering of the longitudinal vortices generated at the wing's leading edge.

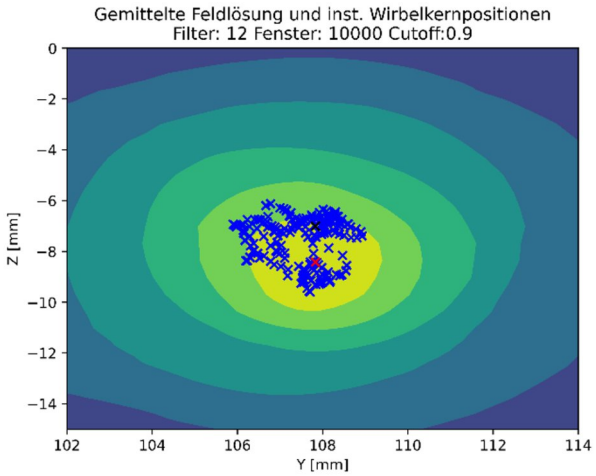

Figure 2: Detected instantaneous vortex centers (blue crosses) in the LES data from an automatic vortex-core detection method.

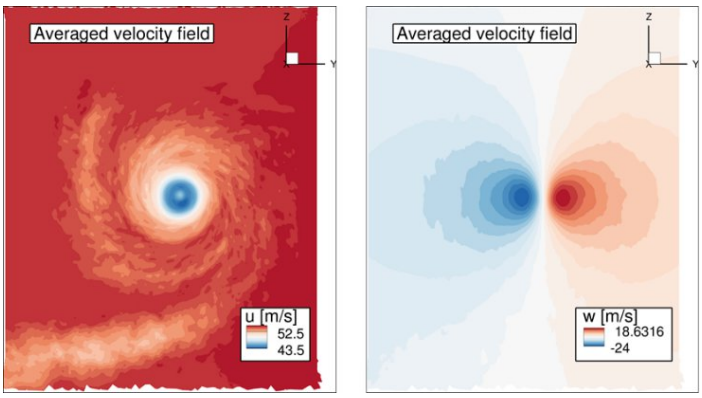

Figure 3: "Wandering-corrected" averaged velocity fields in the vortex computed with LES.

Figure 3), a turbulent (random), and a wandering (coherent) component. This triple decomposition allows analyzing the different phenomena independently and making meaningful comparisons to reference data. For instance, overall satisfying agreement with the PIV (particle-image velocimetry) measurements of the turbulent kinetic energy (TKE) in an accompanying wind-tunnel experiment could be shown, see Figure 4.

Another goal was to use the highly-accurate flow data from LES to improve the modelling approaches that are used in engineering simulation methods. These more efficient, yet less accurate approaches like hybrid RANS/LES employ turbulence modelling (instead of resolving) at least in parts of the flow domain. In the present investigation, the switch from modelled to resolved turbulence is realized by the injection of unsteady fluctuations from a synthetic turbulence generator (STG), which itself involves modelling approximations with significant uncertainties.

To reduce these uncertainties, the LES of the Delta wing was used to obtain reference data at the location of synthetic-turbulence injection and, thus, to replace modelling approximations by more accurate parameters. Thanks to this, both the turbulent frequency spectrum

and the so-called integral turbulent length scale could be prescribed more accurately at the interface from modelled to resolved turbulence in the hybrid simulation [4]. The hybrid simulations were also conducted on SuperMUC-NG, but required much less computational effort and were therefore not limited to one single simulation.

As visible in Figure 4, in particular the modification (i.e. reduction) of the turbulent length scale helps to significantly improve the prediction of turbulent kinetic energy with the engineering approach, when compared to the LES and the wind-tunnel experiment (PIV). This will allow engineers of future aircraft to accurately predict the stabilizing effect of vortices on the flow on the wing, making more efficient and climate-friendly designs possible.

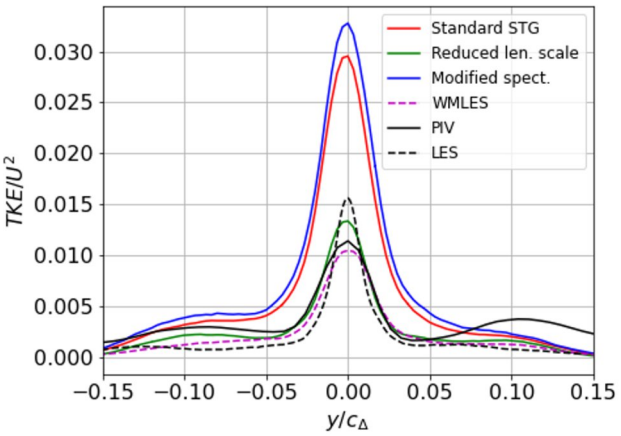

Figure 4: The turbulent kinetic energy (TKE) in the vortex core from different simulation methods, compared to experimental data (PIV). The green line shows the improvement of an efficient engineering method, by using input from the high-resolution LES.

Ongoing Research / Outlook

While the main part of the project work, i.e. the generation of highly-accurate LES data of the vortex on SuperMUC-NG and the derivation of an improved engineering simulation method has been concluded, the analysis of this “data treasure” to gain further knowledge on the flow physics of strake vortices is ongoing.

References and Links

[1] Devenport, W. J. et al. Journal of Fluid Mechanics, Vol. 312, 1996, pp. 67-106.
[2] Francois, D et al. Journal of Aircraft. Vol. 58, No. 6, 2021 DOI: 10.2514/1.C036273.
[3] Ströer, P. et al. SuperMUC-NG report-book 2022, pp. 114-115.
[4] Soliman, E. et al. STAB Workshop 2023, Göttingen.

Rocket engine components in close-up

Research Institution
Raumfahrtmobilität und -antriebe, TU München

Principal Investigator
Andrej Sternin

Researcher
Tiziano Santese, Oskar Haidn

Project Partner
–

SuperMUC Project ID: pr53we

Introduction

This report focuses on the examination of rocket engines, specifically targeting the cooling channels and combustion chamber. These components are critical for engine performance and durability, with the cooling channels thermal management and the combustion chamber facilitating energy release and thrust. To understand the complex physics within these parts, we conducted Direct Numerical Simulations (DNS) under conditions mimicking real engine environments. DNS offers a detailed view of the flow dynamics. The insights from our simulations aid in developing numerical models that accurately describe the observed phenomena, improving our understanding of fluid flow, heat transfer, and combustion dynamics. This knowledge is crucial for predicting engine performance and guiding design optimizations.

Results and Methods

Duct with ribbed wall

We conducted three-dimensional DNS simulations to investigate the characteristics of flow over a roughened wall in a duct, focusing on the variation of the Reynolds numbers from 4400 to 13200 and examining the influence of the side wall. Our findings offer valuable insights into the complex behavior of such flows and have important implications for a wide range of engineering applications. First and foremost, we observed a significant increase in the friction factor with respect to the smooth case, indicating heightened resistance to flow due to surface roughness. This increase was due to a remarkable deviation of the flow patterns with respect to the smooth case: three recirculation zones in the streamwise direction were identified which highlight the intricate separation phenomena caused by the presence of a roughened wall, Fig.1a. Additionally, the presence of two major vortices on the plane perpendicular to the flow reveals the intricate vortical structures formed in this configuration, Fig.1b. Furthermore, the influence of the Reynolds number on these vortices seems to be not well-scalable with respect to the bulk velocity, as usually happens for smooth ducts. Indeed, these vortices present an increasing near-wall intensity as the Re is increased, Fig.1c. Further analysis is conducted by investigating the streamwise vorticity distribution, which has its peak on the bottom wall and in the bottom corners but assumes important values also further away along the side wall. The influence of the side walls has been further investigated: an increased turbulent production rate is found along the side walls, where the secondary motions are stronger, and the dissipation rate is not enough to balance the TKE, resulting in non-equilibrium turbulence. Also, turbulent normal and shear stresses are found to be highly sensible to the Reynolds number and to the side wall, increasing when moving towards the side wall and when the Re is increased; in all cases, the symmetry (or antisymmetry) of the stresses is lost, due to the presence of the riblet which causes a significant peak of the stresses near the ribbed wall, which once again, becomes even stronger when Re is increased. Lastly, a convergence analysis is carried out for the first time, giving meaningful insights on the convergence behavior of a ribbed wall simulation against a smooth wall one; the convergence is found to happen quicker, probably due to the stronger intensity of the vortex which can be captured with less convective timesteps.

Combustion chamber

The approach of utilizing Direct Numerical Simulation (DNS) for analyzing combustion chambers presents a complex and costly challenge. It necessitates not only a deep understanding of design principles for accurately dimensioning and configuring the DNS setup for turbulent combustion studies but also necessitates the development of an effective strategy for processing and interpreting the resulting data. To manage computational demands, the size of the DNS domain is significantly reduced compared to actual combustion chamber geometries, typically by a factor of at least 20. Despite this reduction, the maintenance of comparable Reynolds, Damköhler, and Karlovitz numbers ensures the relevance and applicability of the results. Figure 2 illustrates the utility of DNS in combustion research. On the left, a

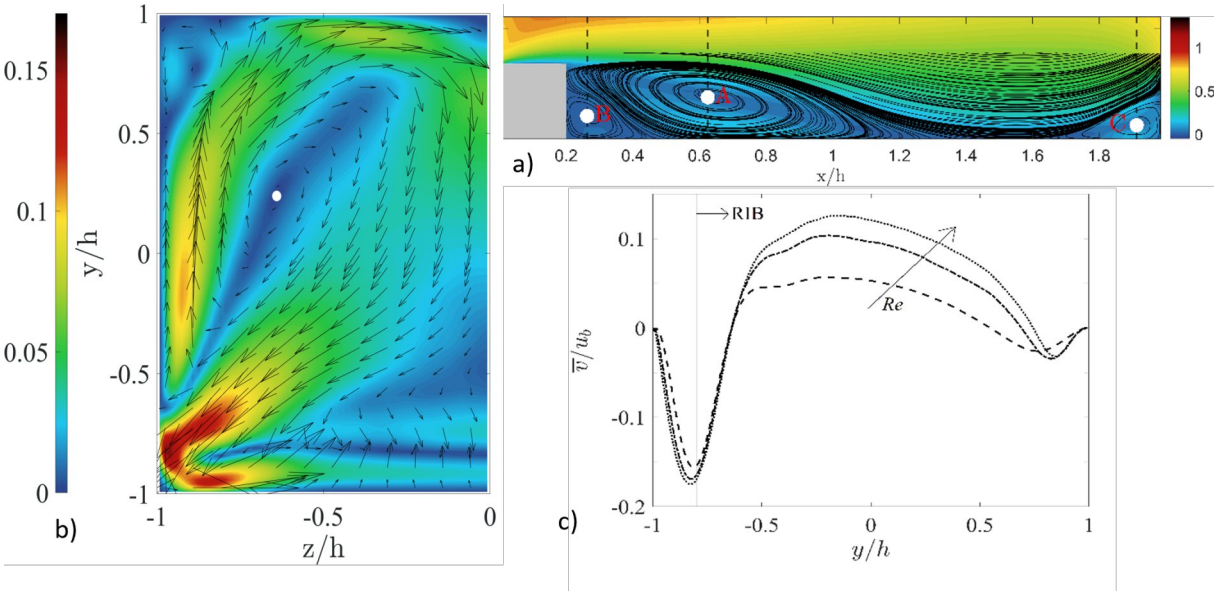

Figure 1: a) three recirculation regions are presented between two riblets, at the center-plane z = 0;b) cross-flow velocity presented in x/h = 0.625, where the center of vortex A is; c)vertical velocity component, presented for three Reynolds number in x/h = 0.625 and z/h = -0.95.

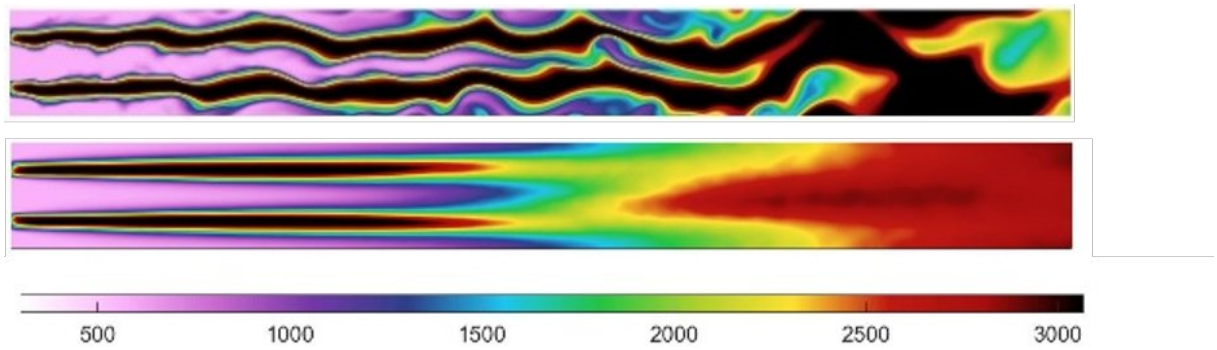

Figure 2: Instant and averaged picture of a diffusion flame from a DNS simulation.

snapshot from a DNS simulation of a diffusion flame vividly captures a moment in the combustion process. On the right, an image generated by averaging data over several hundred time points is presented. This averaging process filters out transient phenomena, highlighting only the repetitive patterns inherent in the combustion process. Through the application of various analytical tools to this data, insightful observations about different physical processes within the flame can be made. For example, the analysis reveals how the flame generates additional turbulence. Furthermore, it uncovers that while the increase in turbulence enhances mixing and therefore the combustion process initially, there comes a point where further increases in turbulence begin to hinder rather than help the process, becoming counterproductive.

Ongoing Research / Outlook

The primary challenge we face in conducting direct numerical simulation studies lies in managing the vast quantities of data. This challenge places significant demands on operating memory, storage capacity, and CPU resources. Despite these hurdles, we are committed to continuous improvement of the tools and methodologies used at every stage of the workflow. Through these enhancements, we aim to progressively increase the size of the simulation domain. This expansion allows us to capture physical phenomena and geometrical features at increasingly larger scales, broadening the scope and applicability of our DNS simulations. By pushing the boundaries of our computational capabilities, we pave the way for more comprehensive and detailed insights into complex processes.

References and Links

[1] T. Santese, D. Martinez-Sanchis, A. Sternin, et. al. “DNS investigation of turbulent flows in rectangular smooth and rough ducts”, Aerospace Europe Conference EUCASS-CEAS, 2023, Lausanne.

[2] D. Martinez-Sanchis, A. Sternin, T. Santese, et. al., “The role of turbulence in the characteristic velocity and length of rocket combustors”, Aerospace Science and Technology, 2023.

[3] A. Sternin, D. Martinez-Sanchis, D. Sternin, “Characterisation and Design of Direct Numerical Simulations of Turbulent Statistically Planar Flames”, Aerospace 9, 2022.

[4] I. Nasser, Y. Torres, T. Santese et al. “A comprehensive investigation of heat transfer in a high aspect ratio cooling channel of a rocket engine using LNG coolant”, Acta Astronautica, 213, 2023.

DNS study of the early flame kernel development under hydrogen engine conditions

Research Institution
Institute for Combustion Technology, RWTH Aachen

Principal Investigator
Heinz Pitsch

Researcher
Hongchao Chu

Project Partner
–

Funding
DFG, FOR2687

SuperMUC Project ID: pr89pa

Introduction

Hydrogen is a promising energy carrier of future energy systems, which can be used as a fuel in internal combustion engines (ICE) to reduce greenhouse gas and pollutant emissions. However, the low volumetric density, low minimum ignition energy, and high flame speeds of hydrogen-air mixtures present specific challenges for hydrogen ICE such as low volumetric power density, pre-ignition, and back-fire. Direct fuel injection can be applied in hydrogen ICE to mitigate such issues. Experimental studies of optical engines have shown that the mixture remains inhomogeneous near top dead center (TDC) for fuel injection during the compression stroke. Such inhomogeneities can affect the flame kernel development, which significantly impacts the entire combustion stroke. The present study investigated the effects of inhomogeneous equivalence ratio on hydrogen early flame kernel development under engine conditions. For this purpose, datasets of direct numerical simulations (DNS) for idealized engine geometries have been established and analyzed. Additionally, as a reference, flame kernel development for conventional spark-ignition fuel, iso-octane, has also been simulated and compared.

Results and Methods

Mixture state and composition in the simulation have been chosen to be representative for direct injection spark ignition (DISI) engine operation at medium load. In order to provide well-defined, but engine-relevant flow configurations, flame kernels have been computed in decaying isotropic turbulence. The turbulent Reynolds number, Damköhler Number, and Karlovitz number are specified as close as possible to engine conditions. To be consistent with realistic engine flame development, an ignition heat source has been employed to initialize the flame kernels. The source term is defined to vary smoothly in time and space such that the low-Mach number flow and ideal gas regimes remain valid. Multiple realizations were simulated by adding the ignition energy at different locations in the isotropic turbulence field in both homogeneous and inhomogeneous mixtures. In addition to flame kernels in turbulence, a laminar hydrogen flame kernel is also simulated. Visualizations of the flame kernels are given in Figure 1. Different realizations of the turbulent kernels are denoted by “k1", “k2", etc. Differences among the homogeneous flame kernels of each fuel are attributed to the different local flow conditions and their impact on the kernel evolution. When comparing the morphology of the inhomogeneous turbulent flame kernels with the corresponding homogeneous counterparts of each fuel, similar global shapes can be identified, indicating similar effects of the local turbulent flow on the corresponding flame kernels. Differences in the local topology between the corresponding homogeneous and inhomogeneous flame kernels can also be observed, which are attributed to the mixture inhomogeneity. In particular, mixture inhomogeneity prevents the breakup of iso-octane kernel 1, which occurs in the homogeneous mixture due to interactions with turbulence. The effects of the mixture inhomogeneity are also illustrated by the stronger variations of the heat release rates of the flame kernels in inhomogeneous mixtures. As expected, cellular instabilities occur for the laminar hydrogen kernel due to differential diffusion effects. The present DNS database has been generated with the CIAO code [1] using a time- and space-staggered finite-difference reactive Navier–Stokes solver on Cartesian grids. Mass, momentum, and kinetic energy are conserved discretely in the low-Mach number formulation. In this study, the momentum equations are discretized with a fourth-order-accurate central difference scheme. For the scalar transport equations, a weighted essentially non- oscillatory (WENO5) scheme is used. This scheme has fourth-order

Case	Core-h [mio]
Generation of flow field	0.8
Iso-octane flame kernels	8
Hydrogen flame kernels	13
Total	21.8

Figure 2: Snapshot of the temperature field (upper) and the mixture fraction (lower) in the methane-based combustion chamber.

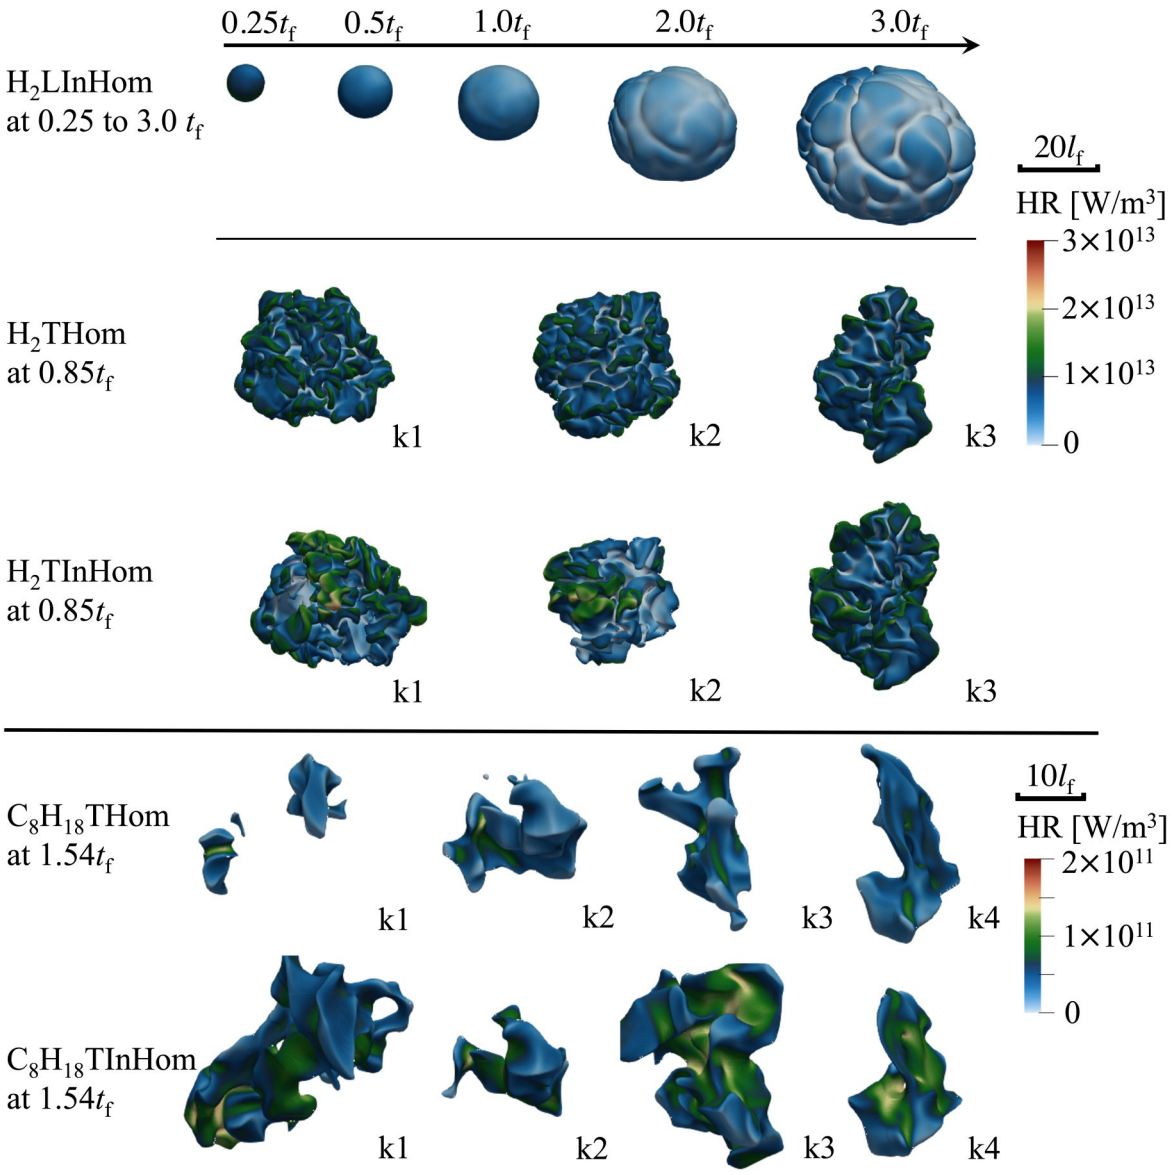

Figure 1: Flame kernel visualizations: iso-surfaces of temperature colored with heat release rate. H_2LInHom: hydrogen kernel in laminar inhomogeneous mixture. H_2THom: hydrogen kernel in turbulent homogeneous mixture. H_2TInHom: hydrogen kernel in turbulent inhomogeneous mixture. C_8H_{18}THom: iso-octane kernel in turbulent homogeneous mixture. C_8H_{18}TInHom: iso-octane kernel in turbulent inhomogeneous mixture. tf: chemical time scale. k1, k2, k3, etc.: different realizations.

accuracy in smooth monotone regions and third-order accuracy near discontinuities. To avoid inconsistencies between discrete mass conservation and scalar convection, a finite-volume mass flux is used in the scalar transport equations. Time integration is performed with a second-order-accurate Crank-Nicolson scheme. For time-advancement of the Navier-Stokes equations, a variant of the fractional step method is used. The pressure correction is obtained by solving the Poisson equation with the algebraic multi-grid solver Hypre Boomer AMG. To efficiently advance the stiff advection-diffusion-reaction equations for species and temperature, the symmetric operator splitting is used. The resulting system of ordinary differential equations for the zero-dimensional homogeneous reactor in each grid cell is solved using a fully time-implicit backward difference method. The Jacobian is evaluated analytically for computational efficiency. The datasets of iso-octane have been computed with a grid of 9,603 cells on SuperMUC-NG with 10,752 cores in the “general” partition and the datasets of hydrogen with 14,403 cells and 21,504 cores in the same partition. A summary of the performed simulations is given in Table 1.

Ongoing Research / Outlook

Current investigations focus on the effects of turbulence on the ignition of the mixture by spark plasma channel. Turbulence is expected to impact plasma channel shape and thus the initial kernel topology, which significantly influences the kernel growth rate. In addition, the heat loss to the spark plug will be considered in the simulations.

References and Links

[1] O. Desjardins, et al., J. Comput. Phys. 227 (2008) 7125–7159.

Lie-Symmetry theory of wall turbulence: V and W moments

Research Institution
[1]TU Darmstadt

Principal Investigator
Martin Oberlack[1]

Researcher
Sergio Hoyas[2], Cat Tuong Nguyen[1]

Project Partner
[2]Universitat Politècnica de València

Funding
ERC starting UltimateRB, ERC advanced MultiMelt

SuperMUC Project ID: pn49ya

Introduction

Wall turbulence is probably one of the open problems in physics with most applications in daily life. Even if the equations determining these flows have been known for almost 200 years, we still lack a complete theory. As wall-bounded turbulence is responsible for the 5% of the CO_2 dumped by humankind into the atmosphere every year, this is a problem of the uttermost importance. Experimental techniques dominated the research of turbulent flows until the eighties of the last century, when 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 is 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 in 1987, the friction Reynolds number has grown steadily. In earlier project, we have run a simulation reaching the 10,000 frontier. This simulation allowed us to study high Reynolds number effects. We published two highly cited papers [1,2], that were extraordinarily well received by our community. However, this friction Reynolds number is still less than the largest flow realization obtained by experimental means. Our objective in this project is to obtain an even higher Friction Reynolds number and scaling laws for the spanwise and wall-normal directions. We are running two simulations for Reynolds numbers 15k and 20k. 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, and statistical quantities such as the mean velocity are sufficient. However, considering turbulence as a statistical process 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).

The PI and his co-workers developed an approach that 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. In his 1905 seminal work on special relativity, Einstein 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 allows 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 10k 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. The code uses a low-storage third-order semi-implicit Runge-Kutta (R-K) time-stepper for time integration. 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 for the spatial derivatives. The use of CFD allows greater flexibility in the choice of the mesh, which consists of 533×10^9 points, needing 11TB of Ram memory. LISO is written in Fortran90, and for SuperMUC-NG, it is built with the Intel ifort compiler. It has been carefully checked and optimized for SuperMUC-NG using a test account and in a SuperMUC extreme scaling workshop. The implementation uses FFTW routines for the Fourier transforms (MKL library) and MPI + OpenMP

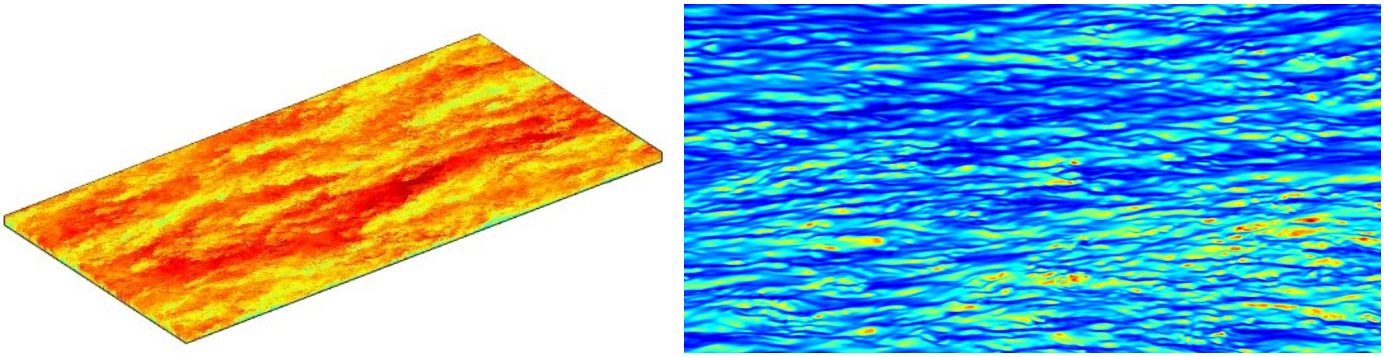

Figure 1: Left, Instantaneous visualization of the streamwise velocity at the logarithmic layer. Right, details of the streaks of the flow at the viscous layer.

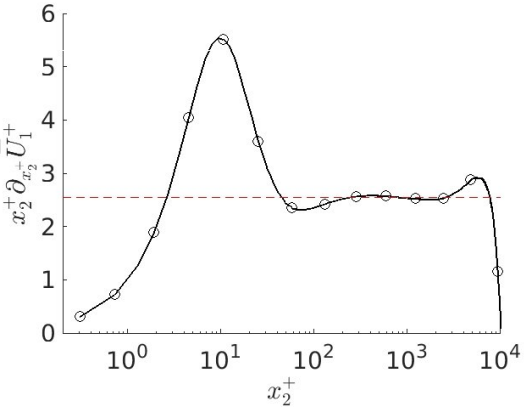

Figure 2: Indicator function. The flat region between 200 and 2,000 indicates the existence of a very long logarithmic layer. The value of 0.39 for the Karman constant can be taken as almost definitive for this range of Reynolds numbers.

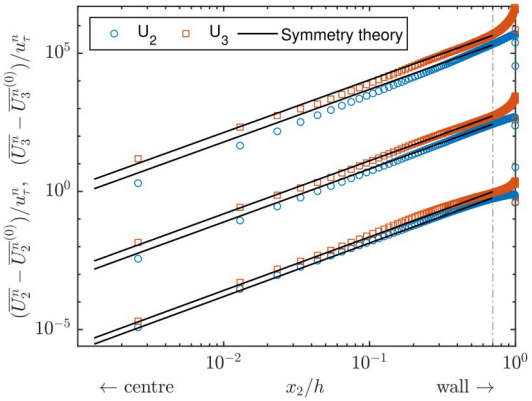

Figure 3: From bottom to top, scaling of moments of the wall-normal and spanwise velocities for moments 2, 4 and 6. Dot line, data coming from DNS. Continuous line, data from theory.

for the parallelization, where Intel MPI is our preferred environment for MPI. I/O routines use the parallel implementation of HDF5 in SuperMUC. It took less than 20sec to save an instantaneous flow realization, with a size of 1.5 TB. This field was distributed in as many as 171 nodes. Due to limitations on the code, the maximum number of processors is 8,192 cores. As we needed to run the code for a considerable time, more than two years, to get enough statistics from the turbulent field, we ran three independent cases simultaneously. In total. In total LISO has been running for about 60M core hours. 200M core hours are at least needed to finish this activity. This simulation has provided many interesting results about the kinematics of boundary layers. One of the most important features it that it allowed us to validate the symmetry-based Turbulence theory, 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 predicts 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 order 7, 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

We sincerely think that Symmetry Theory can be a game changer in the theory of turbulent flows. Thus, we plan to investigate further the theory of symmetry for what we need more statistics from the flow. We have applied for a large project, where we think we will be able to unveil more conservation laws of turbulent wall-bounded flows.

References and Links

[1] Oberlack, Martin, et al. "Turbulence Statistics of Arbitrary Moments of Wall-Bounded Shear Flows: A Symmetry Approach." Physical Review Letters 128.2 (2022): 024502.
[2] Hoyas, Sergio, et al. "Wall turbulence at high friction Reynolds numbers." Physical Review Fluids 7.1 (2022): 014602.
[3] Hoyas, Sergio and Oberlack, Martin. "Simulations at f." The near-wall region of wall-bounded turbulence up to friction Reynolds number of 150007. *Submitted*.

Processes and Upscaling of Nanoparticle Spray-Flame Synthesis

Research Institution
TUM - Chair of Aerodynamics and Fluid Mechanics

Principal Investigator
Marcus Giglmaier

Researcher
Sebastian Klukas

Project Partner
–

SuperMUC Project ID: pn34gi

Introduction

The SuperMUC-NG resources are applied to explore nanoparticle spray-flame synthesis. This method is highly efficient for producing nanoparticles with precise control over their size and composition, which enables the tailored creation of materials with specific properties for various applications. Nanoparticle spray-flame synthesis also allows for scalable production, making it suitable for industrial applications and advancing the development of materials in fields like electronics, medicine, and energy. This report covers spray-flame calculations and an analysis of two upscaling approaches.

Results and Methods

The laboratory sized SpraySyn burner [1] is utilized for analysis due to the extensive experimental measurements conducted on it. The spray-flame synthesis setup comprises a central injection of a liquid ethanol/precursor mixture in a coaxial nozzle. Pure oxygen is inserted through an outer annular gap at a much higher speed than the entrained liquid. The emerging turbulent gas jet introduces high shear forces at the liquid-gas interface, which eventually leads to ligament formation and break-up into small droplets. A surrounding pilot flame supports spray evaporation, ignition, and stabilization of the spray-flame. Nanoparticles form within the spray-flame.

Numerical boundary conditions are determined by the SpraySyn configuration [1]. Given that the dense spray obstructs measurements near the injection nozzle, interface resolving simulation results of primary atomization are used to extract spray boundary conditions for subsequent Eulerian-Lagrangian spray-flame calculations.

To model spray-flame synthesis with an iron nitrate precursor, a specific OpenFOAM solver has been developed. An exemplary simulation result is visualized in Fig. 1. The implementation features a Eulerian-Lagrangian approach to reflect the reactive multiphase flow. A flamelet progress variable (FPV) combustion model is implemented to accurately predict turbulence-chemistry interactions. Population balance modeling is applied to track the inception and growth of nanoparticles. The three-dimensional geometry of the SpraySyn burner is discretized by about $4.5 \cdot 10^6$ cells for RANS and $13.5 \cdot 10^6$ cells for LES simulations. Obtaining averaged results, including spray variables, requires about 500,000 core hours for LES simulations. For RANS calculations approximately half the computation time is needed. The influence of different RANS turbulence models, in comparison to an LES result is depicted in Fig. 2.

Burner scale-dependent effects on the synthesis of iron-oxide nanoparticles are investigated. For that purpose, the laboratory-scale SpraySyn burner is compared to the pilot plant-scale IUTA burner (*Institut für Umwelt & Energie, Technik & Analytik*) by numerical means. Key findings include the high-temperature particle residence time (HTPRT) being almost identical across scales despite differences in burner design and geometry. Nonetheless, the pilot-scale burner generates larger nanoparticle agglomerates caused by higher particle concentrations. The pilot flame design of the IUTA burner is integrated into a two-step simulation approach, which treats the premixed methane/oxygen pilot in a separate simulation. This leads to two distinct numerical domains, which are illustrated in Fig. 3. Fig. 4 compares simulation results of the laboratory and pilot plant burner against each other. It shows that the distribution of the initially formed monomer particles and coagulated particles is different across the burners. In the SpraySyn burner particles accumulate around the central axis, whereas the IUTA burner generates particles along a straight line at an angle away from the central axis. This is caused by the different distribution of the spray's liquid volume fraction, Fig. 4a.

An alternative approach for scaling up nanoparticle synthesis is characterized by injecting the precursor liquid into a steadily anchoring hydrogen/air flame [2]. To explore this configuration a numerical study of

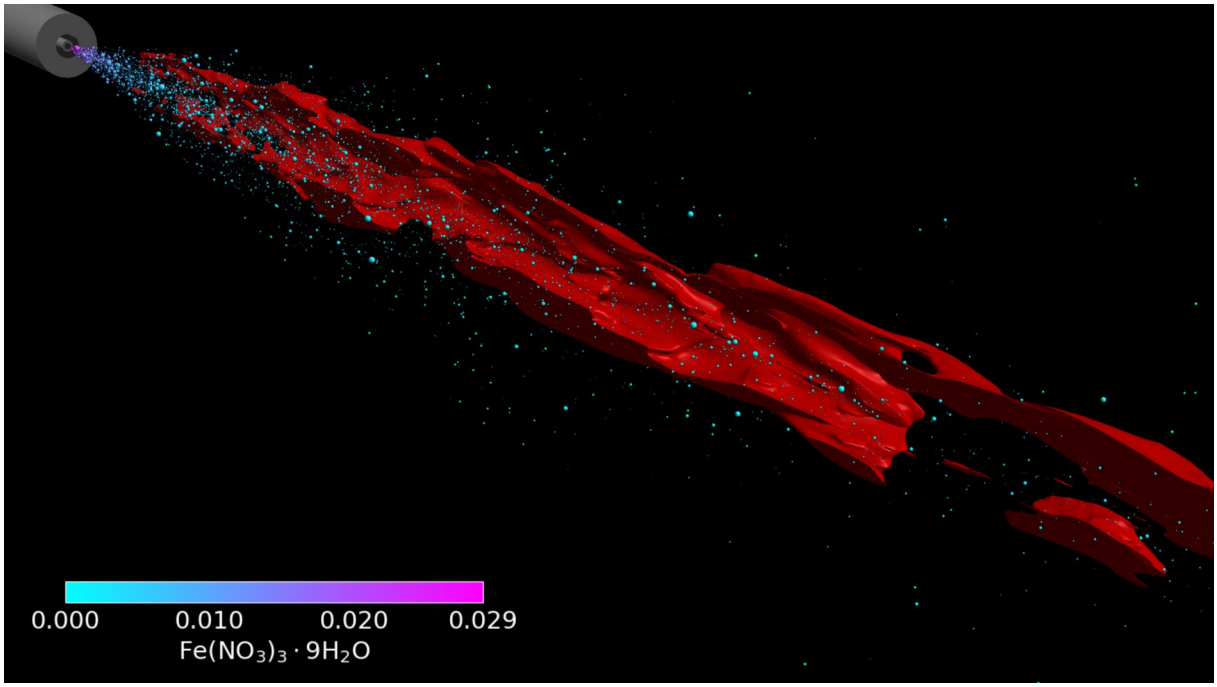

Figure 1: 3D snapshot visualization of SpraySyn LES results. The stoichiometric mixture fraction is indicated by the red isocontour. The color of spray droplets shows the concentration of remaining iron nitrate precursor.

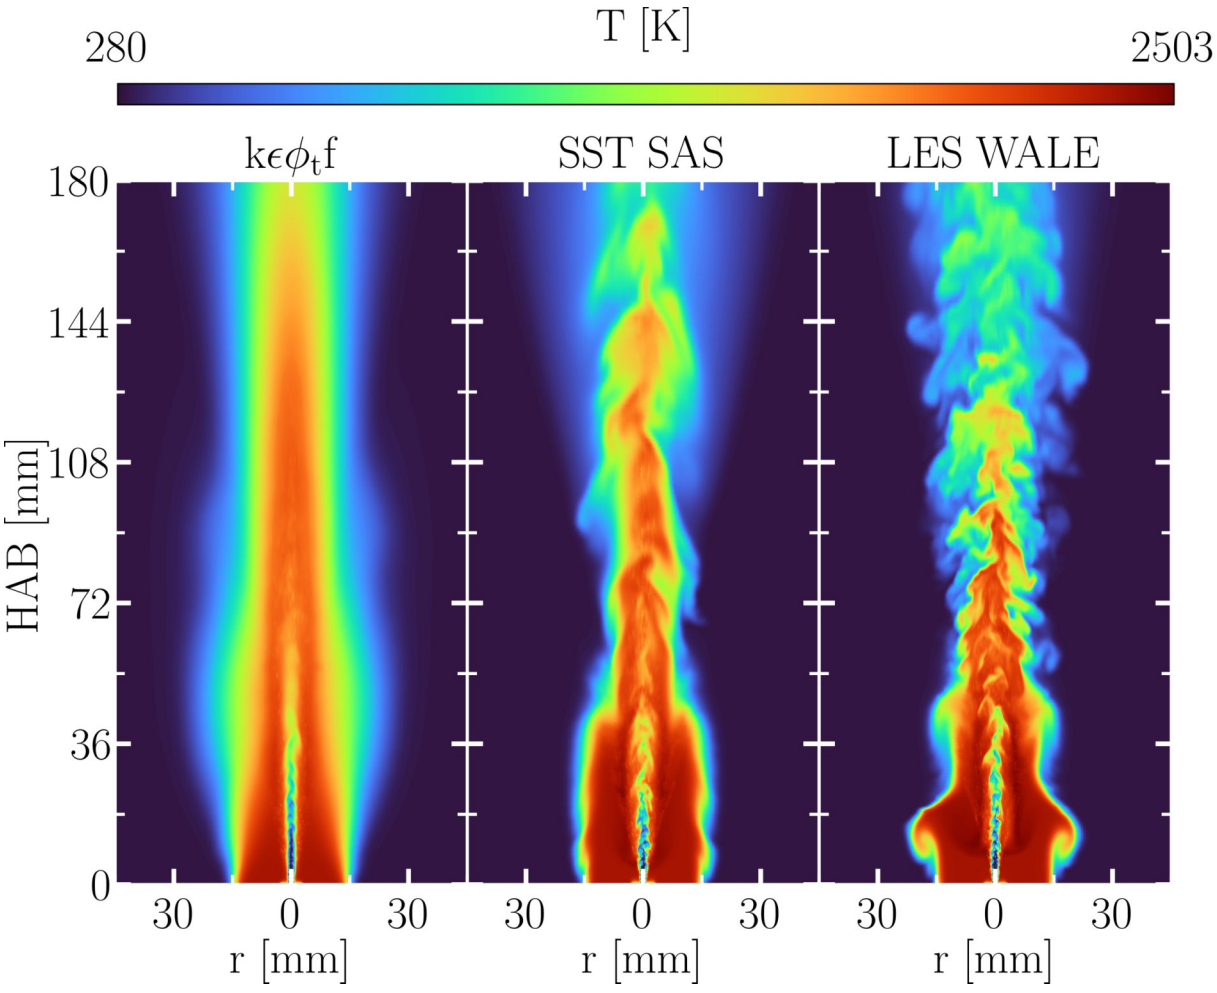

Figure 2: SpraySyn simulation results for different RANS turbulence models in comparison to an LES outcome. Temperature color maps are displayed.

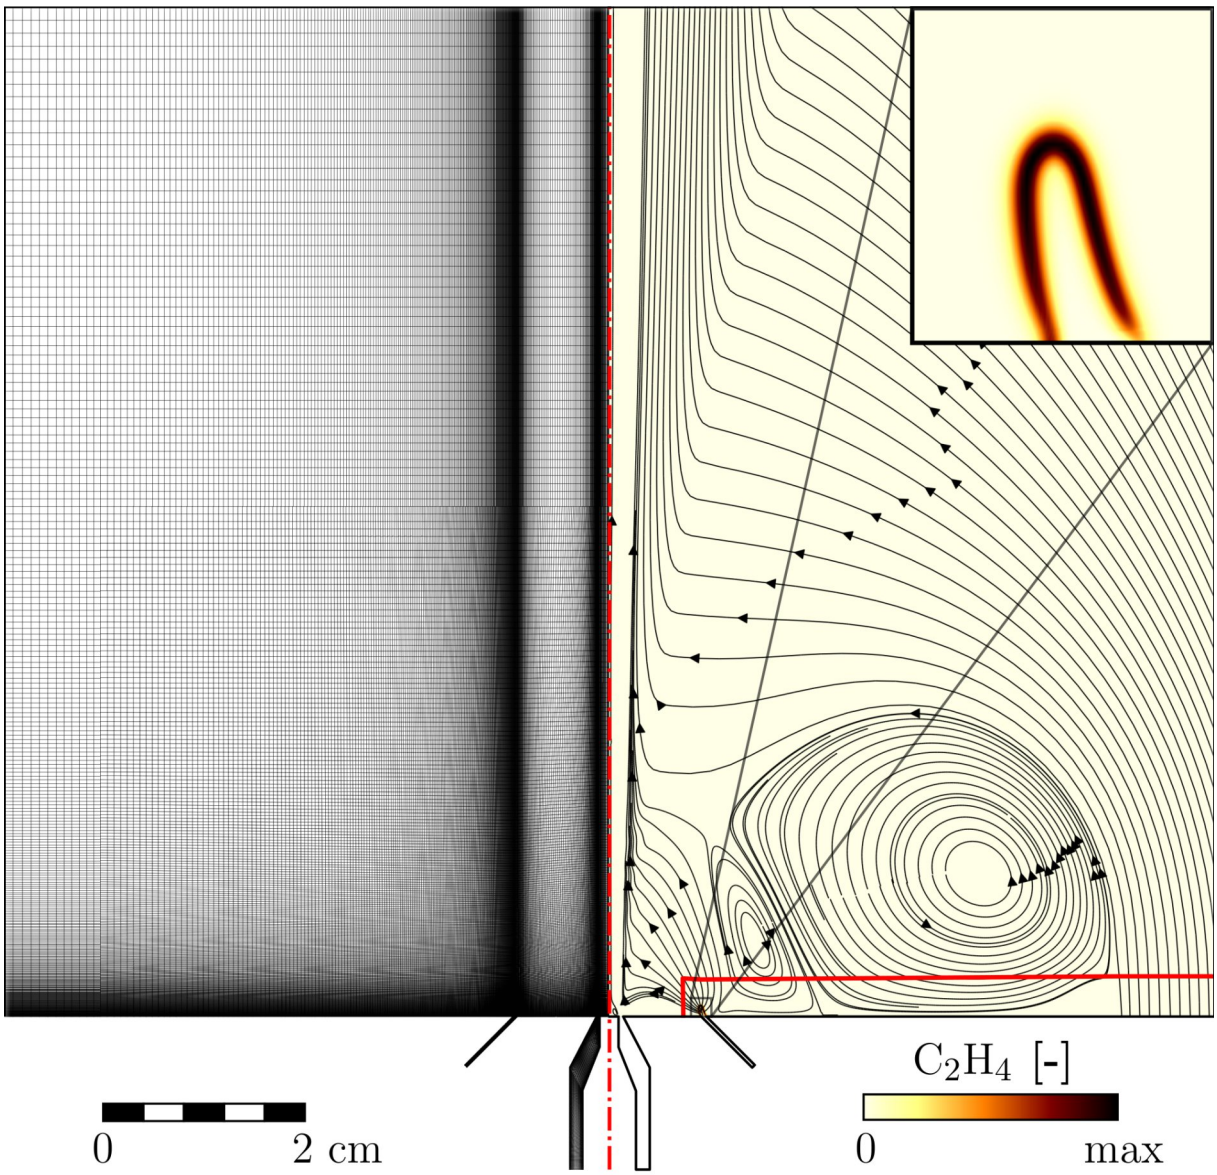

Figure 3: Visualization of the grid design (left), velocity streamlines and pilot flame front location (right) of the IUTA burner. Two interacting vortices between pilot and co-flow inlet are evident. The pilot flame front location is signified by a quickly recombining speciesand confirms that the reaction zones of the spray-flame and pilot flame do not interfere with each other. Accordingly, the simulation domain separation location is visualized by the solid red line. Adopted from [3].

turbulent premixed hydrogen/air flames anchoring on externally heated walls, Fig. 5, was conducted and compared to experimental measurements. It aims to bridge the knowledge gap between laminar flame studies and real-world applications, emphasizing the importance of flame behavior in turbulent conditions. This research might also be interesting for the development of efficient and safe hydrogen energy systems, considering the growing interest in hydrogen as a clean energy carrier. The findings are conclusively published in [2].

Ongoing Research / Outlook

Based on the numerical investigations of scale-dependent effects in nanoparticle spray-flame synthesis and hydrogen pilot flames for large-scale applications, a novel pilot plant-scale spray-flame synthesis burner is being developed in conjunction with the project partners at IUTA.

References and Links

[1] F. Schneider et al., SpraySyn - A standardized burner configuration for nanoparticle synthesis in spray flames, Review of Scientific Instruments 90 (8), 5090232 (2019).
[2] S. Klukas et al., Anchoring of turbulent premixed hydrogen/air flames at externally heated walls, International Journal of Hydrogen Energy 45 (56) (2020) 32547–32561.
[3] S. Klukas et al., Iron oxide nanoparticle synthesis: Simulation-based comparison of laboratory-and pilot plant-scale spray-flame synthesis, Applications in Energy and Combustion Science 18 (2024) 100263.

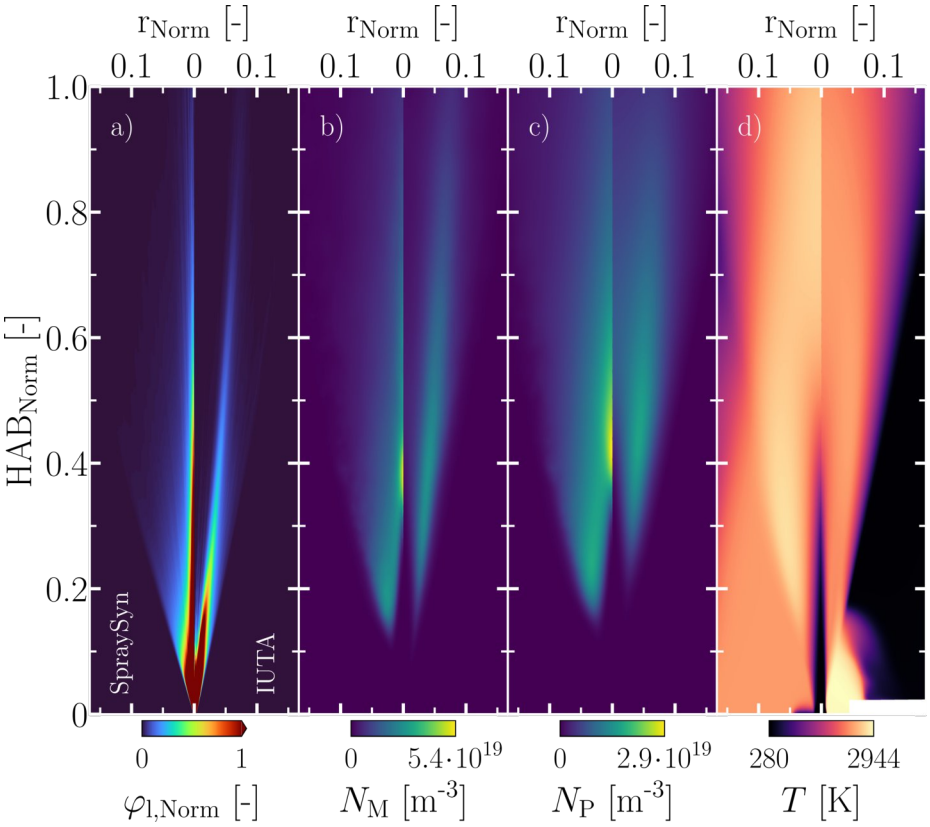

Figure 4: Color maps of normalized liquid volume fraction, $\varphi_{l,Norm}$, particle concentrations, N_M and N_P, and temperature, T, for the laboratory-scale SpraySyn (left) and pilot plant-scale IUTA burner (right). Spatial normalization ensures uniform spray-flame lengths across burners. Adopted from [3].

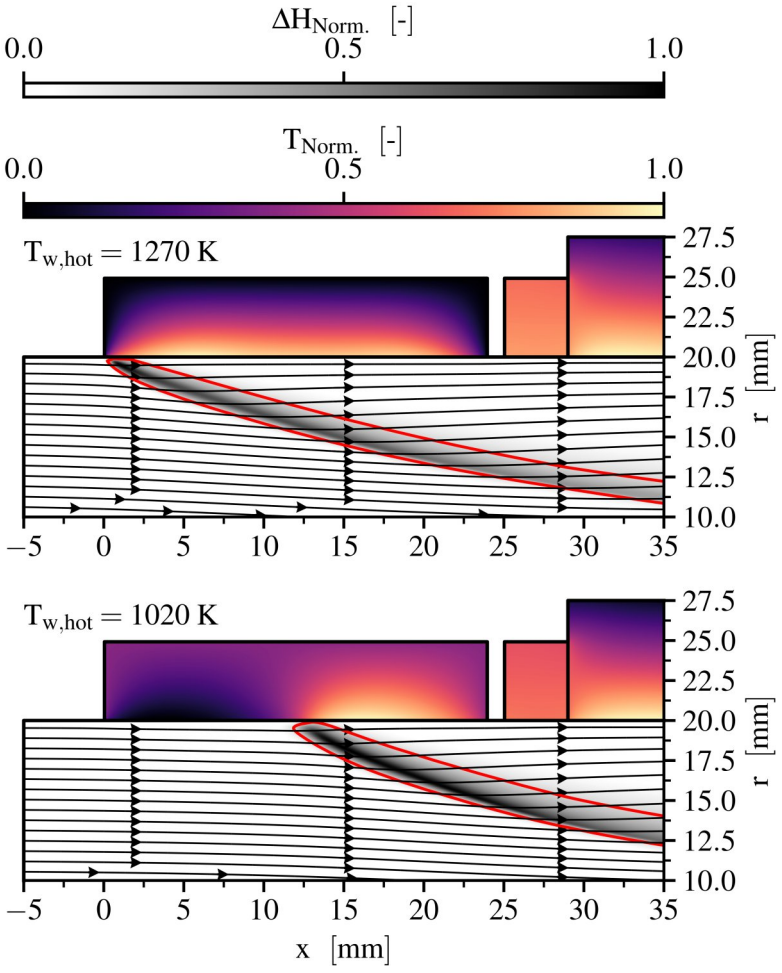

Figure 5: Influence of hot wall temperature on flame-wall interaction and anchoring position. Steady flame anchoring is visualized for $T_{w,hot}$ = 1,270 K (top) and $T_{w,hot}$ = 1,020 K (bottom). RANS simulation results with the k-ω BSL turbulence model are shown. The velocity field is plotted by streamlines while red isolines replicate one percent of the normalized heat release rate $\Delta H_{Norm.}$ distribution. Solid temperature contours are normalized separately for the heated and non-heated wall segments. Adopted from [2].

Rocket engine components in close-up

Research Institution
[1]SusatainableFuture Mobility, TU München
[2]Space Propulsion and Mobility, TU München

Principal Investigator
Andrej Sternin[1,2]

Researcher
Daniel Martinez-Sanchis[1,2], Oskar Haidn[2], Agnes Jocher[2]

Project Partner
–

SuperMUC Project ID: pn57ge (Gauss Large-Scale project)

Introduction

This project is engaged in an extensive endeavor to enhance the physical comprehension and mathematical description of high-pressure combustion processes involving hydrocarbons and hydrogen. The objective is to refine numerical tools to advance the design of combustion devices for space propulsion applications. The primary methodology employed in this research is Direct Numerical Simulations (DNS), a subset of Computational Fluid Dynamics (CFD) simulations. DNS offers a precise and direct analysis of physical quantities with negligible uncertainties. However, the computational demand for this method is substantial, often requiring several million core hours. Consequently, High-Performance Computing (HPC) clusters, such as the LRZ, are essential to provide the necessary computational resources. The execution of these simulations not only consumes significant computational resources but also generates vast amounts of data. This necessitates the application of sophisticated mathematical treatments, including advanced averaging techniques and pattern analysis. The development of big-data tools has been imperative to manage and interpret the detailed results from these simulations. Although the simulation conditions were initially inspired by aerospace applications, the methodologies and derived mathematical models have broad applicability across various combustion-related industries, particularly within the energy sector. The insights gained from our simulations hold potential for widespread industrial application. Subsequent sections of this report will detail the achievements resulting from this project.

Results and Methods

In this section, we detail the methodologies devised or refined to manage extensive simulations and the resultant large-scale data.

Simulation Perimeter:
The simulation domain is designed to replicate the geometry of the injection region within a rocket engine, as illustrated in Figure 1. It includes features such as the recess in the injection plate and the post tip area between oxygen and hydrogen injections. Multiple injection elements are positioned adjacent to each other to study their interactions. A sufficient flow length is incorporated to observe the majority of the mixing process. The scale of these features is approximately 1:20 relative to those in actual rocket engines. This configuration is projected in depth, rendering the case pseudo-2-dimensional. Unlike real rocket engines, which have a circular injection geometry, this simplification to a plate-like form is intended to facilitate detailed studies by maintaining a manageable complexity. This approach helps avoid the compounding of variables that could obscure the analysis of physical interactions. The simplification extends to equalizing the velocities of injected gases to eliminate phenomena caused by initial shear stresses.

Mesh Generation:
To adhere to the highest scientific standards, the mesh is entirely composed of homogeneous hexahedrons with identical grid spacing in all directions. The size of these cubes is determined by the lesser of two values: the smallest Kolmogorov eddies expected within the domain and the thermal and reactive gradients that must be resolved. For this purpose, the OpenFOAM mesher "blockMesh" is deemed suitable. It eliminates the need for a preliminary step of geometry generation, instead directly creating the mesh. Given that OpenFOAM was not originally designed to manage meshes comprising billions of cells, a parallelization of the entire meshing procedure was necessary. This involved the customization of meshing, refineMesh, and mapping processes through the development of proprietary scripts.

Initialization:
The gases introduced into the simulation domain are set at a temperature of 300 K, preventing autoignition. Unlike real rocket engines, which employ various techniques for starting and ignition, this aspect is not pertinent to our study. Our focus is solely on the stationary burning conditions to generate relevant data. Consequently, the resources expended to achieve these conditions are

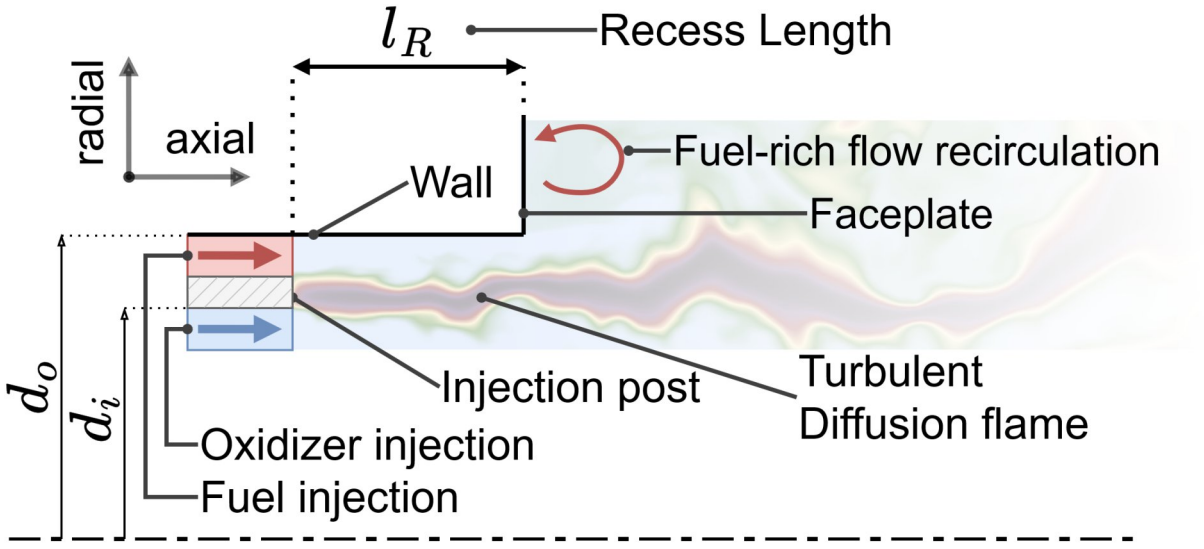

Figure 1: Injection geometry and most relevant phenomena [1].

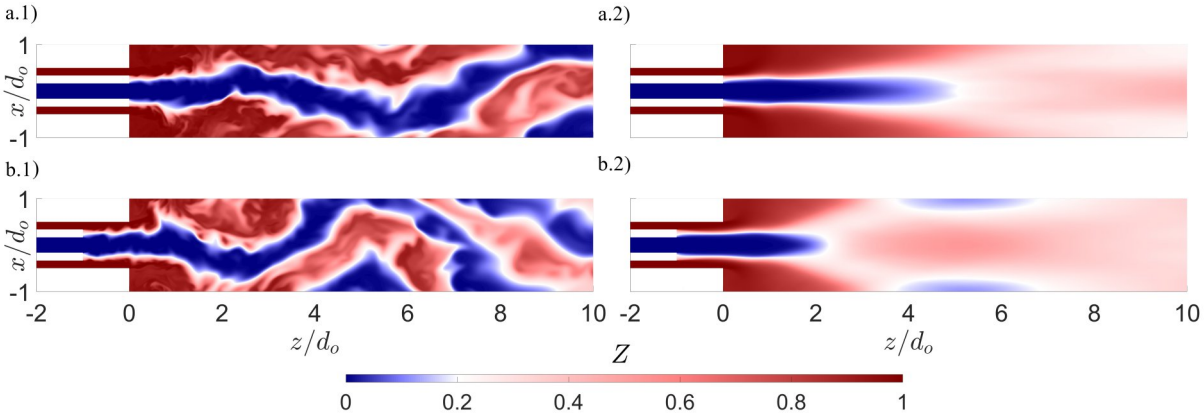

Figure 2: Instant and averaged picture of mixture fraction with two different injection geometries from DNS [1].

considered non-recoverable, potentially amounting to several million core hours. To minimize these costs, a sophisticated initialization procedure was developed. Initially, a low-resolution solution is generated using the same geometry, where stable flow conditions are achieved without chemical reactions. Subsequently, the area between the injected gases is heated to induce ignition. Once a stable flame pattern is established, the initial heat source is removed. The entire field is then mapped onto a mesh with twice the resolution, and this process of mapping and refining continues until the desired cell size is reached, enabling the generation of relevant data.

Data Generation:
To optimize storage space, data generation was tailored based on spatial and temporal independence, recognizing that writing out all data points would introduce redundancies due to the proximity of positions sharing common eddies and phenomena. Although this additional data can be useful, it is not as valuable as statistically independent data. Analysis procedures, including Right-tailed significance tests and turbulence theories, guide the selection of specific parts of the simulation domain and appropriate data saving frequencies. The resultant data fully utilizes available storage while maximizing utility.

Postprocessing:
To obviate the need for downloading the entire dataset from the cluster, partial processing is conducted in situ. Initially, data from all ranks are compiled into common files for preliminary averaging. This averaged data is then downloaded for further analysis. The complete datasets are retained on the cluster, along with the simulation case and its configurations, to facilitate future simulation runs. All potentially useful data for future research is transferred to the Scientific Storage (DSA), ensuring its accessibility for years to come.

Ongoing Research / Outlook

As this project is currently in progress, we anticipate a substantial volume of data to become available for postprocessing. This presents a significant opportunity to employ machine learning tools to aid in the analysis and identify patterns that may be overlooked by human observation.

References and Links

[1] D. Martinez-Sanchis, A. Sternin, O.J. Haidn, A. Jocher, Physics of Fluids 2024.
[2] D. Martinez-Sanchis, A. Sternin, T. Santese, et. al., Aerospace Science and Technology, 2023.
[3] D. Martinez-Sanchis, A. Sternin, O.J. Haidn, A. Jocher, Physics of Fluids 2023.

Identification of entropy waves in a partially premixed combustor

Research Institution
[1]Associate Professorship of Thermo-Fluid Dynamics, Technical University of Munich

Principal Investigator
Wolfgang Polifke[1]

Researcher
Alexander J. Eder[1]

Project Partner
[2]CAPS Laboratory, ETH Zurich, Zurich
[3]CERFACS, Toulouse

Funding
DFG transfer project NoiSI (PO 710/23-1)

SuperMUC Project ID: pn57so

Introduction

Unsteady combustion in hydrogen-powered gas turbines may generate coherent temperature fluctuations in the burnt gas, also known as entropy waves. These inhomogeneities are convected through the combustion chamber and attenuated by turbulent mixing and spatial variations in the velocity field. In aero-engine combustors, where convective dispersion is slow relative to the residence time, entropy waves often reach the turbine stage with high amplitude fluctuations, where they are accelerated. This acceleration results in a mode conversion into pressure oscillations (i.e. acoustic waves). On the one hand, the downstream traveling acoustic waves add to the overall noise emissions of an aircraft as so-called indirect combustion noise. On the other hand, the upstream propagating acoustic waves might interact constructively with the unsteady combustion (i.e. the heat released by the flame) and cause thermoacoustic combustion instability, particularly at low frequencies (known as "rumble"). This coupling is achieved through multiple connected mechanisms involving fluctuations of entropy, pressure, velocity and fuel concentration, and can lead to reduced combustor lifetime or even system failure. In this work, lean partially premixed (or technically premixed) turbulent flames are considered, where the fuel is injected close to the burner outlet resulting in spatial mixture inhomogeneities at the flame. Fluctuations in the air-fuel ratio can contribute to combustion instability only if driven in a coherent manner by the combustion process and pressure oscillations, thus closing a positive feedback loop. Coherent fluctuations in fuel concentration originate from pressure and velocity fluctuations at the location of the fuel injection, which are convected to the flame. These locally richer and leaner pockets of fresh mixture cause local changes in the adiabatic flame temperature, i.e. entropy waves. In the absence of air-fuel ratio fluctuations, entropy waves may occur due to unsteady wall heat loss (dominant effect) or differential diffusion, as shown in our previous study on entropy waves in a fully premixed non-adiabatic combustor with hydrogen enrichment [1]. These mechanisms are coherent with velocity fluctuations.

The two excitation mechanisms mentioned above and their interaction result in a frequency-dependent response of entropy to flow perturbations, known as entropy transfer function (ETF). The model of a partially premixed combustion system with a large pressure drop across the fuel injectors located in acoustically compact distance to the burner outlet reduces to a single-input single-output (SISO) model structure. For the configuration considered here, the forcing of air mass flow determines the overall ETF in partially premixed flames. In this work, high-fidelity large eddy simulation (LES) is combined with system identification (SI) [2] to determine the frequency response of entropy waves.

Results and Methods

Large eddy simulations of the swirl-stabilized, turbulent, partially premixed ETH combustor shown in Fig. 1(a) are performed using the Fortran-based, massively parallel, explicit, cell-vertex AVBP code [3] developed by CERFACS, which solves the multi-species compressible Navier-Stokes equations. Subgrid Reynolds stresses are modeled using the SIGMA turbulent closure without wall functions. A second-order Lax-Wendroff scheme is used to discretize the convective terms. The time step is fixed to 1.3×10^{-8} s, assuring an acoustic Courant-Friedrichs-Lewy (CFL) number below 0.8. The computational domain shown in Fig. 1(b) consists of the burner (axial swirler, fuel injectors and mixing duct) and 300 mm of the combustion chamber and is spatially discretized with 72.8 million tetrahedral elements. The grid is shown in Fig. 1(c) and refined in the injectors (80 µm), fuel jets (80 µm), swirler (250 µm), flame region (250 µm), mixing duct (300 µm), and combustor walls (400 µm). Inlet and outlet boundaries are modeled with the Navier-Stokes characteristic boundary conditions (NSCBC) imposing mass flow rates and ambient pressure, respectively. An analytically reduced chemistry (ARC) mechanism (20 species, 166 reactions) for methane-hydrogen-air combustion is employed. Turbulence-combustion interaction is modeled with the dynamic thickened flame (DTFLES) model, where the flame is thickened based on a flame

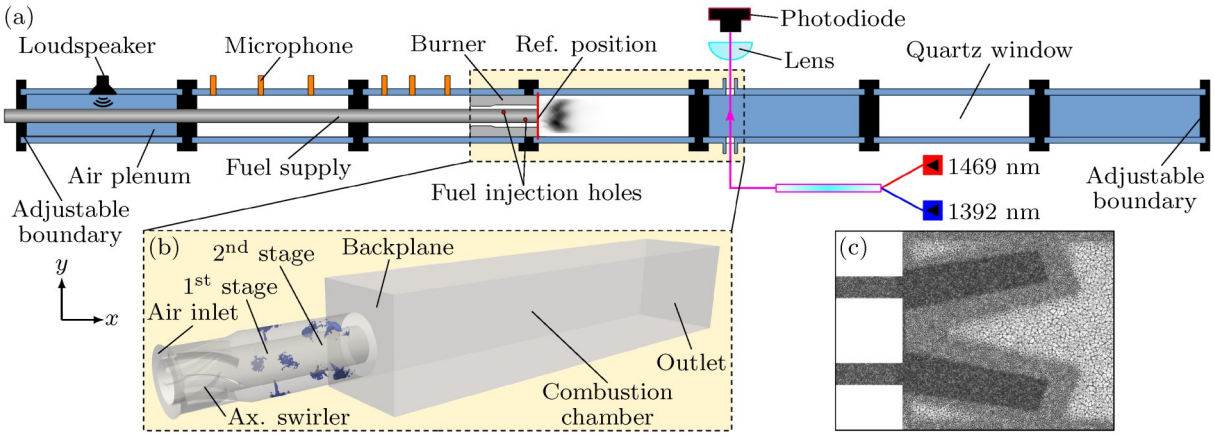

Figure 1: (a) Schematic of the swirl combustor equipped with laser spectroscopy and microphone arrays for entropy wave measurement. (b) Computational domain investigated in LES with instantaneous iso-contours of hydrogen (colored in blue) to visualize the fuel injection. (c) Grid refinement in the flame region. Adopted from [4].

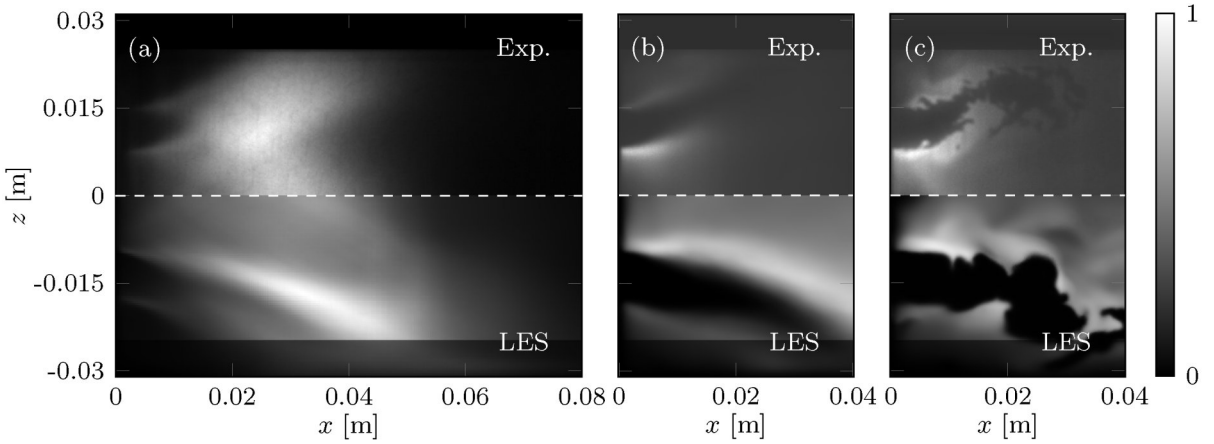

Figure 2: Comparison of flame shape between experiment (top) and LES (bottom), with (a) line-of-sight integration of time-averaged heat release rate, (b) time-averaged OH* mass fraction, and (c) instantaneous OH* mass fraction. Adopted from [4].

sensor. The compressible LES resolves all physical mechanisms needed to describe the generation and transport of entropy waves. There are three main challenges in this project that required the use of a supercomputer. First, the data-driven system identification methodology – which estimates the entropy transfer functions from LES data – requires relatively long time series (in the order of 100 ms) to accurately predict the respective response over a wide range of frequencies. Second, the enrichment of the fuel with hydrogen introduces the numerical difficulty that the higher laminar flame speed of hydrogen (compared to methane) reduces the flame thickness and thus requires fine spatial discretization in the flame region (i.e., a significant increase in the number of required mesh cells). Third, due to the partially premixing (fuel and air are mixed shortly before combustion), very high velocities (unity Mach) occur in the present burner due to tiny fuel injection holes (8 mm), which requires fine temporal discretization. A first comparison of time-averaged and instantaneous flame images from LES and optical measurements, respectively, is shown in Fig. 2. The mean fields in LES are time-averaged over a duration of 35 ms. Although OH* (hydroxyl radical) light intensity is not a quantitative indicator for heat release rate in partially premixed flames, it may serve as a qualitative measure to compare the shape and characteristics of the flame from experiment and LES. As shown in Fig. 2(a), the high volume fraction of hydrogen in the fuel causes a short flame with dominant M-shape that exhibits pronounced reactions in the inner and outer shear layer, which is well captured in LES. Furthermore, time-averaged and instantaneous OH-PLIF images (experiment) are compared with OH mass fractions (LES), see Fig. 2(b) and 2(c), respectively. Both flames show an accumulation of OH in the inner shear layer. As noticeable in Fig. 2(a) and 2(b), the flame length is overestimated compared to experiments, which may be partially attributed to the simplifications of the thermal boundary modeling and associated inaccuracies in the absolute wall temperature and heat distribution. Nevertheless, the agreement between experiment and LES is considered acceptable for the subsequent entropy response analysis.

Ongoing Research / Outlook

With the allocation of computing time through the Gauss Centre for Supercomputing e.V., we are able to perform several large eddy simulations to generate time series data. This data is then post-processed using state-of-the-art system identification techniques, such as the finite impulse response filter, to identify the entropy response of partially premixed flames. First results are presented in [4]. More information about the project can be found under [5].

References and Links

[1] A. J. Eder et al., (2023). J. Eng. Gas Turbines Power, 145(11), 111001. doi: 10.1115/1.4063283.
[2] Polifke, W. (2014). Ann. Nucl. Energy, 67C, 109-128. doi: 10.1016/j.anucene.2013.10.037.
[3] www.cerfacs.fr/avbp7x/index.php
[4] A. J. Eder et al., (2024) Proc. Combust. Inst., *accepted*.
[5] https://www.epc.ed.tum.de/tfd/forschung/noisi-combustion-noise-and-dynamics-of-partially-premixed-flames/

Earth, Climate and Environmental Sciences

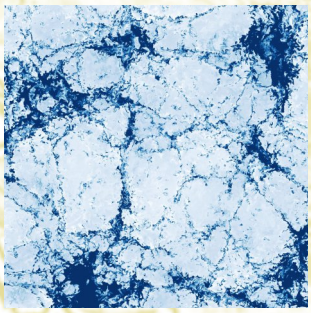

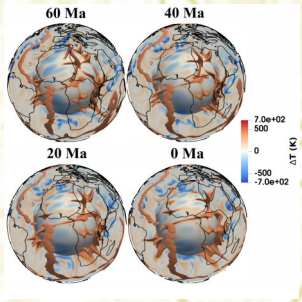

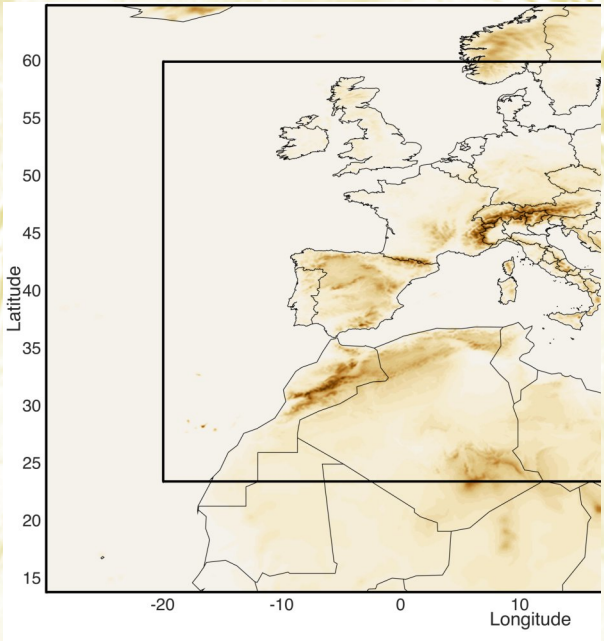

Resolving Cloud-Turbulence Interactions in the Climate System

Research Institution
University of Hamburg

Principal Investigator
Juan Pedro Mellado

Researcher
Raphael Pistor

Project Partner
–

Funding
DFG EXC CLICCS

SuperMUC Project ID: pn49de (Gauss Large-Scale project)

Introduction

Stratocumuli are low level clouds at the top of the atmospheric boundary layer, at altitudes of about 1 km and with thicknesses of about 100 m. They are key elements of the climate system. On the one hand, they extend over thousands of kilometers in the eastern boundaries of the subtropical oceans, e.g., off the coasts of California, Peru, and Namibia, favored by the temperature contrast between the cold upwelling water and the warm subsiding air. On the other hand, they reflect more incoming solar radiation (higher albedo) than the underlying surface of the ocean, while they emit similarly in the long-wave range. This combination of large coverage and net cooling effect substantially affects the Earth's radiative budget. How will stratocumuli change with climate warming? Despite continuous efforts, we still do not know [1].

A major difficulty is to understand how stratocumuli break and how this depends on the environmental conditions. Stratocumuli are embedded in the marine boundary layer and therefore depend on the sea surface temperature and the free tropospheric conditions. At the same time, the cloud-top region plays a crucial role. It is here where turbulent convection starts, which helps to maintain the cloud, but also where the cloud mixes with the warm air above it, contributing to cloud break-up The cloud-top region is very thin and experiences rapid changes; for instance, the temperature can vary 10 K within only 10 m. This implies that simulating cloud-turbulence interaction needs meter-scale resolution or less [2]. Simulation techniques such as direct numerical simulation (DNS) now reach this resolution [3].

This LRZ project focuses on the cloud-top region and on one cloud-turbulence interaction: droplet sedimentation. Because of the gravitational force, cloud droplets tend to fall. Their typical fall speed in stratocumuli is on the order of millimeters per second, much smaller than the turbulent velocity fluctuations and comparable to the rate at which the cloud top, in the mean, moves up and down. Our previous work suggested that insufficient resolution and numerical artifacts mask the effect of sedimentation in stratocumuli, which challenges the current consensus on the role of these clouds in the climate system [4,5]. This work, however, neglected the finite size of the cloud and the vertical transport across the boundary layer, which connects the sea surface with the cloud-top region. In this project, we have addressed this gap and considered the whole boundary layer, as shown in Figure 1.

Results and Methods

There are 3 major findings so far. First, we have found that the effect of sedimentation is of order one, as shown in Figure 2, which confirms the hypothesis that we need 1 m resolution or less to resolve cloud-turbulence interaction in stratocumuli. The second main result is that, in contrast to previous studies, the turbulent entrainment intensifies, although the mean entrainment velocity diminishes. The reduction in mean entrainment velocity is caused by the sedimentation flux, which opposes the turbulent flux. This is important for parametrizations in global climate models. Last but not least, simulations show a tendency towards Reynolds number similarity in cloud-relevant properties such as the integrated water content, shown in Figure 2. This allows us to extrapolate simulation results to atmospheric conditions more reliably.

To obtain these results, we have performed DNS of the marine boundary layer in domains of 6 km and with environmental conditions corresponding to the DYCOMS-II field campaign. We varied two parameters, namely, the sedimentation strength and the Reynolds number. The Reynolds number is equivalent to the grid spacing, and we have achieved resolutions of 1 m. The grid size in this last case spans 6,144 points in each horizontal direction and 1,536 points vertically, making these simulations the largest worldwide for studying this specific problem. The computational time used in all

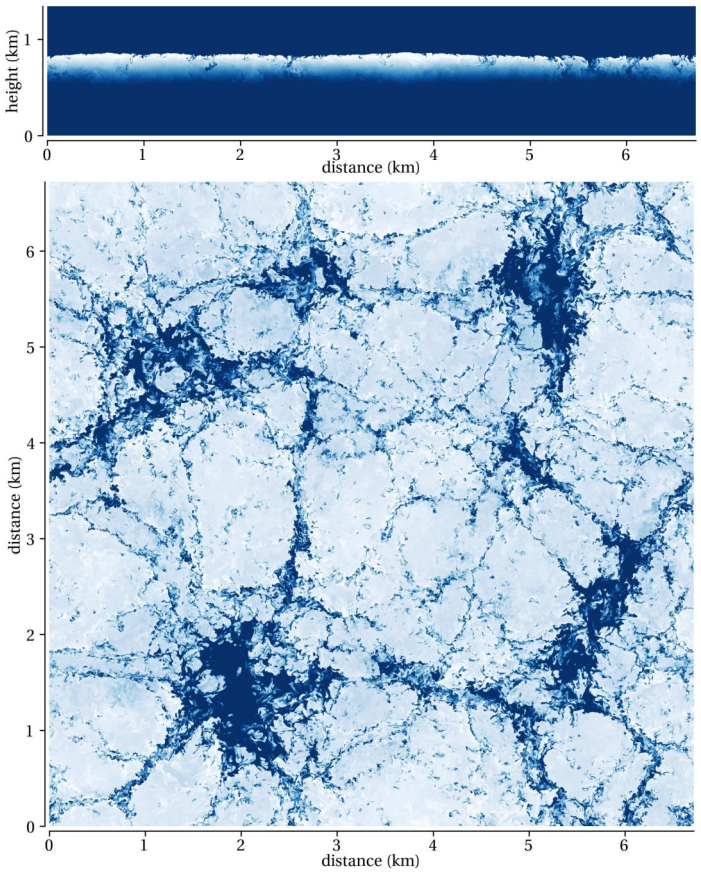

Figure 1: Vertical cross-section (top) and horizontal cross-section through the cloud (bottom) of the liquid water content: white, cloud; blue, dry air.

these simulations is 45 million core hours in SuperMUC-NG, 30 million corresponding to the high-Reynolds number simulation.

We rely on DNS because it solves the Navier-Stokes equations directly, without turbulence models, which removes the uncertainty associated with these models. For this problem, we use the anelastic approximation to the Navier-Stokes equations to remove the acoustic waves, which is a good approximation because of the small air velocities [4]. Grid sensitivity studies allow us to obtain grid independent results. Reynolds-number sensitivity studies (considering 4 m, 2 m and 1 m grid spacing, see Figure 2) allow us to estimate the role of limited scale separation in the simulations compared to atmospheric conditions.

The numerical algorithm is based on sixth-order compact finite differences to discretize the spatial derivatives, and low-storage, fifth-order Runge-Kutta schemes for the temporal integration. The pressure-Poisson equation is solved by a Fourier decomposition in the two horizontal directions. The code is written in Fortran 2008 and uses MPI for the domain decomposition in the two horizontal directions. The code is publicly available [6].

Ongoing Research / Outlook

The capability for such high-resolution studies relies exclusively on supercomputers like SuperMUC-NG, encompassing both computational and data aspects.

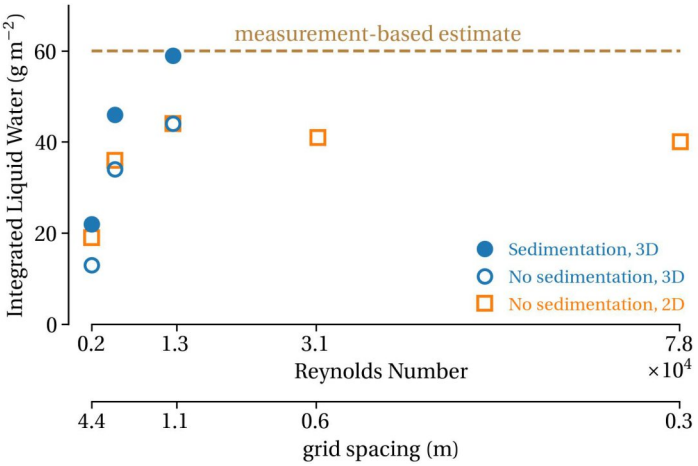

Figure 2: Dependence of the water content on sedimentation strength and Reynolds number (grid spacing). Measurement-based estimate corresponds to the DYCOMS-II field campaign. 2D simulations were used to explore Reynolds number similarity.

Each 3D field occupies 216 GiB of disk space, and, to describe the system, we need 5 fields along with their temporal evolution. The post-processing also needs an easy access to this large amount of data to calculate averages, spectra or visualizations.

The current challenge is to accelerate the code. As illustrated in Figure 2, we want to increase the Reynolds number to confirm the tendency towards Reynolds-number similarity, which would allow us to extrapolate results from simulations to atmospheric conditions more reliably. Ideally, this means reducing the grid spacing by a factor 2, which implies an increase in computational resources by one order of magnitude.

Accelerating the code will also allow us to study the effect of different environmental conditions and thereby how stratocumuli will change in future climates, as the sea surface temperature and the free troposphere in the subtropics change. Part of this work is already ongoing as a continuation of this project, in particular, testing the hypothesis that insufficient resolution in the cloud-top region has masked cloud-adjustment mechanisms associated with entrainment and sedimentation.

To accelerate the code, we have considered serial and parallel optimization. Our reformulation of the pressure-Poisson solver reduces the number of linear systems to solve and increases stability with non-uniform grids. We have also reformulated the I/O to single precision to save disk space. Currently, we are considering the MPI communication in single precision as well. We also want to reformulate the calculation of the evolution equations to increase locality in memory.

References and Links

[1] L. Nuijens, P. Siebesma, Curr Clim Change Rep, 5 (2019) 80–94.
[2] J.P. Mellado, Annu Rev Fluid Mech, 41 (2017) 145–169.
[3] J.P. Mellado et al., J Adv Model Earth Syst 10 (2018) 1421–1438.
[4] A. Lozar, J.P. Mellado, J Atmos Sci 74 (2017) 751-765.
[5] B. Schulz, J.P. Mellado, J Adv Model Earth Syst 11 (2019) 1830-1846.
[6] https://github.com/turbulencia/tlab

So2Sat - 10^{16} Bytes from Social Media to Earth Observation Satellites

Research Institution
[1]Data Science in Earth Observation, Technical University of Munich

Principal Investigator
Xiaoxiang Zhu[1]

Researcher
Yuanyuan Wang[1], Dawood Wasif[1], Yilei Shi[2]

Project Partner
[2]Chair of Remote Sensing Technology, Technical University of Munich

Funding
ERC

SuperMUC Project ID: pr53ya (Gauss Large-Scale project)

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 *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 is 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

Since we employed machine learning for 3D building height estimation, one important question is the uncertainty of the machine learning model predictions. However, uncertainty quantification becomes unclear in black-box like machine learning models. In order to benchmark the performance of uncertainty quantification from machine learning models, a dataset not only containing the ground truth of the prediction, but also the ground truth of the variance of the prediction shall be curated. In such context, we developed a synthetic building segmentation aerial imagery dataset [2] that is first of its kind and serves the following two-fold purposes: 1) Provides realistic "error-free" synthetic dataset that not only contains perfect labels (i.e., the building segmentation masks) but also mirrors the complexity of real-world scenarios; 2) With its variants with various noise perturbations (both in image and labels) enables a systematic comparison and benchmarking of uncertainty quantification methods by utilizing computed pseudo ground truth uncertainty. The simulation concept is shown in Figure 1, where we use 3D building models as input and simulate noisy 2D aerial images together with other attributes like building mask. Large amounts of image rendering was done on SuperMUC-NG. We used the software Blender and 3D building models of Berlin as the input data. Blender were containerized using charliecloud for smooth execution on SuperMUC-NG. Since the rendering of individual images is independent, we employ task parallelization, reducing the communication among nodes. A job distribution script were implemented to distribute the task to individual nodes in a single slurm job requesting multiple nodes. This also reduces the total number of slurm jobs we submit. We simulated variations of our synthetic dataset by changing 3D viewpoints (or perspectives) while generating the rasterized images. From an implementation perspective, this is achieved by varying the camera angle in Blender by a value sampled from a Gaussian distribution with fixed degrees of standard deviation. We used fixed standard deviation values of 1, 2, 4, and 8 degrees and for each distribution, we sample and render 50 variations of each of the 400 patches of our dataset in Blender to simulate diverse viewpoints for capturing the 3D environment. In total 80,000 large images were synthesized. Those images forms a dense sample of the image data noise. An example of such variations is shown in Figure 2. The visual different is subtle, as the angle variation has a maximum of a standard deviation of 8 degree.

Figure 1: The concept of the aerial image simulation. We use 3D building models (left) as input, and simulate 2D aerial image, associated with many other attribute like building mask, building height, which can be used for uncertainty quantification of many downstream tasks.

Figure 2: Image rendered from 3D building model with different camera viewing angles. 2 samples from the batch with standard deviation of 8°.

Ongoing Research / Outlook

A main technical challenge is the very small number of synthetic aperture radar images with poor orbit configurations for many cities, causing low accuracy of height reconstruction. Therefore, we have being looking for alternative data and method for height reconstruction. During the project, we have acquired numerous high resolution optical images apart from the synthetic aperture radar images, which could be potentially employed for height reconstruction. Therefore, we have been developing an algorithm for height reconstruction from optical image. Different from the well-established stereo matching method, we chose height reconstruction from single optical image. This is because stereo matching has strong requirements on camera position (orbit) and orientation, whereas single image height reconstruction eases such strong constraints on orbit and sensor orientation. Single image height reconstruction is an emerging field, especially since the last few years thanks to the strong predictive power of large machine learning models. Since we want to predict a single height from a single image pixel, it can have infinitely many solutions. Therefore, it requires well-designed networks for enhanced representations to improve performance. Moreover, the height distribution is long-tailed with the low-height pixels, e.g., large area of flat ground, as the head. If directly trained using such data, the networks are usually biased and tend to underestimate building heights. To solve the problems, we propose HTC-DC Net [3] following the classification-regression paradigm, with the head-tail cut (HTC) and the distribution-based constraints (DCs) as the main contributions.

References and Links

[1] www.so2sat.eu
[2] D. Wasif, Y. Wang, M. Shahzad, R. Triebel, and X. X. Zhu, “Towards a Benchmark EO Semantic Segmentation Dataset for Uncertainty Quantification,” presented at the IEEE International Symposium on Geoscience and Remote Sensing, 2023.
[3] S. Chen, Y. Shi, Z. Xiong, and X. X. Zhu, “HTC-DC Net: Monocular Height Estimation from Single Remote Sensing Images.” 2023. http://arxiv.org/abs/2309.16486

Computational Hydrometeorology – with Advanced Processing Tools to Enhanced Realism (CHAPTER)

Research Institutions
[1]Ludwig-Maximilians-Universität München

Principal Investigator
Dieter Kranzlmüller[1]

Researcher
Antonio Parodi[2], Stephan Hachinger[3], Maximilian Höb[1,3], Martina Lagasio[2], Anna Napoli[2,4], Agostino N. Meroni[2,5], Nazario Tartaglione[2,6]

Project Partner
[2]CIMA Research Foundation, Savona, [3]Leibniz Supercomputing Centre
[4]Università di Trento, [5]Università degli Studi di Milano-Bicocca,
[6]Istituto Superiore per la Protezione e Ricerca Ambientale (ISPRA), Roma

Funding
EU Horizon 2020 GAs # 825532 (LEXIS – Large-Scale Execution for Industry and Society),
777533 (PROCESS – PROviding Computing solutions for ExaScale ChallengeS)

SuperMUC Project ID: pn29co

Introduction

Severe weather events are responsible for hundreds of deaths and millions of Euros of damage every year worldwide, with a particularly large impact on mountain and coast regions. With climate change progressing at a larger speed then ever before, more extreme events are expected. Mountain areas, coastal zones, floodplains, and areas alongside large rivers are examples for zones particularly vulnerable to the consequences.

Hydrometeorological simulations help predict hazards to local populations and ecosystems in this context, such as flash floods. In recent years, the field has vastly advanced, enabling forecasts in mountain and coastal regions which are traditionally hard to simulate with their orographic details. Improvements are due to better simulation methods, due to large-scale assimilation of weather station data, due to higher resolution and – last not least – due to better initial and boundary conditions. The latter are provided by global climate models, e.g. of the European Centre for Medium-Range Weather Forecast (ECMWF) or the American National Oceanic and Atmospheric Administration (NOAA).

Leaps forward in computational (hydro-)meteorology and climate science will not only be enabled by progress in the simulation techniques themselves, but also critically depend on IT infrastructure. Novel frameworks and platforms, as developed in the EU projects PROCESS and LEXIS, automatize simulation workflows and data handling, and help data analysis and visualisation. The LEXIS platform enables data-driven workflows across heterogeneous and geographically distributed HPC and Cloud-Computing resources, contributing to HPC-Cloud-Big Data convergence. PROCESS has focussed on preparing tools for the upcoming epoch of Exascale computing, in particular an interactive execution environment and scalable data analysis tools. Both projects, as well as new services at LRZ, have been aiming at automating data management within workflows, taking into account aspects of FAIR [1] research data management (RDM).

In this setting, the CHAPTER project has had the goal to compute a unique high-resolution model reflecting hydrometeorological conditions in East, West, Central and Southern Europe (to Northern Africa) over the last 40 years. The resulting dataset, to be made public, can be used to physically explain extreme events and to drive a wealth of application models simulating river flood waves, agricultural yields or wind power. Finally, the model will also facilitate a deep and uniform scrutiny of deficits and achievements of current modelling techniques, and of the consequences of climate change. To achieve all this, we have performed a dynamical downscaling of the quite recent ERA5 data set [2] of ECMWF (providing us initial and boundary conditions at a resolution of 31 km) onto a fine spatio-temporal grid (3 km, 1 h) which resolves clouds. For the computations, we have used the Numerical Weather Prediction (NWP) model WRF-ARW [3], proven and widely accepted within the community. We thus have simulated accurate surface and troposphere meteorological fields, which we now provide from 1981 to 2022 for our core domain d02 (cf. Figure 1).

Throughout the project, we have explored novel IT as described above. Besides downscaling (Subproject/

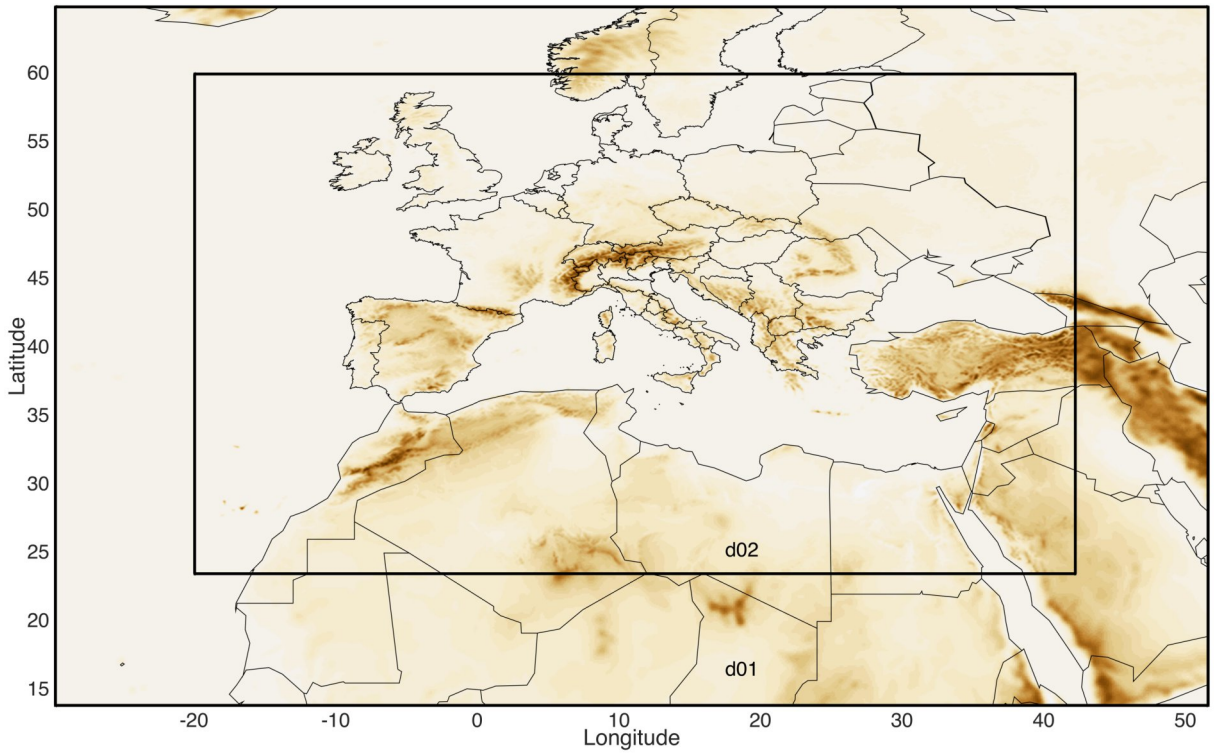

Figure 1: Nested (two-)domain setup for our project, foll-owing WRF domain design rules (d02: 3 km resolution).

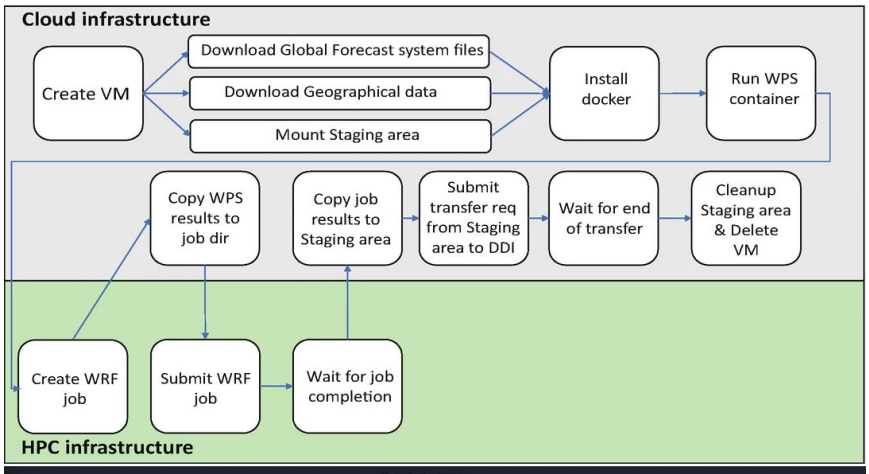

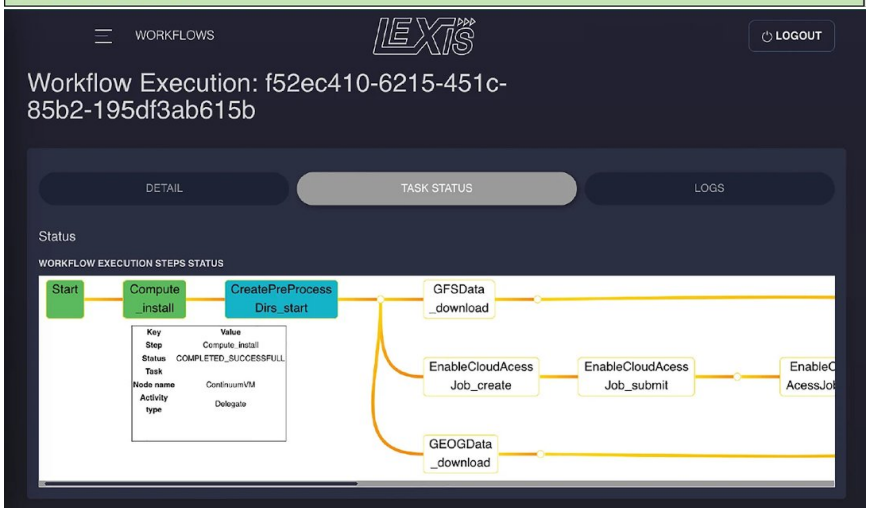

Figure 2: Typical WRF workflow (upper panel), here run (first steps, lower panel) on LEXIS. Fig. adapted from [4].

SP 2), we have focused on testing WRF speed with AVX (SP 1), and on analysis and FAIR management of data (SP 3).

Results and Methods

Exploration of vectorisation speed-ups (SP 1)
Once our basic WRF setup (for Subprojects 1 & 2) had been devised, we ran small test cases to explore whether WRF speed (in our case) profits from compilation with the AVX options of the WRF build system. The results will be reported; unfortunately we found no significant speed-up.

Scientific production: a 40-year weather data suite for Europe at 3 km / 1 h resolution (SP2 core); first analysis (part of SP 3)
CHAPTER has produced a competitive cloud-permitting atmospheric reanalysis-based dataset for central Europe and the Mediterranean. To this end, WRF-ARW has been run for downscaling on daily basis with a spin-up period of 6 hours, from 18 to 00 UTC from the day before. All relevant processes are simulated by WRF-ARW at a resolution of 9 km (d01, cf. Fig. 1) and 3 km (d02), respectively, in this typical nested-domain setup.

The WRF model physical setup is derived to a large extent from studies of Pieri et al. [4] and von Hardenberg et al. [5] at cloud-resolving grid spacing, which was found very well suited for severe weather phenomena. The schemes used for planetary boundary layer turbulence closure, radiation, convection, and for the soil model have been carefully chosen, as will be described in upcoming publications. No cumulus scheme was activated in the innermost domain at 3 km grid spacing, because the grid spacing allows us to resolve the convection dynamics. The Thompson microphysics scheme [6] has been adopted.

The dataset produced by CHAPTER comprises data in the PB range, which we have been extensively validating, with the aim of public release of interesting parts in mind. On the partners' side, the validated data will be used to understand extreme-weather events and to model flash floods or forest fires, atmospheric dispersion or air quality, just to name a few applications.

The CHAPTER data will show its strength in pursuing long-term aims, such as:

- understanding the mechanisms of past natural disasters and phenomena,
- discovering climate-change signals in the phenomena investigated, and
- understanding how current models in NWP, hydrology, atmospheric transport, and similar fields have to be set up in order to reproduce observations and thus to produce predictions useful for administrative and political decision making.

WRF workflow control (part of SP2)
During the project, we explored options for workflow and data flow control. For the SP2 runs, we used a traditional script-controlled workflow as LEXIS-based workflow descriptions and system adaptors still underwent changes at the beginning. However, we extensively explored the possibilities of LEXIS orchestration with WRF workflows focused on different target areas/topics [7]. These include automatised system control from the invocation of WPS (WRF preprocessor system) through the actual WRF run (Figure 3) to postprocessing or application model runs [7].

LRZ Data Science Archive (DSA) & FAIR RDM (part of SP3)
CHAPTER has been probably the first PB-range test case for the LRZ Data Science Archive (DSA). We plan to publish output compilations – according to the FAIR principles [1] – via a new LRZ FAIR data portal demonstrator. Data access will likely rely on GLOBUS sharing.

Ongoing Research / Outlook

CHAPTER has not only produced a dataset which will be a great basis for hydrometeorological studies and applications. It has also shown that science can heavily profit from innovative IT services for workflow control and data handling. With publication (of data and papers) ongoing, the fruits of this work just begin to be reaped.

References and Links

[1] Wilkinson, M.D., et al.; Sci Data 3, 160018 (2016)
[2] Hersbach, H., et al.; Q. J. R. Meteorol. Soc. 146(730), 1999 (2020)
[3] Skamarock, W.C., et al.; Tech. Rep. NCAR/TN-556+STR (2019)
[4] Pieri, A.B., et al.; J. Hydrometeorol., 16(4), 1872 (2015)
[5] von Hardenberg, J., et al; In: Engineering Geology for Society and Territory-Volume 1, pp. 209-213. Springer (2015)
[6] Thomspson G., et al.; Mon. Weather Rev. 136(12), 5095 (2008)
[7] Hachinger, S., et al.; in: Curry, E., et al. (eds): Technologies and Applications for Big Data Value. Springer (2022)

SeisSol: Earthquake Simulation with Poroelasticity

Research Institution
[1]Technical University of Munich

Principal Investigator
Michael Bader[1]

Researcher
Sebastian Wolf[1], Alice-Agnes Gabriel[2], Martin Galis[3]

Project Partner
[2]UC San Diego
[3]Comenius University in Bratislava

Funding
KONWIHR, ChEESE (EuroHC JU, grant No 101093038)

SuperMUC Project ID: pr83no

Introduction

SeisSol [1] is an HPC software for the simulation of earthquake source mechanisms and resulting ground motion. It employs the Discontinuous Galerkin method with ADER time-stepping (ADER-DG) to solve the elastic wave equation in various types of elastic media, including isotropic and anisotropic and viscoelastic materials. It also features coupling to acoustic materials for the simulation of earthquake-tsunami scenarios. For earthquake sourcing, SeisSol simulates the dynamic coseismic breaking of pre-existing faults due to frictional failure as a physics-based model ("dynamic rupture simulation") to excite seismic waves.

For the GCS project pr83no, we enhanced SeisSol's capabilities to simulate small induced earthquakes, as they may occur during the preparation and operation of Geothermal energy production plants. For the simulation of such scenarios, the interaction of the solid phase (rock) and the injected fluid (water phase) has to be accounted for. We therefore implemented and optimized the simulation of poroelastic materials [2] in SeisSol, we refactored and extended SeisSol's routines for dynamic rupture simulations, and we performed first benchmark simulations combining these new features.

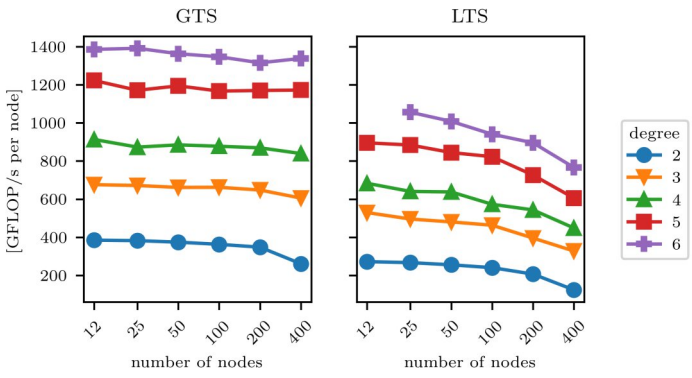

Figure 1: Strong scaling results for the poroelastic version of SeisSol comparing different convergence orders and time stepping schemes.

Results and Methods

Poroelasticity
Poroelastic media consist of an elastic porous skeleton, which is filled with a viscous fluid. The deformation of the skeleton drives fluid flow and vice versa. The interaction of both phases renders the governing equation stiff and requires a more complicated solution procedure for the ADER-DG scheme – using a locally implicit space-time-predictor. Here, a lot of medium-sized linear systems (with ~1,000 unknowns) have to be solved. Since all these linear systems share the same sparsity pattern, a specialized solution algorithm has been developed. This algorithm takes the block-wise sparsity pattern into account and reduces the required floating point operations by up to a factor of 25 compared to straightforward LU decomposition. The code generator YATeTo [3] is able to map the solution algorithm to optimized implementations of matrix-matrix multiplication (GEMM) and employs hardware-aware backends for optimal performance. Figure 1 shows a strong scaling study for a wave-propagation scenario in poroelastic media. We compare global (GTS) and local time stepping (GTS) for different convergence orders. Scenarios computed with higher convergence order achieve a better performance due to the higher computational intensity. Also, global time stepping achieves a higher performance compared to local time stepping due to the simpler update scheme. Despite the reduced performance, local time stepping reduces the time to solution significantly. The parallel efficiency plot highlights the good scaling behaviour of SeisSol. With a node-level performance of up to 1.4 TFLOP/s per node, SeisSol reaches approximately 40% of the available peak performance of SuperMUC-NG, when using polynomials up to degree 6 as basis functions.

Dynamic Rupture Simulation
The dynamic rupture source model combines frictional failure along fault surfaces with the radiation of seismic waves. Both aspects are linked and drive each other. This source mechanism allows researchers to gain insight into the hidden nucleation processes of earthquakes.

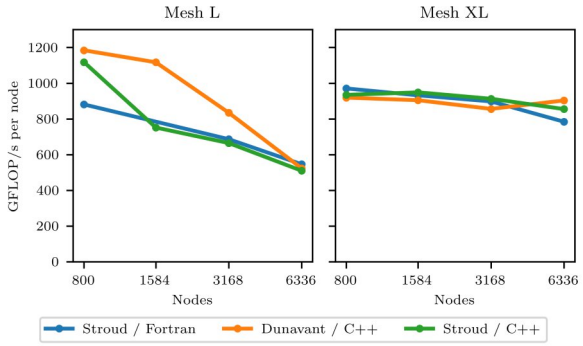

Figure 2: Strong scaling results for a complex scenario with two different mesh resolutions up to the full SuperMUC-NG.

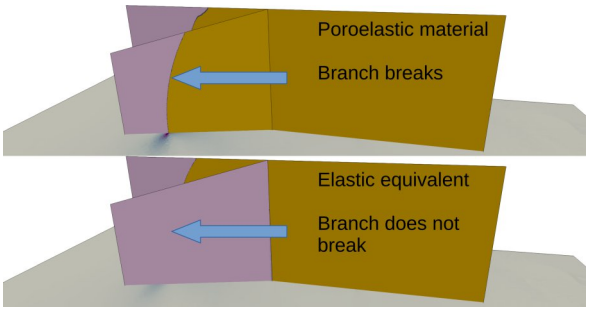

Figure 3: Comparison of the poroelastic (top) and elastic (bottom) version of the same fault branching scenario. Purple colours indicate parts of the fault that are still intact, yellow indicates where the fault is sliding.

SeisSol's early roots date back almost 20 years. Over these years, it has almost entirely been re-implemented from Fortran to C++ – with the dynamic rupture implementation being one of the last "legacy" parts. In particular to enable GPU support for upcoming machines, like SuperMUC-NG Phase 2, the respective routines have now been redesigned, as well.

SeisSol uses OpenMP (on CPU machines) or SYCL/CUDA (on GPU architectures) for intra-node parallelism. SYCL has been chosen as it allows a performance-portable implementation without code duplication for AMD, Intel and NVidia GPU architectures. On CPUs, padding of the vector of degrees of freedom enables the efficient use of AVX-512 vector instructions. The refactoring thus yields a better performance in comparison to the legacy implementation also on CPU-based supercomputers.

As part of the redesign, we also re-examined the numerical treatment of dynamic rupture simulations. The latter requires the evaluation of a friction law on the prescribed fault surface, which is discretized by a triangular mesh. The friction law is evaluated at quadrature points following a quadrature rule on the triangular surface cells. In addition to the originally used Stroud rule, we implemented the Dunavant rule, which achieves the same accuracy with less integration points, thus reducing the computational effort. The figure shows a strong scaling study from 800 to 6,336 nodes of SuperMUC-NG comparing the legacy Fortran version of SeisSol and the refactored version with both quadrature rules. While on smaller scales, the improvements are more clearly visible and the C++ version outperforms the legacy Fortran version, the picture is more complicated at large scale. We observe large variations in performance between different runs of the same configuration, which we attribute to effects such as power capping or dynamic changes of the clock frequency.

Results

For earthquake simulation, the poroelastic material model has to be combined with a respective dynamic rupture source mechanism: the fluid pressure weakens the fault and facilitates rupture. After a careful verification of the combined model, a study on a branched fault has been conducted. The branch lies in a region, where the waves emerging from the main fault induce an increased fluid pressure, which weakens the fault and facilitates rupture on the branch. The best available elastic equivalent model does not capture the physics of the poroelastic model correctly, because the fluid pressure is not accounted for properly. Thus, the branch does not break in the elastic case. This scenario shows that proper treatment of poroelastic effects is required for earthquakes in fluid-rich Earth. SeisSol now enables users to conduct respective large-scale studies.

Ongoing Research / Outlook

SeisSol now supports the most common material models (acoustic, elastic, poroelastic, viscoelastic) for modeling seismic wave propagation in solid Earth. Many interesting research questions circle around coupling different material models, such as elastic-acoustic coupling for the simulation of earthquake-tsunami events or to simulate acoustic noise patterns of induced earthquake. Our long-term goal is to expand SeisSol's support for elastic-acoustic interfaces towards coupling of all available material models with each other, while also supporting all exascale-relevant compute architectures.

References and Links

[1] https://github.com/SeisSol/SeisSol
[2] S. Wolf et al., 'An efficient ADER-DG local time stepping scheme for 3D HPC simulation of seismic waves in poroelastic media', Journal of Computational Physics, vol. 455, pp. 1–29, Apr. 2022, DOI: 10.1016/j.jcp.2021.110886.
[3] C. Uphoff and M. Bader, 'Yet Another Tensor Toolbox for Discontinuous Galerkin Methods and Other Applications', ACM Trans. Math. Softw., vol. 46, no. 4, p. 34:1-34:40, Oct. 2020, DOI: 10.1145/3406835.

The Influence of a Thin, Weak Asthenosphere in High Resolution Mantle Circulation Models

Research Institution
Geophyics Section, Department for Earth and Environmental Sciences, Ludwig-Maximilians-Universität München

Principal Investigator
Hans-Peter Bunge, Bernhard Schuberth

Researcher
Hamish Brown, Ingo L. Stotz

Project Partner
–

Funding
DFG research training group UPLIFT 440512084 (GRK 2698) and DFG STO1271/2-1

SuperMUC Project ID: pr48ca

Introduction

The idea of a thin, weak layer beneath the Earth's rigid outer shell is over a century old, long predating the widespread acceptance of plate tectonics. This layer, termed the asthenosphere, was originally proposed in order to account for Earth's internal support of topographic loads such as mountain ranges (isostasy). Our view of this layer expanded following the plate tectonic revolution of the late 1960's, after which it came to be viewed as a lubricating layer to facilitate plate motion; passively flowing in response to the overlying plates. In the past two decades, it has become clear that the asthenosphere is a much more exciting and active part of the convecting mantle than previously imagined. Upwelling mantle plumes from deep within the Earth (found, for instance, beneath Hawaii) feed this layer with hot material. This creates high pressure regions within the asthenosphere, while sinking tectonic plates produce low pressure regions. The resulting pressure gradients can induce strong pressure-driven flow, in addition to the component of flow driven by the overlying plates. This fast asthenospheric flow can allow for the lateral transport of plume-sourced material over significant distances in relatively short geological timescales. As a result, surface uplift and volcanism may be induced in regions where we would not have previously expected it, which could even have implications for past and future changes in sea level. In addition, the flow can actually alter the motion of the overlying tectonic plates if pressure-driven flow velocities become high enough. The asthenosphere is therefore a crucial layer in the Earth system; providing a link between deep mantle dynamics and surface processes.

While the importance of the asthenosphere is now well understood, simulations which accurately include such a layer remain a computational challenge. Fluid dynamical simulations of past mantle convection have now become a ubiquitous tool in the field of global geodynamics. However, resolving the sharp viscosity contrasts and short flow wavelengths associated with the asthenosphere require a grid-point spacing of ~10 km at the very least. In a global model, this equates to ~670 million grid points. Solving the conservation equations for the velocity, temperature, and pressure at each grid point means solving a system of equations with several billion unknowns during each time step. A simulation covering a significant time span of geological history (say, 400 million years) may require ~10,000 such time steps. With these numbers in mind, the necessity of modern HPC systems in solving these problems becomes abundantly clear. Few geodynamics groups around the world have access to vast computational resources such as those found at the LRZ. Therefore, as it stands, the vast majority of published models remain at a grid spacing of ~25 km or coarser. It is due to this that influence of a thin, low viscosity channel in global models of past mantle convection has so far remained elusive.

Results and Methods

During this project, we have generated high-resolution simulations of past mantle convection which feature a low viscosity asthenospheric channel in the uppermost mantle. Simulations were carried out using the parallel finite element mantle convection code TERRA. Tectonic plate histories were assimilated at the surface of the simulation as boundary conditions on surface velocity, allowing us to study the resulting mantle flow. An Earth-like initial temperature field was first generated by simulating 1 billion years of convection with free-slip velocity boundary conditions at the surface. Following this, tectonic velocities were assimilated from 400 Ma (million years ago) to the present day. The generation of the initial state alone requires ~300,000 core hours on the SuperMUC-NG, while the remainder of each simulation requires an additional ~250,000 core hours. Even when parallelised amongst 2,048 cores, this is equivalent to 11 days of compute time for a single simulation. The simulation has a resolution of ~10 km, allowing us to create a thin (~100 km), low viscosity

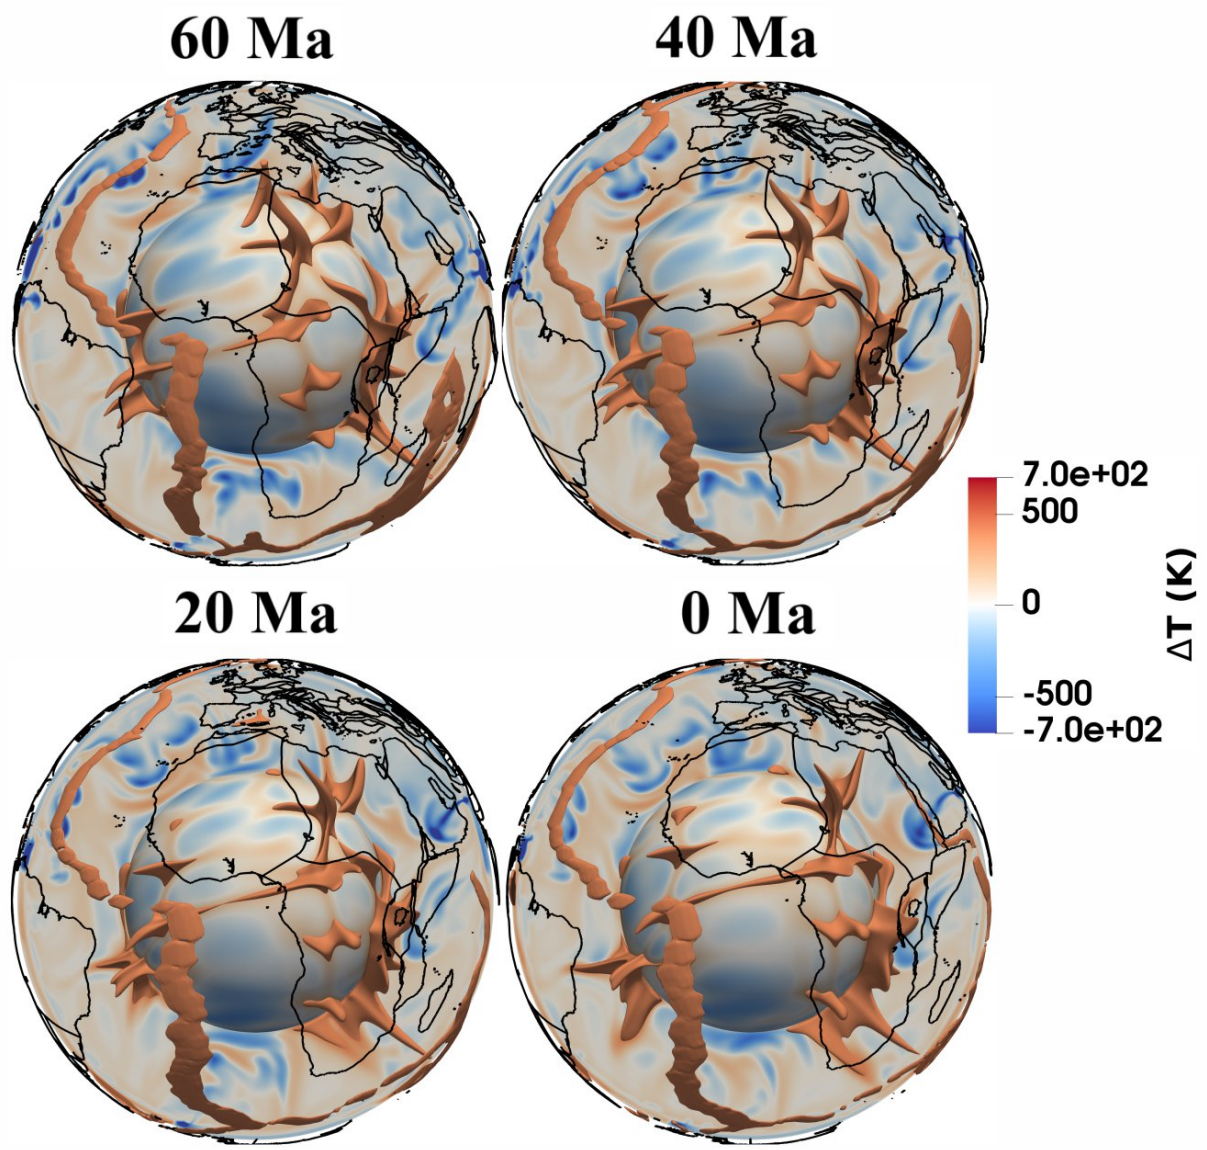

Figure 1: Evolution of mantle temperature anomalies beneath Africa from 60 million years ago (Ma) to the present day (0 Ma). Upwelling mantle plumes are shown by 3-D contours of temperature deviation at 400 K.

(~ $5 \cdot 10^{19}$ Pa.s) channel beneath a stiff outer shell, the tectonosphere. In addition to radial viscosity layering, viscosity was made temperature-dependent, meaning that cold, sinking tectonic plates can retain their strength while passing through the asthenosphere.

Figure 1 shows four transects of the African hemisphere, spanning from 60 Ma (shortly after the extinction of the dinosaurs at 66 Ma) to the present day (0 Ma). A contour of temperature deviation has been extracted at 400 Kelvin in order to highlight hot, upwelling mantle plumes. Old, sinking tectonic plates can be seen as cold (blue) regions in the background. For comparison, we also ran simulations at a more "standard" resolution of ~ 25 km grid-point spacing. Models such as this are far more limited in their ability to simulate asthenospheric dynamics, and instead can only feature a thicker (~ 200 km) and higher viscosity (~ $5 \cdot 10^{20}$ Pa.s) asthenosphere. A comparison between the two is shown in Figure 2. Temperature deviations within the asthenosphere of each model at 200 km depth are shown at the same time steps as Figure 1. The difference in the character of asthenospheric temperature anomalies is immediately clear. Note the long-wavelength character of the hot regions in the standard resolution model (left), as plume material cannot spread out effectively when it reaches the uppermost mantle. Meanwhile, the lower viscosity channel allowed by the high resolution model (right) allows for the fast transfer of plume material in response to pressure gradients and overlying plate velocities. Hot material is carried towards subduction zones (blue lines) where it builds up, an effect which is not reproduced in the standard resolution model. The differences between the asthenospheric flow patterns between these two models has broader implications for their prediction of surface uplift (dynamic topography) and volcanism. As such, high resolution models as the one shown in Figure 1 ought to become the new "standard"—something which can only become possible with computational resources such as the SUPERMUC-NG.

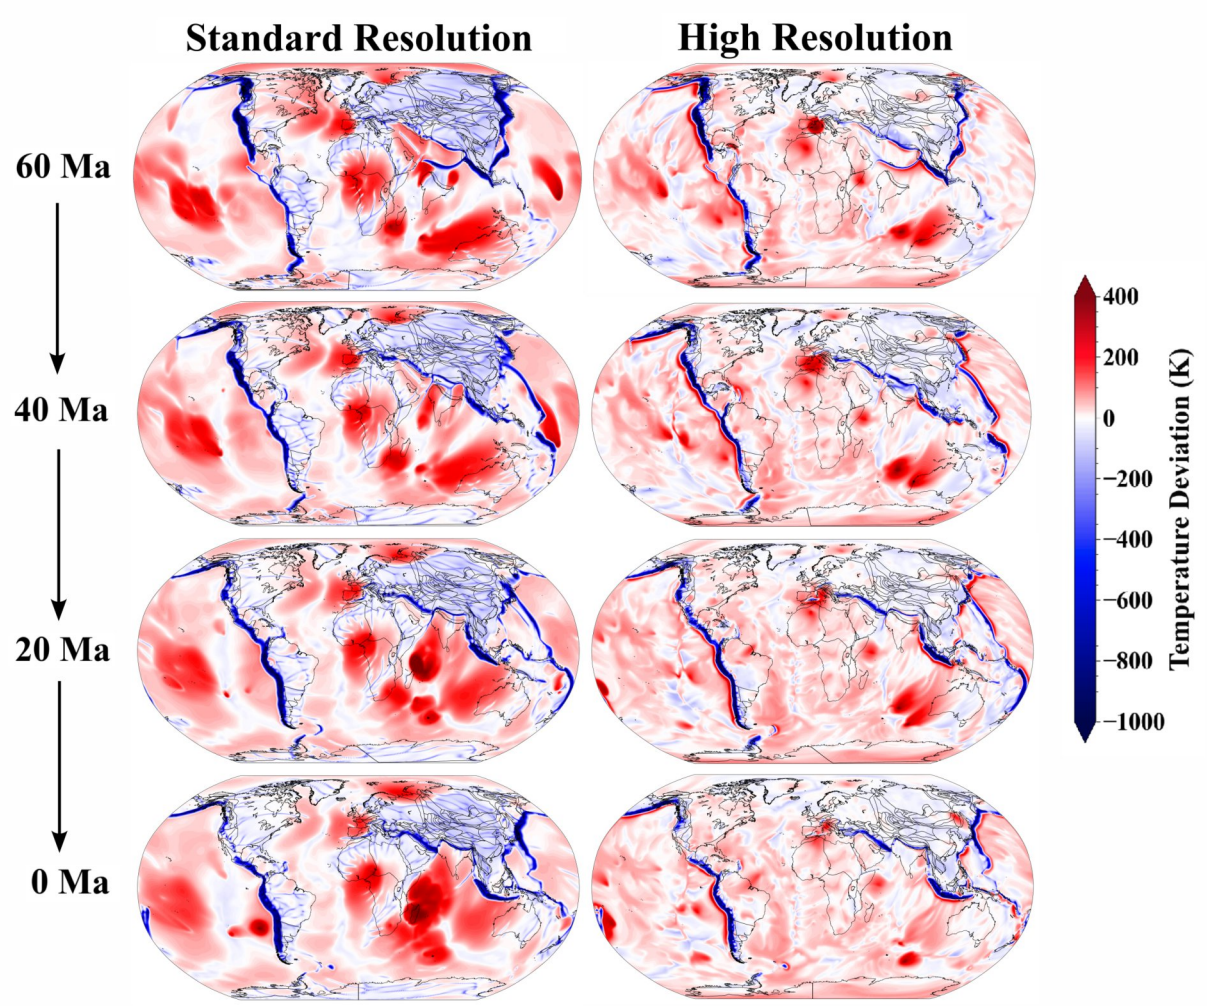

Figure 2: Comparison between temperature anomalies within the asthenosphere at 200 km depth of two mantle circulation models. Left column: a "standard" resolution Earth model; right column: a high resolution model made possible using the SUPERMUC-NG.

Ongoing Research / Outlook

While the results presented here have provided invaluable insight into asthenospheric dynamics, the grid spacing of ~10 km is still not fine enough to simulate viscosity structure of the real asthenosphere inferred from observations. For this, a grid point spacing of ~1–5 km would be required globally, leading to sparse linear systems with more than a trillion unknowns. This necessitates scalable codes which can take advantage of exascale-capable modern HPC systems. To this end, our TerraNeo project has focussed on the development of a new mantle convection code, capable of extreme-scale simulations. The project is a collaboration between FAU Erlangen-Nürnberg, LMU Munich, TU Munich and the Leibniz Supercomputing Center (LRZ). The TerraNeo application is built upon the HyTeG [1] (Hybrid Tetrahedral Grids) framework, an open-source C++ framework aimed at efficient and scalable multigrid solvers for Finite Element simulations. The extreme scalability and performance of HyTeG's matrix-free multigrid methods has already been demonstrated on the SuperMUC-NG. The solver was able to solve the Stokes equations with more than a trillion (10^{12}) unknowns in less than a minute when parallelised amongst 147,456 processes [2]. The TerraNeo application has now been equipped with the essential physical features, such as tectonic plate velocities and temperature-dependent viscosity, in order to generate the first ever simulations of global mantle circulation with ~1 km resolution within the coming year.

References and Links

[1] https://i10git.cs.fau.de/hyteg/hyteg
[2] N. Kohl and U. Rüde, "Textbook efficiency: massively parallel matrix-free multigrid for the Stokes system", SIAM Journal on Scientific Computing, Vol. 44, no. 2, C124–C155, 2022

Elementary Particle Physics

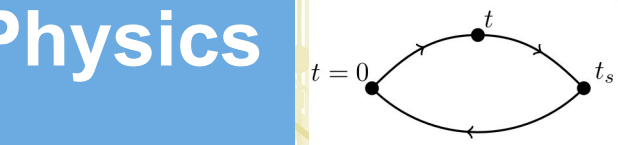

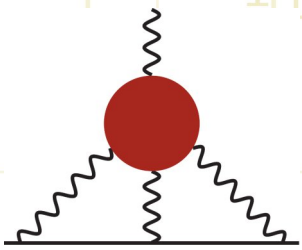

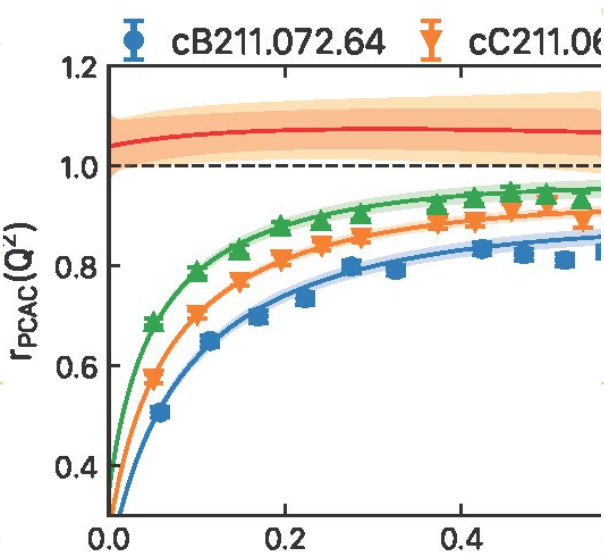

LPC $P_1^z/P_2^z = 4/2$ (200

LPC $P_1^z/P_2^z = 4/3$ (2005.14572)

Bernstein (2003.06063)

Hermite (2003.06063)

$h_1(P^+ = 2.27\text{Ge}$

$h_1(P^+ = 1.74\text{Ge}$

$h_1(P^+ = 1.74\text{Ge}$

$h_1(P^+ = 1.25\text{Ge}$

Impact of magnetic fields on dense strongly interacting matter

Research Institution
[1]University of Bielefeld

Principal Investigator
Bastian B. Brandt[1]

Researcher
Szabolcs Borsányi[2], Gergely Endrődi[1], Jana N. Günther[2], R. Kara[2], Adeilton Dean M. Valois[1]

Project Partner
[2]Bergische University, Wuppertal

Funding
"Netzwerke 2021"

SuperMUC Project ID: pn36ri

5

Introduction

The strong interactions as part of the Standard Model of particle physics are described by Quantum Chromodynamics (QCD). Due to its strong coupling at typical energy scales in today's Universe, predictions for strongly interacting matter, such as the one of the quark-gluon plasma, appearing in collisions of heavy nuclei at the Relativistic Heavy Ion Collider (RHIC) and future efforts, cannot be obtained using perturbative methods. The numerical treatment of QCD, discretized on a spacetime lattice – lattice QCD – has proven to be a viable tool to investigate the properties of QCD in the strongly coupled regime. Due to the strong coupling, quarks and gluons are confined into hadrons at low energies, while asymp-totic freedom implies the existence of a phase with quasi-free quarks and gluons at high energies, i.e., high temperatures or quark densities, for instance. The phase transition between these two phases is the subject of intense theoretical and experimental research and also of particular interest for astrophysics and cosmology.

For simulations in lattice QCD, the regime at non-zero baryon density remains inaccessible due to the complex action problem and can only be accessed indirectly through extrapolations from vanishing or imaginary chemical potentials. In particular, this also concerns the possible existence of a critical endpoint, separating the crossover line of the QCD transition from a hypothetical first order phase transition at large chemical potentials. In recent years, several different extrapolation schemes have been explored [1, 2], showing agreement with direct simulations at small volumes.

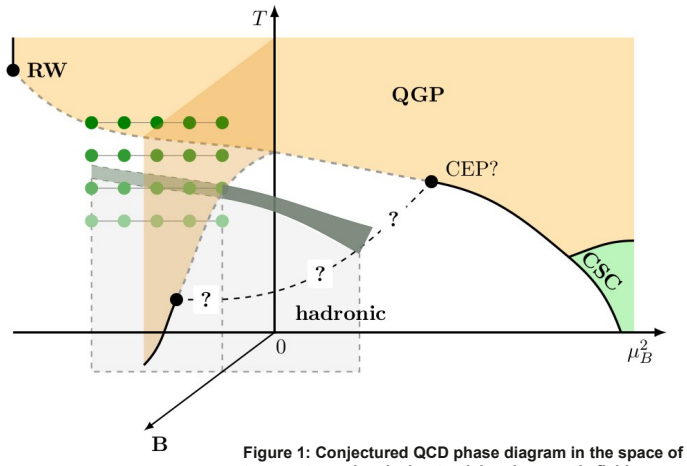

Figure 1: Conjectured QCD phase diagram in the space of temperature, chemical potential and magnetic field.

Apart from a non-zero baryon density, a broad range of physical systems also feature strong magnetic fields, such as magnetars, off-central heavy-ion collisions and, potentially, even the early Universe. For a phenomenological description of such systems, incorporating the effects due to the magnetic field – an external field on typical QCD timescales – is mandatory, but a systematic study is still missing. The properties of QCD subject to an external magnetic field at vanishing quark densities are by now well understood, including the strengthening of the crossover and the existence of a critical endpoint in the phase diagram. Naturally the question arises whether the two critical endpoints (if both exist) are connected by a single critical line. Consequently, studying the impact of the combinationof magnetic fields and quark densities on the equation of state and the phase structure is important not only for the description of the aforementioned physical systems, but also might provide a different angle on one of the long-standing questions concerning the phase diagram at vanishing magnetic field.

Results and Methods

In our study we are investigating the properties of QCD subject to both non-zero chemical potential and external magnetic fields using Taylor expansion and analytic continuation from imaginary chemical potentials. In this first step, our main interest is in the equation of state (EoS) in the phenomenologically relevant setup for heavy-ion collisions where the freeze-out conditions are

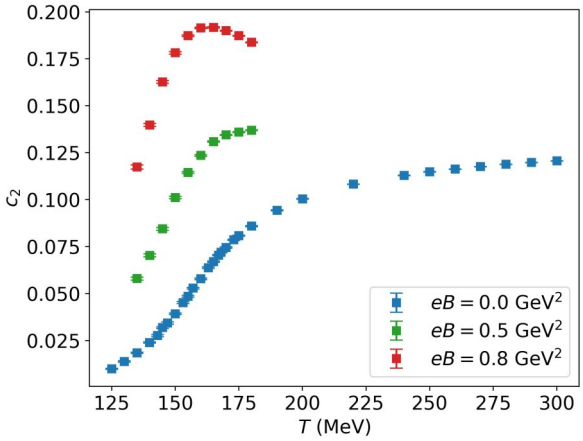

Figure 2: Coefficient between the strangeness and baryon chemical potential that ensures strnageness-neutraliy in our simulations.

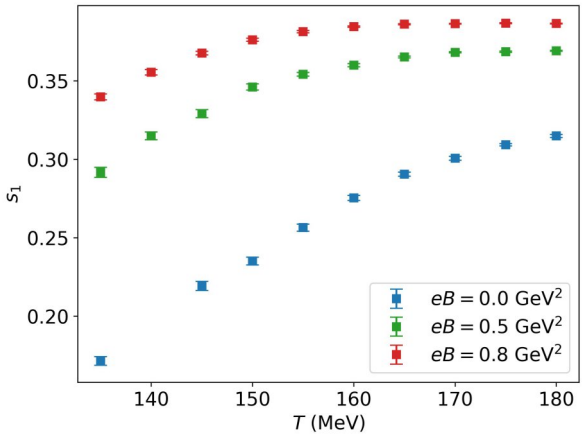

Figure 3: Leading-order Taylorcoefficien of the Equation of State as a function of temperature for three values of the magnetic field.

met. The conditions are: 1) a fixed ratio between the baryonic density and the charged-particle density constraint given by the number of valence quarks present in the colliding ions; 2) the strangeness neutrality, due to the absence of valence strange quarks in the ions. These conditions can be implemented by tuning charge and strangeness chemical potentials when changing the baryon chemical potential. The tuning coefficients can be related to chemical potential susceptibilities [3], which we computed from our simulations at vanishing chemical potential. In the initial stage of the project, we obtained, for the first time in the literature, the proportionality coefficients, between these chemical potentials, labelled q1 and s1 for the charge and strangeness chemical potentials, respectively, as a function of temperature for different magnetic field strengths [4]. As an example, we show the s1 coefficient, ensuring strangeness-neutrality in Figure 2. For this success the usage of SuperMUC-NG has been essential, since a large number of gauge field configurations for 30 different ensembles with sizeable lattices (8×32^3) had to be generated for each value of the magnetic field. In addition, the computation of the observables is numerically expensive, due to large numbers of random vectors necessary to evaluate the operator traces accurately. For an efficient use of SuperMUC-NG our code has been specifically optimized for such large scale cluster architectures.

In terms of physics, the results from Figure 2 can be understood by the suppression of the hadronic contributions of protons and lambda-baryons to the partition function due to the presence of the magnetic field. This leads to a decrease of q1 and an increase for s1, respectively.

Using the tuning coefficients, we were able to compute the leading-order Taylor coefficient for the extrapolation of the Equation of State (labeled c2) to real and imaginary chemical potentials for the three different magnetic field values as functions of temperature. The results are shown in Figure 3. Even these preliminary results can be used to benchmark future heavy-ion collision models that take into account both finite densities and magnetic fields.

Ongoing Research / Outlook

Currently we work on the fine tuning of charge and strangeness chemical potentials to match the experimental setup of heavy-ion collisions, in order to start the first simulations at non-zero imaginary chemical potentials. Those are needed for the final and precise extrapolation to real chemical potentials. In this process the explicit isospin breaking and the strangeness neutrality conditions have to be taken into account, which are expected to lead to further deviations from the tuning at vanishing magnetic field [5]. Eventually, the calculation of high-order fluctuations at imaginary chemical potentials will for the first time lead to an accurate calculation of the EoS for the phenomenologically relevant heavy-ion setup, discussed above, featuring finite baryon density and non-zero magnetic fields.

References and Links

[1] S. Borsányi, J,.N. Günther, R. Kara, Z. Fodor, P. Parotto, A. Pásztor, C, Ratti, and K. K. Szabó, Phys. Rev. D, 105, no.11 (2022), 114504, 2022.
[2] S. Borsányi, Z. Fodor, J. N. Guenther, R. Kara, S. D. Katz, P. Parotto, A. Pásztor, C. Ratti, and K. K. Szabó, Phys. Rev. Lett., 126 (2021), 232001.
[3] A. Bazavov et al., Phys. Rev. Lett. 109 (2012) 192302.
[4] S. Borsányi, B. B. Brandt, G. Endrődi, J. Guenther, R. Kara, and Valois, A. D. M. QCD. 12 2023.
[5] S. Borsányi et. al. Phys. Rev. Lett. 125 (2020) 052001.

The angular momentum structure of the proton and other hadrons

Research Institution
[1]Regensburg University

Principal Investigator
Andreas Schäfer[1]

Researcher
Michael Engelhardt2, Sebastian Lahrtz[1], Maximillian Schlemmer[1], Enno Scholz[1], Wolfgang Söldner[1], Hai-Tao Shu[1], Lisa Walter[1], Christian Zimmermann[3]

Project Partner
[2]New Mexico State University, Las Cruces
[3]University of Marseille

Funding
DFG Scha-458/21, Scha-458/22

SuperMUC Project ID: pn73vu, pn49ge

5

Introduction

Quantum Chromodynamics (QCD) is without any reasonable doubt the correct fundamental theory of the strong interactions between quarks and gluons. Thus it describes exactly all properties of particles like the proton, which is a bound state of quarks and gluons, so-called hadrons. However, while it is rather straightforward to calculate all properties of atoms, i.e. bound states of electrons and nuclei, from Quantum Electrodynamics, the theory which describes the electromagnetic force, this is not the case for QCD. In fact, the QCD-equations are so difficult to solve that many or even most properties of hadrons are only poorly understood at the QCD level. Simulations on powerful computers like SuperMUC-NG, using Lattice QCD provide some, but only very limited insight, because they require physical time to be analytically continued to imaginary time which precludes, e.g., the possibility to describe any real time dynamics. Quantum computers might improve this situation some day but, at present, this day still seems to be far away. Another very severe technical problem is the so-called operator mixing which further reduces the number of calculable quantities dramatically. Thus, people explore already for many years alternative formulations, which might allow to broaden the range of calculable quantities, so far, unfortunately without decisive success. In our research projects we explore two such ideas. In this report we will focus on just one of them which goes under the name of LaMET (Large Momentum Effective Theory).

Research in this field proceeds in two steps. First, one develops complete parameterizations of those hadron properties which can be experimentally observed in terms of a hierarchy of functions (typically with rather exotic names, e.g. transversity, worm-gear functions, distribution amplitudes, or Boer-Mulders function.) This hierarchy is artfully constructed such that it groups these functions according to the difficulty of calculating them. It turns out that at each level of this hierarchy there are angular momentum dependent and independent functions and proceeding from the simple to the more difficult ones one typically calculates all of them before continuing to the next, more difficult level. This is what we did by calculating first the quark transversity and, after this was successful, the quark Boer-Mulders function in the proton. The Boer-Mulders function is a so-called Transverse Momentum Dependent distribution (TMD) which in addition depends on the transverse momentum (relative to the total proton momentum) of the hit quark. TMDs have very complicated properties closely related to the properties auf gauge links, which are the most basic, gluonic building blocks of QCD. The extraction of TMDs is, e.g., the primary goal of the new accelerator center EIC (Electron Ion Collider) to be built in the US, as well as its Chinese competition, the EicC (Electron ion collider in China). Without going into any detail let us add that in all probability experiments alone will not allow to reach this goal. Instead, to do so will require complementary information from the lattice of the type we are providing. Let us finally stress that we perform only pioneering work aiming primarily at clarifying the many open questions of TMD physics. LaMET requires taking several non-trivial limits (most importantly the proton momentum must be larger and nobody knows as of today how large to reach a certain accuracy) and our simulationstest whether these can be controlled with real world resources. The upshot of this report is that they can and that we do. This success paves the way for future high-precision simulations, which, unfortunately, will require far more computer time.

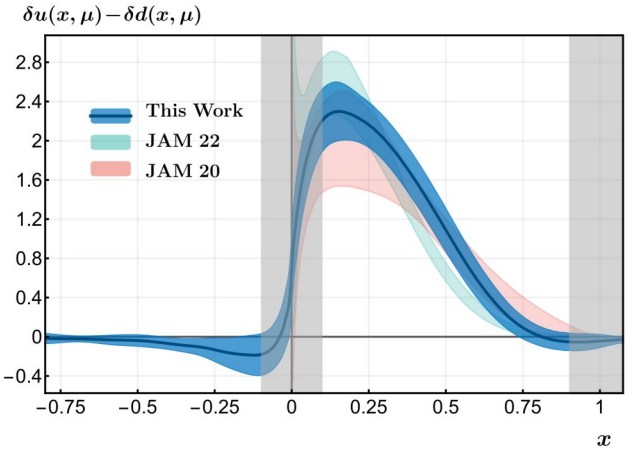

Figure 1: Isovector quark transversity distribution in the nucleon, which is given by the difference of transversities for up and down quarks (*u* minus *d*). *x* is the momentum fraction carried by the hit quark. JAM 20 and JAM 22 are phenomenological fits from 2020 and 2022 of the JAM collaboration.

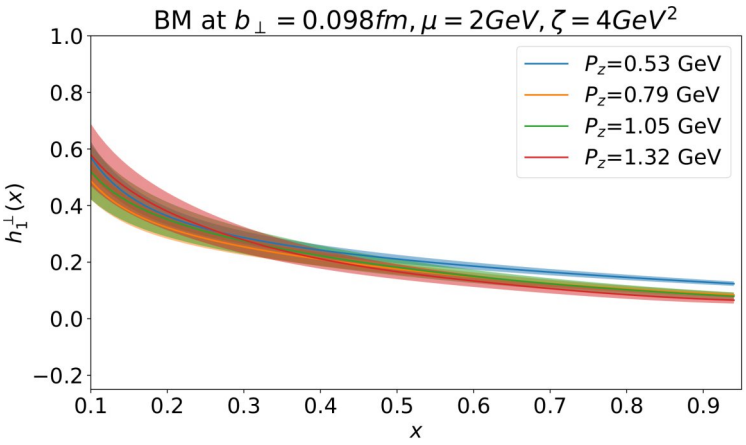

Figure 2: First results for the isovector quark Boer-Mulders function in the pion for different pion momenta which indicate convergence with pion momentum.

Results and Methods

Work on QCD is dominated by two techniques namely, analytic, perturbative calculations (pQCD) and numerical lattice QCD (LQCD). LaMET combines both. The quantities one has to calculate are non-local in time and space and, as explained above, time looses its originalmeaning when analytically continued. (In fact it turns into an inverse temperature.) Therefore, the calculation of functions which are non-local in time is not feasible in standard LQCD. In LaMET one calculated instead correlators which are only non-local in space and matches them to the required objects by continuum perturbation theory. More precisely, one calculates the perturbative matching function with which the LaMET lattice result, called e.g. quasi parton distribution function must be convoluted. This strategy was pioneered by X. Ji (University of Maryland) and is meanwhile employed by several large collaborations, primarily in the US and China. (We are members of the Lattice Parton Collaboration based in China.) While all of this might sound rather natural and straight forward, the technical problems are hellishly difficult. Still, the cooperative effort of all these groups has succeeded to overcome most problems such that TMDs and the similarly complicated Generalized Parton Distributions (GPDs) can meanwhile be calculated.

Figure 1 shows our results for the isovector quark transversity distribution in the nucleon which we analyzed first, leading to the publication [1]. Transversity is no TMD but a comparably simple observable which has an approximate probabilistic interpretation. It parameterizes the probability that a quark in a transversely polarized nucleon has momentum fraction *x* and a spin orientation parallel to that of this nucleon minus the probability that the quark spin is oriented in the opposite direction. Negative values of *x* indicate the corresponding probability for an antiquark with positive momentum fraction. The grey bands indicate regions with large and poorly known systematic errors. Ours was a pioneering calculation demonstrating that such calculations are possible with moderate present day computer resources. Future work with larger resources should markedly improve their precision.

Figure 2 shows our first results for the isovector quark Boer-Mulders function of the pion which is a TMD and thus much more difficult to compute. This function contains far more information than, e.g., the transversity distribution and is correspondingly more difficults to obtain. However, the figure demonstrates that we succeed in doing so. The large amount of information it contains shows up in the many variables on which the Boer-Mulders function depends, namely five, which makes it difficult to plot. In Figure 2 three of these are chosen fixed (top line), one is indicated by the color code, and one is plotted on the abscissa. The main message of this figure is that although the pion momentum P_z is significantly smaller than one would like it to be convergence towards the limit of very large P_z seems already to be observed.Presently we perform runs with finer lattices and larger P_z to corroborate this finding. Let us stress that reaching larger momenta is the most crucial requirement to obtain higher precision and that this is to a large extend made possible by an algorithmic techniques we developed in [2], which is meanwhile used by all groups working in this field.

Ongoing Research / Outlook

Our future work has two main objectives namely to extend the range of applications of LaMET and to improve the precision reached. The latter has two aspects. The purely statistical precision is rather straight forward to improve as so far we have used only rather limited computer resources. In contrast, identifying and solving the remaining conceptual problems of LaMET is a very demanding task. The most difficult remaining such problem is to develop methods allowing to reach substantially larger hadron momenta at acceptable cost beyond what was suggested in [2]. Though all groups in this field try to do so, progress so far was rather incremental than decisive.

References and Links

[1] F. Yao et al., PRL 131 (2023) 261901; doi:10.1103/PhysRevLett.131.261901.
[2] G.S. Bali et al., Phys. Rev. D 93 (2016) 094515; doi:10.1103/PhysRevD.93.094515.

Nucleon observables as probes for physics beyond the standard model

Research Institution
[1]NIC, DESY, Zeuthen

Principal Investigator
Karl Jansen[1,2]

Researcher
C. Alexandrou[2,3], S. Bacchio[2], M. Constantinou[4], J. Finkenrath[5], B. Kostrzewa[6], G. Koutsou[2], M. Petschlies[6], F. Pittler[2], G. Spanoudes[3], C. Urbach[6]

Project Partner
[2]The Cyprus Institute, [3]University of Cyprus, [4]Temple University, PA, [5]University of Wuppertal, [6]University of Bonn

Funding
"NiceQuarks", RIF, EXCELLENCE/0421/0195, "3D-nucleon", RIF, EXCELLENCE/0421/0043

SuperMUC Project ID: pr74yo

Introduction

In high-precision, low-energy particle physics experiments, small but significant discrepancies have been found when compared to expectations from theory. This has substantially increased the interest in precision nucleon structure measurements. This so-called *precision-frontier* of particle physics is serving a complementary role to new physics searches in the *high-energy frontier*, such as those pursued by experiments at the Large-Hadron Collider (LHC) at CERN. A major challenge in such searches is determining the precise contribution from the strong force component of the Standard Model. Experimentally, the proton, which is stable and abundantly available, is an optimal probe for studying strong interaction phenomena and is thus the target of several ongoing experiments such as at MAMI in Mainz, Jefferson Lab, and Fermi Lab as well as at the planned Electron-Ion Collider (EIC) in the US, where an extensive program to map out the proton's rich sea structure is foreseen. Theoretically, strong interaction phenomena are governed by Quantum Chromodynamics (QCD) and at energy scales relevant to proton structure, the only known way of studying QCD from first principles is via large scale simulations using the lattice formulation. Our project on SuperMUC-NG targets such simulations, using the so-called twisted mass fermion formulation, which has the advantage that observables, such as nucleon structure observables we are targeting here, converge faster to the continuum limit. Using the resources provided at LRZ, our Extended Twisted Mass Collaboration (ETMC) [1] has simulated QCD gluon configurations which include dynamical mass degenerate up and down, as well as the strange, and charm quark flavors ($N_f = 2+1+1$) with masses tuned to their physical values and at four values of the lattice spacing parameter. These state-of-the-art ensembles are being used as building blocks for a wide range of lattice QCD studies in addition to nucleon structure which we highlight here.

Results and Methods

Physical point $N_f = 2+1+1$ ensembles
Lattice QCD refers to the discretization of the Lagrangian of QCD on a 4-dimensional Euclidean grid that allows generating representative configurations via a Markov chain Monte Carlo process. The generated gauge configuration ensembles are then used to calculate hadronic quantities of interest. Continuum physics requires extrapolating using ensembles with multiple values of the lattice spacing and volume to the continuum and infinite physical box length limit. Several algorithmic improvements we have implemented during our allocations on SuperMUC-NG include better integration schemes that scale more favorably with the lattice volume and advanced linear solvers [2, 3]. These algorithmic improvements complement the increased availability of computational resources to allow for the simulation of ensembles with larger volumes and with finer values of the lattice spacing, as we document in terms of SuperMUC-NG core-hours over the years in Figure 1.

Nucleon axial form-factors
The ensembles generated on SuperMUC-NG have been used to obtain the axial and pseudo-scalar form factors of the nucleon and in particular to obtain the form factors at the continuum limit directly at the physical point. The axial structure of the nucleon is important for both understanding strong interaction dynamics and in revealing new physics. However, unlike the Electromagnetic form factors of the nucleon [4], the axial form factors are not well known experimentally making lattice QCD an important tool for their study.

Our recent analysis [5] using the three coarser values of a shown in Figure 1 is shown in Figure 2, including a continuum limit to $a \rightarrow 0$. For the axial form factor, the value at $Q^2 = 0$, known as the axial charge (g_A), is well-known experimentally via a dedicated process and thus

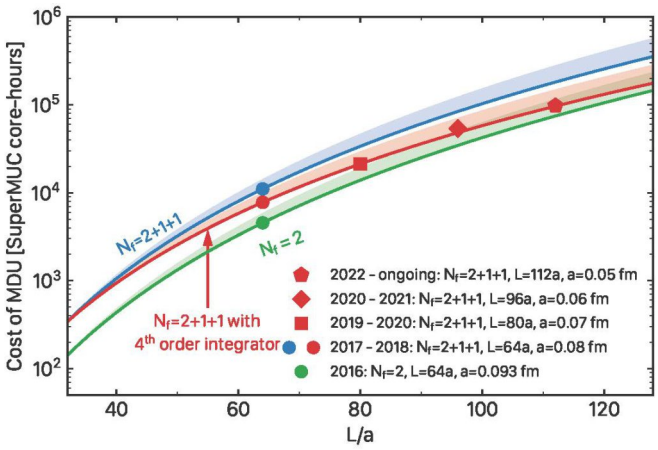

Figure 1: Simulation points carried out on SuperMUC since 2016. The red curve shows improvements achieved by our project to reduce computational cost while increasing the lattice volume.

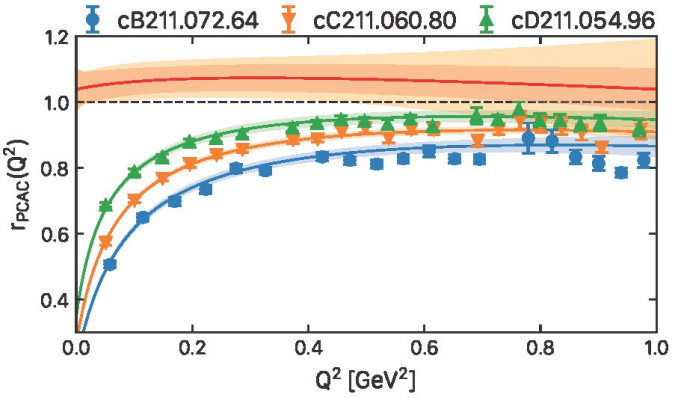

Figure 3: The PCAC ratio for the ensembles with a≈0.08 fm (blue circles), 0.07 fm (orange down-pointing triangles), and 0.06 fm (green up-pointing triangles) and after continuum extrapolation with statistical errors (red band) and with both statistical and systematic errors (orange band) considered.

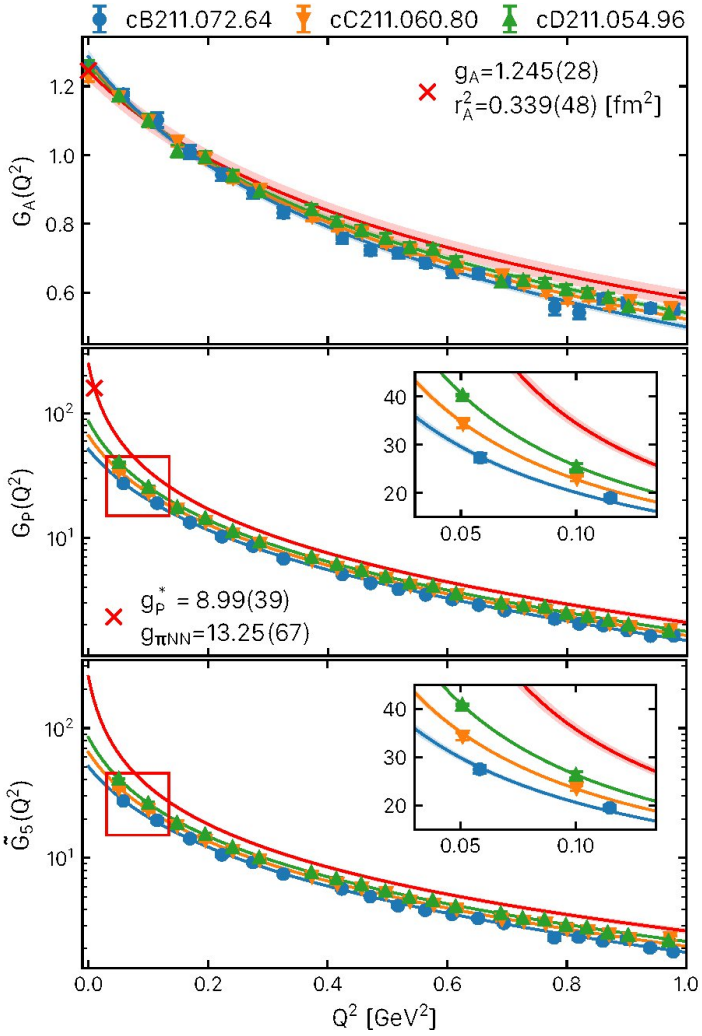

Figure 2: Axial (top), induced pseudoscalar (center), and pseudoscalar (bottom) form factors of the nucleon computed on the three ensembles with a≈0.08 fm (blue circles), 0.07 fm (orange down-pointing triangles), and 0.06 fm (green up-pointing triangles) and after continuum extrapolation (red curve).

its reproduction by our simulation is a confirmation of the validity of our lattice setup.

In QCD, the partial conservation of the axial current (PCAC) leads to a so-called PCAC relation which relates the three form factors shown in Figure 2 in a non-trivial way. Confirming that this relation holds in lattice QCD is an important verification of the extracted form factors and their subsequent use in phenomenology. In Figure 3, we show the ratio r_{PCAC}, a rearrangement of the three form factors in the PCAC relation that should yield r_{PCAC}= 1 if the relation holds. We see that individually the ensembles at finite lattice spacing do not yield 1 and thus to not fulfil the relation. However, at the continuum limit, when also considering all systematic error, the relation is restored.

Ongoing Research / Outlook

Our analysis for the extraction of the full set of nucleon structure observables on the ensembles generated on SuperMUC-NG is ongoing. In particular, the axial form factors presented here have been carried out using three N_f = 2+1+1 ensemble at the physical point. While this allows for a first continuum extrapolation with all ensembles simulated using physical quark mass parameters, the analysis on all four ensembles shown in Figure 1 will allow us for a more robust continuum limit. As the lattice spacing becomes smaller and lattice thus the number of lattice grid points increases, large-scale resources such as SuperMUC-NG become all the more important for such studies.

References and Links

[1] ETMC: https://github.com/etmc; http://www-zeuthen.desy.de/~kjansen/etmc/
[2] C. Alexandrou, S. Bacchio, J. Finkenrath, Comput. Phys. Commun., 51-64 (2019), doi:10.1016/j.cpc.2018.10.013 [arXiv:1805.09584 [hep-lat]].
[3] C. Alexandrou et al., Phys. Rev. D98, no.5, 054518 (2018) doi:10.1103/PhysRevD.98.054518 [arXiv:1807.00495].
[4] C. Alexandrou, et al., "Nucleon electromagnetic form factors using Nf=2+1+1 twisted mass fermion ensembles at the physical mass point", PoS LATTICE2022 (2023) 114, doi:10.22323/1.430.0114.
[5] C. Alexandrou et al., Phys. Rev. D109, no. 3 034503 (2024) doi:10.1103/PhysRevD.109.034503, [arXiv:2309.05774].

A strong side of weak decays: How beauty and charm resonate in pions and kaons

Research Institution
[1]Helmholtz-Institut für Strahlen- und Kernphysik, Rheinische Friedrichs-Wilhelms-Universität Bonn

Principal Investigator
Marcus Petschlies[1]

Researcher
Constantia Alexandrou[2], Luka Jevšenak[3], Luka Leskovec[3], Stefan Meinel[4], John Negele[5], Srijit Paul[6], Ferenc Pittler[2], Andrew Pochinsky[6]

Project Partner
[2]Cyprus Institute & University of Cyprus, [3]Jozef Stefan Institute, University of Ljubljana,
[4]University of Arizona, [5]Center for Theoretical Physics, Massachusetts Institute of Technology,
[6]School of Physics and Astronomy, University of Edinburgh

SuperMUC Project ID: pr27yo

Introduction

The B and D mesons, which are built each from a pair of bottom and light or charm and light quarks, are not present in ordinary matter but can be produced at high-energy particle colliders. However, they play an eminent role in our present-day efforts to probe our understanding of the fundamental strong and electroweak forces in Nature, which we describe in the Standard Model of elementary particle physics. Both mesons are unstable under the electroweak interaction, and the transition amplitudes for their decays to states including two pions or a pion and a kaon (denoted for brevity as e.g. $B \to \pi\pi l\bar{\nu}$ and $B \to K\pi ll$) are an essential ingredient to the investigation of semileptonic B and D decays in the effective electroweak sector of the Standard Model. These flavor changing processes, involving transitions at the quark level of heavy bottom and charm quarks to lighter up and strange quarks (for short, e,g. $b \to ul\bar{\nu}$ and $b \to sll$, plus leptons and neutrinos), are excellent candidates at the high-precision frontier for testing the Standard Model, as well as in the search for physics beyond the Standard Model: the transition $B \to \pi\pi$ pertains to the precision estimate of the element V_{ub} of the Cabibbo-Kobayashi-Maskawa (CKM) quark mixing matrix, a fundamental parameter of the Standard Model.

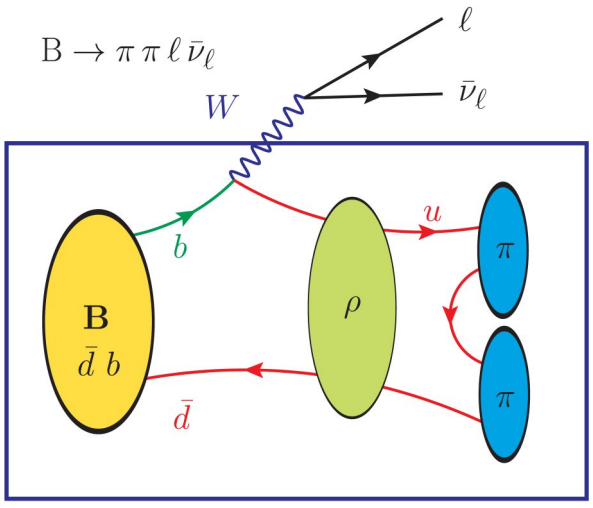

Figure 1:Illustration of the B-decay to a π π state via theρ resonance.

The *flavor anomalies* in $b \to sll$ observed as tensions between measured B-meson decays and their Standard Model prediction allow for exciting investigations of New Physics scenarios.

The decay is $B \to \pi\pi l\bar{\nu}$ illustrated in Figure 1. The transition matrix elements factorize in an electroweak part, and in one part that is governed by the strong interaction and determined numerically in pure Quantum Chromodynamics (QCD). The latter factor is pictorially defined by the blue box frame. Being subject to the strong force, the theoretical physics investigation of this process requires the High Performance Computing methods of Lattice QCD and thus the computational power of the supercomputer SuperMUC-NG. Studying the process $B \to \pi\pi$ becomes even more challenging, when the elastic pion-pion scattering in the final state, which features the famous ρ resonance, is taken into account.(Or the κ and K^* resonances for the transition $B \to K\pi$.)

The rigorous calculation of such matrix elements with an unstable final state must include the proper conversion of our Lattice QCD data obtained in a finite box of size L^3 to the in finite volume of the physical world. Thus the study of $\pi\pi$ and $K\pi$ scattering in Lattice QCD forms an integral part of our computational project.

Results and Methods

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 the P -wave (ρ resonance) channel and (2) for $K\pi$ I = 1/2 with focus on

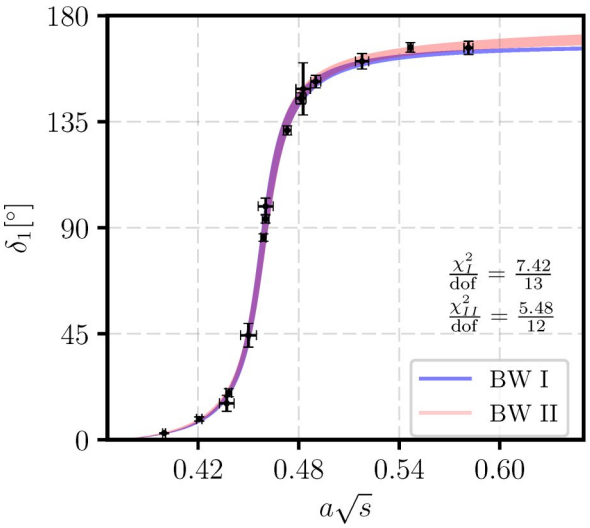

Figure 2: phase shift fit results for pion -pionP-wave at 320 MeV pion mass (top), kaon- pion S-wave (center) and P-wave (bottom) elastic scattering at 180 MeV and 320 MeV pion mass, determined from Lattice QCD [1,2].

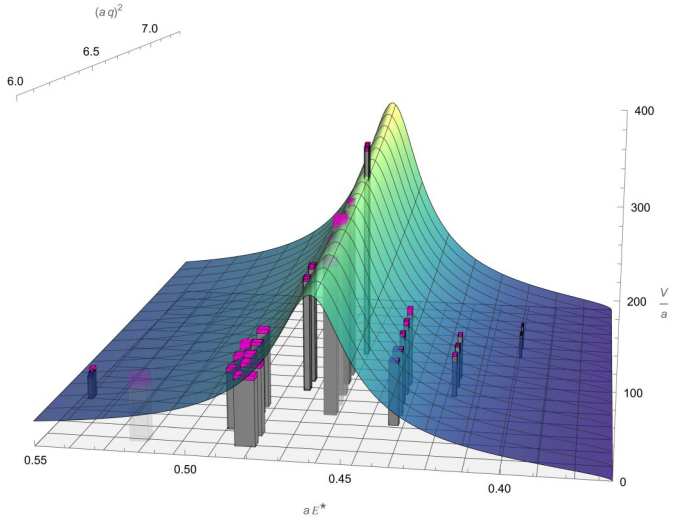

Figure 3: Form factor as a function of momentum transfer q^2 and invariant mass of the 2-pion system E^*. All data is shown in lattice units.

S-wave (κ resonance) and P-wave (K^* resonance). Step (3) is the production of 3-point correlation 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) with the Lellouch-Lüscher method and its generalization by Briceno, Hansen and Walker-Loud. The required Wick contractions to construct 2-point correlation functions and matrices thereof for variational analysis of the Lattice QCD spectrum are calculated based on SuperMUC-NG partitions with up to 2,592 cores and approximately 25,000 core-hours per job, and using the C/C++ QLUA software suite [3] and QOPQDP-multigrid solver for numerically solving the large sparse linear system of the Dirac equation on the lattice. With the help of SuperMUC-NG resources we started our investigation with the transition of a B meson to $\pi\pi$ via coupling to a vector and axial vector current. The phase shift δ_1 for elasticpion-pion scattering n the top right picture with its prominent S-shape indicates the famous ρ resonance.

Our analysis of the transition amplitude $B \to \pi\pi$ via vector current is shown in the bottom right picture. The pion mass here is 320 MeV withbox size of our lattice 3.6 fm and lattice spacing 0.11 fm. We show the original lattice data as columns with size of the uncertainty as magenta caps. The heat-map surface shows the best fit to the data and beautifully displays the resonant enhancement around the ρ pole.

Ongoing Research / Outlook

With pion-pion and kaon-pion scattering amplitude information in hand, we simulate the beauty and charm meson 3-point functions for the transitions $B/D \to \pi\pi/K\pi$ and extract the infinite-volume matrix elements in analogy to the case of $\pi\gamma^* \to \pi\pi$ in [4]. In the near future, we focus on the analysis of the plethora of data produced with the resources of SuperMUC-NG. Follow-up investigations target simulations ever closer to the conditions of the physical world, including lighter pions, larger lattice volume, smaller lattice discretization, and probing such intricate effects as coupled decay channels (mixing of 2-kaon and 2-pion states, and of kaon-pion and kaon-eta states), towards a rigorous contact with experiment.

References and Links

[1] C. Alexandrou et al., Phys. Rev. D 96 (2017) no.3, 034525.
[2] G. Rendon et al., Phys.Rev.D 102 (2020) 11, 114520.
[3] USQCD software Qlua package, https://usqcd.lns.mit.edu/w/index.php/QLUA
[4] C. Alexandrou et al., Phys. Rev. D 98 (2018) no.7, 074502.
[5] L. Leskovec et al., PoS LATTICE2022 (2023) 416; L. Leskovec PoS LATTICE2023 (2024) 119.

Form factors of pseudoscalar mesons with Lattice QCD

Research Institution
Institute for Nuclear Physics, University of Mainz

Principal Investigator
Georg von Hippel

Researcher
Konstantin Ottnad

Project Partner
–

Funding
DFG Project HI 2048/1-2

SuperMUC Project ID: pn49yu (Gauss Large-Scale project)

Introduction

With the discovery of the Higgs particle at the Large Hadron Collider (LHC), the Standard Model of Elementary Particle Physics has been vindicated as the correct description of all fundamental forces in Nature (except only gravity). However, many aspects of the Standard Model are still not fully understood. Besides the nature of the elusive neutrinos, this includes in particular the dynamics of the strong interaction, which binds quarks into composite particles called hadrons (which include the familiar protons and neutrons making up atomic nuclei).

The main difficulty with the strong interaction is that it is strong – analytical methods such as perturbation theory, so successful in quantum electrodynamics, will not work when the coupling constant governing the strength of the interaction is large, as is the case in Quantum Chromodynamics (QCD), the theory of the strong interaction, at low energies.

There are two main approaches to dealing with this issue and obtain theoretical predictions for low-energy QCD: lattice QCD and effective field theories, such as Chiral Perturbation Theory (ChPT). The latter describes the low-energy regime of QCD in terms of light hadronic degrees of freedom, particularly the pions, and contains a large number of so-called low-energy constants that need to be determined from experiment or predicted from the underlying theory of QCD. One problem in determining some of these constants experimentally is that they parameterize the dependence of the properties of the pions on the masses of the light quarks, which of course are fixed in Nature and cannot be varied as part of an experiment. This is where Lattice QCD comes to the rescue, since it permits simulating at different quark masses, allowing an interpolation to determine the low-energy constants of ChPT.

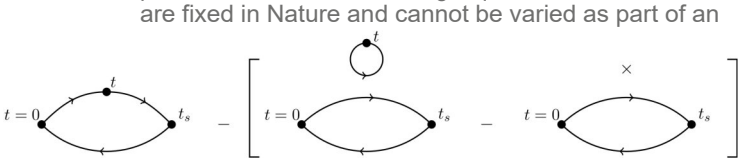

Figure 1: The quark-flow diagrams contributing to the scalar form factor of the pion. The quarks are to be considered moving in a sea of strongly interacting gluons. The hard-to-compute quark-disconnected contributions are bracketed.

This determination is the main purpose of this project studying the structural properties of the pions and their heavier pseudoscalar siblings, such as the kaons, which can be encoded in so-called form factors describing how charges are distributed within the hadrons.

Results and Methods

Lattice QCD proceeds by replacing continuous spacetime with a discrete hypercubical lattice, where quarks live on the sites of the lattice and the gluon fields mediating the strong interaction on the links connecting neighboring lattice sites. This discrete formulation enables the numerical treatment of the theory by means of large-scale simulations.

A lattice QCD simulation proceeds in several steps: first, a sufficiently large ensemble of gluon fields that are statistically distributed in accordance with the fundamental equations of QCD is generated by a long-running Markov-Chain Monte Carlo simulation. Second, one computes how quarks propagate in each of these gluon field backgrounds by solving the Dirac equation in the presence of this gluon field, and constructs the relevant operators for the quantities of interest from the resulting quark propagators. Finally, the observable quantities so constructed are analyzed using statistical methods in order to obtain the actual physical results.

The first two steps require HPC resources due to the very large size of the equation systems that need to be solved repeatedly, with typical problems involving matrices with dimensions of order $10^8 \times 10^8 \sim 10^9 \times 10^9$. Very efficient Krylov-space solvers and effective preconditioners are required for this purpose even on HPC hardware.

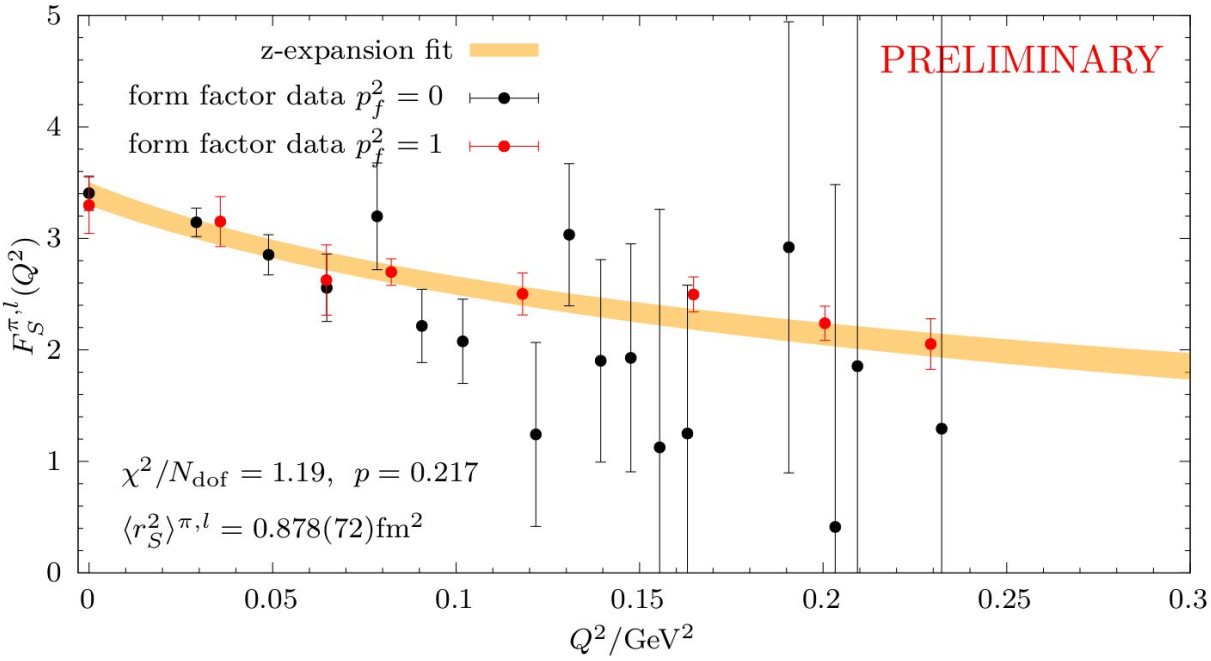

Figure 2: The scalar form factor of the pion at physical quark masses as a function of the relativistic four-momentum transfer. The red and black points correspond to different reference frames used in the computation, and the yellow band shows a fit to the momentum dependence.

In the case of the scalar form factor of the pion, an additional computational complication arises from the fact that there is a contribution from a quark-disconnected diagram, in which a quark propagates from a point back to the same point (cf. Figure 1). Considered naively, this would require what is known as an all-to-all propagator, i.e. the complete inverse of the Dirac operator, which is computationally infeasible to compute. Fortunately, it turns out that by using a stochastic estimator and leveraging a number of computational tricks, viz. combining the so-called one-end-trick (which allows computing the difference of two matrix inverses efficiently) with a generalized hopping parameter expansion (which is very accurate for heavy quark masses) and hierarchical probing (which exploits the spatiotemporal structure of the Dirac operator in order to evaluate its trace more efficiently), the quark-disconnected contribution can be evaluated with good precision at a feasible computational cost.

Using these methods, we have been able to obtain precise results for the scalar form factor of the pion at a range of light quark masses (cf. Figure 2 for a representative example [1]). The range of momentum transfers covered and the statistical accuracy reached are unprecedented for this quantity.

The quark-disconnected loops determined as part of this project can also be used in the context of other projects, such as determining the transition form factors governing the decay of a neutral pion into two photons [2], the hadronic contribution to the so-called anomalous magnetic moment of the muon [3], or the so-called sigma term determining the contribution of the quark masses to the mass of the nucleon [4].

Ongoing Research / Outlook

We are currently in the process of analyzing our simulation results so as to determine also the vector form factor of the pion and the corresponding form factors of the kaon. The next step in the analysis will involve extracting the values of the low-energy constants of ChPT from our results using fits to the quark-mass dependence of the radii of the hadrons determined from our form factors. We expect that the resulting values of the low-energy constants will be the most precise determination in existence.

Since the statistical error of the quark-disconnected part is still the dominant source of overall uncertainty in our results even after we have treated the one-point loop with all of the advanced computational methods at our disposal to the point where it cannot be further improved except by adding additional gluon fields, we anticipate further improving the precision of the two-point loop also entering the disconnected part by adding further evaluations of the quark propagator on the existing gluon field backgrounds to increase statistics.

Looking further afield, we plan to eventually determine the transition form factors governing the decays of the pseudoscalar mesons with the same kind of precision as we have determined the form factors. This will help to better understand the next-to-leading hadronic contributions to the anomalous magnetic moment of the muon, helping to address one of the great current puzzles of elementary particle physics.

References and Links

[1] K. Ottnad and G. von Hippel, arXiv:2401.06914.
[2] J. Koponen et al., arXiv:2311.07330.
[3] M. Cè et al., Phys. Rev. D 106 (2022) 114502.
[4] A. Agadjanov et al., Phys.Rev.Lett. 131 (2023) 26.

The Static force and other operators with field strength tensor insertions

Research Institution
[1]TUM NAT Physics Department, Garching

Principal Investigator
Nora Brambilla[1]

Researcher
Alexei Bazavov[2], Rafael Delgado[3], Andreas Kronfeld[4], Viljami Leino[5], Julian Mayer-Steudte[1], Peter Petreczky[6], Sebastian Steinbeißer[7], Antonio Vairo[1], Johannes Heinrich Weber[8]

Project Partner
[2]Michigan State University, [3]Universidad Politecnica de Madrid, [4]Fermi National Accelerator Laboratory, [5]Helmholtz Institut Mainz, [6]Brokkhaven National Laboratory, [7]Leibniz-Rechenzentrum, Garching, [8]Humboldt-Universität zu Berlin

Funding
DFG "ORIGINS", EXC-2094-390783311, DFG No. BR 4058/2-2, STRONG-2020 No. 824093

SuperMUC Project ID: pn49lo

Introduction

Low energy correlation functions arising due to factorization in the construction of nonrelativistic effective field theories are a key for a number of contemporary frontier problems such as: the strong coupling constant (α_S) of the quantum chromodynamics (QCD), spectrum and decay properties of quarkonia and exotics X,Y,Z, and the nonequilibrium evolution of heavy probes in the Quark Gluon Plasma (QGP). While the appropriate effective field theory description is known up to high orders, these results require non-perturbative calculation of low energy correlators using lattice QCD. Depending on the application, we require calculation of the static QCD force or a similar correlator with inserted chromoelectric or chromomagnetic fields components.

The static energy of an infinitely heavy quark-antiquark pair is the QCD representative of the well-known central potential problem. The static energy is a theoretically clean observable, which is known to next-to-next-to-next-to-leading order (N^3LO). The static energy is a function of the distance r between the static quark and antiquark, which is a well-defined quantum number that can be made small enough to serve as the (inverse) high energy scale. Other than that, the static energy only depends on the QCD Lambda parameter (Λ_{QCD}), which is directly related to αs through the QCD beta function. Among the lattice determinations of α_S, the static energy is one of the most precise results. However, the static energy suffers from a renormalon ambiguity; nevertheless, the derivative of the static energy, which defines the static force, is a physical quantity and might be the preferred choice of the object to investigate. We measure the static force directly through a generalized Wilson loop with a chromoelectric field insertion, avoiding the step of performing a numerical derivative, where additional systematic uncertainties arise. We apply a novel method called gradient flow to handle lattice artifacts located in the discretization of the field strength components, and to improve the signal-to-noise ratio. Furthermore, we can show the renormalizing effect of gradient flow on observables with field insertions by relating the direct force measurement with the numerical derivative of the static energy. This makes the gradient flow method extremely valuable for computations with field components needed in nonrelativistic effective field theories.

In the zero-separation limit, we can reduce the correlator to an adjoint Wilson line. Connecting two chromoelectric fields to such a line, we can study a multitude of physical phenomena. At finite temperature, this operator describes the diffusion of heavy quarkonium in QGP. Likewise, at zero temperature, it enters as a missing non-perturbative piece to the description of the annihilation rate of a P-wave spin-triplet quarkonium into light hadrons.

Results and Methods

The quantities we are interested in, and which encode the underlying physics, are formulated through expectation values of observables in a path integral formulism. With classical computational methods, this can be solved only with Monte Carlo methods on supercomputers. To produce reliable results, we need to generate Markov chains with low autocorrelation and high statistics, at different resolutions of the space-time grid.

For our simulations, we have used the publicly available code of the MILC collaboration [1], which is a hybrid MPI-OpenMP code written in C and in steady development since the 80s.

We have generated a range of pure gauge ensembles in both zero and finite temperature over which we run the gradient flow method. We use the state-of-the-art

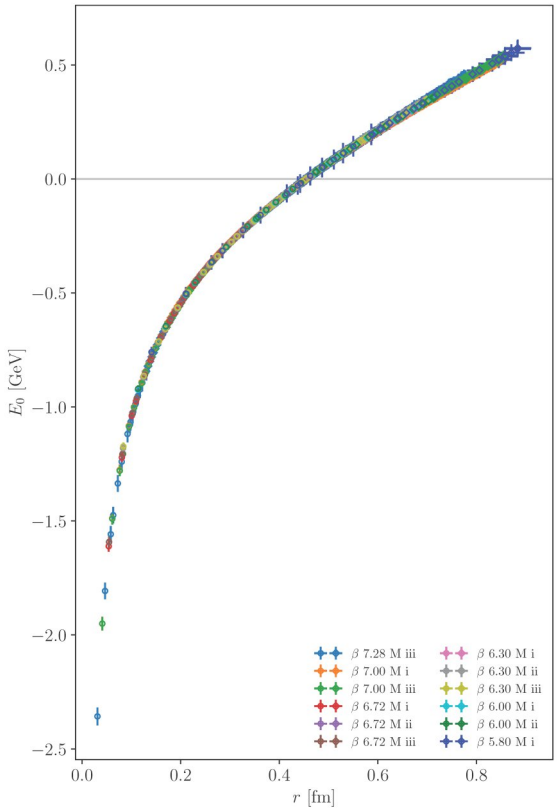

Figure 1: Result of the static potential with dynamical quarks.

discretization of the gradient flow, based on Runge-Kutta methods, to improve the convergence of the algorithm.

We have used Supermuc-NG for the simulations and measurement. To generate the ensembles, we used 0.4 M core hours, and 16 M core hours to perform the gradient flow with measuring our observables at several flow times. In a previous Supermuc project with ID pn56bo, 30 M core hours were used to extract the static potential at ensembles with dynamical fermions. Another 8 M core hours were used to perform perturbative calculations to control discretization artifacts. We store our data to data science storage (DSS) associated to the project ID pn56bo. Parts of the data can be reused for further projects by us and also by our collaborators using different HPC systems.

Results with dynamical fermions, but without gradient flow, were published last year [2]. Fig. 1 shows the result of the static potential with dynamical quarks. The gradient flow study without dynamical quarks was preliminary published in proceedings [3], and recently in a preprint version [5]. The recent study provides a theoretical prediction for ΛQCD which can be seen in Fig. 2 where we compare the result with previous determinations of Λ_{QCD} with different methods. Furthermore, we have shown the renormalization property of gradient flow explicitly, which prepares the application of the gradient flow method for further studies with chromo field insertions.

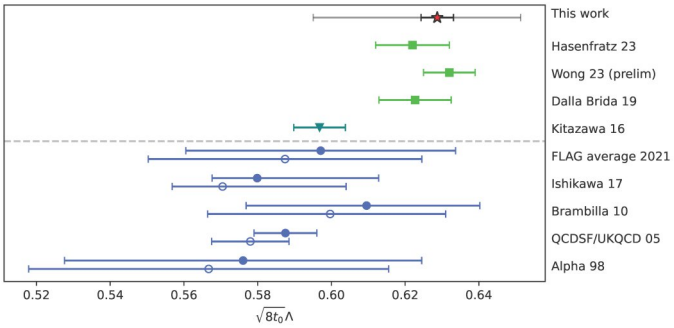

Figure 2: Literature comparison of the theoretical prediction of Λ_{QCD} for zero dynamical fermions.

Ongoing Research / Outlook

So far, the computations were carried out with pure gauge ensembles. The study showed the concept, and we aim for finer lattices to approach closer to the perturbative regime and extend the calculations to ensembles with dynamical quarks in the sea, which will complement our ongoing effort of measuring the strong coupling from a simulations with physical charm quark mass. The former aim requires the production of fine lattices with novel methods, for the latter aim we can use ensembles from previous projects.

Furthermore, we apply the gradient flow method to adjoint correlators, where chromo field components are connected by adjoint Wilson lines. The gradient flow renormalizes the field insertions and improves the signal-to-noise ratio. The connecting adjoint Wilson line by itself comes with an additional linear divergent self-energy contribution on the lattice. This divergence is made visible with the gradient flow. We are developing different ways to remove this divergence in order to measure all the operators where adjoint Wilson line manifests. After removing the divergence, fixing the associated constant is equivalent to choosing a scheme. A proceeding about the current status was published recently [5].

Especially, this continues our long-standing progress of determining the heavy quark diffusion, by also measuring the heavy quarkonium diffusion coefficient. For this observable, the divergence can possibly be removed in an easier way, by normalizing the operator with a Polyakov loop, which has same divergence structure.

For other observables, such as the P-wave decay constant, we are still in progress of developing state of the art methods for the divergence removal.

References and Links

[1] MILC code: http://www.physics.utah.edu/~detar/milc/
[2] Brambilla et. al., Phys. Rev. D 107, 074503 (2023)
[3] Mayer-Steudte J., arxiv: 2401.08244.
[4] Brambilla N., Leino V., Mayer-Steudte J., Vairo A., arxiv:2312.17231.
[5] Leino V., arxiv: 2401.06733.

Non-perturbative Heavy Quark Effective Theory

Research Institution
[1]Universität Münster, Institut für Theoretische Physik
[2]Neumann Institute for Computing NIC, DESY Zeuthen

Principal Investigator
Jochen Heitger[1], Rainer Sommer[2]

Researcher
Alessandro Conigli[3,4], Julien Frison[2], Patrick Fritzsch[5], Antoine Gérardin[6], Gregorio Herdoiza[7], Simon Kuberski[8], Carlos Pena[7], Hubert Simma[2]

Project Partner
[3]Helmholtz-Institut Mainz, GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt
[4]Johannes Gutenberg-Universität Mainz, [5]School of Mathematics, Trinity College Dublin
[6]Centre de Physique Théorique, Aix-Marseille Université, [7]Instituto de Física Teórica, UAM-CSIC Madrid
[8]Theoretical Physics Department, CERN

Funding
GRK 2149, GRK 2575, EuroPLEx H2020-MSCAITN-2018, STRONG-2020, HiCoLat, AMX-18-ACE-005

SuperMUC Project ID: pr84mi (Gauss Large-Scale project)

5

Introduction

The Standard Model of particle physics stands as one of the most remarkable theoretical constructions of scientific theory. Developed over several decades, it provides a comprehensive framework for understanding the fundamental constituents of matter and their interactions. This theory describes the structure of matter at length scales below the diameter of atomic nuclei with astonishing precision. It also successfully predicts particle decays and scattering cross sections of high-energy processes up to the energies reached at the Large Hadron Collider (LHC) at CERN in Geneva. Despite the great success in explaining a wide range of phenomena, there are still various open problems in fundamental physics that cannot be understood within the framework of the Standard Model, such as the existence of dark matter or the matter-antimatter asymmetry.

The hunt for new physics beyond the Standard Model is mainly performed within two complementary approaches. Direct searches aim at probing the existence of new particles, interactions and phenomena potentially detectable by high-energy particle colliders such as the LHC. On the other hand, indirect searches of new physics exploit precise measurements to probe new phenomena that may manifest as deviations from the Standard Model predictions. In this context, a key role is played by the “strong” interaction part of the theory, known as quantum chromodynamics (QCD), and describing the nuclei structure in terms of their constituents, quarks and gluons. In fact, at low energies, the dynamics due to the strong interaction cannot be determined analytically. Instead, we need to solve QCD numerically on a finite space-time grid on supercomputers such as SuperMUC-NG. This is done by stochastically evaluating the Feynman path integral defining the theory and represents the very heart of *lattice QCD*, a systematically improvable, non-perturbative approach to extract predictions of the theory.

Small deviations between experiment and theory were found in quantities related to bound states of heavy quarks. So far, these tensions are not significant enough to claim discovery of physics beyond the Standard Model. To improve methodologically compared to previous routes on the theory side, we study QCD contributions to decays of B-mesons, i.e., bound states of a heavy b-quark (with mass of about 5 GeV) and a light antiquark. Even on modern supercomputers, the properties of B-mesons in physically large volumes are difficult to investigate via lattice QCD, since they require fine resolutions and thus very large system sizes not yet affordable. We hence rely on Heavy Quark Effective Theory (HQET), an expansion in the inverse heavy quark mass, which – when properly implemented on the lattice, including next-to-leading order terms – can lead to new systematic precision [1].

Results and Methods

A lattice QCD calculation involves two steps: first, ensembles of gauge field configurations are generated according to the distribution dictated by the Boltzmann weight in the path integral. The second step is to calculate physical quantities on these configurations (“measurements”). In the current stage of our project, we have mainly focused on the latter, through the measurement of B-physics observables in both large and small volume ensembles, using in-house codes based on the publicly available openQCD package [2]. Utilising SuperMUC-NG, physical quantities were calculated in the relativistic regime and the static limit of HQET, where the infinitely heavy b-quark mass decouples from the theory. Measurements were performed on approximately 35 ensembles in both small and large volume, at different valence quark mass values from the charm region to the

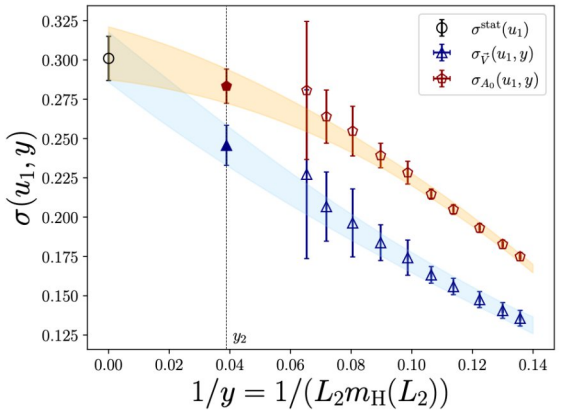

Figure 1: SSFs' interpolation to the target b-quark mass between relativistic and static data for QCD matrix elements entering B-meson leptonic decays.

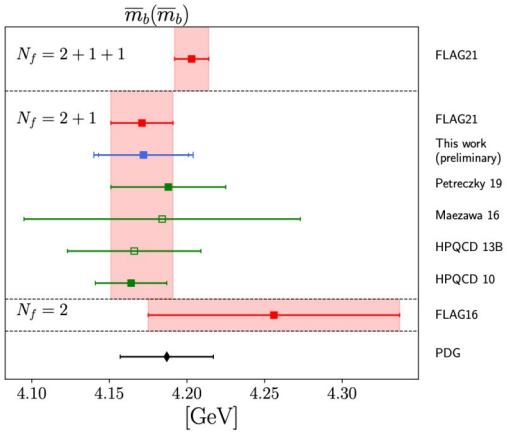

Figure 2: Comparison of our result for the b-quark mass in (2+1)-flavour QCD with FLAG averages and results from other collaborations [6].

bottom scale and beyond. In the large volume, we employed gauge configurations generated within the CLS simulation effort [3]. Here, for the lattice with finest resolution and 192×96^3 sites, we exploited up to 6,144 cores and observed excellent scaling in the entailed numerical solution of the light and heavy quarks' Dirac equation. Such challenging calculations are only possible on massively parallel computers such as SuperMUC-NG, allowing to collect the large statistics at fine resolutions required to achieve precise results.

A strategy for lattice B-physics, complementary to [1] and detailed in [4], is based on quantities called Step Scaling Functions (SSFs), constructed from the above relativistic and static measurements as suitable ratios such that intricate renormalisation and matching factors occurring in the static theory cancel.These SSFs admit a continuum limit and a simple dependence on the (inverse of the) heavy quark mass, thus enabling interpolations between the static and relativistic sets of data. Combining multiple SSFs provides the link from finite-size volumes to large, physical ones. We have presented in [5] our preliminary results for the b-quark mass and B-meson leptonic decay constants obtained from data generated on SuperMUC-NG. These new measurements will also play a crucial role for calculating QCD matrix elements of semi-leptonic decays of the B-meson. Here, an improvement in the precision of the theory predictions is essential to understand the noticed tensions with the Standard Model.

In Figure 1 we illustrate the interpolation in the (inverse of the) dimensionless heavy quark mass proxy m_H of the SSFs for the pseudo-scalar and vector B-meson decay constants between relativistic data and the static point in the continuum. The vertical line indicates the target b-quark mass scale, and the filled symbols display the result of the interpolation. By its very construction, this methodology allows to constrain B-physics observables at the physical b-quark scale without relying on dangerous extrapolations. The precision may be further improved by prolonging the Monte Carlo chains of our small-volume simulations, by adding data at masses closer to the target b-scale and by carefully studying the systematics in the HQET expansion used to parametrise the data.

These promising results demonstrate the applicability of our strategy to B-physics and its great precision potential by interpolating relativistic and static data to the physical b-scale. Figure 2 shows a comparison of our b-quark mass result with FLAG averages [6] and other studies.

Ongoing Research / Outlook

We plan to extend the computations to more large-volume ensembles with close-to-physical light sea quark masses, in order to reach the physical point by joint extrapolations to the continuum limit and to physical pion and kaon masses. By this we expect to make precise theory predictions with reliable error estimates and good control of systematic uncertainties. In addition, we are implementing our strategy to also determine semi-leptonic form factors of B-mesons. Their accurate knowledge is crucial for new physics searches in heavy flavour physics.

References and Links

[1] J. Heitger and R. Sommer, JHEP 0402 (2004) 022.
[2] M. Lüscher and S. Schaefer, openQCD, https://luscher.web.cern.ch/luscher/openQCD/
[3] https://wiki-zeuthen.desy.de/CLS/
[4] R. Sommer et al., PoS LATTICE23 (2023) 268.
[5] A. Conigli et al., PoS LATTICE23 (2023) 237.
[6] FLAG21: http://flag.unibe.ch/2021/

Charmonium and mixing with light hadrons

Research Institution
[1]Bergische Universität Wuppertal

Principal Investigator
Francesco Knechtli[1]

Researcher
Jochen Heitger[2], Roman Höllwieser[1], Tomasz Korzec[1], Jan Neuendorf[2], Michael Peardon[3], Juan Andrés Urrea Niño[1]

Project Partner
[2]Universität Münster
[3]Trinity College Dublin

Funding
DFG Research Unit FOR5269, NRW-FAIR network NW21-024-A

SuperMUC Project ID: pn29se (Gauss Large-Scale project)

Introduction

Quantum Chromodynamics (QCD) is the sector of the standard model describing the strong nuclear force, which binds quarks and gluons inside hadrons. The theory confines these constituents, which are never observed directly in experiment. In 2023 QCD celebrated its 50th anniversary. Yet a fundamental question like the existence of states made of gluons only and called glueballs, still awaits an answer from experiment and theory.

In this project we study charmonium, a system of mesons containing a charm quark-anti-quark pair. These studies underwent a revolution after a number of entirely unexpected narrow resonances called the X, Y and Z states were discovered by the Belle and BaBar experiments at the start of the new millennium. In spite of almost two decades of theory investigations, no clear picture of the internal mechanics of these resonances has emerged. Excitations of constituent gluons are likely to be an important part of some candidate hybrid states. Extensive experimental studies of charmonium and the XYZ states are underway at LHCb at CERN and there are significant plans at new experiments such as Belle II at the SuperKEKB B-factory and BESIII at BEPC and a future Super-c-tau factory.

Lattice QCD provides a framework to study charmonium from first principles. A rich spectrum is seen, including candidate hybrids in recent calculations. However these use an approximation where states are assumed to be stable and effects from charm-quark annihilation and mixing with glueballs and light hadrons are neglected. In lattice QCD the masses of charmonia are extracted from the exponential fall-off of the correlator in the Euclidean time separation between source and sink of the meson. Charm annihilation contributions are quark-line disconnected diagrams where the charm quark and anti-quark annihilate at the source and/or at the sink. Through charm annihilation charmonium states can mix with states made of light quarks or states made of gluons only. They constitute the significant remaining uncertainty in precise calculations of the low-lying charmonium states. Charm-annihilation effects are suppressed by the Okubo-Zweig-Iizuka (OZI) rule, but these dynamics are poorly understood.

In this project we have developed a technique to precisely calculate the masses of charmonium, *D*-mesons and multi-hadron operators like two-pion operators. We obtained first results on the spectrum of charmonium including the mixing with light hadrons.

The investigations reported here are part of the research programme of project 1 “The confined gluon: precision spectroscopy with charm quarks” in the Research Unit “Future methods for studying confined gluons in QCD” (FOR 5269) funded by the Deutsche Forschungsgemeinschaft (DFG) and of work package 1 of the NRW-FAIR network.

Results and Methods

Our improved distillation method [1] is based on optimal profiles for each particle and state. The profiles are wave functions of the mesons in distillation space. The latter is spanned by the eigenvectors corresponding to the lowest eigenvalues of the 3D gauge Laplacian defined on the time-slices of the lattice. Quark propagation from a time-slice to another is restricted between the distillation spaces and can be computed exactly. We performed these computations on a close-to-physical ensemble of gauge configurations that has been previously generated by us [2] using resources from a dedicated project on the JUWELS supercomputer. These are lattices of size 96×32^3 including the dynamics of three light degenerate quarks (whose mass sum is equal to its value in nature) and of a physical charm quark. The lattice spacing is $a = 0.054$ fm. We use 100 Laplacian eigenvectors for the light quarks and 200 eigenvectors for the charm quark.

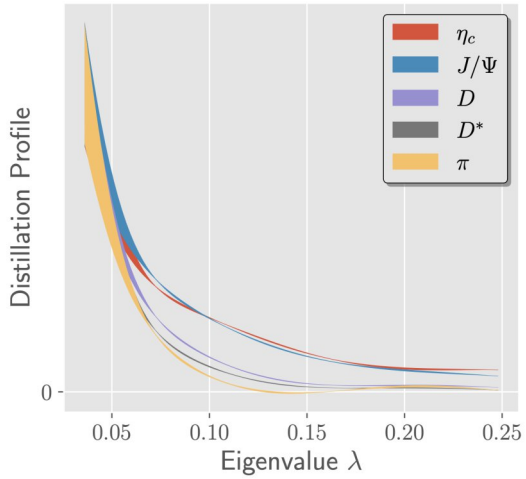

Figure 1: Optimized profiles in distillation space for different particles.

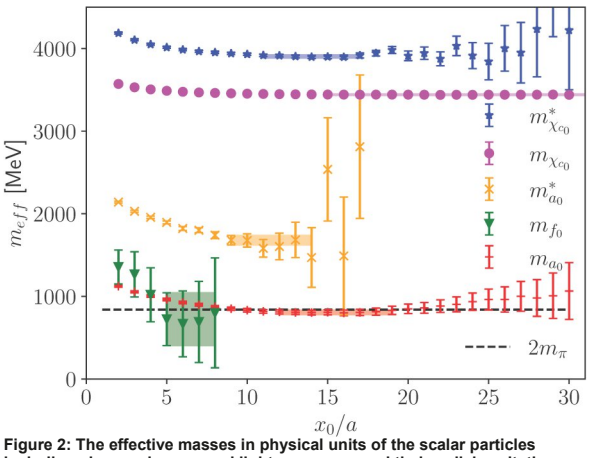

Figure 2: The effective masses in physical units of the scalar particles including charmonium χ_{c0} and light mesons a_0 and their radial excitations (marked with a star in the legend). The f_0 meson has light-quark disconnected contributions

We developed a software package for the measurement of correlation functions using the improved distillation method. Our programs are based on "QCDlib", a library written by us in C+MPI, that facilitates massively parallel QCD calculations. The computationally most intensive part is the calculation of the quark propagators. They require solutions of the Dirac equation for each Laplacian eigenvector on the right hand side. The solutions are obtained by calling the package openQCD [3] using a deflated SAP GCR solver with AVX instructions enabled for SuperMUC-NG. The computations for light and charm quarks on 4,000 gauge configurations required 60 millions core-h.

Charmonium and hybrid charmonium resonances above the open charm threshold can decay into a pair of *D*-mesons with back-to-back momentum. As a first step towards the study of these decays, optimal creation operators for a *D*-meson with zero and non-zero momentum were constructed [4]. In Figure 1 the optimal profiles as a function of the eigenvalues of the 3D Laplacian operator are displayed for different particles at zero momentum. The profiles of charmonium are wider than the ones of the *D*-meson, which themselves are wider than the ones of the pion. Lighter particles seem to have narrower profiles and therefore a better overlap with more smeared sources.

The charm quark anti-quark pair in charmonium can annihilate and mixing with other states with the same quantum numbers but different quark content (hidden flavour mixing) or no quarks (glueballs) is allowed. In [5] we investigate the mixing between charmonium and the other flavorless mesons. The Okubo-Zweig-Iizuka (OZI) rule suppresses strong decays of charmonium to light hadrons, which must proceed via exchange of virtual gluons. Very little is known from first principles about this suppression and its dynamical origin. In Figure 2 we show first results for the scalar mesons. We plot the effective masses of the ground and excited states for charmonium χ_{c0} and for the light scalar meson a_0 neglecting the effects of quark-annihilation. The light f_0 particle includes the effects of light quark-annihilation. Its mass is very close to twice the pion mass (dashed line). The horizontal bands are the plateau averages of the effective masses from which we extract the masses of the particles. In [5] small effects of charm-light flavour mixing were observed for the case of the pseudoscalar particles for larger values of the light quark masses.

Ongoing Research / Outlook

Currently we are implementing operators for multi-hadron states to be combined with the optimal distillation profiles. As seen in Figure 2 we need to include two-pion operators in our analysis. Operators for two *D*-mesons will be needed to study decays of hybrid charmonium states. Also states made of gluons only, the glueballs can mix with charmonium states. The correlation functions of glueballs suffer from a severe exponential increase of the statistical noise over the signal with time. To tackle this problem the Research Unit FOR 5269 is developing new numerical algorithms for calculations of glueballs.

References and Links

[1] F. Knechtli et al., Phys. Rev. D 106 (2022) no.3, 034501.
[2] R. Höllwieser et al., Eur. Phys. J. C 80 (2020) no.4, 349.
[3] http://luscher.web.cern.ch/luscher/openQCD/
[4] J. Neuendorf et al., arXiv:2401.05828.
[5] J.A. Urrea-Niño et al., arXiv:2312.16740.

Isospin breaking effects in QCD

Research Institution
[1]Juelich Supercomputing Centre, Forschungszentrum Juelich

Principal Investigator
Kalman Szabo[1]

Researcher
Zoltan Fodor[2], Laurent Lellouch[3], Finn Stokes[4], Balint Toth[2]

Project Partner
[3]Department of Physics, University of Wuppertal, [4]School of Physics, University of Adelaide
[5]Department of Physics, Aix Marseille University

Funding
ERC-MUON

SuperMUC Project ID: pn56bu (Gauss Large-Scale project)

5

Introduction

In this project we calculate some of the most important isospin breaking (IB) effects in the strong interactions. Nowadays large scale computations of the strong interaction, lattice QCD, can easily reach precisions well below 1%. At this level of precision it is neccessary to take into account isospin breaking arising from

1. the presence of the electromagnetic interaction, these are called electromagnetic isospin breaking effects and
2. the mass difference between up (u) and down (d) quarks, as called as the strong-isospin-breaking (SIB) effects.

One prominent example, where such precision is required, is the magnetic moment of the muon. Its experimental value, announced by the Fermi National Laboratory (Fermilab) in 2021, is in a 4.2 σ tension with the traditional, data-driven theoretical determination of the same quantity, indicating new physics and new unknown interactions. On the other hand, our ab initio theoretical determination [1], published on the same day as the Fermilab announcement, agrees with its value on the 1.5 σ level, confirming the standard model of particle physics. On August 10, 2023, Fermilab has announced an updated experimental result [2], which increases the discrepancy with the data driven theoretical determination to 5 σ. In order to distinguish between the possible scenarios, the uncertainties of the lattice determination have to be on the same level as that of the experimental value. Another prominent example where lattice QCD can give valuable input is related to the unitarity of the Cabibbo-Cobayashi-Maskawa (CKM) matrix. Recently, tensions have emerged between these different determinations, commonly referred to as the “Cabibbo anomaly” [3]. Lattice QCD can provide accurate values for the pion and kaon decay constants, which together with the experimental values of the meson decay rates, can be used to give strong constraints on the CKM matrix elements. In this project we compute the strong-isospin-breaking contribution to the decay constants. Combining it with the isospin-symmetric component yields the full QCD results for the decay constant.

Results and Methods

Results on the muon magnetic moment
This project aims at further reducing the theoretical uncertainty by reducing the statistical and systematic errors on the strong-isospin-breaking contributions to the muon magnetic moment. In our earlier work [1], these contributions were obtained using an extrapolation technique: we performed computations with valence quark masses that are some multiple κ of the sea quark mass and then performed an extrapolation to the target point at $\kappa = 1$. We measured the contractions at five different values of κ and used the three lightest of these, $\kappa = 3, 5, 7$, to perform a linear extrapolation. Our goal is to eliminate this extrapolation from the computations and do the measurements directly at $\kappa = 1$. The standard measurement approach is not feasible here: the statistical error is about a factor 7 larger at $\kappa = 1$, than the statistical error of the extrapolation. In order to reach a similar precision, we extend a technique called Low Mode Averaging (LMA), that we were already using in our previous work [1] to reduce the noise on other contributions of the muon magnetic moment. On SUPERMUC-NG we re-measured the strong-isospin-breaking contribution directly at $\kappa = 1$, on all the ensembles that were used previously for this quantity in our work [1]. The comparison of the two approaches is shown in Figure 1. Using the new measurement technique, the systematic error related to the κ-extrapolation is completely eliminated, while the statistical errors are also reduced by approximately 30%.

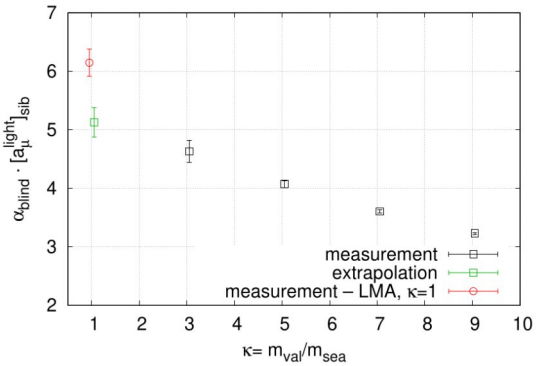

Figure 1: Comparison of our results for the previously used extrapolation and the recently implemented direct thechniques (the latter involves a blinding factor).

Results on the decay constants

The goal of this subproject is to provide the strong-isospin-breaking corrections to the meson decay constants, to leading order in the light quark mass difference. To this order, the strong-isospin-breaking of the pion decay constant is zero. In order to work with dimensionless quantities, we take the combination $f_K w_0$, where w_0 is the gradient-flow based scale setting parameter, whose value we take from our previous work [1]. In order to account for the systematic uncertainties, we perform many different analyses, where we vary the continuum fit form, the number of coarsest lattice spacings omitted from the fit, the effective mass fit ranges for the meson propagators, etc., which gives a total of 9,792 systematic combinations. The combined distribution of the so obtained systematics and the statistical uncertainties, is shown in Figure 2. The major source of uncertainties comes from the continuum limit procedure, where we extrapolate away the details of the lattice discretization. This extrapolation step is shown in the figure below, the value at zero lattice spacing should be taken as our result. The orange points were computed on the SUPERMUC-NG in recent grant periods. Note that the results are not finalized yet, and therefore they are still blinded. This SIB contribution is then combined with the isospin-symmetric values obtained in another project, to yield the full QCD results for the decay constants.

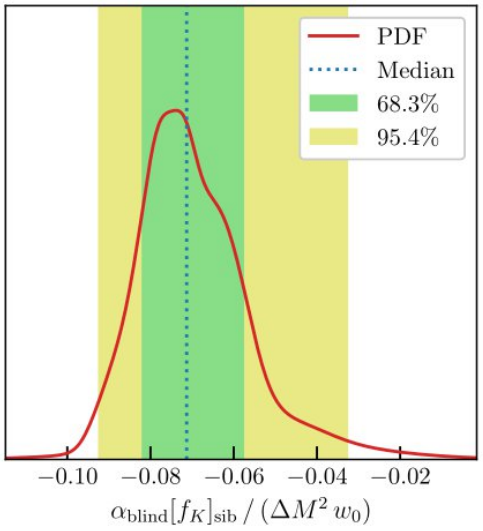

Figure 2: Distribution of the strong isospin breaking of the kaon decay constants. The dashed line gives the central value, the green and yellow bands determine the uncertainty.

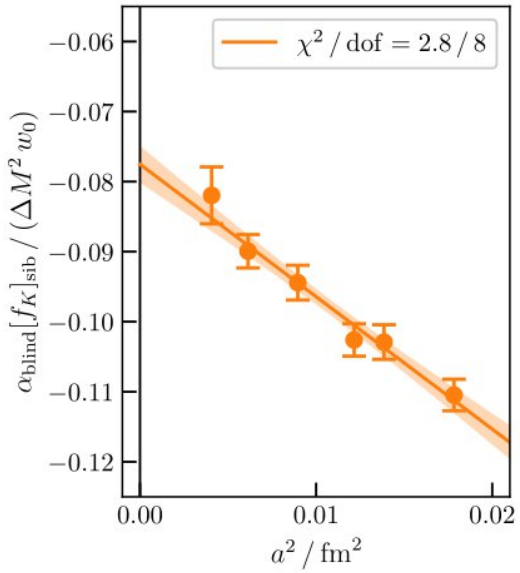

Figure 3: Continuum extrapolation of the strong isospin breaking of the kaon decay constant.

Ongoing Research / Outlook

We are currently finalizing the analyses and preparing for publication in case of both physical observables: the isospin-breaking of the muon magnetic moment and the decay constants. Combining these with the results of state-of-the-art isospin-symmetric lattice QCD simulations, we will be able to set the most stringent constraints on beyond the standard model physics in several experiments.

References and Links

[1] Borsanyi et al Nature 593 (2021) 7857, 51-55.
[2] Albahri et al Phys.Rev.D 103 (2021) 7.
[3] A Crivellin arxiv:2207.02507.

Flavor-singlet meson physics from lattice QCD

Research Institution
[1]Juelich Supercomputing Centre, Forschungszentrum Juelich

Principal Investigator
Kalman Szabo[1]

Researcher
Antoine Gerardin[2], Letizia Parato[3], Willem Verplanke[2]

Project Partner
[2]Aix-Marseille University
[3]ETH Zürich

Funding
ERC-MUON

SuperMUC Project ID: pn73xi (Gauss Large-Sacale project)

Introduction

The muon, a short-lived cousin of the electron, has provided a longstanding discrepancy between the standard model of particle physics and experimental measurements. This suggests that a not-yet-known particle or force perturbs the muon, altering its magnetic moment. The discovery of such “new physics” would have profound consequences on our understanding of Nature. In 2021, the attention of the world was drawn to this discrepancy when the announcement of independent confirmation of the experiment by Fermilab [1] coincided with the publication of our ab-initio calculation (Nature [2]). Our result dramatically updated the theoretical prediction on the magnetic moment of the muon, bringing it significantly closer to the experimental value: it may be possible to explain the Fermilab measurement without any new physics, even with the latest Fermilab update [3]. The experimental uncertainties are expected to be reduced by a further factor of 2 by the Fermilab experiment and a new experiment at J-PARC, Japan, is under construction. The theory error is completely dominated by effects of the strong interaction which can be separated into two distinct contributions: the hadronic vacuum polarization (HVP) and the hadronic light-by-light scattering (HLbL). To fully benefit from the upcoming, more precise experimental results, a precision of 0.2% for the HVP and below 10% for the HLbL is required. In the present work we focus on the latter. We present a first ab-initio calculation of the π^0, η and η' transition form factors performed with physical light-quark masses. We provide a complete parametrization of the form factors that includes both single and double-virtual kinematics. In a second step, our parametrizations of the transition form factors are used to compute the dominant pseudoscalar-pole contributions to the hadronic light-by-light scattering in the muon magnetic moment.

Results and Methods

Lattice QCD
The ab-initio calculation means in the above context lattice quantum field theory. A space-time grid is introduced and 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.

Masses, mixing angles, decay constants
Mesons have played an important role in the development of high-energy physics. These particles were proposed first by Yukawa to describe the nuclear force and subsequently found in experiments. A large variety of such particles were discovered, and eventually they were organized into multiplets by Gell-Mann. These developments ended in the discovery of color and quarks. From all these mesons the so-called flavor-singlet mesons are special. They behave strikingly differently from naive expectations. Let us just mention the anomalous electromagnetic decay of the π^0 or the anomaly in the η-η' mass splitting. The explanation of these anomalies sparked many important theoretical developments. We now know the masses of these particles to a relatively good precision from experiments and we also have a solid qualitative understanding of the mass generating mechanisms. There are several observables beyond the masses of these particles, which have a huge phenomenological importance. In some cases the experimental determination is not precise enough or not yet possible. Here enters the ab-initio approach of lattice QCD to these problems. As the first target, we are computing the π^0, η and η' transition form factors (TFFs). These describe the interaction of these flavor-singlet mesons with two

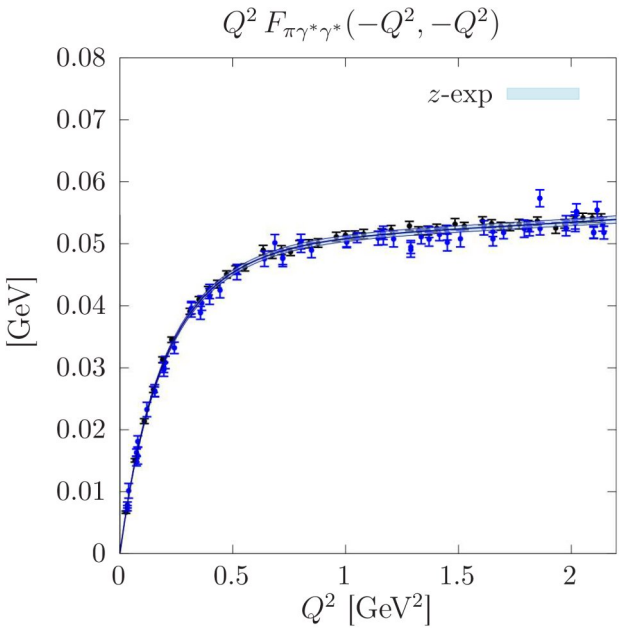

Figure 1: Transition form factor (TFF) of the π⁰ in the doubly virtual regime.

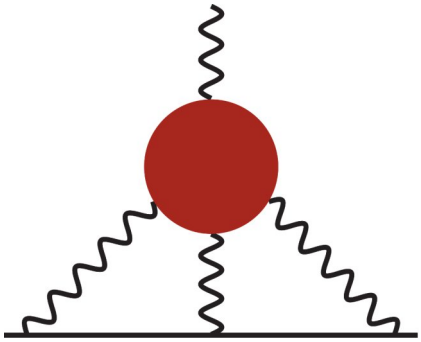

Figure 2: The hadronic ligh-by-light (HLbL) contribution to the muon magnetic moment: The external magnetic field (depicted by the vertical wavy line in the top of the figure) interacts with the muon (shown as the horizontal straight line) through the hadronic blob in the middle and the exchange of three photons.

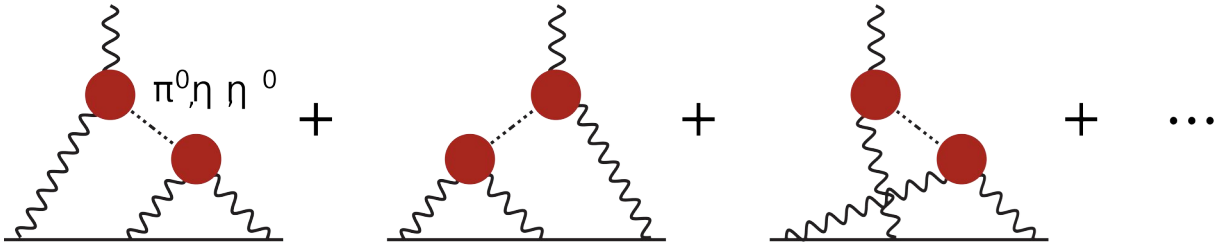

Figure 3: The largest contributions to the HLbL are given by the processes with the flavor-singlet mesons π0, η or η' as intermediate states.

photons. We have given a complete parametrization of these TFFs in [4]. As an example, the doubly virtual TFF of the π^0 is shown in Figure 1.

Hadronic light-by-light scattering

Our second target is to compute the hadronic light-by-light (HLbL) scattering contribution to the magnetic moment of the muon, g_μ. This contribution is illustrated in Figure 2. There is large attention surrounding the muon magnetic moment, with the hope of a discovery of a new physical interaction. The experimental and theory determinations of g_μ have been choosen to one of the **10 Breakthroughs of the Year in 2021** by the Science Magazine [5]. The current picture on the hadronic contributions to the magnetic moment is far from being settled. The HLbL contribution is small, but currently its error is one of the largest in the error budget of the theory determination of g_μ. Independent determinations of the HLbL are therefore highly welcome. Our strategy relies on the so-called dispersive framework: one considers contributions from all-possible intermediate states in the corresponding scattering amplitude (see Figure 3). It turns out that by far the largest contributions among all possible intermediate states are given by the flavor-singlet mesons: π^0, η and η'. These contributions can be given in terms of the transition form factors (TFF). There is very little experimental information on these amplitudes and an ab-initio computation, that we are carrying out, is invaluable. Using our parametrization of the TFFs given in [4], as our final result for the flavor-singlet pole contribution to the magnetic moment of the muon we obtain

$$[(g_\mu\text{-}2)/2]_P^{\text{HLbL}} = 85.1(4.7)(2.3) \times 10^{-11},$$

where the first error is statistical and the second is systematic.

Ongoing Research / Outlook

Future plans include adding large-volume ensembles at smaller lattice spacings for the η and η' transition form factors. These would help to better constrain the normalization of the form factor and to clarify the tension on the $\eta \to \gamma\gamma$ decay rate. The main advantage of large physical volume is to provide a better sampling of the low-virtuality region close to the origin. Reducing the error on the decay rate has a strong impact on the π^0 and η-pole contributions and, to a lesser extent, on the η'-pole contribution.

References and Links

[1] Abi et al, Phys.Rev.Lett. 126 (2021) 14, 141801.
[2] Borsanyi et al, Nature 593 (2021) 7857, 51-55.
[3] Aguillard et al, Phys.Rev.Lett. 131 (2023) 16, 161802.
[4] Gerardin et al, arXiv:2305.04570.
[5] www.science.org/content/article/breakthrough-2021

Life Sciences

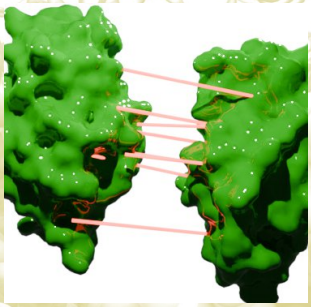

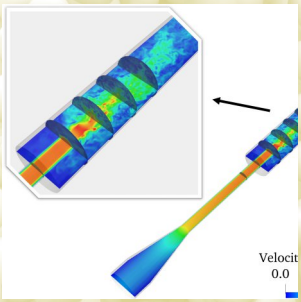

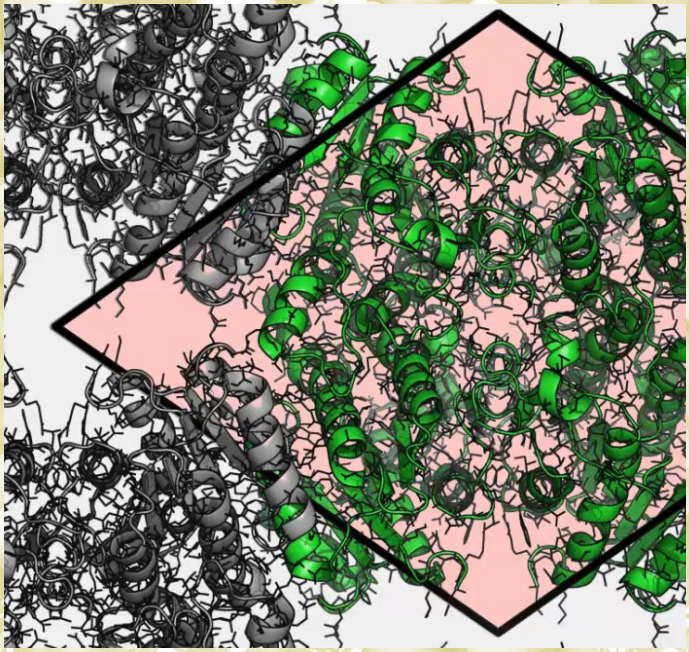

Lattice Boltzmann simulation of hemodynamics in cerebral aneurysms

Research Institution
[1]Laboratory of Fluid Dynamics and Technical Flows, Otto-von-Guericke-University Magdeburg

Principal Investigator
Dominique Thévenin[1]

Researcher
Feng Huang[1], Seyed Ali Hosseini[2]

Project Partner
[2]Department of Mechanical and Process Engineering, ETH Zürich

Funding
China Scholarship Council (grant number 201908080236)

SuperMUC Project ID: pn73ta

Introduction

Intracranial aneurysms, affecting approximately 3% of the western population, pose a serious threat due to the risk of rupture, leading to irreversible disabilities or death. Leveraging the computational power of modern supercomputers, our project employs the lattice Boltzmann method (LBM) to delve into the complexities of hemodynamics within intracranial aneurysms, utilizing our in-house numerical solver, ALBORZ.

Our research addresses the challenge of complex geometry in patient-specific aneurysms by introducing a curved boundary condition, enhancing the accuracy of LBM simulations. Additionally, we propose a robust modified central Hermite polynomial-based multiple relaxation time lattice Boltzmann model for modeling blood flow, challenging assumptions about high shear rates within aneurysm sacs.

A crucial focus of this project lies in the exploration of unsteady flow and flow fluctuations within patient-specific aneurysms. Flow instability emerges as a promising hemodynamic metric for assessing rupture risk, and our study investigates these fluctuations using both Newtonian and non-Newtonian fluid models.

Our findings reveal intriguing insights, particularly in ruptured aneurysms where hydrodynamic instability during pulsatile inflow contributes to high-frequency fluctuations around the rupture location. Importantly, we observe minimal differences between Newtonian and non-Newtonian outcomes in unruptured cases but significant variations in ruptured cases.

This work significantly contributes to understanding hemodynamics in intracranial aneurysms, shedding light on the role of non-Newtonian behavior and flow fluctuations in assessing the rupture risk of these critical medical conditions.

Results and Methods

To validate the accuracy and reliability of our numerical solver (ALBORZ), we initially selected a patient specific aneurysm [2]. Figure 1 presents a qualitative comparison between the hemodynamic simulations based on the lattice Boltzmann (LB) approach and in-vitro phantom measurements. The agreement with particle image velocimetry (PIV) data indicates the solver's capability to accurately capture the flow structure, even in relatively under-resolved simulations.

For benchmarking, we employed the US Food and Drug Administration (FDA) suggested idealized medical nozzle device. The simulation, conducted on 1,500 cores in parallel, utilized a discretized domain of 155×155×3,093 equidistant grid points (74.3 million points). Figure 2 displays the results from the FDA nozzle simulation at Re = 6,500, revealing local turbulence downstream of the sudden expansion region, consistent with both PIV and numerical observations [3].

In addressing the challenging task of assessing the risk of rupture in intracranial aneurysms, we conducted direct numerical simulations on patient-specific cases. Both Newtonian and non-Newtonian fluid models were considered. The domains for the aneurysms were specified as 21.2 × 37.6 × 19.7 mm and 32.5 × 18.4 × 11.1 mm, with corresponding core hours of 8,760 and 3,810. Figure 3 provides a qualitative comparison of flow streamlines within the aneurysm sac, offering insights into dynamic behavior during different phases of the cardiac cycle. This initial qualitative analysis revealed a stronger flow with larger variations in time between ruptured and unruptured cases.

In summary, our validation process and simulation results contribute to the robustness of ALBORZ and pave the way for a detailed investigation into the elusive criterion for assessing rupture risk in intracranial aneurysms.

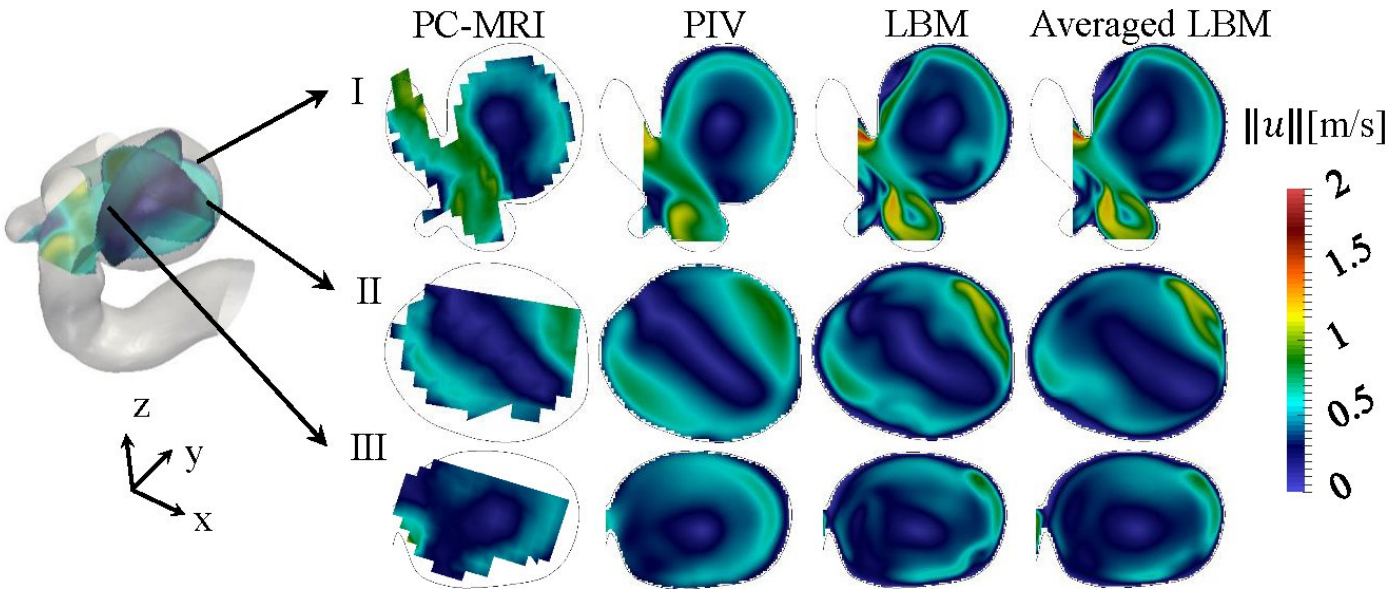

Figure 1: From left to right: Qualitative comparison of the velocity fields acquired using PC-MRI (1st column), stereoscopic PIV (2nd column), instantaneous (3rd column) and time-averaged LBM (last column) simulation in the aneurysm sac on three different planes, i.e. from top to bottom: planes I, II and III.

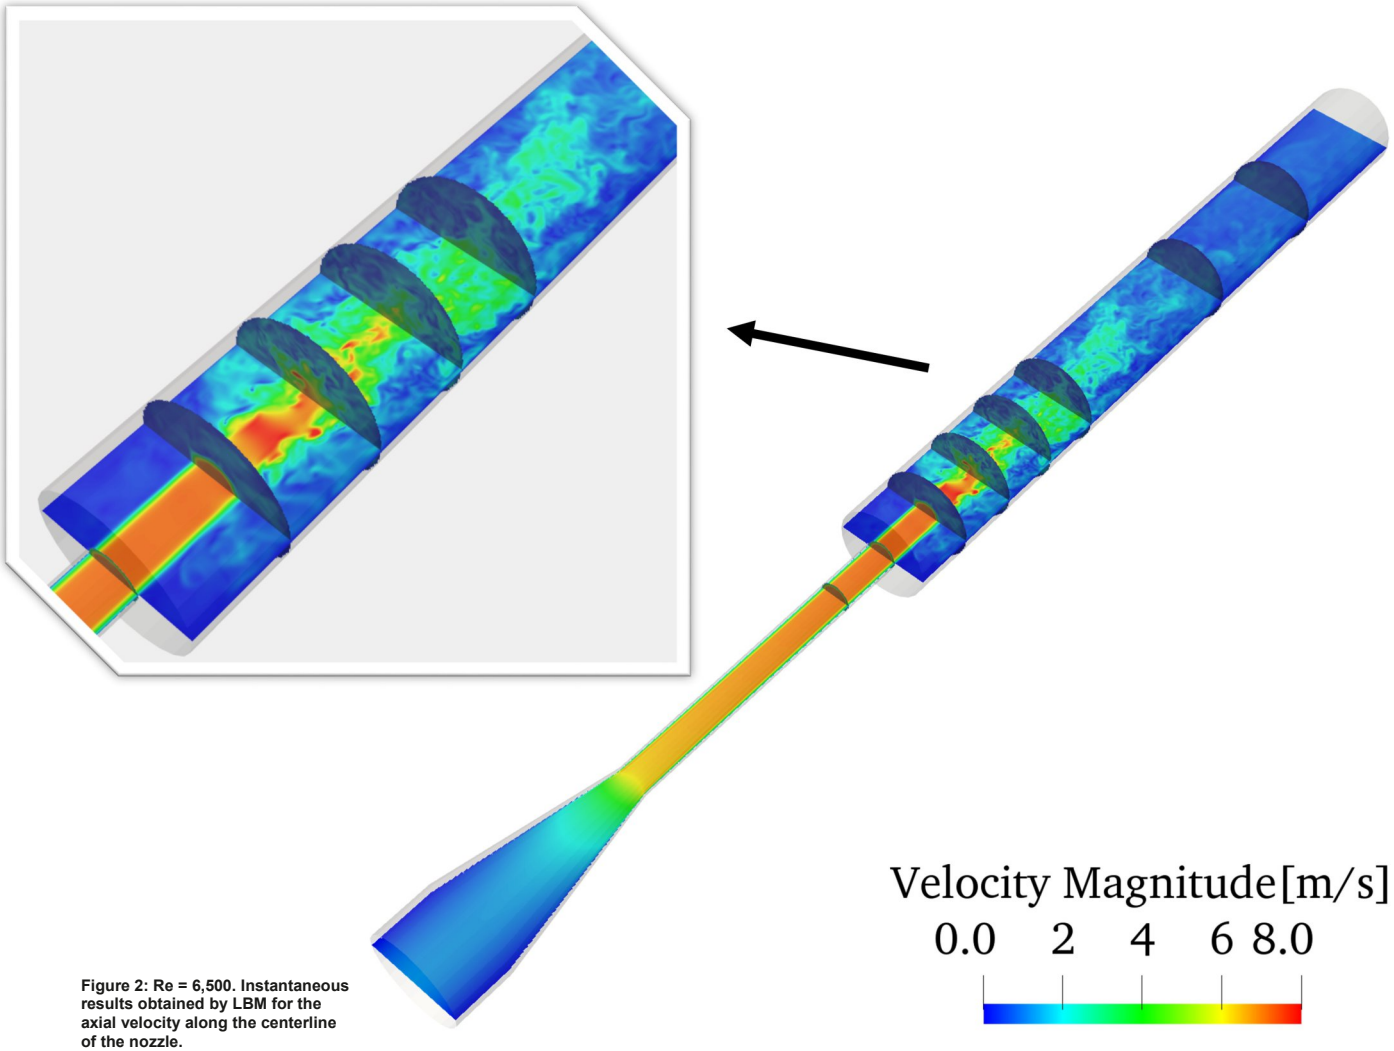

Figure 2: Re = 6,500. Instantaneous results obtained by LBM for the axial velocity along the centerline of the nozzle.

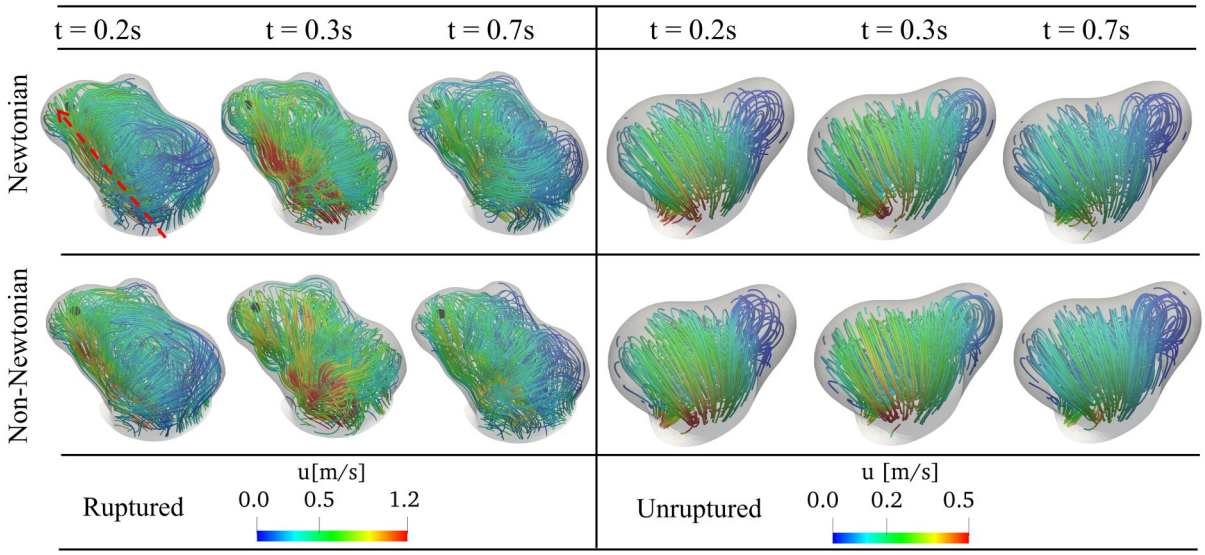

Figure 3: Flow streamlines colored by flow velocity inside the aneurysm sac.

Ongoing Research / Outlook

In the coming phases of our research, we aspire to elevate our numerical procedure to simulate an even broader spectrum of patient-specific cases with heightened accuracy. This evolution entails advancements in resolution, incorporation of realistic geometries, and a nuanced consideration of additional physiological factors. By doing so, we aim to extend our analyses to a more extensive dataset, cultivating a comprehensive understanding of aneurysm behavior that transcends the limitations of smaller datasets.

Furthermore, our future study will extend to encompass numerical simulations involving stents placed within aneurysms. This expansion is motivated by the need to evaluate the effectiveness of existing and emerging medical devices for aneurysm treatment. By simulating the interaction between stents and blood flow, we aspire to assess device stability, discern hemodynamic changes, and explore opportunities for optimization.

In summary, our future research directions aim to not only enhance then accuracy and scope of our simulations but also to explore new frontiers in medical device evaluation within the context of intracranial aneurysms.

References and Links

[1] https://www.lss.ovgu.de/lss/en/
[2] S.A. Hosseini et al., Computers in Biology and Medicine 131(7):104251 (2021).
[3] F. Huang et al., Computer Methods and Programs in Biomedicine 221(1-3):106863 (2022).

Free Energy Calculations for Drug Discovery

Research Institution
[1]Ludwig Maximillian University, [2]University College London

Principal Investigator
Dieter Kranzlmüller[1], Peter V. Coveney[2]

Researcher
Shunzhou Wan[2], Agastya P. Bhati[2], Alexander D. Wade[2], Shantenu Jha[3]

Project Partner
[3]Department of Computer Engineering, Rutgers University

Funding
EU H2020 CompBioMed2, EU H2020 EXDCI-2, DOE INCITE

6

SuperMUC Project ID: pn98ve

Introduction

Free energy calculations are an essential tool by means of which to evaluate the binding affinity of drugs to their target proteins and their applications are becoming increasingly widespread in the pharmaceutical industry and also clinical settings. Both relative and absolute binding free energy (RBFE and ABFE, see Fig. 1) calculations have been used in pharmaceutical companies in the last few years, albert the former is much common than the latter. Despite of the increasing use of the approaches, insufficient attention has been paid to the reliability, reproducibility and uncertainty quantification of the methods used. We have been using ensemble approaches for free energy calculations which are successful in generating reproducible and reliable predictions in conjunction with statistically robust uncertainty quantification [1]. We have shown recently that the ensemble-based alchemical approach TIES (thermodynamic integration with enhanced sampling) [2] generates relative binding free energies with chemical accuracy in all cases for a large demonstration dataset [3]. Here we further assess the performance of non-equilibrium alchemical methods for both RBFE and ABFE calculations, and to examine the impact of simulation parameter settings on their predictions. The study of binding affinities is of immense interest in pharmaceutical industry. To access the compound screening for a drug discovery project, we use the IMPECCABLE (Integrated Modeling PipelinE for COVID Cure by Assessing Better LEads) workflow we have developed at UCL. Our studies deploy both physics-based (PB) techniques including molecular dynamics (MD) based methods, and machine learning (ML) techniques to screen a large set of candidate compounds. It involves performing computer-intensive molecular dynamics simulations and ultra-high throughput docking calculations for improved inhibitor pose selection. IMPECCABLE workflow can analyse several million compounds to provide more accurate binding affinity ranking. We have also applied BAC (binding affinity calculator) to proteins with sequential variations, which has potential to offer a reliable way to select subgroups of patients for a given drug treatment, based on their genetic profiles.

Results and Methods

We adapt TIES [2], the equilibrium method we have developed, to perform non-equilibrium simulations for the calculation of RBFE and ABFE corresponding to an alchemical transformation. A thermodynamic cycle approach (Fig. 1a) is invoked for both the equilibrium (Fig. 1b) and non-equilibrium (Fig. 1c) methods, in which the compound pair is transferred alchemically from one to another in RBFE, or from something to nothing in ABFE, both in aqueous solvent and bound to a protein. We evaluate the non-equilibrium method by applying it to large datasets, comprising 503 ligand pairs bound to a diverse set of fourteen pharmaceutically relevant protein targets for RBFE [4], and 219 complexes of 31 different proteins and 186 compounds with a variety of scaffolds for ABFE [5]. These are the largest datasets used so far for validating the RBFE and ABFE methodologies to our knowledge. For the RBFE calculations, we typically perform 260 ns simulations for each ligand pair, and extend up to 40.4 µs for a subset of the dataset. For the ABFE, we use 1.76 µs for a single free energy prediction, and an overall simulation time of ~2 ns for the entire molecular systems. Because of the high demanding of computational resources for these studies, we have used SuperMUC-NG and other supercomputers we have access to. All of the simulations in these studies make use of ensemble approach; all replicas in an ensemble can be executed concurrently in the wallclock time of a single short run. Because of the multiple sequential steps in non-equilibrium approach, it introduces longer wall clock times, as well as higher computational complexity, than the equilibrium approach. The detailed and systematic analysis of the settings of the simulation parameters in non-equilibrium simulations enables us to make definitive recommendations for the implementation of the methods to minimise the loss of accuracy and precision of the results obtained while keeping computational costs minimised. We have made recommendations on the ensemble size, the number of non-equilibrium transitions from on endpoint to the other, and the simulation lengths of these transitions for the application of non-equilibrium method to RBFE [4] and ABFE [5] simulations. We have evaluated and applied the entire

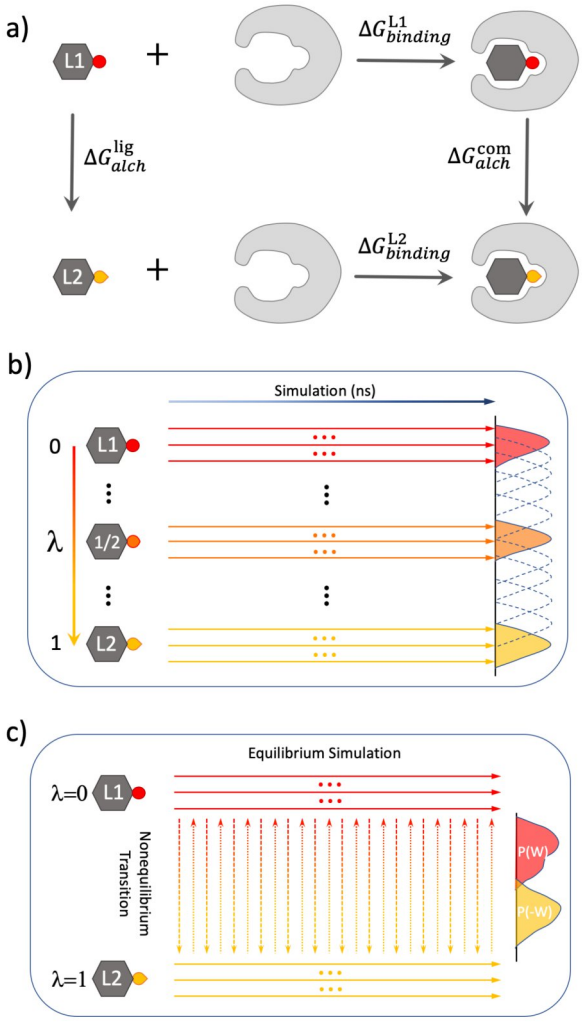

Figure 1: Free energy calculations. a) Thermodynamic cycle for the calculation of binding free energy changes ($\Delta\Delta G_{bind}$). The horizontal arrows correspond to the binding free energy of two ligands, L1 and L2, to the binding site of a protein. For ABFE, one of the two ligands is nothing. The vertical arrows correspond to the alchemical transitions of L1 to L2 in bulk solvent (solvent leg) and in the bound state (complex leg). The free energy changes in the alchemical process can be calculated using b) an equilibrium or c) a non-equilibrium method. Ensemble simulations are performed whenever required.

IMPECCABLE workflow, as well as its individual components, using different protein targets and diverse sets of compounds. The workflow performs a scalable concurrent multi-stage task-execution which requires to use RADICAL-Cybertools (RCT) middleware. We have prepared RCT's virtual environment outside of the system and then moved it to SuperMUC-NG login node. In this way, RCT could be launched from a login node via the pre-set environment, without the need for outbound Internet access. On SuperMUC-NG, we automated the launching of both services on a dedicated compute node, which was provided by a special service queue with unlimited walltime, while the workers for RP were provided by the regular batch system. We have used the IMPECCABLE to find small molecules which have better bind affinities with the target proteins. After a few iterations between the generation of compounds from a generative ML method by AstraZeneca and the binding free energy prediction from ESMACS approach by UCL. We are getting a converging behaviour over the course of iterations, i.e. a downward trend in terms of the free energy. The results we have got so far are

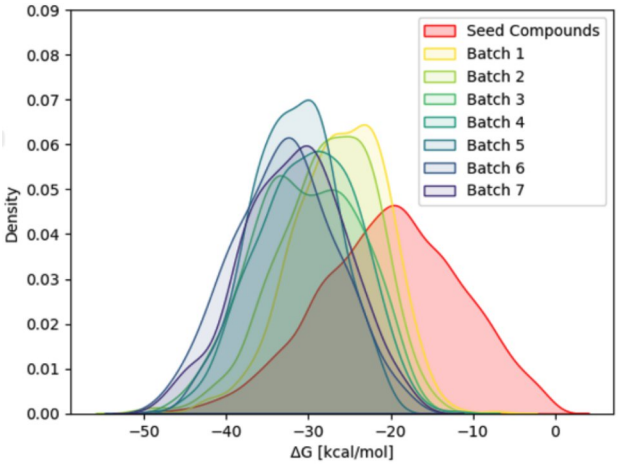

Figure 2: Distribution of calculated *ΔG* from ESMACS for each iteration. The *ΔG* distribution of seed compounds used to train the initial surrogate model is shown in red.

encouraging: we have already seen some improvement in the generated compounds when comparing with those from the previous iterations (Fig. 2).

Ongoing Research / Outlook

Drug resistance is often mediated by mutations in target proteins. We are investigating the relative binding affinity changes associated with amino acid mutations in proteins. We are involved in joint co-design efforts with Arm, Atos, Leibniz Supercomputing Centre (LRZ), UCL and Oxford University in the domain of genomics with a specific application to prediction of mutations in Mycobacterium tuberculosis (TB) which confer drug resistance and on how to treat patients on an individual basis in the face of this. We are using TIES-PM in conjunction with NAMD to calculate free energy changes associated with hundreds of amino acid mutations in RNA polymerase and pncA. Combining the most advanced high-throughput molecular dynamics simulations and associated free energy calculation approaches with rapid sequence analysis technology, BAC and IMPECCABLE have potential to offer real-time support for making clinical decisions on patient stratified drug selection. In collaboration with our Oxford colleague Phil Fowler, we have identified tens of amino acid mutations in RNA polymerase, and hundreds in pncA. These mutations are a subset selected from existing experimental data collected by the Oxford group. We have started the TIES-PM simulations for the dataset, and completed a subset of the identified mutations. The work now underway will predict the free energy changes for all of these mutations.

References and Links

[1] M. Vassaux et al., J. Chem. Theory Comput. 2021, 17, 5187-5197.
[2] A.P. Bhati et al., J. Chem. Theory Comput. 2017, 13, 210-222.
[3] A.P. Bhatiet al., J. Chem. Theory Comput. 2022, 18, 2687-2702.
[4] S. Wan et al. J. Chem. Theory Comput. 2023, 19, 7846–7860.
[5] Bhati, A.P.; Wan, S.; Coveney, P.V. 2024, Equilibrium and Non-Equilibrium Methods for Accurate, Precise and Reliable Absolute Binding Free Energy Calculations at Large Scale. *Preprint*.

AVOCADO - AdVanced fOrmulation using Coarse grAineD mOdelling

Research Institution
Department for Pharmacy, LMU Munich

Principal Investigator
Benjamin Winkeljann

Researcher
Jonas Binder, Tim Sarter, Katharina Steinegger, Wolfgang Frieß, Olivia M. Merkel

Project Partner
–

Funding
ERC-2022-COG-101088587 to O.M.M

SuperMUC Project ID: pn25ma

Introduction

Biologics offer significant therapeutic advancements, but their complex nature poses formulation challenges due to their sensitivity to environmental factors. Molecular dynamics (MD) simulations can accelerate formulation development by investigating biophysical behavior and guiding the design of stable and effective formulations. However, MD simulations face limitations in computational resources and force field accuracy. This report details the use of supercomputing capabilities for large-scale MD simulations, exploring various formulation parameters using coarse-grained (CG) models. Experimental data was used to cross-validate simulation results, ensuring the reliability of insights and demonstrating the potential to streamline the development of optimized biologic formulations.

Results and Methods

Subproject 1: Investigating nucleic acid delivery and interaction with polycationic vectors.
Polymeric nanoparticles for RNA drug delivery commonly range in size from around 10 to 200 nm and self-assemble in an incubation time of minutes to hours. To retrieve comprehensive insights on the structural organization of those particles via MD simulations, sufficient simulation box sizes and extended time spans are prerequisite. The computational power to assess such systems was provided by the SuperMUC-NG.

This setup enabled a screening for influencing factors on particle structure: Firstly, we identified the hydrophilicity (represented by the percentage of lipophilic moieties (% OA)) of our polymers as the main parameter for particle morphology (Figure 1). Secondly, changes in ionic strength and buffer pH were recognized to influence polymer self-interactions and aggregation tendencies of the particles. Identification of these factors allows correlation to wet lab results and trends seen in in vitro experiments. Additionally, it will lead to optimization of simulation protocols to efficiently characterize further particle systems.

Subproject 2: Protein Aggregation at Moving Interfaces.
We developed a CG MD setup designed to replicate the proposed mechanism of protein aggregation at moving interfaces such as tubing during transfer and filling operations. Given the necessity to simulate a high number of large systems for statistically significant results, employing job farming was essential for the success of this project. The majority of these simulations were non-equilibrium MD simulations, demanding increased computational power. Our simulation outcomes correlated with experimental data on subvisible particles and turbidity, thereby validating our model. Using this model, we were able to gain mechanistic insights. For instance, we found that compression velocity does not affect the aggregation behavior of preformed protein films, but rather influences their regeneration. With the possibility of studying the impact of different variables upon compression and dilation at the interface on a molecular level, our model contributes to the understanding of the mechanisms of protein aggregation at moving interfaces.

Subproject 3: Rational design of the antibody ion cloud targeting protein-self-interactions.
Although CG-MD simulations have been shown to provide valuable insights for investigating nanoscale phenomena, commonly used CG force fields, such as Martini 2, often suffer from overestimation of protein-protein interactions. This inaccuracy limits their ability to reliably predict important solution properties (e.g., B_{22}

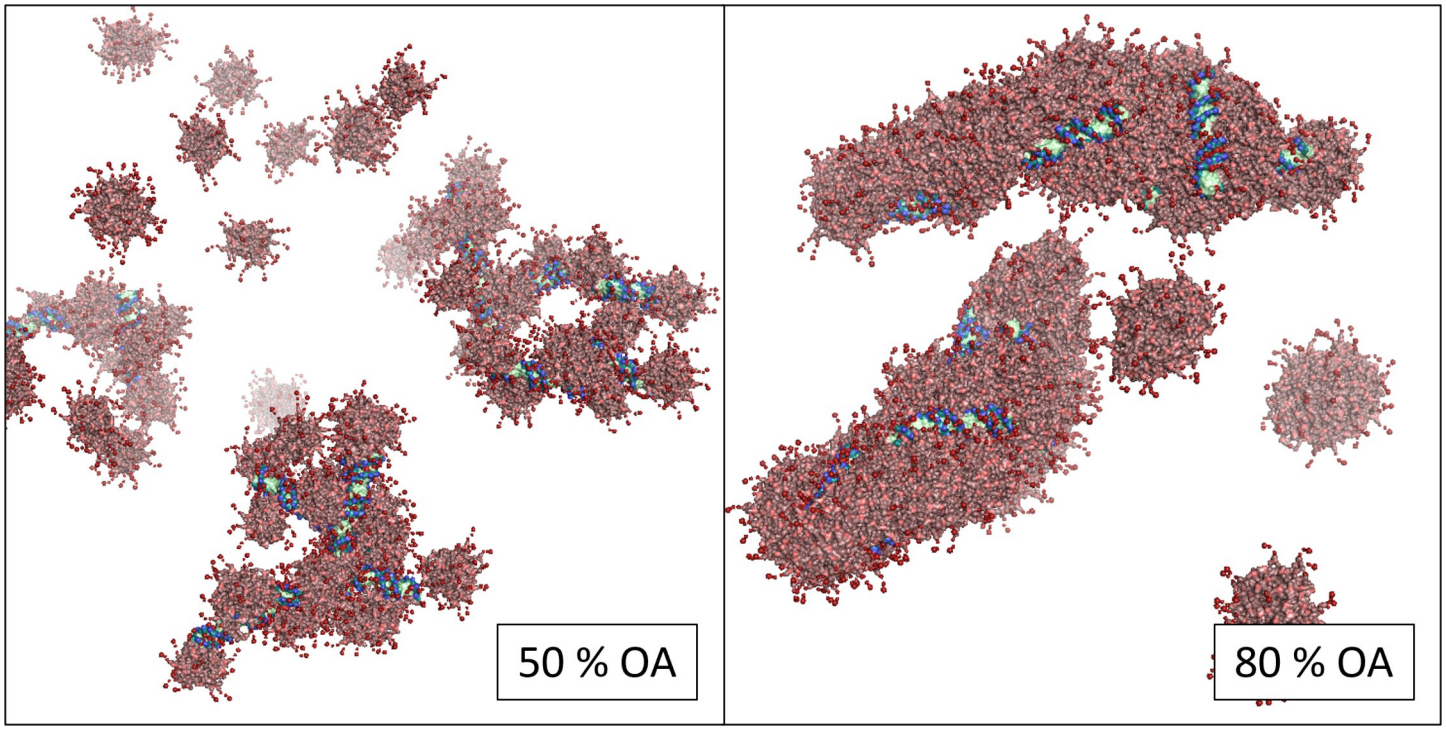

Figure 1: Influence of the lipophilicity of the polymers (%OA) on nanoparticle morphology.

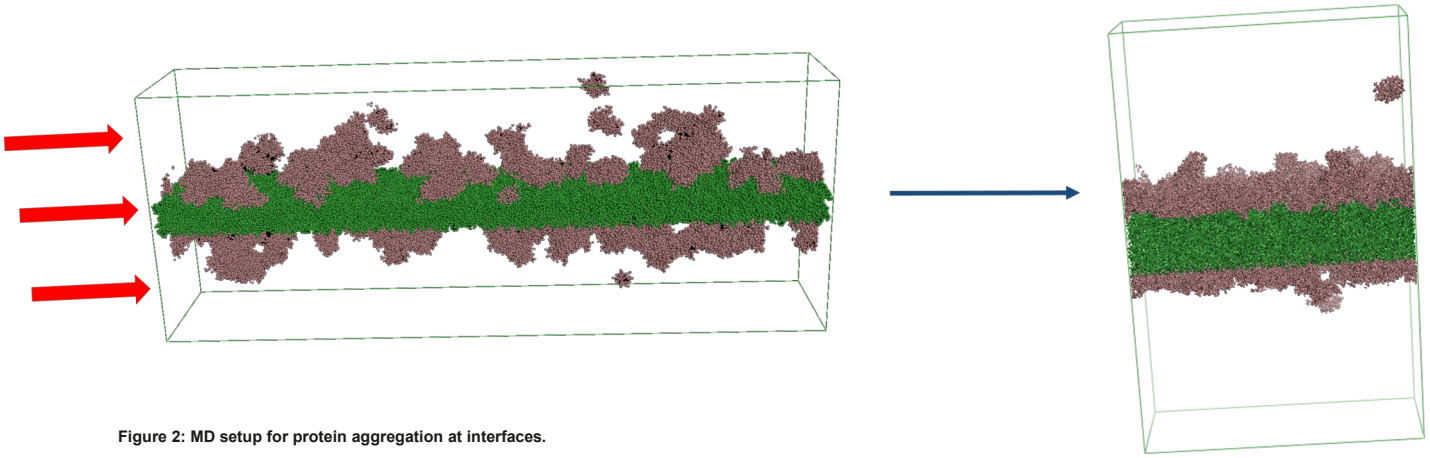

Figure 2: MD setup for protein aggregation at interfaces.

value, viscosity, aggregation propensity), which are critical for computer-aided drug and formulation design. While the newer Martini 3 force field improves modeling, it still exhibits discrepancies in protein-protein interactions, highlighting the need for further refinement to better match experimental data. We developed a novel approach combining metadynamics and umbrella sampling with the Martini 3 force field to investigate protein-protein interactions. By calculating potentials of mean force across various proteins and salt concentrations, we identified systematic overestimations in the force field. To enable this systematic approach, a large number of simulations were necessary, which was enabled by a job farming approach on SuperMUC-NG.

Subproject 4: Simulation of true-sized polyethyleneimine (PEI) for the complexation of RNA drugs.

A CG model for PEI was developed within the Martini 3 forcefield, and MD simulations were executed to investigate the complexation between PEI and siRNA. A code was developed that allows for the generation of input files for MD simulations of PEI at its true size (several kilodalton) and variable degrees of branching. We further proposed a methodology for systematically mapping large molecular structures through fragmentation. Electrostatic interactions were identified as central to the process, with molecular weight and branching of PEI playing pivotal roles. The importance of a combined

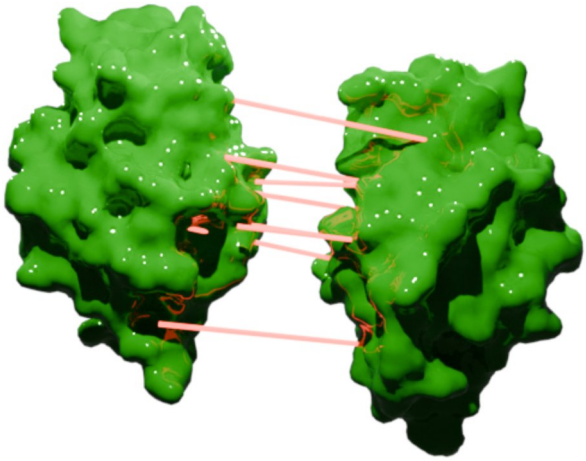

Figure 3: Visualization of the molecular interaction between two lysozyme molecules.

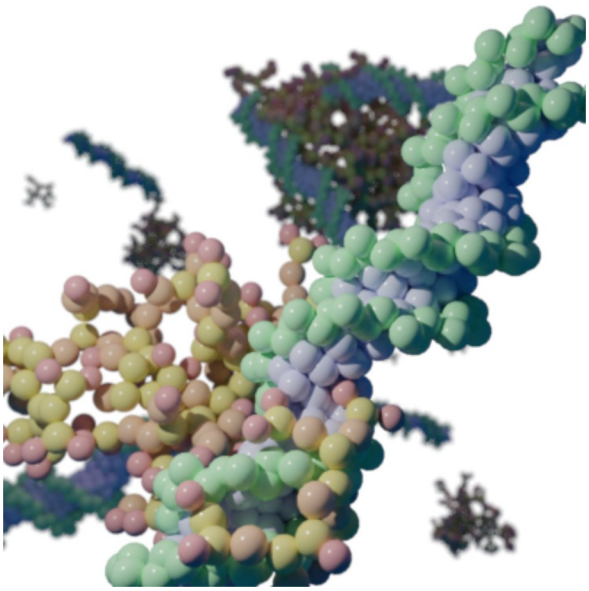

Figure 4: Complexation of siRNA with PEI molecules.

computational and experimental approach in understanding these complex systems is highlighted. On the computational side, we employed metadynamics, much like Subproject 2 and 3. This approach is computationally very intensive, requiring large numbers of cores to calculate the complex free energy landscapes along predefined collective variables.

Ongoing Research / Outlook

While our initial simulations on supercomputing systems have yielded promising results that were validated by experimental data, MD's potential is limited by computational constraints and force field accuracy. To fully explore the complexities of biologic formulations, more simulations covering longer timescales, diverse parameters, and refined coarse-grained models are essential. This demands a significant increase in computing resources and core hours to advance biologic formulation strategies and ultimately benefit patient outcomes.

References and Links

[1] Tim Sarter and Wolfgang Frieß, Molecular Pharmaceutics (2024). https://github.com/jonbind/closingthegap
[2] Jonas Binder, Joshua Winkeljann, Katharina Steinegger, Lara Trnovec, Daria Orekhova, Jonas Zähringer, Andreas Hörner, Valentin Fell, Philip Tinnefeld, Benjamin Winkeljann, Wolfgang Frieß, and Olivia M. Merkel, Molecular Pharmaceutics (2024). https://doi.org/10.1021/acs.molpharmaceut.3c00747
[3] https://github.com/jonbind/closingthegap

Capturing stroke in the circle of Willis

Research Institution
[1]Leibniz Supercomputing Centre

Principal Investigator
Dieter Kranzlmüller[1]

Researcher
Jon McCullough[2], Peter Coveney[2], Elisabeth Mayer[1], Salvatore Cielo[1], Johannes Gunther[3]

Project Partner
[2]University College London
[3]Intel Germany

Funding
CompBioMed2 (EC 823712), CompBioMedX (UKRI EP/X019446/1), SEAVEA (UKRI EP/W007711/1)

SuperMUC Project ID: pn72qu

Introduction

The work of this project 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-scale MRI data. Over the duration of our efforts using SuperMUC-NG we have studied several topics that have advanced our capabilities for the study of personalised vasculatures including methods for coupled arterial-venous simulations, efficient elastic wall modelling, visualisation of large datasets and improved computational performance at scale. In our most recent efforts on SuperMUC-NG, we have focussed on simulating stroke in high-resolution models of the circle of Willis. This vascular structure is located within the brain and plays a crucial role in distributing blood flow throughout the cerebral tissue and, ideally, provides alternative pathways for this in the case of a blockage. However, the presentation of this structure of vessels varies widely between individuals, with only a relatively small proportion possessing the complete structure. This variation means that many people will be at higher risk of negative outcomes from stroke due to physiological limitations in how blood is distributed in the brain. 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

In our study of stroke in the circle of Willis [3], we sought to investigate how a nominally identical stroke would present in individuals with different structures of the circle of Willis. Based on the observed distribution of circle of Willis structures in a population, we studied three forms that represent a combined 31% of the population – the full circle of Willis (Fig. 1) and those missing the left or right posterior connecting arteries. The impact of a blockage of the basilar artery was investigated by changing the transient inflow to this vessel and comparing the observed outlet flows feeding brain tissue to those seen with healthy inflow. We particularly observed that, due to the choice of infarct in the basilar artery, outflow from the posterior cerebral arteries was reduced by the largest amount - up to 50% in the full geometry and almost 70% when a vessel is missing (Fig. 2). As circle of Willis physiology can vary significantly between individuals, our results highlight the benefit of personalised models when assessing stroke risk. SuperMUC-NG was first utilized to generate high-resolution domains of these domains – these consisted of around 170 million lattice sites resolved at 25μm spacing. Each individual simulation consisted of 3 million iterations, this equates to 3 seconds of physical time and captures 3 full heartbeat cycles. Each was run on at least 150 nodes and typically took 8-10 hours to complete. In this work, we also took advantage of our previous work in generating tools for visualizing HemeLB data on SuperMUC-NG using Intel OSPRay [4]. Using this platform, as well as ParaView, we were able to effectively represent the 3D nature of our simulations to highlight the spatial variations in velocity (Fig. 3) and wall shear stress (Fig. 4) variations within our various scenarios of interest. We believe that further development of efficient, immersive visualizations of biomedical simulations will be crucial in advancing end-user comprehension and uptake of modelling results.

Figure 1: The full circle of Willis used in our study. A stroke was introduced by reducing flow in the basilar artery (vessel 1). See [3] for further detail on the cases studied.

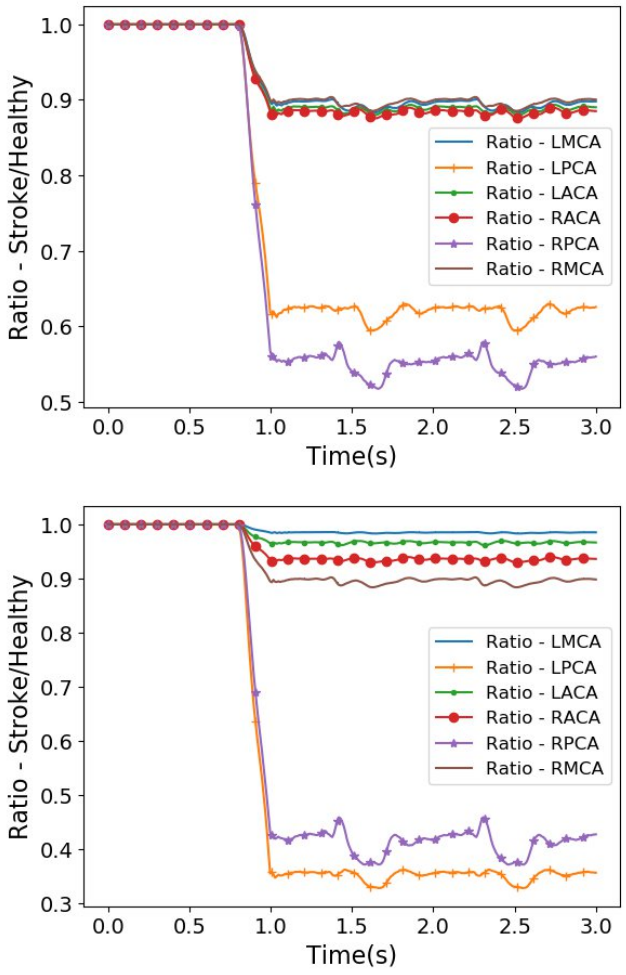

Figure 2: Changes in flow rate seen at the outlets of the circle of Willis due to a complete blockage of the basilar artery: (top) full domain, (bottom) domain missing left posterior connecting artery.

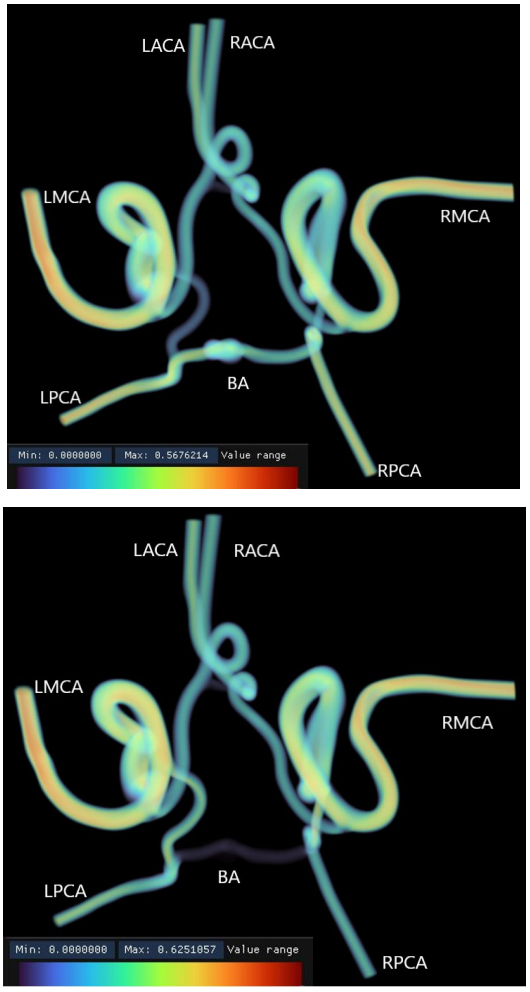

Figure 3: Velocity magnitudes observed within the full circle of Willis: (left) healthy flow and (right) following a blockage of the basilar artery.

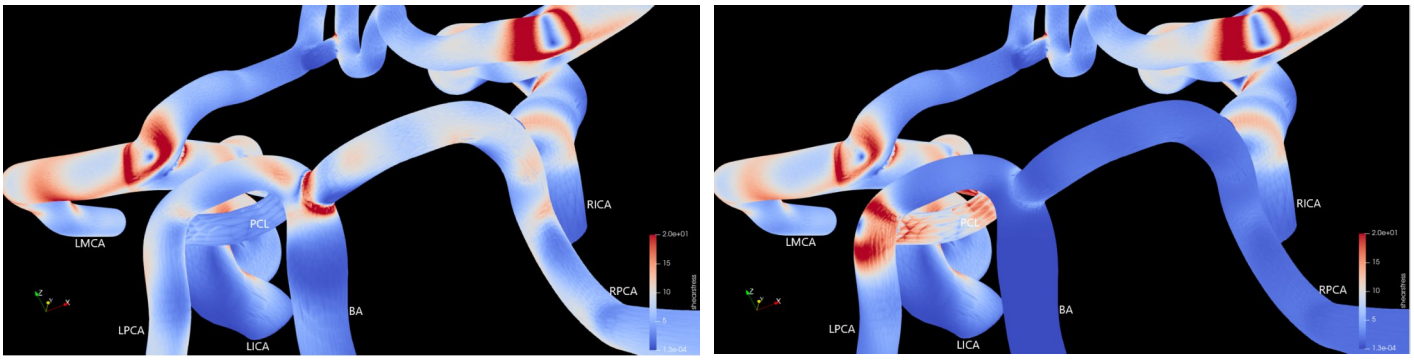

Figure 4: Wall shear stress magnitudes observed at the bifurcation of the basilar artery within the domain missing the right posterior connecting artery: (left) healthy flow and (right) following a blockage of the basilar artery.

Ongoing Research / Outlook

As this work continues to progress, we will continue to develop our flow algorithms and domains to facilitate simulation of full human vasculatures. Significant effort has been expended in porting the HemeLB code to multiple GPU architectures [5] including for Intel such as will be deployed on SuperMUC-NG Phase 2. The use of GPUs is anticipated to better enable simulation and visualisation approaches to be deployed in clinical decision making for personalised healthcare.

References and Links

[1] http://hemelb.org
[2] https://www.compbiomed.eu/
[3] J. W. S. McCullough and P. V. Coveney, Sci. Rep., 13, 21665 (2023). DOI: 10.1038/s41598-023-48776-0
[4] E. Mayer, J. McCullough, J. Gunther, S. Cielo, P. Coveney, SC21 SciViz Showcase presentation (2021), https://www.youtube.com/watch?v=Sd8_cujMP_4
[5] I. Zacharoudiou, J. W. S. McCullough, P. V. Coveney, Computer Physics Communications, 282, 108548 (2023) DOI:10.1016/j.cpc.2022.108548

Enabling genomic-assisted breeding of cattle and chickens in developing countries

Research Institution
[1]Ludwig-Maximilians-Universität München, [2]Queen Mary University of London

Principal Investigator
Laurent Frantz[1,2]

Researcher
Said I. Ng'Ang'A[1,2], Steven Fiddaman[3], Christophe Klopp[4], Philippe Bardou[4], Mathieu Charles[4], Michèle Tixier-Boichard[4], Johannes Geibel[5], Christian Reimer[5], James Alfieri[6], Jacqueline Smith[7], James Ward[8], David McHugh[8], Olivier Hannotte[9], Richard Crooijmans[10], Martijn Derks[10], Gary Vaugh-Smith[11]

Project Partner
[3]University of Oxford, [4]Institut national de la recherche agronomique (INRA), [5]Friedrich-Loeffler-Institut, Greifswald
[6]Texas A&M, College Station, [7]Roslin Institute, Edinburgh, [8]University College Dublin,
[9]The International Livestock Research Institute, Ethiopia, [10]Wageninigen University, The Netherlands
[11]Silverstreet Capital LLP, London
Funding
European Research Council (ERC), Deutsche Forschungsgemeinschaft (DFG), European Commission (Horizon 2020), Science Foundation Ireland (SFI), French National Research Agency (ANR), Global Challenges Research Fund (GCRF), Bill & Melinda Gates Foundation, Silverstreet Capital LLP

SuperMUC Project ID: pn29qe

Introduction

Intensified production has led to remarkable gains in chicken production (5x), growth rate (3x), and milk yield (2x). However, this relies on narrowed genetic diversity in primarily European commercial breeds, increasing their vulnerability to future biotic (disease) and abiotic (climate) stresses. Local breeds, critical for adaptation, are experiencing alarming declines (34% endangered, 5% extinct in 15 years) [1]. This is driven by the adoption of high-yielding European breeds, particularly in Africa. While initially productive, these breeds are often poorly adapted to local environments and lack disease resistance. Our research leverages genomics to understand the impact of industrialization and globalization on livestock genetics. We analyze thousands of diverse genomes to reconstruct history, characterize adaptations, and assess threats like inbreeding. Additionally, ancient DNA analysis reveals how livestock species have evolved, informing sustainable management strategies that balance productivity with genetic diversity and adaptability. This approach, combined with our large computationally intensive genomic dataset, offers the potential to develop genomic-based breeding tools for underutilized breeds. These breeds, while lower-performing than their highly selected counterparts, possess local adaptations that enhance their long-term sustainability. While European breeds offer short-term benefits for non-European smallholder farmers, like in Africa, their vulnerability to drought and local diseases makes them risky investments in the long term, especially considering the vital role livestock plays as savings in communities as in Sub-Saharan Africa [2]. Therefore, we aim to pave the way for genomic-assisted breeding (genomic selection) of local breeds. This technology, successfully used in Europe to improve productivity while limiting inbreeding, requires high-quality reference imputation panels representative of complex livestock (chicken and cattle) ancestries. This will enable the deployment of genomic breeding tools, promoting sustainable livestock management in developing regions.

Results and Methods

This project has leveraged a state-of-the-art pipelines deployed on the SuperMUC-NG to assemble large genomic datasets for chickens and cattle.

Chicken Dataset [3]:
- 4,392 genomes encompassing domesticated chickens, red junglefowl subspecies, and congeneric Gallus species.
- Diverse geographical origins representing commercial birds, experimental lines, and local breeds.
- ~50 ancient genomes reveal past genetic diversity and inform domestication history.

Cattle Dataset:
- 128 newly generated genomes from 32 diverse African populations.
- Integrated with publicly available data for a total of 3,527 genomes.

These comprehensive datasets enable deeper understanding of genetic diversity, adaptation, and evolutionary history in both species, informing sustainable management strategies.

Data processing pipeline
To ensure uniformity across all samples, the entire dataset was subjected to a state-of-the-art genomic pipeline on the SuperMUC-NG computing cluster. The processing of these genomes required ~25 M core hours.

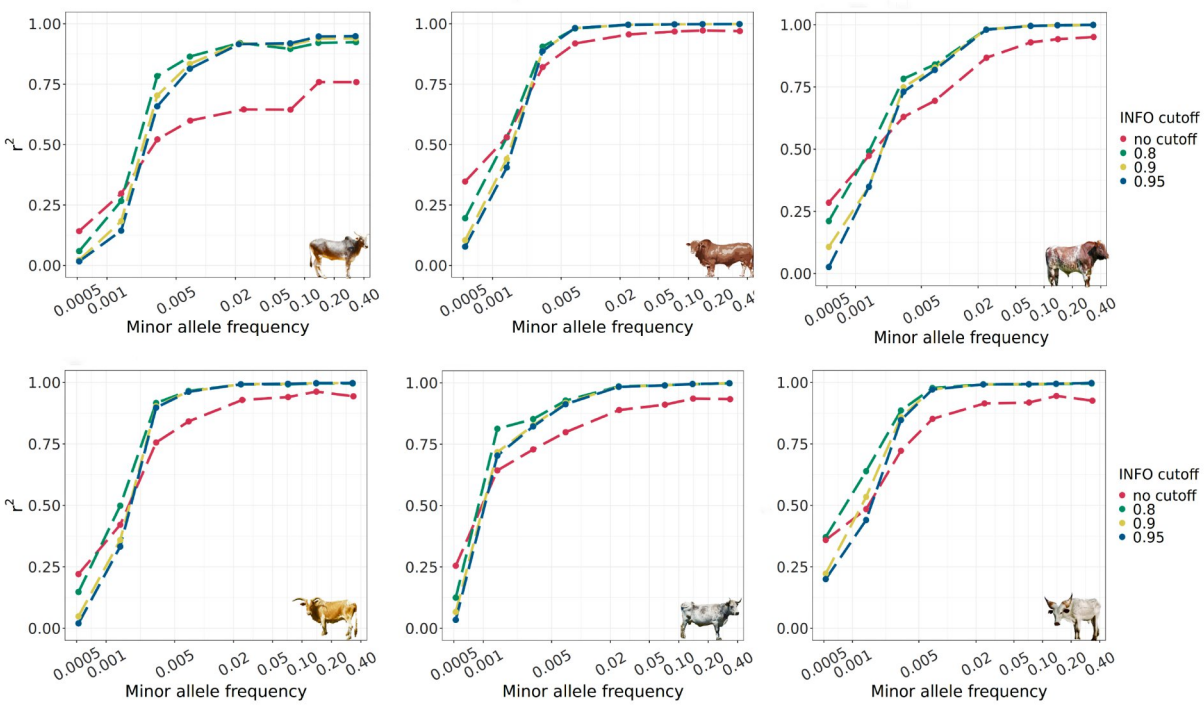

Figure 1: Accuracy of imputation in African cattle breed of 0.5x coverage. The y-axis represents r², with 0 indicating poor imputation quality compared to a 20*x* coverage genome (≈ €350) and 1 indicating equivalent quality. The x-axis displays allele frequency, ranging from 0.001 to 0.5. The findings demonstrate that our imputation pipeline can elevate the quality of a 0.5*x* genome to match that of a 20*x* genome even for rare mutation (found in 0.2% of cattle individuals). The analyzed breeds encompass diverse geographic origins and genetic backgrounds: Hariana (South Asia), Afrikaner (South Africa), Shorthorn (Europe), Kilimanjaro Zebu (East Africa), Creole (South America), and Kuri (West Africa).

Imputation pipeline

Large-scale genomic datasets, processed on the SuperMUC-NG, were merged into reference panels containing tens of millions of genetic variants. These panels enable the deployment of low-coverage genome sequencing with imputation, a cost-effective tool for genomic-assisted breeding (genomic selection) in developing countries. Briefly, imputation computationally predicts missing genotypes based on known variants in reference panels.

Evaluating Imputation Accuracy:

We assessed the performance of our reference panel for imputation. Traditionally, reliable genome data requires 10x coverage, exceeding €200 per sample. Using our imputation pipeline and reference panel, we achieved similar quality with 0.5x coverage. This will reduce the cost of generating genome data to less than €20 per sample. This significant cost reduction empowers animal breeders in developing countries to generate data for improving local breeds. This opens avenues for:

- Sustainable breeding: Selecting animals with desired traits based on their genetic makeup, leading to improved productivity and adaptation.
- Preserving diversity: Characterizing and conserving genetic diversity within local breeds, ensuring their resilience to future challenges.
- Empowering local communities: Enabling farmers to make informed breeding decisions and enhance their livelihoods.

By facilitating cost-effective genomic data generation, this approach holds immense potential for advancing sustainable livestock management and empowering developing communities.

Ongoing Research / Outlook

African cattle genomic data

A manuscript detailing our research on imputation accuracy in African cattle is currently under preparation (Figure 1). Additionally, we have generated over 1,000 low-coverage genomes from diverse African cattle populations, which are currently undergoing analysis in conjunction with ancient genomic data. This combined approach aims to elucidate how African cattle have evolved to thrive in a wide range of ecological contexts.

Chicken genomic data

Our research efforts in chickens parallel those undertaken for African cattle. We are currently: 1) Developing an imputation panel. This panel will enable us to assess the effectiveness of imputation in reducing genotyping costs for chickens, paving the way for more affordable genomic studies in this species. 2) Preparing a large-scale chicken genomics paper. This paper, incorporating ancient genomic data, will investigate how chickens, as tropical birds, have adapted to the diverse environments they inhabit today. These findings will contribute to the development of improved breeding strategies and conservation efforts for this vital livestock species.

References and Links

[1] FAO (2013) FAOSTAT database collections. Food and Agriculture Organization of the United Nations.
[2] Menjo D, Bebe B, Okeyo A, Ojango J (2009) Survival of Holstein-Friesian heifers on commercial dairy farms in Kenya. Appl Anim Husb Rural 2:14–17.
[3] Fiddaman SR, Klopp C, Charles M, et al (2023) Chicken Genomic Diversity consortium: large-scale genomics to unravel the origins and adaptations of chickens. Cytogenet Genome Res.

A Computational Approach for Rational Design of Protein Variants to Improve Crystallization Behavior

Research Institution
[1]Chair of Theoretical Biophysics, Technical University of Munich
[2]Chair of Biochemical Engineering, Technical University of Munich

Principal Investigator
Martin Zacharias[1]
Dirk Weuster-Botz[2]

Researcher
Daniel Bischoff[2], Brigitte Walla[2], Julian Mentges[2], Christina Wegner[3]

Project Partner
[3]Karlsruhe Institute of Technology

Funding
DFG: Research project WE2715/14-2 within the framework of priority program SPP 1934

SuperMUC Project ID: pn68jo

Introduction

Proteins are essential macromolecules in biological systems. Research on the rational design of proteins promises advances in the understanding and improvement of catalytic, sensory, regulatory, and structural functions. Similar to the multifaceted role of proteins in living organisms, the process of protein crystallization has a fundamental role in current research areas, including protein structure elucidation through diffraction, unveiling binding mechanisms of protein interfaces, its involvement in pathologies, as well as its use as a method for protein purification used in industrial downstream processes [1]. However, obtaining high-quality protein crystals remains a challenging problem due to the vast variety of protein crystallization conditions parameters. Thus, the present research project focuses on the computational study of the protein self-assembly process that results in the formation of protein crystals. In contrast to commonly used experimental control mechanisms, this computational study investigates the particular protein-protein interfaces formed in the proteins crystallized phase and the discovery of appropriate amino acid axchanges (mutations) that lead to improved crystallization behavior. This is achieved through comparative analysis of the proteins' solvated and crystallized state on an atomistic scale using Molecular Dynamics (MD) simulations. These MD simulations enable a physics-based representation and time-resolved propagation of solvated and crystallized proteins on a microsecond timescale. Despite the difficult process of setting up accurate protein crystal MD simulations, the finally obtained data sets contain valuable information about thermodynamic properties that support the discovery of mutations for further free energy calculations [2], which in turn result in a quantitative estimate of the predicted crystallization behavior. The accuracy of the predictions is then experimentally tested through site-directed mutagenesis followed by the protein variants crystallization and comparison to the unmodified protein.

Results and Methods

Since the associated time-scale of the crystallization process is not accessible by current MD simulations, the employed alternative predictive model for protein crystallization is based on relative changes in the Gibbs free energy during the crystallization process of different protein variants. By using a thermodynamic cycle as a theoretical model for protein crystallization processes, this approach allowed the derivation of the relative crystallization free energy as an estimator for the crystallization behavior. MD simulations required for this quantity can be performed in the proteins solvated state, as well as its crystallized state [2]. Since these free energy calculations require a comprehensive simultaneous sampling of many parallel alchemical states that need to communicate with each other, the use of the SuperMUG-NG high-performance computer was essential to enable the presented research.

While we were able to demonstrate a promising correlation of the free energy based estimator with the experimentally observed crystallization behavior, the correlation quality was dependent on the knowledge of the protein variants' crystal system, which provided valuable information for the protein crystal simulations in the form of positional restraints. In absence of these specialized positional restraints, the initial approach required the use of positional restraints of the protein crystal simulations with respect to the unmodified protein crystal, resulting in notably worse correlation with the experimental observations. These positional restraints are required, since MD simulations of the crystal phase are difficult to stabilize, as quantified by the root-mean-square deviation to the initial crystal structure obtained through diffraction experiments. Thus, further improvements in the setup of the protein crystal simulations were required to improve the predictive estimates for the most common situation where the protein mutant crystal system is not already known.

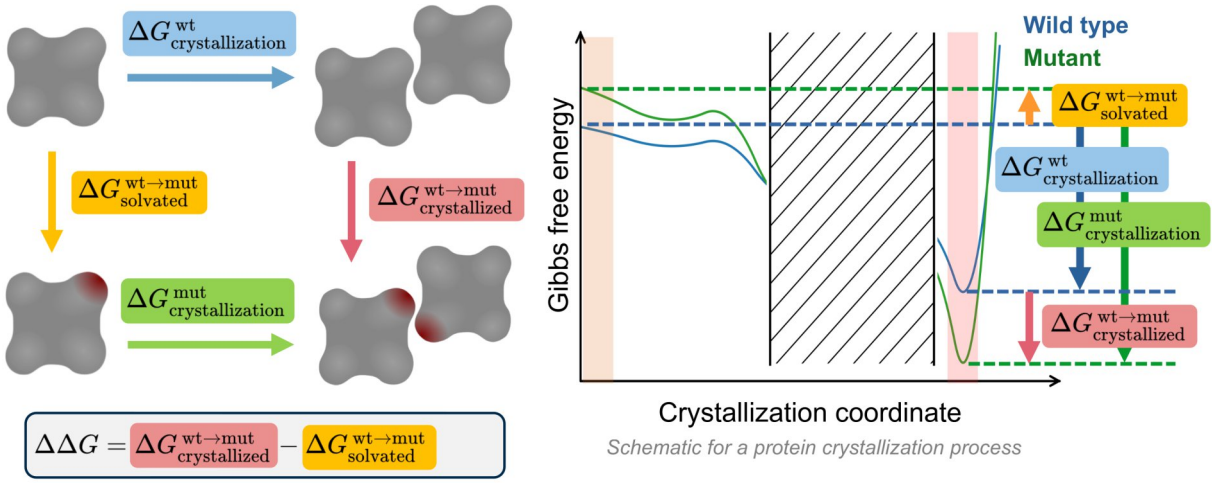

Figure 1: (Right) Schematic representation for a potential protein crystallization process depicting a Gibbs free energy landscape on an imaginary reaction coordinate. (Left) Thermodynamic cycle that demonstrates, how relative free energy differences can be obtained for a crystallization process although the process itself (horizontal arrows) cannot be simulated. Left part of the figure is adapted from Hermann et al. [2].

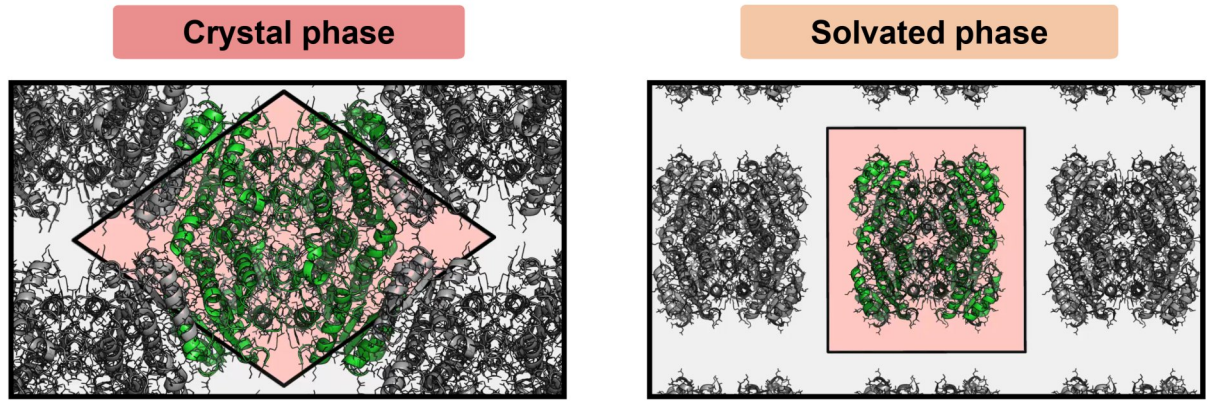

Figure 2: Depiction of the crystallized phase of the protein (left) and the solvated phase (right). The unit cell indicated in red represents the simulated volume that encodes the translational symmetries of an equivalent infinite system. The ability of MD simulations to simulate an infinite periodic system efficiently is particularly suited for protein crystal simulations.

To this end, a large-scale parameter study was conducted with the aim of finding those parameters that influenced the protein crystal simulations stability the most. Among the studied parameters were the size of the simulated protein crystal in units of simulated unit cells (e.g. supercells), the solvent content of the simulation, as well as the different compressibilities for anisotropic pressure coupling. Since the simulation of larger systems comprised of multiple unit cells is computationally very expensive, the use of the SuperMUC-NG high-performance computing platform in combination with the excellent implementation of distributed parallelism provided by the GROMACS MD simulation package enabled efficient screening of a large number of parameter combinations. The use of supercells provided no additional advantage. Similarly, altering the compressibility tensor of the system had no major impact on the protein crystals stability. However, it was found, that the most important factor for stable protein crystal simulations is the exact composition of the solvent content of the protein crystal simulation.

Ongoing Research / Outlook

With the optimized parameter set for protein crystal MD simulations, ongoing work focuses on the use of the predictive quality of the free energy-based estimator for three novel and technically relevant proteins. However, even with the mentioned improvements, the computational demand for a complete screening of mutations at protein crystal contact interfaces exceeds the resource and time constraints for the present project. Therefore, current work focuses on a combination of individual estimates for enthalpic and entropic contributions of single amino acid mutations that are used as prefilters to narrow down the design space for protein variants. The best scoring mutants are then used as candidates for the free energy-based estimator obtained from the optimized MD simulations. Mutants that pass this final criterion result in predictions that are proposed as protein mutants with improved crystallizability. The final predictions are currently experimentally tested.

References and Links

[1] B. Walla et al., Crystals 11.6 (2021) 975.
[2] J. Hermann et al., Crystals 11.6 (2021) 588.

Unraveling the dynamics of guanylate-binding proteins and their membrane interactions

Research Institution
[1]Heinrich-Heine-Universität Düsseldorf
[2]Forschungszentrum Jülich

Principal Investigator
Birgit Strodel[1,2]

Researcher
Jennifer Loschwitz[1,2], Wibke Schumann[1,2], Bastian Bundschuh[1,2]

Project Partner
–

SuperMUC Project ID: pn98zo

Introduction

Toxoplasma gondii is an intracellular parasite that can be lethal to immune-suppressed people or unborn children. To fight off the infection, guanylate-binding proteins (GBPs) are vital. GBPs initiate membrane disintegration by binding to parasitophorous (mouse) or lipopolysaccharide (human) membranes, but the exact molecular mechanisms are unknown. GBPs bind guanylate triphosphate (GTP) which enables oligomerization and large-scale hinge motions [1]. These motions resemble the dynamics occurring in a related group of proteins, called dynamins, whose membrane-remodeling capacities are well studied. This project investigates the similarities and differences in the movements of GBPs and dynamins. To this end, we perform atomistic and coarse-grained molecular dynamics simulations of GBPs.

Results and Methods

We simulate protein motions through molecular dynamics (MD) simulations using the software GROMACS [2] on SuperMUC-NG. Despite employing parallel hardware, it still takes several days to simulate the nano- to microsecond timescales. To extend the time scale, enhanced sampling methods are employed. These can work by selectively heating the system, as in replica exchange MD, or constraining the system in certain ways, thereby forcing it to perform specified motions, as in umbrella sampling MD. Furthermore, the coarse-graining (CG) approach can be used in which four atoms are grouped into one bead, which reduces computational cost by decreasing the resolution and therefore opens up the possibility of simulating larger systems.

One of our research aims is to elucidate how the hinge motion present in dynamins and GBPs can achieve the transition between open and closed protein forms. We already demonstrated that hGBP1, mGBP2, and mGBP7 exhibit a large hinge motion of the M/E domain [3-6]. The bacterial dynamin-like protein (BDLP) has been structurally resolved in both open and closed states. We first explored the space between these extreme states, then simulated 87 intermediate conformations, for 500 ns each, and determined the energetic barrier of the full conformational change. An energy barrier, accessible by GTP hydrolysis, of about 30 kcal/mol was revealed (Figure 1). We simulated the stability of open and closed BDLP monomers and dimers in solution and inserted them in a lipid membrane via the paddle domain. The large systems were modeled at the CG level for 1 µs each and continued after backmapping to the all-atom level for an additional 100 ns. It was found that the open BDLP is only stable when dimerized and membrane-bound, leading to the proposal of a BDLP cycle (Figure 2). This study was published in [6].

We also focused on the membrane binding of mGBP2 and mGBP7 in mono- and dimeric states. Through unbiased CG-MD simulations, we investigated alternative binding poses of mGBP2 and mGBP7 monomers and evaluated the membrane binding capabilities of dimers (Figure 3). As lipid membrane composition we selected DOPC/cholesterol (75%/20%) with 5% cardiolipin (CL) or POPS, and a control without the latter to mimic experimental conditions (Figure 4). The monomer prefers membrane-binding via anchoring the geranylgeranyl group (mGBP2) or the C-terminal (CT) tail (mGBP7) into the membrane. Electrostatic interactions between C-terminal R/K residues, which are independent of the lipid composition, support the membrane binding.

In addition to these expected binding modes, two alternative binding poses were observed, via the G domain or the tip of the M/E domain, which might be better compatible with polymerization and membrane fusion. The statistics indicates that the alternative binding poses are more likely than the insertion of the GG anchor or CT tail for the dimers. This can be explained by the orientation of the proteins within the dimer, where only one of the GG anchors or CT tails is available for membrane binding. We conclude that the G domain dimer must undergo domain motions or even assemble differently, which is in principle possible as demonstrated by an integrative approach involving MD simulations (done within this project), XL-MS and SAXS experiments [7], to enable membrane-bound mGBP2/7 multimers.

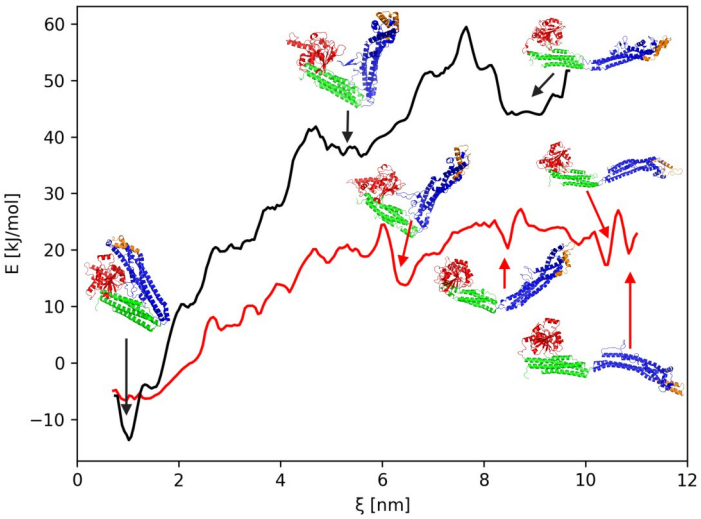

Figure 1: Free energy profile of the closed-to-open transition of apo- (black) and holo-BDLP (red). Allosteric effects triggered by GTP binding lower the free energy for opening.

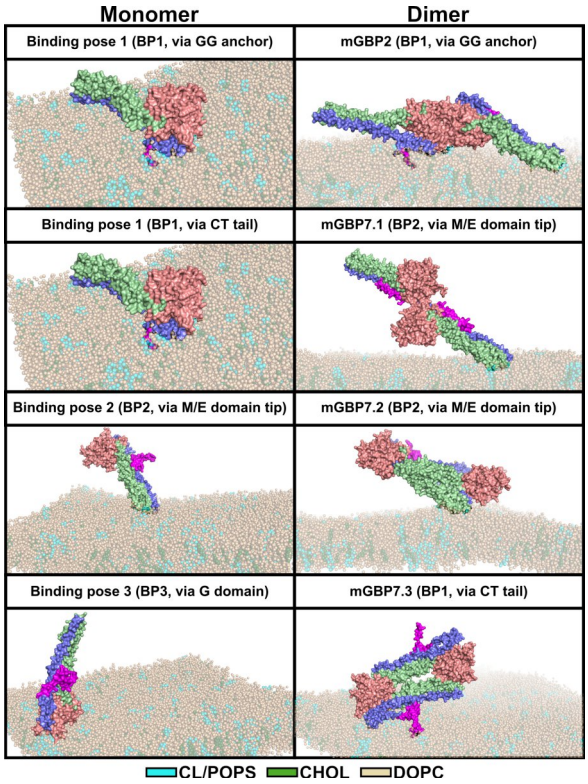

Figure 3: CG-MD Simulation results for mGBP2/7-membrane systems. Membrane binding poses uncovered via the geranylgeranyl anchor and the CT tail for the monomers and dimers.

Ongoing Research / Outlook

In the future, we will extend our knowledge on BDLP to analyze how GTP binding changes the energetic profile of the close ↔ open transition of hGBP1. Moreover, we will analyze which other factors, like dimerization and membrane binding stabilize the open form of hGBP1. In our expectations, this should lead to an hGBP cycle, that explains the defense mechanisms against pathogens, like T. gondii. Another focus in the coming period will be on the binding modes of mGBP2/7 on different membranes in combination with PIP lipids (Figure 4), to uncover determinants for mGBP2/7-membrane interactions and prediction of membrane-bound mGBP2/7 multimer

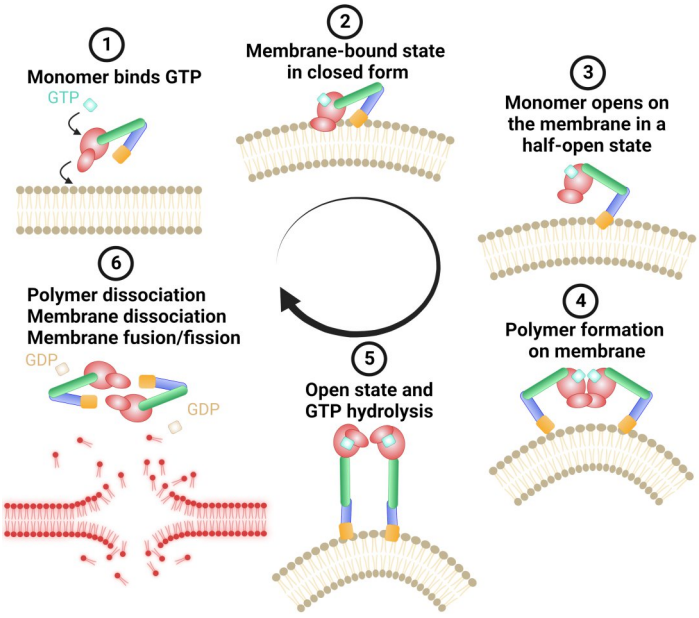

Figure 2: The BDLP cycle starts with the binding of GTP (1) and adsorption to the membrane (2). The half-open state is reached when the paddle is impeded by the membrane (3) and stabilized upon dimerization (4). The process can be repeated for thousands of BDLPs, raising their density and inducing membrane curvature (5). The cycle ends with the hydrolysis of GTP, detachment of the BDLP polymers, and thus membrane fusion or fission (6).

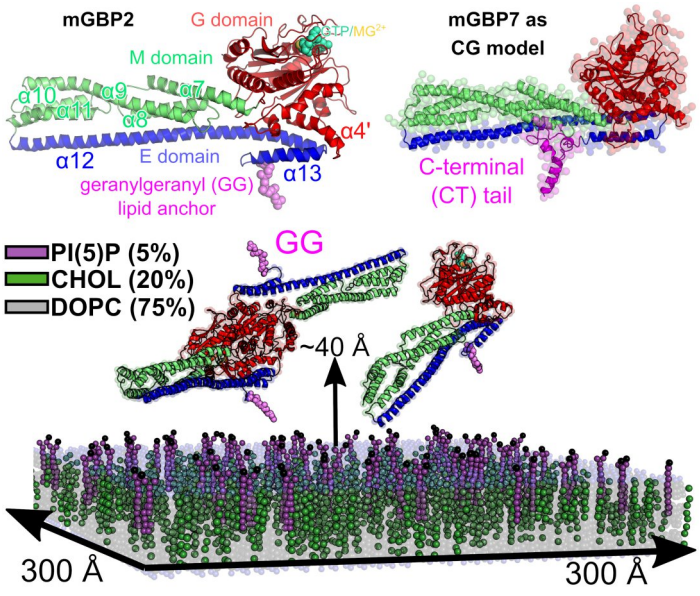

Figure 4: mGBP2/7 on a PI(5)P membrane. While mGBP2 has a geranylgeranyl (GG) lipid anchor, mGBP7 exhibits an additional 49 more residues helical structure, called CT tail.

structures. We aim to consolidate all structural, biochemical, and biophysical knowledge of GBPs to develop a mechanism of the whole GBP cycle, including defining details between species. Thus, the overall goal is to provide a mechanistic understanding of the GBPs' roles in pathogen defense. The current compute project has been granted until 06/2024 and 95% of the allocated resources have been used to date.

References and Links

[1] E. Kravets et al., Elife (2016) 5:e11479.
[2] M. Abraham et al., SoftwareX, (2015) 1–2:19–25.
[3] Legewie, et al., Biochemical journal (2019) 476:3161-3182.
[4] Barz et al., PLOS Computational Biology (2019) 15:e1007193.
[5] Loschwitz et al., Scientific Reports (2023) 13:679.
[6] Schumann et al., bioRxiv (2023) 10.1101/2023.01.16.524228.
[7] Schumann et al., Protein Science (2023) 32:e4818.

Plasma Physics

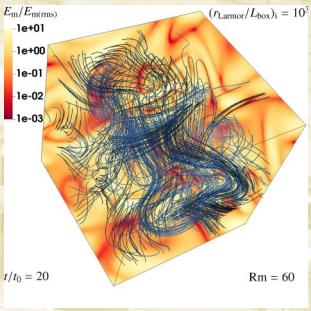

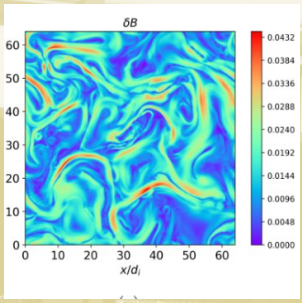

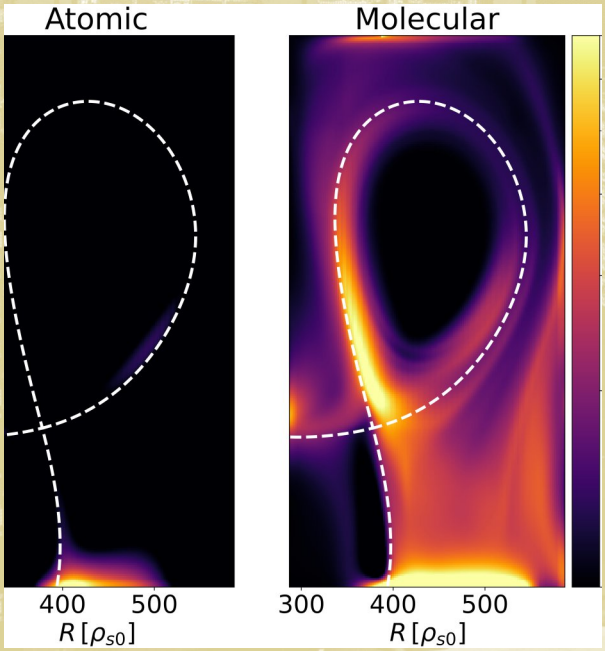

First-principles turbulence simulations of the tokamak boundary

Research Institution
[1]École Polytechnique Fédérale de Lausanne (EPFL), Swiss Plasma Center
[2]École Polytechnique Fédérale de Lausanne (EPFL), SCITAS

Principal Investigator
Paolo Ricci[1], Gilles Fourestey[2], Christian Theiler[1]

Researcher
Davide Galassi[1], Kyungtak Lim[1], Davide Mancini[1], Diego Sales de Oliveira[1]

Project Partner
–

SuperMUC Project ID: pn49ja (PRACE project ID 2021250044)

Introduction

Based on practically inexhaustible resources, nuclear fusion is one of the few options for a sustainable energy source with no release of greenhouse gases. Magnetic confined fusion in tokamaks is nowadays the most developed path towards a fusion reactor. In this concept, the fusion fuel, which is in the plasma state, is magnetically confined in the center of a donut-shaped vacuum vessel. This is the configuration adopted for the ITER experimental reactor, under construction in southern France, and is being explored for DEMO, the first demonstration fusion power plant. The success of the fusion program is closely tied to the plasma dynamics of the tokamak boundary. The plasma-solid interface is set in this region, with the magnetic geometry diverting particles and energy flows that escape the confined plasma to the so-called divertor targets. Plasma detachment, an important process to protect the divertor targets in future fusion reactors, also takes place in the tokamak boundary. In detached conditions, the plasma in front of the divertor targets is cooled by radiative losses with neutrals and/or impurities, mitigating the heat load on the solid plates. Although of crucial importance, the understanding of plasma dynamics at the tokamak boundary is still incomplete. The reason is largely due to the complex dynamics emerging from the nonlinear interaction of turbulence with the multiphysics process, for example, large scale flows and intricate atomic processes. Turbulence appears in the form of coherent structures elongated along the magnetic field-lines, transporting particles and energy in the crossfield direction. Global, self-consistent, and realistic turbulence simulations are key to address the complex and inherently 3D nature of the plasma dynamics at the tokamak boundary in order to provide a first-principles understanding. Thanks to progress in numerical schemes and computational power, such simulations are now possible.

Results and Methods

Global simulations of plasma turbulence of the tokamak boundary are carried out with the GBS code [1,2,3]. GBS evolves in time the drift-reduced Braginskii equations, a set of fluid equations for electrons and ions in a strong magnetic field. The boundary conditions to describe the interaction of the plasma with the solid wall were rigorously derived at the Swiss Plasma Center. The code solves at the same time a kinetic model for the evolution of neutrals, therefore, computing self-consistently, sources and sinks of plasma particles. There is no separation between equilibrium and fluctuation quantities, ensuring all turbulence scales are addressed, only limited by time and grid resolutions. GBS uses a spatial discretization based on a non-field aligned fourth-order finite difference scheme on a regular grid, allowing simulations with magnetic geometries imported directly from the experiment. Time stepping is performed by using an adaptive scheme based on the standard 4^{th} order Runge-Kutta advancement. The source code is written in Fortran 2003 language and parallelized using a 3D domain decomposition with the Message Passing Interface (MPI) library. An iterative solver for the Poisson and Ampere equations is implemented using the PETSc library, which improves the scaling of GBS and includes electromagnetic effects. A tokamak simulation using 128 nodes requires approximately 107 corehours to simulate 50 time units, the typical characteristic time of the global evolution of the plasma in the tokamak boundary. The numerical output for each run is stored in an HDF5 file, containing the snapshots of the plasma and neutral fields every 50,000 time steps, or 0.1 time unit, and whose size is ~10 GB. For the neutral part, a memory of 33 GB/node is required for the 128 nodes simulation and 40 GB/node for the 16 nodes simulation, which is largely within the SuperMUC-NG capacity.

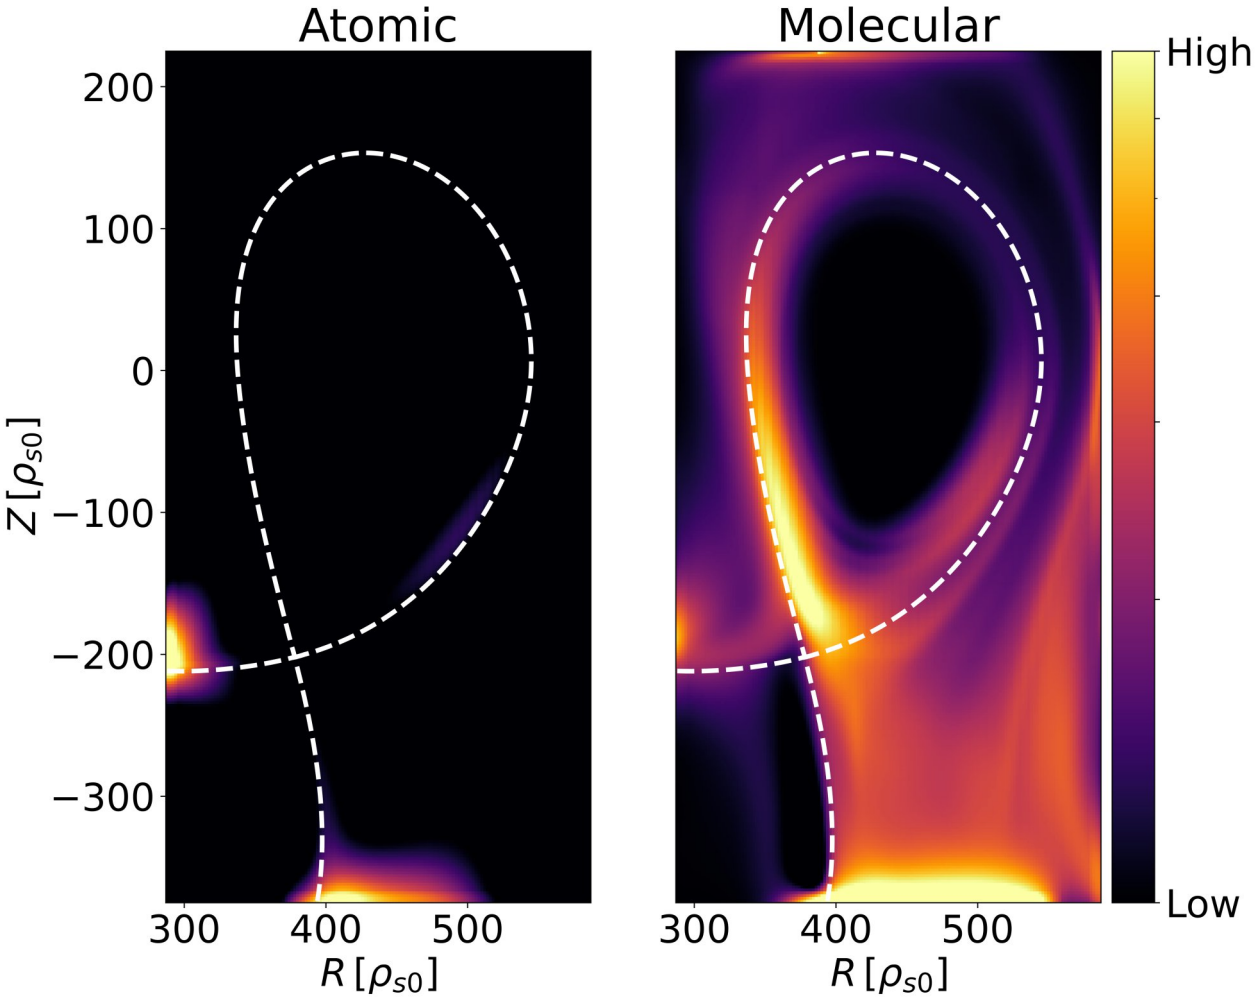

Figure 1: Cross-section of TCV tokamak from GBS simulations showing the volumetric losses obtained with the atomic (left) and molecular (right) kinetic models. The dashed line separates the confined and boundary regions.

Simulations of the TCV-X21 [4] scenario including atomic and molecular deuterium were performed. The TCV-X21 is a diverted L-mode deuterium plasma scenario, specially designed for the validation of turbulence simulations developed at TCV (Tokamak à Configuration Variable), located at the Swiss Plasma Center and one of the main devices of the EUROfusion consortium. The simulation with atomic deuterium reproduced well the density profile at upstream locations, while broader profiles than in the experiment were predicted at the target plates. The molecular dynamics revealed to be essential in the plasma dynamics near the targets. Molecules lead to a large momentum and energy dissipation and, therefore, to a large mitigation of the heat flux. This large dissipation proved essential to reproduce the detachment of the high-field-side divertor observed in experiments. To our knowledge, these turbulence simulations are the first to include the self-consistent treatment of a multispecies plasma with kinetic neutrals [5].

Ongoing Research / Outlook

The multiscale nature of the physics in the tokamak boundary requires simulations with a high number of degrees of freedom and spawning time series of a few times larger than the turbulence characteristic times. SuperMUC-NG is one of the few supercomputers in the world able to perform large-scale simulations as the ones in this project, in a human time scale of a few months. High and low density TCV plasmas depicting an elongated divertor leg and double-null configurations, with four divertor targets, are being simulated and will be rigorously validated. These plasma configurations show improved power exhaustion handling compared to standard divertor geometries and are being considered for DEMO. Along those lines, simulations investigating the interplay between magnetic geometry and turbulence are currently undergoing. These simulations will represent an important further step in the development of reliable tools for understanding the plasma dynamics of the tokamak boundary, allowing prediction and optimization of future fusion reactors.

References and Links

[1] The GBS code for tokamak boundary simulation: https://gbs.epfl.ch
[2] P. Ricci et al Plasma Phys Contr Fusion, 54 (2012) 124047.
[3] M. Giacomin et al. J Comput Phys, 463 (2022) 111294.
[4] D.S. Oliveira et al Nucl. Fusion, 62 (2022) 096001.
[5] D. Mancini et al Nucl. Fusion 64 (2024) 016012.

Studying multi scale processes in space plasmas via numerical simulations: a turbulent example

Research Institution
[1]Institute for Theoretical Physics I (TP1), Ruhr University Bochum

Principal Investigator
Maria Elena Innocenti[1]

Researcher
J. Dargent[1], A. Micera[1], K. Schoeffler[1], E. Boella[2], G. Lapenta[3], G. Arrò[4]

Project Partner
[2]University of Lancaster, [3]University of Leuven, [4]Los Alamos National Laboratory

Funding
SFB 1491, DFG 497938371

SuperMUC Project ID: pn49yi (Gauss Large-Scale project)

Introduction

Our work aims at addressing the fundamental links between the very large and very small scales in heliospheric plasmas.

The heliosphere is the cocoon in interstellar medium dominated by activity originating on our Sun. Earth itself, located 1 Astronomical Units (AU) away from the Sun, is part of the heliosphere. Almost the entire heliosphere is filled by plasma, from the terrestrial magnetosphere a few hundred kilometers on top of our heads, to the solar wind, up and including the solar atmosphere.

Plasmas are ionized gasses characterized by collective behavior. This means that plasma particles, electrons and protons, do not simply follow electric and magnetic fields passively, but contribute to modify them self-consistently. Another fundamental characteristic of plasmas is their multi-scale nature: large and small scale processes are tied non-trivially, and interact with each other in fundamental and sometimes unexpected ways. This is both fascinating and very challenging to investigate via numerical simulations. If one chooses (as we do) to describe via simulations all the processes that can occur in a plasma, without approximations, one needs to resolve the very small scales where waves and particles interact. If, at the same time, one wants to capture the large-scale drivers and consequences of these same wave/ particle interactions, one has to simulate extremely large domains. To give an example of what “small” and “large” means: in the solar wind, “small” scales (the electron skin depth) are of the order of the kilometer. (Relative) “large” scales (hundreds or thousands of proton skin depths), instead, are 10,000 or 100,000 times bigger! [1]. From these numbers, one can see that the computational cost of simulations quickly skyrockets, hence the need of using some of the best HPC systems in the world, such as SuperMUC-NG.

One of the most fascinating multi-scale space plasma processes which we study is turbulence. Turbulence is a term used to indicate a seemingly random and chaotic behavior typical of fluids and plasmas. Most heliospheric plasmas live in a turbulent state. We call turbulence a multiscale process because its very defining property is the ability to move energy across scales: the energy injected into the system from large scale structures ends up in small scale structures. Simply put, in turbulence, very big eddies produce very small structures. The latter are typically associated with extremely energetic phenomena and are responsible for dissipation, i.e. for the conversion of turbulent kinetic energy into heat. Understanding how turbulence dissipates energy is one of the most important unsolved problems of modern physics. The study of turbulence is particularly challenging when dealing with plasmas, and this is because their dynamics is coupled to electromagnetic fields. These are generated and sustained by the plasma itself, and eventually become turbulent too, producing small scale fluctuations that strongly interact with charged particles, accelerating and heating them, thus causing dissipation. Hence, in order to solve the problem of dissipation in plasmas, it is necessary to understand the nature and properties of the small scale electromagnetic structures produced by the turbulence.

When studying turbulence in the heliosphere we have the great advantage of decades of in situ spacecraft observations, the most recent ones coming from the Parker Solar Probe and Solar Orbiter missions, which are flying towards the Sun right now. Spacecraft observations have revealed the presence of a wide variety of small scale electromagnetic structures in the solar wind. We use these observations to set up our simulations and as a benchmark for our results.

Among the different structures observed by satellites, there are magnetic holes, small scale magnetic field dips, associated with intense particle heating [2]. The origin

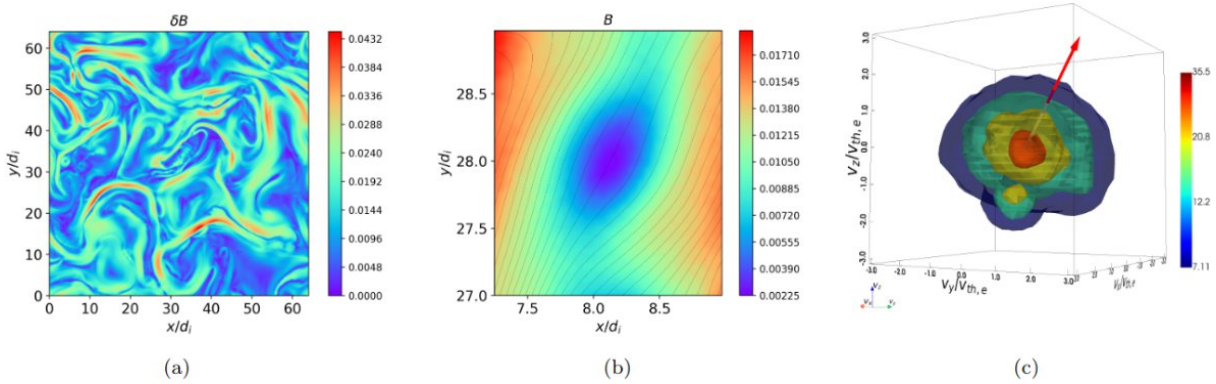

Figure 1: Turbulent magnetic field (a), magnetic hole (b) and electron velocity distribution inside the structure (c).

of magnetic holes is still under debate and we investigate their formation and properties by means of numerical simulations of plasma turbulence.

Results and Methods

We performed large scale particle-in-cell (PIC) simulations of plasma turbulence with parameters inspired by satellite observations. The PIC method consists in following the evolution of a huge amount of charged particles (electrons and protons) under the influence of self-generated electromagnetic fields. In order to obtain results comparable with satellite observations, very large simulation domains, containing a huge amount of particles (of the order of tens of billions of particles) need to be employed. These requirements make our simulations extremely expensive from a computational standpoint, and such numerical experiments can be performed only on large HPC systems, such as SuperMUC-NG.

As the turbulent plasma evolves in our simulation, we observe the formation of a wide variety of small scale structures, including magnetic holes. Thus, we find that plasma turbulence naturally produces magnetic holes, which explains why such structures are frequently observed in many astrophysical turbulent environments, such as the solar wind. Panel (a) of Figure 1 shows a snapshot of the turbulent magnetic field obtained in our simulation, with a close-up on a magnetic hole in panel (b). Here, the magnetic hole is represented by the oval-shaped dark shaded area, which indicates a drop in the magnetic field intensity. Many more magnetic holes with similar properties are found in the simulation [3, 4].

In addition to understanding whether magnetic holes are produced spontaneously by the turbulence, we also studied their role in the turbulent dynamics, and how they contribute to dissipation. As mentioned, dissipation happens when the energy associated with plasma motion (kinetic energy) is converted into heat. We investigated this aspect by analyzing the velocity distribution of particles inside magnetic holes. Panel (c) of Figure 1 shows such distribution for electrons, measured within the magnetic hole in panel (b). The velocity distribution is essentially a 3D histogram of particle velocities, and the wider the distribution is, the hotter particles are (i.e. their temperature is higher). The red arrow in the figure represents the magnetic field direction inside the magnetic hole. Normally, the velocity distribution has a spherical, symmetrical shape, but turbulence drastically alters the nature of the distribution, as observed in the figure. We see that the distribution is asymmetrical, being wider in the plane perpendicular to the magnetic field direction, and narrower in the parallel direction. This shape implies that the electron temperature parallel to the magnetic field is smaller than the perpendicular temperature. Thus, inside magnetic holes, electron heating is more efficient in the direction perpendicular to the magnetic field, and the heating is said to be anisotropic.

The strong electron heating observed inside magnetic holes is caused by the fact that these structures are able to trap particles. This indeed causes dissipation. To understand why, let us assume there is a group of electrons moving all together toward a magnetic hole. This population of electrons has a certain kinetic energy, associated with their collective drift motion. However, when the electrons meet the magnetic hole, they get stuck and start wandering inside the structure. Thus, their collective drift is converted into a localized trapped motion inside the magnetic hole. As more electrons get trapped, the velocity distribution gets wider, so the plasma gets hotter. In other words, the kinetic energy of electrons is dissipated into heat as they get trapped, and this is how magnetic holes contribute to dissipation in turbulence.

Ongoing Research / Outlook

Energy dissipation in turbulence is only one of the multiscale plasma processes which we investigate with our simulations. Others are magnetic reconnection, which converts energy stored in the magnetic fields into particle acceleration and heat, and a host of so-called “kinetic instabilities”, rearrangement of the energy stored in plasmas. They are of particular relevance when occurring in the solar wind because, even if they occur at the very small, kinetic scales, they influence energy evolution over the entire extension of the heliosphere, from the Sun up to its boundaries, 100 astronomical units away from it.

References and Links

[1] D. Verscharen et al., Living Reviews in Solar Physics, 16(1), 5.
[2] S. Y. Huang et al., ApJL, 2017, 836, L27.
[3] G. Arrò et al., A&A, 2022, 668, A33.
[4] G. Arrò et al., ApJ, 2023, 958, 11.

Turbulent Magnetic Field Amplification in Weakly Collisional Plasmas

Research Institution
Research School of Astronomy and Astrophysics, Australian National University

Principal Investigator
Christoph Federrath

Researcher
Radhika Achikanath Chirakkara, Amit Seta

Project Partner
–

Funding
Australian Research Council Discovery Project DP230102280; UA-DAAD

SuperMUC Project ID: pr32lo

Introduction

The gas between galaxies is a hot, diffuse, nearly collision-less plasma, which hosts dynamically strong magnetic fields of micro-Gauss strength. It is one of the cosmic mysteries that these fields are so strong, considering that the magnetic fields in the early Universe, after the Big Bang, were extremely weak, many orders of magnitude smaller than observed today. This is even more surprising because the plasma of the intra-galaxy-cluster medium is nearly collision-less, and therefore does not obey the more conventional physics of collisional gases, the so-called "magneto-hydrodynamical" (MHD) limit. Instead, those plasmas cannot be treated as a gas/fluid, and can only be described by kinetic theory, in which individual particle motions and trajectories, particularly those of the ions, must be followed with special numerical techniques, so called "particle-in-cell" (PIC) methods.

In this project, published in [1], we have developed a new hybrid PIC code to model the magnetic field amplification by turbulent motions in the intra-cluster medium. We know that turbulence is present in these plasmas, yet how efficiently they can amplify fields under nearly collision-less physical conditions remained elusive. Now we have developed the methods to tackle this complex problem with high-resolution, detailed simulations of "turbulent dynamo" amplification in the weakly collisional regime.

The turbulent dynamo is the most efficient physical process to amplify tiny seed magnetic fields exponentially fast to powerful levels where the field can strongly influence the dynamics of the plasma, accelerate charged particles to relativistic speeds, slow down star formation, and drive powerful jets [2]. The way the turbulent dynamo works is by the so-called "stretch-twist-fold-merge" mechanism. Starting with a very weak magnetic-flux tube, the turbulence, through its differential complex motions (so-called "vorticity"), stretches and twists the flux tube until it folds onto itself and merges to a then much stronger flux tube. Thus, in this dynamo process, the magnetic field grows at the expense of turbulent energy.

While we have known this process to work well in ordinary, collisional gases that can be described by the equations of MHD, it was less clear whether the turbulent dynamo would work in nearly collision-less plasmas, such as those in the intra-cluster medium. Here we show that it does, and we determine the plasma conditions under which it can or cannot.

Results and Methods

We have developed a state-of-the-art hybrid PIC module for the FLASH code [3], with a new plasma cooling technique that allows us to maintain a steady turbulent Mach number [1]. This will enable exploration of much more compressible plasma dynamos than ever before, even into the supersonic regime. However, as a first application of the new method, we focus here on the collision-less turbulent dynamo in the subsonic, kinematic regime for different values of the magnetic Reynolds number (Rm), and initial Larmor ratio, that is the ratio between the Larmor radius of the charged particles (ions) and the system size. We use resolutions up to 180^3 grid cells and up to 200 charged particles per cell. Thus, the total number of resolution elements is as high as 10^9, which poses a significant computational challenge, requiring machines such as SuperMUC-NG. Our code is fully parallelised with the MPI and I/O is handled with the parallel HDF5 library.

We investigate both the "un-magnetised" regime of the dynamo, where the magnetic field is relatively weak, and the Larmor radius is larger than the system size, as well as moderately "magnetised" cases, where the Larmor radius is of the order of the system size. Almost independent of the level of magnetisation, we find a new critical

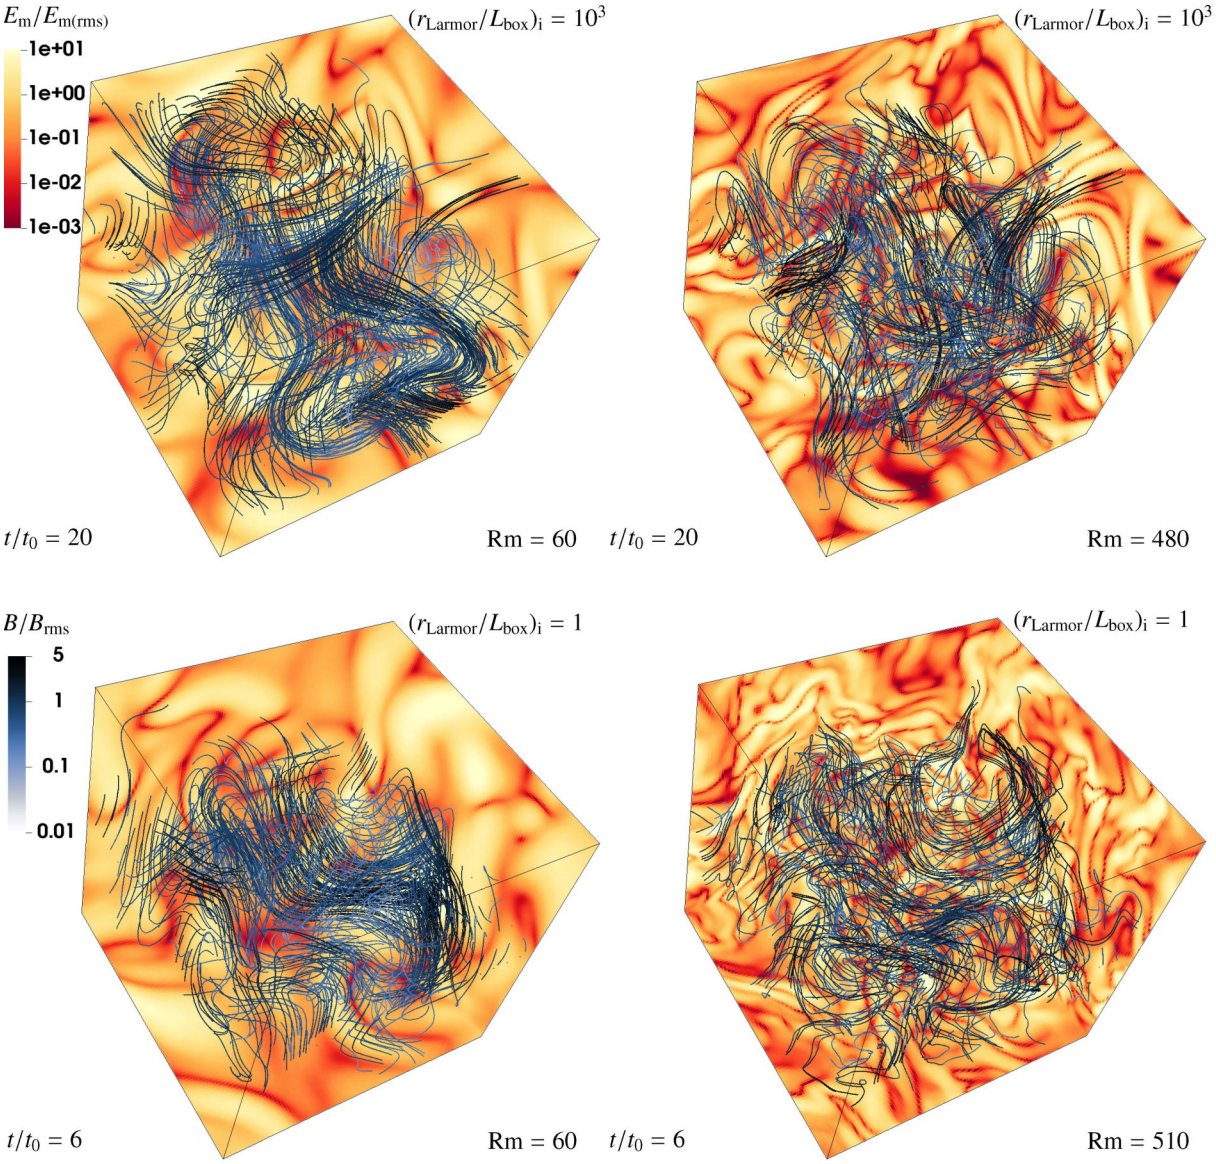

Figure 1: Slices of magnetic energy, with magnetic field streamlines coloured by magnetic-field strength in simulations of turbulent magnetic field amplification in weakly collisional plasmas. The top panels show initially 'un-magnetised' simulation models, where the Larmor radius is 1,000 times greater than the system (box) size, while the bottom panels show 'magnetised' models with the Larmor radius comparable to the system size. The left-hand panels show runs with high resistivity and therefore low magnetic Reynolds number (Rm), while the right-hand panels show the same for high Rm. We find that only models with high Rm yield magnetic field amplification through the turbulent dynamo, which manifests in complex, highly tangled magnetic field structures (see right-hand panels). This outcome is largely independent of the level of initial magnetisation. This is the first time that it is shown that the turbulent dynamo can operate under such conditions and amplify initially very weak magnetic fields to significant levels, such as observed in galaxy clusters today.

magnetic Reynolds number for the dynamo action of Rm_{crit} = 100-130, with higher magnetisation requiring magnetic Reynolds numbers that correspond to the upper limit of this critical range. Thus, the turbulent dynamo is easily excited in weakly collisional plasmas with $Rm > Rm_{crit}$, showing similar characteristics to those in the MHD dynamo (see Figure 1).

Ongoing Research / Outlook

Recent studies suggest that the intra-galaxy-cluster medium may have sonic Mach numbers of order 1 and the plasma is therefore strongly compressible. Thus, after investigating the turbulent dynamo under mildly compressible, subsonic plasma conditions, the next step is to expand this study into the compressible, and ultimately supersonic regime of turbulence.

References and Links

[1] Achikanath-Chirakkara et al. 2004, MNRAS 528, 937.
[2] Federrath 2018, Physics Today 71, 38.
[3] Dubey et al. 2008, ASP Conf Ser. 385, 145.

Math and Computer Science

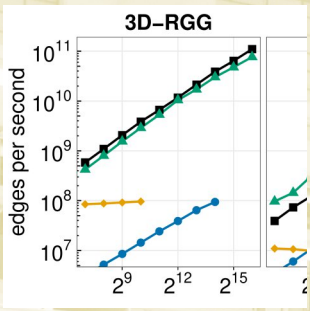

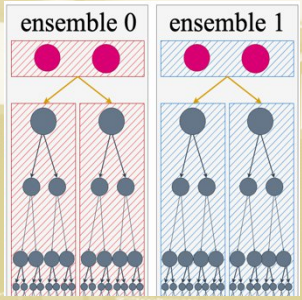

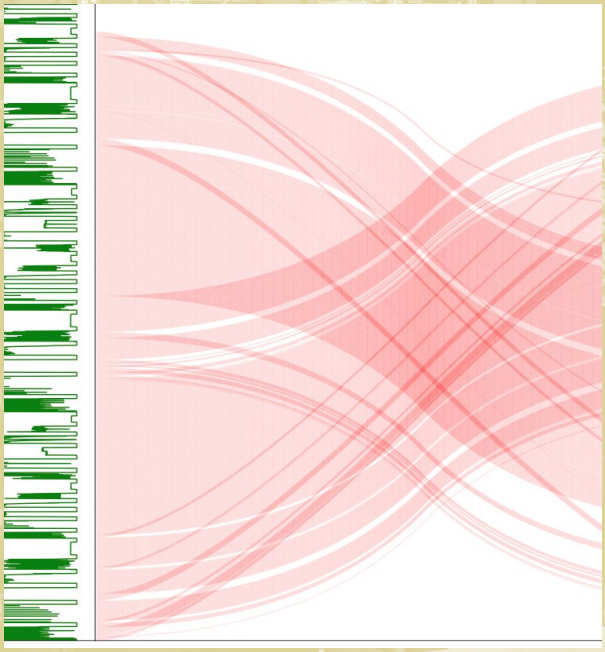

Massively Scalable Discrete Algorithms for the Basic Toolbox

Research Institution
[1]Karlsruhe Institute of Technology

Principal Investigator
Peter Sanders[1]

Researcher
Demian Hespe[1], Lukas Hübner[1,2], Sebastian Lamm[1], Matthias Schimek[1], Dominik Schreiber[1], Daniel Seemaier[1], Tim Niklas Uhl[1]

Project Partner
[2]Heidelberg Institute of Theoretical Studies

Funding
ERC 882500

SuperMUC Project ID: pn72pu

Introduction

Algorithms are at the heart of the information age – they control most nontrivial aspects of computer applications. However, algorithm development (and therefore the progress of the information age) is facing a major crisis by the collision of two revolutionary developments. On the one hand, we observe an explosion of the amount of data to be processed (*big data*). On the other hand, the performance of a single device for executing a sequence of program instructions (a *processor core*) is stagnating. The widening gap between required and available performance can only be closed by decisively improving the algorithms necessary for computationally-demanding applications. The algorithms must become more efficient for large inputs, and they must be able to efficiently use many processors at once (*parallel processing*). In our research, we therefore target the scalability of some of the most frequently used building blocks of algorithms (*the basic toolbox*). More concretely, with our previous and ongoing work we address discrete algorithms that play a role in almost all nontrivial computer programs such as sorting and basic algorithms for graphs. With the increase of the number of processors in current supercomputing systems the probability of hardware failures rises. Therefore, we also look at scalable solutions to fault-tolerance. We use the methodology of algorithm engineering (AE) that integrates the robustness and scalability guarantees of algorithm theory with careful, practice oriented modeling, implementation, and experimental evaluation for the development of massively scalable algorithm for the discrete basic toolbox. The result of these efforts are algorithms that scale up to two orders of magnitude better than the state of the art. Most of them are available as open-source software. Although we are not directly targeting the numerical applications that are so far the main users of supercomputers, such applications will also benefit because our algorithms help with load balancing, collective communication, and data analysis.

Results and Methods

We now discuss our work on scalable basic toolbox graph algorithms [2,3] in greater detail. In addition to this sub-project, we also successfully continued our work on malleable job scheduling, scalable SAT solving and fault-tolerance [1,4,5]. Most state of the art HPC platforms are primarily designed with numerical computations in mind that feature fairly regular data access and communication patterns with often only a small number of communications partners. However, more data intensive and irregular workloads are becoming increasingly popular, i.e., from materials science applications. Graphs are a common representation for the data sets used by these applications, which highlights the need for truly efficient and scalable algorithmic routines for all kind of graph computations. To that end, we investigated algorithms for basic graph problems like connected components, minimum spanning trees, and triangle counting in this project as these are useful for graph data analysis at large scale and serve as building blocks for more complex graph data processing tasks. Designing scalable graphs algorithms on current massively parallel HPC platforms is especially challenging. Graph algorithms often involve highly irregular access patterns with small amounts of data being transferred and (depending on the graph structure) many communication partners. This poses a challenge to current communication hardware and software, such as the commonly used communication libraries implementing the MPI standard. Therefore, massive scalabiltity can only be achieved by using highly-tuned communication primitives, which are not covered by MPI implementations. In the following, we will present our algorithms for the computation of minimum spanning trees (MST) as an example for distributed graph algorithms and briefly sketch our key techniques for achieving scalability up to large processor configurations. We assume each processor holds a subset of the input graph's vertices and the incident edges. Therefore, each processor only has a local view on the input. Our

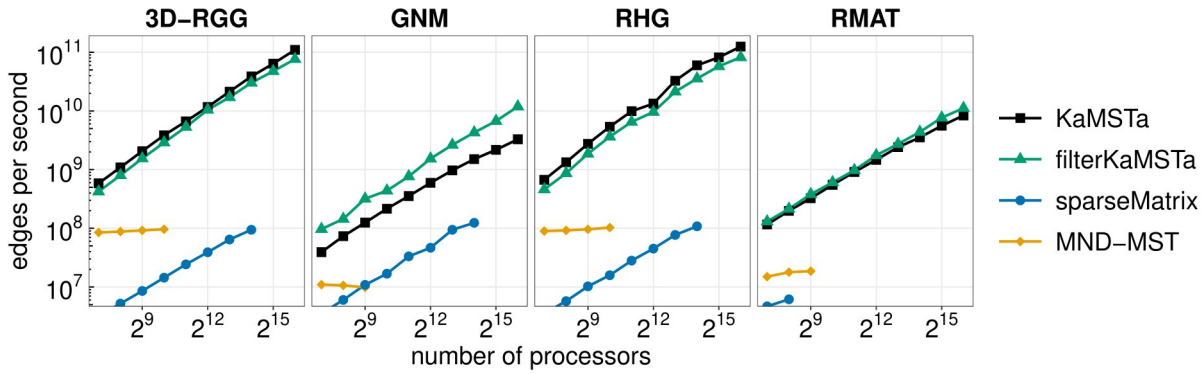

Figure 1: Throughputs of multiple MST algorithms (ours are KaMSTa and filterKaMSTa) in a weak scaling experiment. Each algorithm is executed with eight threads per MPI process.

two distributed algorithms are based on Boruvka's sequential MST algorithm and scale to graphs with hundreds of billions of edges on tens of thousands of processors outperforming the previous state of the art distributed MST algorithms by up two orders of magnitude. One key concept of our algorithms is the exploitation of graph locality in a local preprocessing routine. Here, we use a property of the MST problem which allows to reduce the input graph to an often much smaller kernel by contracting local MST edges without any communication. Since our algorithms additionally follow a hybrid parallelization approach which combines inter-process communication via MPI and shared-memory parallelism, a processor is able to inspect a larger part of the graph without dedicated communication, enhancing the effectivity of this local preprocessing routine compared to an MPI-only setting. Within our MST algorithms we often have to communicate a large number of small messages to varying communication partners. It turned out that the standard algorithms of the used MPI implementations incur a prohibitively large latency overhead for such communication patterns. By delivering messages indirectly in a two phase approach, we are able to reduce the latency considerably. For load balancing, we rely on our previously developed efficient and distributed sorting algorithms. In our second MST algorithm, we additionally adapted an (edge-)filtering technique proposed by Osipov et al. to the distributed setting which allows us to discard many non-MST edges early which can help to reduce communication volume and local work. We compared the scaling behavior of our algorithms (*KaMSTa* and *filterKaMSTa*) with two state of the art distributed MST algorithms: *sparseMatrix* by Baer et al. and *MND-MST* by Panja et al.. For our weak-scaling experiments (see Figure 1), we use different families of random graphs with varying characteristics such as locality and degree distribution. Our algorithms exhibit good scalability on up to 65,536 cores. Furthermore, we clearly outperform our competitors on all instances. Especially on graphs with high locality, we are up to two orders of magnitude faster. To save computation budget, we did not run our competitors on more than 16,384 cores. Note that our competitors could not successfully compute an MST for all tested input configurations leading to further missing data points.

Ongoing Research / Outlook

All of our sub-projects are actively being worked on. We have devised a highly efficient and decentralized approach to *malleable job scheduling*, where individual tasks receive a *fluctuating* amount of computational resources during their execution. On up to 6,400 cores, our approach achieves scheduling latencies in the range of a few milliseconds and near-optimal resource utilization [4]. Within this system, we explore *propositional reasoning*, which is a crucial building block for applications like, e.g., designing electronic circuits, proving the correctness of software, or computing complex timetables. This line of work has resulted in an award-winning distributed solver, which we continue to advance in terms of scalability and reliability. Regarding distributed graph algorithms, we are currently working on (asynchronous) algorithms for shortest path computations which we will then use as a key building block for graph analysis tasks, such as analyzing the structure of large synthetic models of the human brain. Furthermore, we aim to improve our collective communication routines to make them more robust to worst-case inputs and incorporate aggregation mechanism which will help to further reduce their communication volume for many tasks. As supercomputers are getting ever larger and more energy efficient, hardware failures are rising. We are developing algorithms, tools, and frameworks to be able to handle hardware failures in a scalable manner and without restarting the whole application [5].

References and Links

[1] https://ae.iti.kit.edu/
[2] Sanders, Peter and Matthias Schimek. *Engineering Massively Parallel MST Algorithms*. IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2023,pp. 691-701.
[3] Peter Sanders and Tim Niklas Uhl. Engineering a Distributed-Memory Triangle Counting Algorithm. IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2023, pp. 702-712.
[4] Peter Sanders and Dominik Schreiber. *Decentralized Online Scheduling of Malleable NP-hard Jobs*. International Conference on Parallel and Distributed Computing (Euro-Par) 2022, pp. 119-135.
[5] Lukas Hübner, Demian Hespe, Peter Sanders and Alexandros Stamatakis. *ReStore: In-Memory REplicated STORagE for Rapid Recovery in Fault-Tolerant Algorithms*. IEEE/ACM 12th Workshop on Fault Tolerance for HPC at eXtreme Scale (FTXS) 2022, pp. 24-35.

A Next Generation Benchmark for Automated Deep Learning

Research Institution
University of Freiburg

Principal Investigator
Frank Hutter

Researcher
Archit Bansal, Danny Stoll, Arber Zela

Project Partner
–

Funding
01IS18046D, BMBF

SuperMUC Project ID: pn68xi

Introduction

Since the release of ChatGPT in December 2022, artificial intelligence (AI) has garnered unprecedented media attention. Mainstream media outlets, such as Tagesschau, now regularly report AI-related news, and public discourse on the subject has permeated schools and universities. This widespread attention is the result of significant developments in Deep Learning (DL), which has proven incredibly successful at solving previously unsolvable tasks and outperforming known solutions to established problems. Deep Neural Networks (DNNs) are very powerful, but designing their architectures and training pipelines for novel tasks has traditionally required extensive trial-and-error by human experts. Thus, DNNs have been hard to adopt in new fields.

To democratise DL and enable inexperienced users to design good DNNs for novel tasks without the need for extensive research into domain-specific DNN model designs, NAS (Neural Architecture Search) treats the design of these network architectures itself as an optimization problem that can be solved algorithmically. NAS has recently shown very promising results in designing novel DNNs that outperform hand-crafted ones. Parallel to NAS, research in automating the training pipelines for DNNs focuses on algorithmic approaches towards choosing the correct values of a number of parameters, called hyperparameters, that alter the behaviour of the model training pipeline. This field of research is called Hyperparameter Optimization (HPO).

Research on both NAS and HPO is extremely compute intensive, since the evaluation of a single NAS or HPO algorithm may involve training hundreds or thousands of DNNs, cumulatively consuming several weeks or months of compute on large compute clusters. This presents a high barrier to entry for researchers who do not have access to such amounts of compute, as well as raises concerns about the carbon footprint of such research. In response, PI Frank Hutter has previously pioneered the development of benchmarks as a one-time investment of resources in order to amortise the costs of all future NAS research, starting a trend amongst NAS researchers of developing a diverse range of benchmarks.

As the impact of architectures and hyperparameters is intrinsically linked, a joint approach towards optimising both architectures and hyperparameters can be more powerful than NAS or HPO alone. Nevertheless, just as NAS and HPO research was previously bottlenecked by an absence of appropriate benchmarks, so too is research into this problem currently blocked by the absence of suitable benchmarks. Therefore, we proposed JAHS-Bench-201 to enable research into what we refer to as Joint Architecture and Hyperparameter Search (JAHS). JAHS-Bench-201 is likely to become the cornerstone for future research in as-of-yet underexplored approaches of JAHS, and it has already gained significant recognition for its contributions as it was accepted as a featured paper (i.e. top-7.5% of accepted papers and top-2% of submitted papers) in the prestigious 36th Conference on Neural Information Processing Systems' (NeurIPS'22) Datasets and Benchmarks Track [1].

Results and Methods

To create JAHS-Bench-201 [2], we designed a 14-dimensional search space and used the SuperMUC-NG cluster to collect the most extensive dataset of neural network performance data available in the public domain. We used this performance dataset to train surrogate models, so researchers can now use these models to query the JAHS-Bench-201 search space for an approximate evaluation without actually training a complete DNN. This technique reduces the associated compute time from GPU-days to CPU-seconds (see Table 1). A single research paper may need to perform

Task	Surrogate-based [CPU-s]	Training-based [GPU-d]
CIFAR 10	2.8	5.0
Colorectal Histology	2.9	0.4
Fashio MNIST	2.7	2.5

Table 1: Comparison of the compute requirements of surrogate-based experiments (values reported in CPU-seconds) vs training-based experiments (values reported in GPU-days).

Property	Reference	Ours
# Architectures	~1,000	15,625
# Hyperparameters	2 x discrete	2 x discrete 2 x continuous
# Fidelities	0	4
# Metrics	9	20
# Tasks	1	3
# Trained DL models	4,800	810,000
# Data points	192,000	161,000,000

Table 2: Comparison of JAHS-Bench-201 to NAS-HPO-Bench-II (October, 2021), which shares a number of common properties with our work but is much more limited in its scope.

many hundreds of thousands of such evaluations, therefore using our benchmark will both help democratise research on JAHS and save substantial compute time & carbon emissions.

For our search space, we chose to combine one of the most popular search spaces in NAS literature - the search space of the benchmark NAS-Bench-201 - with some common choices for hyperparameters in the HPO literature. We complement the main search space with fidelity parameters that represent low-cost approximations such as reduced network size or reduced image resolutions, as NAS and HPO algorithms use these to gain speedups. As we include cheap low-fidelity evaluations, in contrast to the predominant use of GPUs in other areas of deep learning, our experiments are actually much more cost-efficient to perform on a large cluster of CPUs; thus, we carried them out on SuperMUC-NG.

Our performance dataset is composed of approximately 161 million data points, each consisting of 20 performance metrics, across three image classification tasks. As we record the metrics at each epoch of model training, we yield an additional fidelity parameter in the form of the number of epochs of training. Therefore, JAHS-Bench-201 supports not just 4 distinct fidelity parameters, but also supports research into the joint impact of varying multiple fidelity parameters simultaneously, which is unprecedented in extant related literature.

We were able to utilise the massive number of parallel compute nodes available on the SuperMUC-NG to run simulations for up to 100k Deep Learning trainings in parallel. The greatest limitation we faced was the availability of RAM on a per-node basis, which bottlenecked the number of parallel computations that we can perform. Nonetheless, the SuperMUC-NG provided us with parallel compute resources well beyond what we see in related work (see Table 2).

References and Links

[1] A. Bansal, D. Stoll, M. Janowski, A. Zela, and F. Hutter. *JAHS-Bench-201: A Foundation For Research On Joint Architecture And Hyperparameter Search*. In Proceedings of the 36th Conference on Neural Information Processing Systems. 2022.
[2] https://github.com/automl/jahs_bench_201

Time Series Mining on High Performance Computing Systems

Research Institution
Technical University of Munich

Principal Investigator
Martin Schulz

Researcher
Amir Raoofy

Project Partner
–

Funding
AZ-1214-16 and AZ-1468-20, Bayerische Forschungsstiftung

SuperMUC Project ID: pn25di

Introduction

Modern monitoring infrastructures are equipped with a vast number of sensors that record the history of operations and conditions of underlying systems over long time periods. These infrastructures produce a large series of real-valued historical data. Mining and unraveling this data is a central piece to modeling and ultimately understanding the behavior of the monitored systems, and the first step for that typically includes the investigation of correlations, similarities, and anomalies in the data. This investigation is essential for creating and improving the statistical and machine learning models used in today's intelligent monitoring systems, e.g., for predicting failures, which is a compute-intensive problem, especially for large-scale datasets. In this context, matrix profile [1] is a well-established indexing approach for the mining of time series data. This approach was introduced in 2016 by Yeh et al. in a series of papers and has been successfully applied to mine similarities and patterns in datasets from various fields, such as seismology and medical science, and used for various data mining and machine learning tasks, such as semantic segmentation, clustering, and anomaly detection. For better illustration, in Figure 1, we show an example use case of the matrix profile to find the similarity correspondence between two input time series acquired from cache miss profiles when running an HPC application. The matrix profile itself is a meta index that encodes similarity in time series data, and its computation corresponds to a nearest-neighbors problem with the well-known classical exact approaches to compute it based on exhaustive search operations for nearest neighbors. These exact solutions are generally inefficient for large datasets as the computational costs scale quadratically with the size of the datasets (i.e., the number of records). These approaches typically rely on mitigating the computational costs for the search operations by extensive arithmetic optimizations of compute kernels [2], or using accelerators for computation [3], and deploying on cloud-based [4] or HPC systems [5]. On the other hand, approximate approaches [6] are drawing increasing attention as they can provide solutions that are much more efficient to compute while also being accurate enough in practice. However, existing approaches still suffer either from excessive computational costs [2] (e.g., it takes days to compute matrix profile for 1 billion records on GPUs [4]), are restricted to specific settings and application scenarios, or lack parallelization. In our work [7], we investigate approximate approaches for the computation of matrix profiles, and focus on the scalability of the family of the iterative approximate nearest neighbor methods exploiting tree-based data structures on the SuperMUC-NG.

Results and Methods

We specifically introduce two optimizations to address scaling challenges in the computation of matrix profile using the iterative approximate tree-based approach:

1) We pipeline the stages of the nearest neighbor computation in consecutive iterations to overlap parts (phases) of each iteration, e.g., the construction of the parallel tree data structure and follow-on search operations. This pipelining mechanism helps to stagger the communication latencies within the phases in iterations on top of actual computation and, therefore, achieves latency hiding through communication/computation overlap of multiple iterations. Figure 2 illustrates an overview of this optimization scheme.

2) We enable control over the granularity of parallelism by constructing a forest of parallel trees on smaller ensembles of resources.

The key idea here is to construct a forest of multiple trees, each running on ensembles of the allocated resources, instead of using all the resources to construct a single tree. The matrix profile computation in each ensemble runs independently in parallel, and therefore, the resources in separate ensembles do not actively communicate across. This leads to a much more scalable

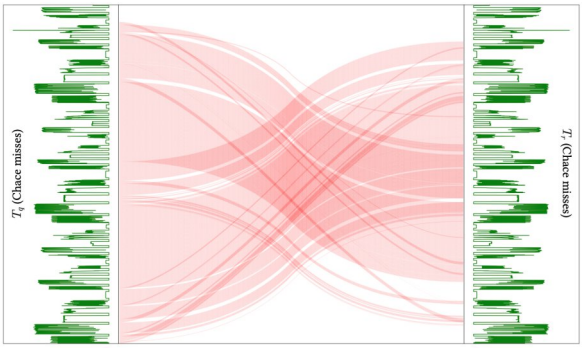

Figure 1: Illustration of similarity correspondence between two time series using the matrix profile.

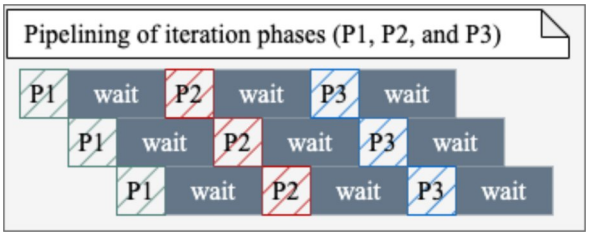

Figure 2: Pipelining of iteration phases of computing matrix profile. Each phase runs in parallel on the entire allocated resources effectively without the penalty of wait latencies on the SuperMUC-NG.

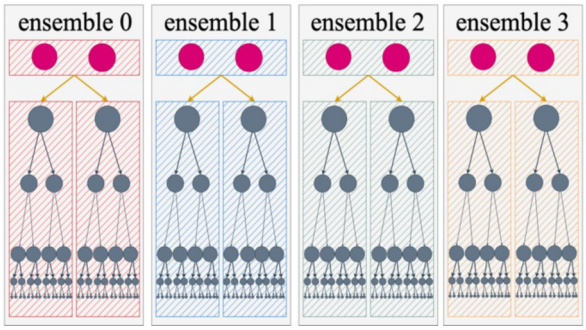

Figure 3: Forest of parallel trees running in ensembles of nodes on the SuperMUC-NG system.

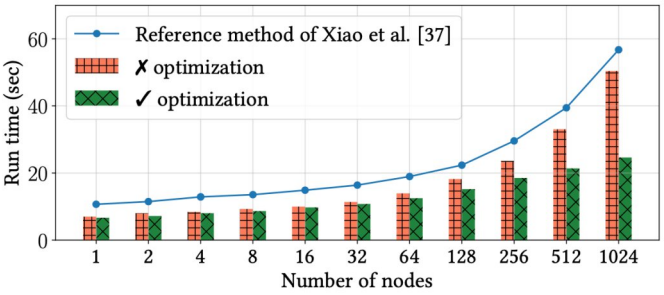

Figure 4: Weak scaling for tree-approach on SuperMUC-NG.

solution. Figure 3 illustrates an overview of this optimization scheme. These approaches are implemented in C/C++ and use pure MPI for parallelization and communication. These optimization schemes are complementary and need to be combined on top of each other and tuned for large-scale runs on the SuperMUC-NG. However, employing them introduces additional memory consumption overhead. Therefore, we conduct extensive parameter studies to find the optimal settings for running matrix profile computation with the highest performance while staying within the memory capacity limits of the SuperMUC-NG nodes. Specifically, the depth of the pipeline (i.e., the number of iterations that can run simultaneously) and the number of ensembles used in the forest are the main tuning parameters. Figure 4 illustrates the results of running a weak scaling experiment on the SuperMUC-NG system starting from 1 to 1,024 nodes. The results suggest improvements surpassing the performance of reference implementations of the approximate approach. Additionally, when the above-introduced optimizations are tuned and enabled simultaneously, we are able to achieve a significant improvement in scalability. Due to its high computational demands, there have only been limited previous efforts to conduct matrix profile computation for large time series. Most notably, the matrix profile computation for a 1-billion-record long time series was deployed on AWS spot instances using the SCAMP [3] algorithm, which took between 10.3 hours and 2.5 days. Therefore, to showcase the capacity of the tree-based approach and illustrate the impact of introduced optimizations for leveraging the compute power of the SuperMUC-NG, we run the matrix profile computation on a billion-record dataset using 48K cores (1,024 nodes) of the SuperMUC-NG. This run takes only 19 minutes to reach 99% accuracy while getting to around 338 Tera Flops of performance.

Ongoing Research / Outlook

The optimization mechanisms presented can be generalized and applied to similar problems (e.g., nearest neighbor computations or randomized iterative algorithms). Further improvements and fine-grained optimization can be investigated to improve the scalability even more. With the presented work, it is feasible to compute the matrix profile for large time series, which previously required days to compute, in just a few minutes. We demonstrated a billion-scale experiment with up to 99% accuracy in 19 minutes on the SuperMU-NG system, which is 3-100 times faster than existing approaches. This provides a superior computation time compared to the alternative approaches by leveraging the power of the SuperMUC-NG system.

References and Links

[1] Yeh, C.M., et al.: Matrix Profile I: All Pairs Similarity Joins for Time Series.
[2] Zhu, Y., et al.: Matrix Profile XI: SCRIMP++: Time Series Motif Discovery at Interactive Speeds.
[3] Zhu, Y., et al.: Matrix Profile II: Exploiting a Novel Algorithm and GPUs to Break the One Hundred Million Barrier for Time Series Motifs and Joins.
[4] Zimmerman, Z., et al.: Matrix Profile XIV: Scaling Time Series Motif Discovery with GPUs to Break a Quintillion Pairwise Comparisons a Day and Beyond.
[5] Raoofy, A., et al.: Time Series Mining at Petascale Performance.
[6] Schall-Zimmerman, Z., et al.: Matrix Profile XVIII: Time Series Mining in the Face of Fast Moving Streams using a Learned Approximate Matrix Profile.
[7] Raoofy, A., et al.: Overcoming Weak Scaling Challenges in Tree-Based Nearest Neighbor Time Series Mining.

Appendices

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 capac-ity of 719 TByte and a peak performance of 26.9 peta-FLOPS (= 26.9 x 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 inter-connect is a fast OmniPath network with a link band-width of 100 Gbit/s. The compute nodes are bundled into 8 domains (islands). Within one island, the Omni-Path 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 Stor-age offers additional 20 PiByte (with 70 GB/s aggre-gated 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 principle of using solid materials of enhanced water vapor

Figure 1: SuperMUC-NG.

adsorption capacity for cooling via the evapora-tion 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 stor-age. 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.

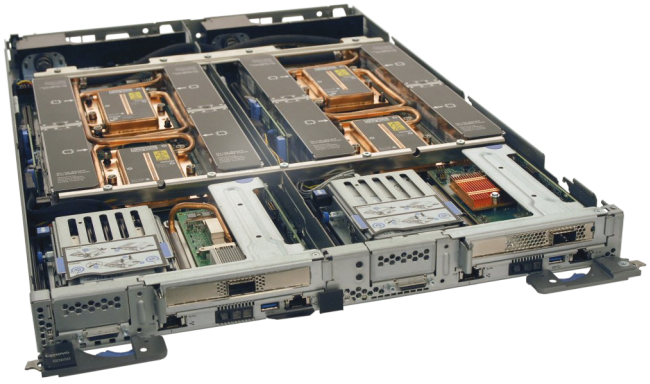

Figure 2: Dual node chassis of SuperMUC-NG.

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)		

Usage of SuperMUC-NG

In the first four and a half years of operation, nearly 10 billion core hours were consumed on SuperMUC-NG, as listed in Table 1. Planned maintenances and occasional hardware failures cover the gap from the achieved 82% productivity to the theoretical maximum, which is right on target for a high-performance machine. During the time of operation, 2.5 million jobs have been processed, from short 30 min one node test runs to large jobs using half the machine for 24 h. During special block operation periods, even jobs on the full machine were run.

High-performance computing is used in a broad range of scientific fields. In total 379 compute time projects have been allocated. The high-level view of Figure 1 (DFG classification level 2) shows that more than half of the compute time of SuperMUC-NG is consumed by 150 physics projects (40% of 379), in which many research groups traditionally use high-performance computing. Thermal and process engineering follows with about a quarter of the resources consumed by 100 projects (26% of 379). The remaining 20% of compute time is shared by 129 (34% of 379) projects in geo sciences, life sciences, computer sciences, mathematics, and further engineering disciplines.

Subdividing scientific disciplines into fields in Figure 2 (DFG classification level 3), computational fluid dynamics simulations (blue) of gas-and liquid flows, and combustion, is the overall largest field, as it takes the largest share of the engineering disciplines. Another small contribution comes from mechanical engineering (dark green) at the end of the overall list. In contrast, physics has three large fields, lead by nuclear and elementary particle physics, e.g., quantum chromodynamics simulations, followed by astrophysics, running simulations of galaxy- and star formation or gravitational waves, and, finally, condensed matter physics, treating quantum effects like superconductivity in solid state matter.

Year	Months	Usage [M core-h]	Productivity	No. of jobs submitted
2019	5	830.1	83.0 %	181,474
2020	12	2,104.1	77.0 %	659,920
2021	12	2,308.5	84.6 %	684,638
2022	12	2,239.8	82.1 %	463,217
2023	12	2,306.3	84.7 %	486,284
Total	**53**	**9,788.8**	**82.3 %**	**2,475,533**

Table 1: Usage, productivity, and number of submitted jobs on SuperMUC-NG. In 2019, SuperMUC-NG started operation in August, therefore only 5 months are reported. In 2022, there are 7 months reported (the reporting time frame for this book ended in July, 2022). Productivity is <100% due to system mainentance.

The remaining disciplines split up in numerous individual fields, whose HPC usage fluctuates over the years indicated by the concentric circles. For the larger disciplines no real trend is deducible from this timeline.

With respect to the origin of the compute time projects, it is not surprising that about 40% of the allocated time has been applied for by researches located in Bavaria. Here, the strong relations of LRZ to the Munich and Bavarian universities play an important role, but also the collaboration with researchers of the Max-Planck facilities in Garching. Another 50% of the compute time is used by researchers in the other federal states of Germany. In 2021, a last allocation of 95 M core-h has been awarded to European researchers via the PRACE (Partnership for Advanced Computing in Europe) initiative.

For detailed and timely reports on the consumption of the SuperMUC-NG compute time, please visit our website [1].

References and Links

[1] https://doku.lrz.de/display/PUBLIC/Usage+Statistics+for+SuperMUC-NG

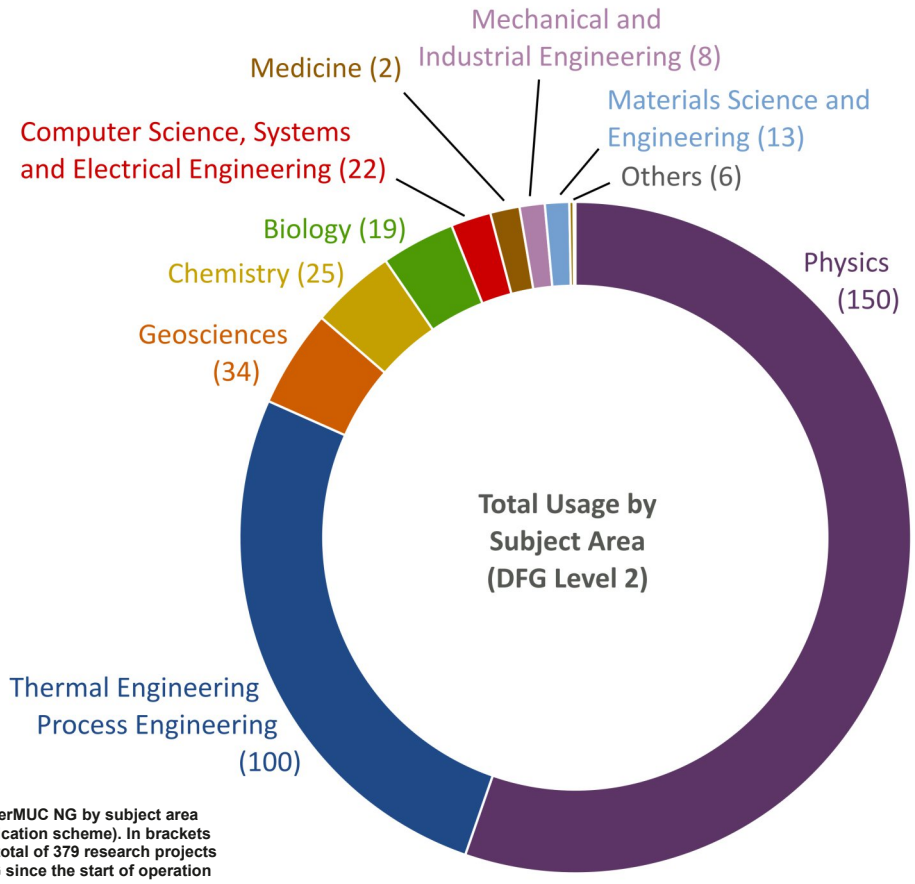

Figure 1: Total usage of SuperMUC NG by subject area (level two of the DFG classification scheme). In brackets number of projects out of a total of 379 research projects performed on SuperMUC-NG since the start of operation in 2019.

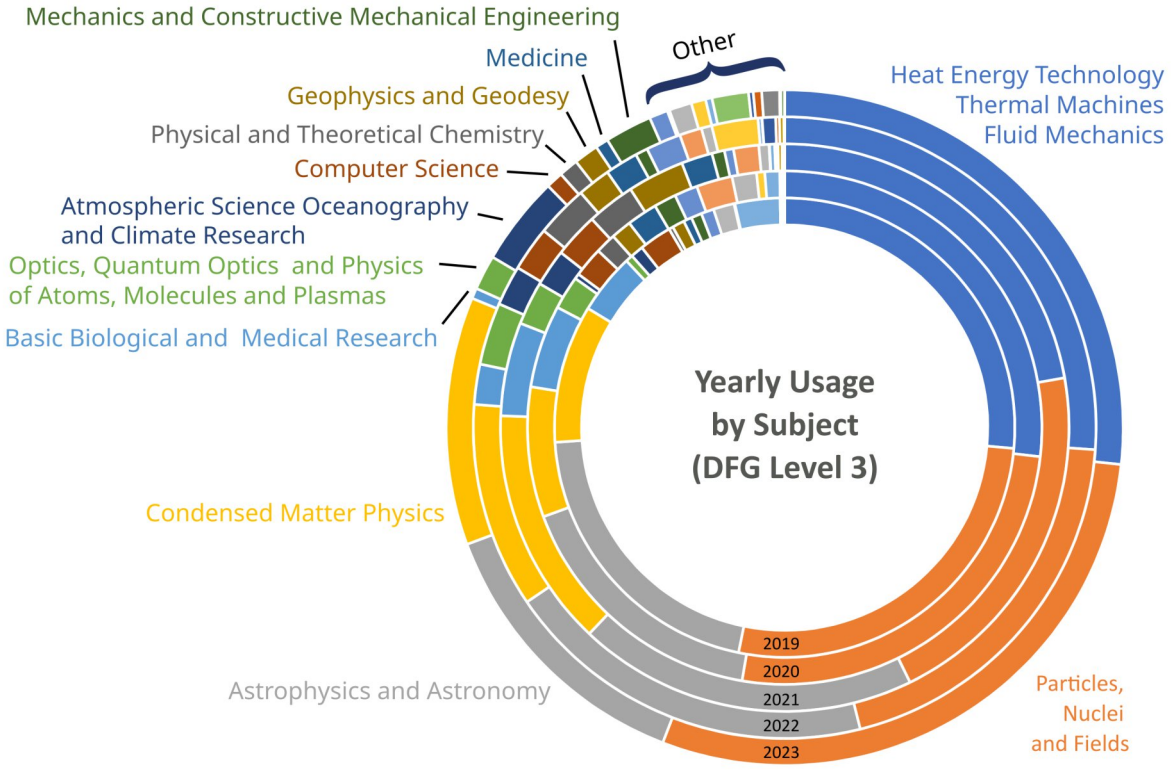

Figure 2: Detailed yearly usage of SuperMUC-NG by subject (level three of the DFG classification scheme).

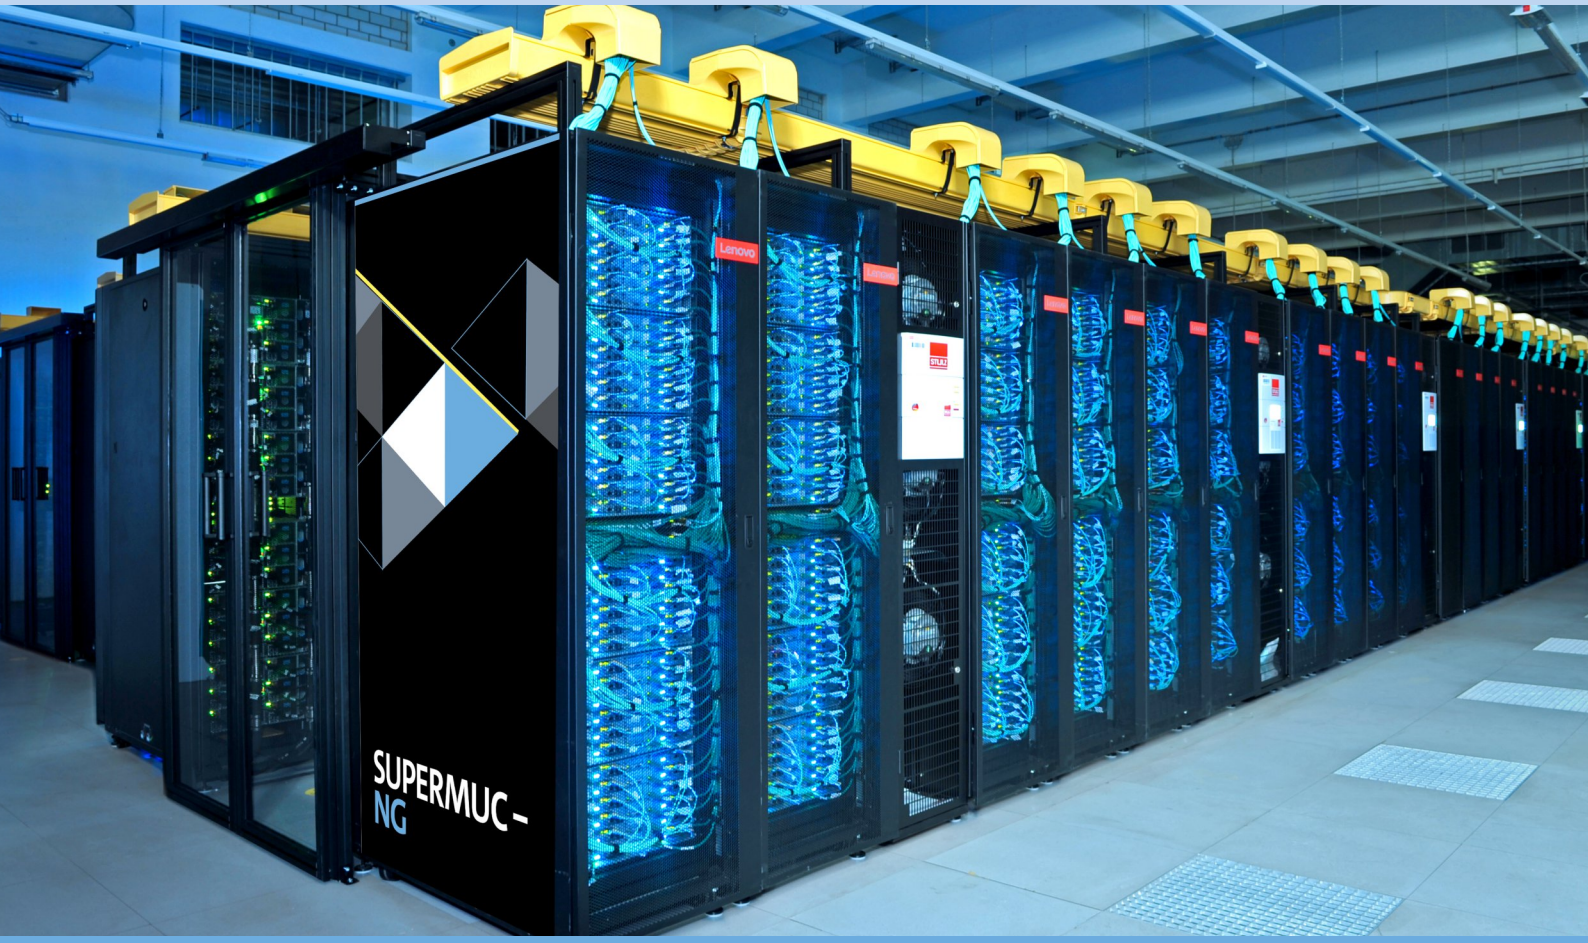

In this book, the Leibniz Supercomputing Centre (LRZ), a member of the Gauss Centre for Supercomputing (GCS), reports on the results of numerical simulations, performed in the time frame January 2022 – December 2023. The 100 project reports give an impressive overview of the utilization of SuperMUC-NG, the Tier-0 system of the Bavarian Academy of Sciences and Humanities.

SuperMUC-NG is based on Lenovo ThinkSystem SD650 DWC compute nodes and equipped with Intel Skylake Xeon Platinum 8174 processors. The machine started user operation in August 2019 and delivered 10 billion core-hours for numeric scientific simulations up to today. A detailed system description can be found in the appendix.

The articles provide an overview of the broad range of applications that use high performance computing to solve 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.

Gefördert durch den Freistaat Bayern

GEFÖRDERT VOM
Bundesministerium
für Bildung
und Forschung

ISBN 978-3-9816675-6-1
https://www.lrz.de/hpcbooks