PRACE

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

PRACE DIGEST 2018

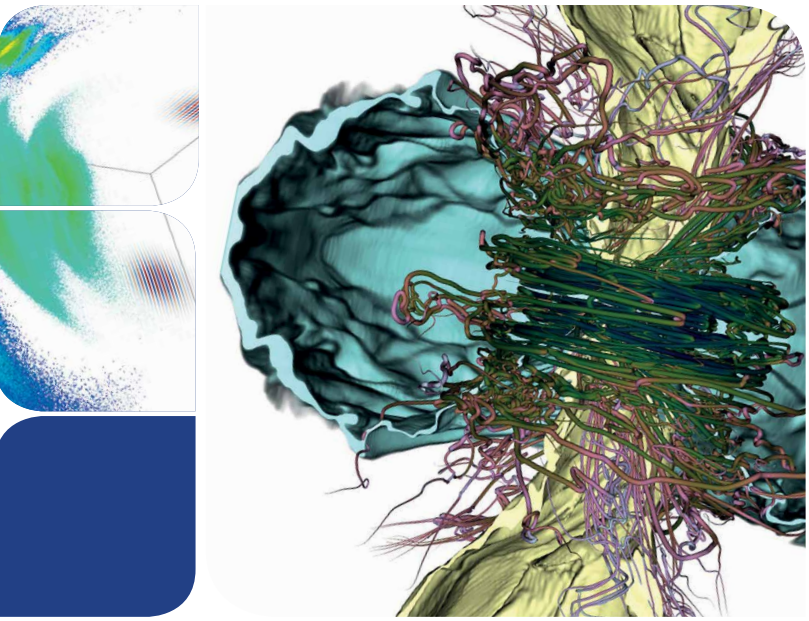

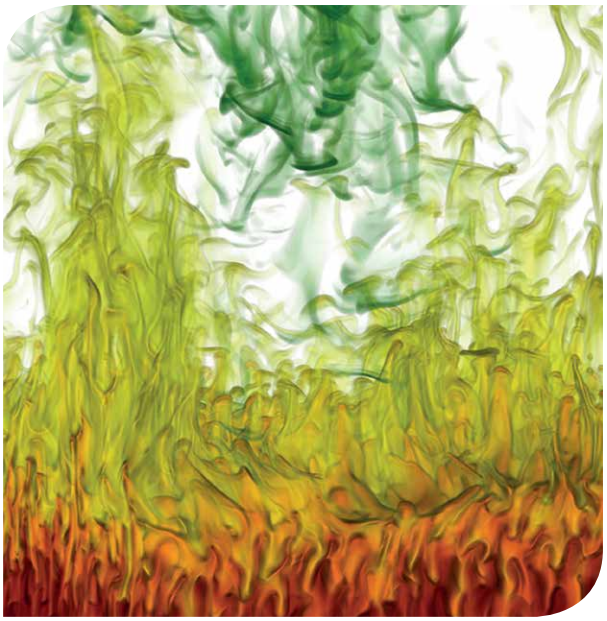

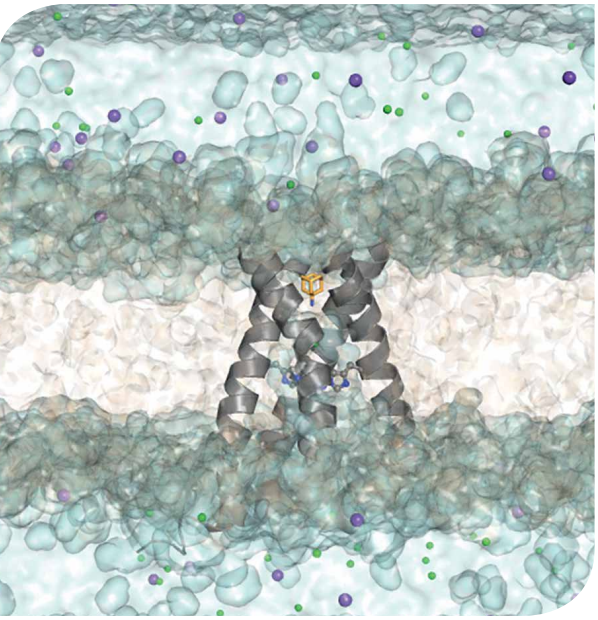

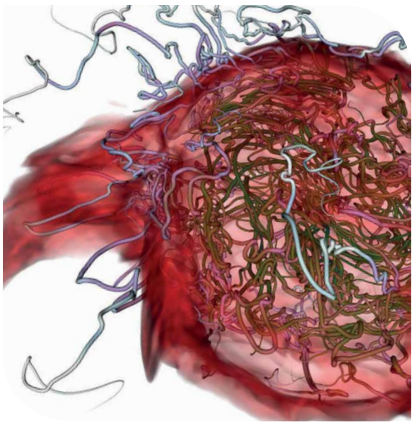

www.prace-ri.eu

PRACE Digest 2018

Neither the Partnership for Advanced Computing in Europe (PRACE) nor any person acting on its behalf is responsible for the use which might be made of the information contained in the present publication. PRACE is not responsible for the external web sites referred to in the present publication. The views expressed in this publication are those of the authors and do not necessarily reflect the official PRACE view on the subject.

© PRACE, 2018
Reproduction is authorised provided the source is acknowledged.

The copyrights of the content and images contained within this report remain with their original owner(s).
The PRACE name and the PRACE logo are © PRACE.

Cover design © Adobe Stock.

PRACE Digest 1/2018
www.prace-ri.eu

Editor-in-Chief
Marjolein Oorsprong – PRACE Communications Officer

Editors
Audrey Gerber, IUCC
Silke Lang, PRACE

Article Editor
Harry O'Neill, Insight Publishers

Communications Assistant
Jeet Landa, Insight Publishers

Graphic Design
Crystal Mather, Insight Publishers

Published by Insight Publishers, Bristol, UK

The Partnership for Advanced Computing in Europe (PRACE) is an international non-profit association with its seat in Brussels. The PRACE Research Infrastructure provides a persistent world-class high-performance computing service for scientists and researchers from academia and industry in Europe. The computer systems and their operations accessible through PRACE are provided by 5 PRACE members (BSC representing Spain, CINECA representing Italy, CSCS representing Switzerland, GCS representing Germany and GENCI representing France). The Implementation Phase of PRACE receives funding from the EU's Horizon 2020 Research and Innovation Programme (2014-2020) under grant agreement 730913.

Published: November 2018
ISBN: 9789082942217

Contents

Editorial

PRACE Digest 1/2018

This eighth edition of the PRACE Digest is launched at SC18 — held this year in Dallas, Texas — a conference where people from all sides of high-performance computing (HPC) are present and connecting. As always, the PRACE Digest aims to cover a variety of HPC applications in different research fields, and this issue includes a number of interesting topics.

Simulation of turbulent flows is a theme that returns in many PRACE projects, and this is because such flows can be found in a diverse range of situations – from the turbulence around airplane wings (page 26) to the sinking salt fingers in the oceans (page 6). All of these flows need a lot of computational power to be modelled in sufficient detail.

Astrophysics is another domain that asks for a continuous stream of large allocations of the most powerful HPC resources. Considering the vastness of the universe and the variety of phenomena in it, there are always more questions to explore and more answers to be found. On pages 14-15, Peter Hauschildt explains how he and his team are preparing their code in 3D to be able to use the next level of systems when they become available.

Looking at the above examples, HPC can seem far removed from our day-to-day lives, but I encourage you to read the article on pages 18-19 to find out how close it can actually get. Molecular dynamics simulations can help us understand how the presence or absence of water influences the effectiveness of drugs. Gaining a full understanding of the complex interactions of water molecules with drugs and the proteins they bind to could lead to much more efficient treatment of diseases.

Of course, this editorial is not meant to give a summary of all articles in the PRACE Digest – some things you will have to discover for yourself!

I will take this opportunity to inform you of the publication of the “Scientific Case for Computing in Europe 2018 – 2026”, compiled and authored by the PRACE Scientific Steering Committee. This edition is the third in a series of publications documenting the HPC requirements of scientists in Europe, and provides a detailed look into what kind of HPC support is needed to do excellent research in various fields. You can read and download the publication here:
www.prace-ri.eu/media/publications/scientific-case/

I wish you pleasant reading.

Serge Bogaerts
Managing Director of PRACE aisbl

Surfing for thermals in turbulent flows

Deciphering the swirling, eddying movement of turbulent fluids has long been a challenge for supercomputers, which are pushed to their limits when solving the over 150-year old Navier-Stokes equations used to describe them. **Professor Roberto Verzicco** of the University of Rome Tor Vergata has been investigating turbulence under both mechanical and thermal forcing.

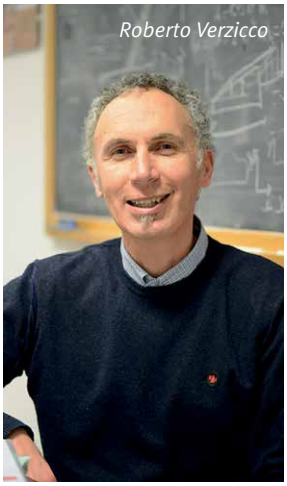
Roberto Verzicco

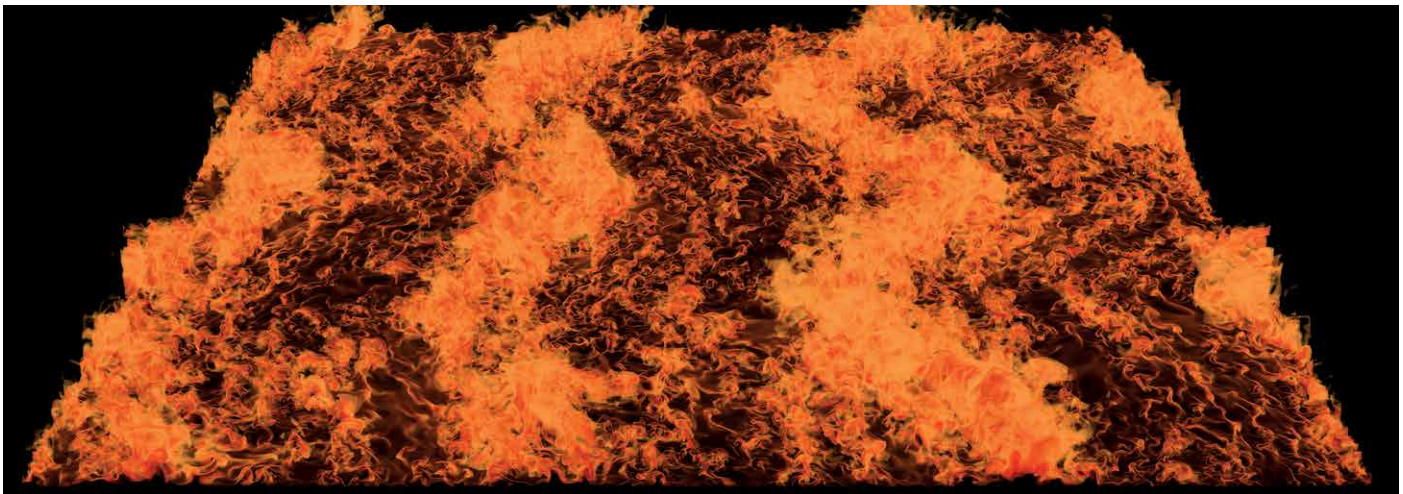
Perspective view of instantaneous temperature isosurface (T=0.8) for sheared thermal convection at Ra=4.6x10^6 and Re_w=8000. Here the flow structures near the bottom plate are presented and the flow is from the near end to the far end. The simulation is run on a mesh of 3.4 Gnodes.

Turbulence is the irregular and chaotic flow of fluid caused by excessive kinetic energy. It exists in various forms, and is famously difficult to solve in mathematical terms. The intricate nature of turbulence has led to it becoming something of a bottleneck in a number of different research areas, including predictive models for engineering applications, such as aeroplane efficiency, and more complex systems, such as the ocean and the Earth's atmosphere.

Roberto Verzicco of the University of Rome Tor Vergata and University of Twente recently led a PRACE project titled "Surfing for Thermals in Turbulent Flows", in which he and his colleagues examined a specific type of turbulence driven both by mechanical forcing (such as a pressure gradient) and thermal forcing (temperature difference). Temperature generates buoyancy by expanding the fluid, which in turn drives momentum. The interplay between the mechanical and thermal forcings gives rise to a very rich physics, which is currently only partially understood.

Using numerical simulations, it is possible to reproduce these systems and analyse the role that each ingredient plays in the features that arise in the flow. Additionally, because they are simulations, it is possible to tinker with the systems in ways which would not be possible in reality, allowing researchers to carry out idealised experiments and explore the systems in radical new ways.

"To capture the whole system, one must simultaneously resolve both the smallest of vortices and the whole system in which they are evolving"

Describing these systems in terms of mathematics involves the solution of coupled non-linear partial differential equations, which is as complex as it sounds, and the reason HPC resources are needed. The datasets produced are enormous, sometimes in the region of petabytes, and visualising them requires specialised HPC graphics clusters.

The main reason turbulence is so computationally expensive to simulate is that as the strength of forcing increases, the grain of scales becomes finer and finer, leading to an almost endless pattern of smaller and smaller vortices. To capture the whole

system, one must simultaneously resolve both the smallest of these vortices and the whole system in which they are evolving. "One can imagine these systems as a room that you want to lay flooring down in," explains Verzicco. "The size of the room stays the same, but the size of the tile decreases and decreases, making the effort required to complete it much bigger. This is why we have to use thousands of processors for tens of millions of CPU hours to develop our databases."

Once the database has been created, it is analysed by testing theories and extracting specific features from it, which can range from essential parameters such as global friction and global heat transfer coefficients, to structures that appear in the flow. Depending on the type of analysis being carried out, this can also involve the use of HPC resources.

A unique point of this project is that the database produced has both Eulerian data in the form of complete snapshots of the flow, as well as Lagrangian data in the form of detailed time evolution of tracer particles injected numerically into the flow. Very little Lagrangian data exists in the literature for this specific type of flow, so these simulations have considerably extended the available data, while unveiling new physics. The results have direct implications on our understanding of transport properties in buoyant flows with applications to engineering, marine and atmospheric science.

One of the main results of the study is an improved understanding of how heat transfer changes as a function of combined thermal and mechanical forcing. The novelty of this study is that the researchers have been able to change each of these parameters individually, giving them greater insight into how each of them affects heat transfer.

Another more specific result of this study involved the researchers looking at flows that mirror those that occur in the oceans. The buoyancy of water in the ocean is determined by both temperature and salinity. As the surface is heated, the water at the top becomes more buoyant, but due to evaporation it also becomes more saline, creating an opposing gradient. This leads to a phenomenon known as salt fingering. When hot, salty water lies over cold fresh water, a thin "finger" of the water above will sometimes begin to sink downwards into the colder, fresher region. Because heat diffuses faster than salt, the finger of water becomes increasingly denser than the water around it and sinks further. Likewise, a small finger of colder, fresher water will be displaced upwards and gain heat by diffusion from surrounding water, which will then make it lighter than the surrounding waters, and cause it to rise further.

These sinking salt fingers play an important role in sustaining life below the surface of the ocean, because the flow created brings down oxygen and other essential nutrients. Verzicco and his colleagues studied the combined effects of salinity and temperature in these flows as part of the project. "The problem is quite a tricky one to study as salt fingers are very thin, so it requires a lot of spatial resolution and adds to the complexity in terms of computational power," says Verzicco.

Instantaneous salinity isosurfaces for the double diffusive convection flow at Ra_S=10^11, Ra_T=4x10^10, Pr_S=21 and Pr_T=7. The density ratio is \Lambda=1.2. The simulation is run on a mesh of 1Gnodes with a refined resolution for the salinity field by a factor 3 in each direction.

Since the end of the project, Verzicco has continued his research into turbulent flows. A quirk of the physics involved in turbulence is that flows which superficially look very different can actually be quite similar when broken down into mathematical equations. For instance, a flow generated by two counter-rotating concentric cylinders (known as a Taylor-Couette flow) is mathematically analogous to the flows that the group studied in the previous project, but is far less computationally expensive to simulate. This has allowed them to increase the complexity of their simulations in other areas, for example by adding gas bubbles into flows of liquid, or droplets of liquid into flows of gas. "Adding this extra level of detail is very rewarding as it more closely reflects the flows found in the ocean and the atmosphere," says Verzicco. "Even small concentrations of bubbles can, for instance, profoundly affect the structure of salt fingers, so this is something we are now working on."

For more information

http://people.uniroma2.it/roberto.verzicco/
https://pof.tnw.utwente.nl/

Resources awarded by PRACE

This project was awarded 68.5 million core hours on FERMI hosted by CINECA, Italy

Publications

Zhu, X., Mathai, V., Stevens, R.J.A.M., Verzicco, R., Lohse, D. "Transition to the Ulti- mate Regime in Two-Dimensional Rayleigh-Bénard Convection", Physical Review Letters, 120(14), 144502, (2018).

Zhu, X., Verschoof, R.A., Bakhuis, D., Huisman, S.G., Verzicco, R., Sun, C. , Lohse, D. "Wall roughness induces asymptotic ultimate turbulence", Nature Physics, DOI: 10.1038/s41567- 017-0026-3, (2018).

Exploring the possibilities of modern laser sources


Luc Bergé

New possibilities for medical imaging and other applications are becoming accessible with the advent of powerful laser sources delivering more intense and ultrashort optical pulses. **Dr Luc Bergé** of CEA has been leading a project which explores the limits of what can be done using state-of-the-art light sources in terms of particle acceleration and secondary radiation.

Modern laser sources deliver optical pulses reaching unprecedentedly high intensity levels by concentrating near-kilojoule energy in ultrashort wave-packets of light, having their duration compressed to a few femtoseconds. The underlying technique used to generate them, known as chirped pulse amplication, was awarded the 2018 Nobel Prize in Physics. Such extreme optical objects question the fundamentals of wave propagation in all materials in general, and in plasmas in particular.

Dr Luc Bergé, director of research at the French Alternative Energies and Atomic Energy Commission, has been leading the CAPITOL project, in which he and his colleagues have been numerically modelling various photon and particle sources generated by intense ultrashort laser pulses interacting with gas or solid targets. There are various potential biological and medical uses for these sources, which range from weakly energetic terahertz radiation for tumour detection, to highly energetic radiations and accelerated ions which can supply low-cost, portable sources for deep-body imagery and cancer surgery.

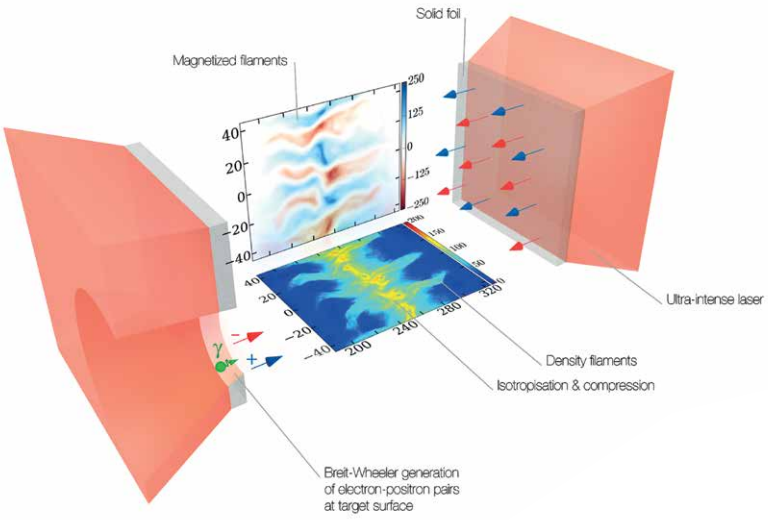

All-optical setup for driving astrophysically relevant collisions of electron-positron pair plasma jets. The pair production results from the decay, via the Breit-Wheeler process, of synchrotron photons radiated by laser-accelerated relativistic electrons in solid foils. The intensity and duration of the two laser pulses are of 10^{24} W cm^{-2} and 60 fs, respectively.

A wide variety of numerical tools were used in the project, including a unidirectional propagation code, a 1D Maxwell-fluid code and multidimensional particle-in-cell (PIC) codes. The PIC algorithm is a time-consuming numerical method that can self-consistently model the kinetic evolution of relativistic plasmas such as those created by intense lasers. Those simulations allowed generic scaling laws to be inferred, which can serve to guide optimisation strategies.

The first task of the project focused on the production of X-ray and gamma-ray sources by ultra-intense lasers. Several radiation mechanisms were explored, mainly in the frame of the "wakefield laser-plasma accelerator". The betatron radiation generated in a laser wakefield accelerator is at the basis of unprecedentedly compact bright X-ray sources, which could lead to important breakthrough in the detection of small breast tumours. "We found that imperfections in the laser beams can be the main cause of the discrepancy between measurements and simulations," says Bergé. "Although the intense, ultrashort pulses exploited for betatron sources are far from ideal Gaussian beams, their imperfections are usually neglected in simulations. In our study, the experiments were modelled in full 3D geometry."

Compared with best-fitting Gaussian beams, this accurate description of the experimental laser profiles revealed decreases of one order of magnitude in the computed photon yields and their growth with the plasma length. The study demonstrated that, for boosting X-ray emission, it is more cost-efficient to increase the laser quality than its power.

A second task concerned laser-induced generation of intense ion beams, addressing the potential of thin foil targets for proton therapy. A more innovative interaction setup employing gas targets was also investigated. These first two actions made extensive use of the PIC kinetic code CALDER in order to quantify the relativistic laser-matter interaction and the related generation of secondary radiation and particles.

The generation of energetic ion beams by intense lasers was investigated within two interaction setups. In the first one, the laser impinged onto an array of metallic micro-pillars supported by

a thin substrate. "In conjunction with measurements performed at the PHELIX laser facility, our PIC simulations confirmed the capability of such targets in enhancing the target-normal-sheath-acceleration (TNSA) of protons," says Bergé. This result was explained by the enhanced production of hot electrons enabled by the modified interaction geometry and the deeper penetration of the laser field into the target structure.

The team also studied ion acceleration from laser-driven collisionless shocks in dilute plasmas. With a few-cycle laser pulse, an electrostatic laminar shock can be triggered in the plasma density down-ramp by the high electron pressure arising from the breaking of the laser wakefield. "We developed a semi-analytical model for the scaling of the hot-electron pressure as a function of the laser-plasma parameters, in good agreement with simulations. We finally examined the formation conditions of turbulent shocks mediated by the Weibel instability in high-velocity plasma collisions, which will be useful for radiation-pressure-induced ion acceleration at next-generation laser facilities."

"Modern laser sources concentrate near-kilojoule energy in ultrashort wave-packets of light"

The final task of the project was to look at the numerical control of broad optical spectra induced by multicoloured femtosecond laser pulses interacting with initially neutral gases, in tightly or loosely focused propagation geometries. The goal was to optimise the terahertz part of the spectra for tumour spectroscopy applications.

Direct comparisons of results provided by CALDER and those obtained from the unidirectional propagation code UPPE3D covered a broad range of laser-plasma configurations and have opened new perspectives in terahertz emission produced over remote distances. The team also performed theoretical investigations including heating by elastic collisions. "Our investigation suggests that the THz energy increases linearly with moderate values of the pump energy, but for higher laser energies, the THz energy yield and the conversion efficiency saturate," says Bergé. Attention has also been given to the transitions between different interaction regimes depending on the laser intensity that can trigger different nonlinear conversion mechanisms, either by the electron currents induced by photo-ionisation (photocurrents) or by the ponderomotive force of the laser. Up to near-relativistic laser intensities, the photocurrents excited by two-colour laser fields remain the dominant mechanism of laser-to-THz conversion.

Investigating laser-matter interaction may require specific electromagnetic-field profiles to be used. A possible method is to prescribe the fields at the box boundaries to achieve the desired fields inside the numerical box. "We showed that the classical paraxial approximation can lead to unexpected field profiles with strong impacts on the interaction results," says Bergé. "We have proposed an efficient algorithm to compute the required laser boundary conditions consistent with Maxwell's equations for arbitrarily shaped, tightly focused laser pulses. It can be parallelised straightforwardly employing domain decomposition, and offers a simple way to simulate more complex pulse configurations or experimental beam profiles."

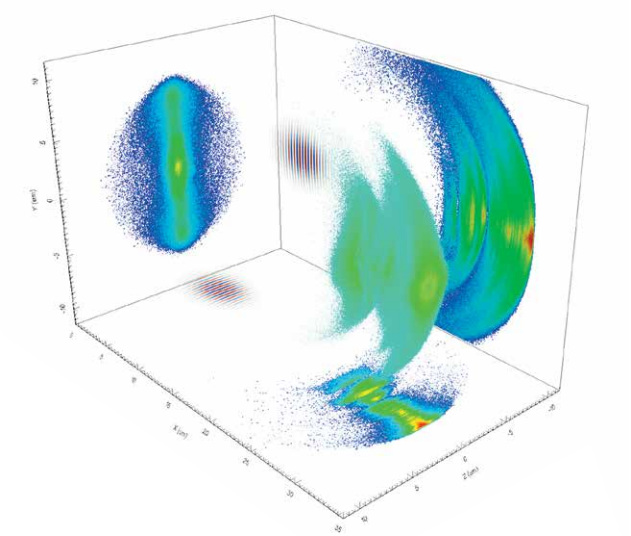

Inverse Compton scattering of a 10^{23} Wcm^{-2}, 15-fs laser pulse off a 3-GeV electron beam generated by a laser wakefield accelerator: 3D distribution of the gamma-ray photons after laser interaction.

As a result of his ongoing research, Bergé was this year awarded the Gentner-Kastler prize, which is given each year to an outstanding physicist in France or Germany whose research has led to top level breakthroughs. Bergé and his team now intend to test new strategies for enhancing the X-ray flux and photon energies, such as plasma density tailoring. Other concepts of micro and nano-structured targets and shaped gas jet profiles will be examined to optimise proton acceleration. Moreover, they are investigating THz pulse generation at relativistic laser intensities, combining both the properties of laser-plasma accelerators and THz emission by coherent transition radiation.

For more information
www.cea-dam.fr

Resources awarded by PRACE
This project was awarded 24.1 million core hours on Curie hosted by GENCI at CEA, France

Publications
J. Ferri et al., "Effect of experimental laser imperfections on laser wakefield acceleration and betatron source", Sci. Rep. 6, 27846 (2016); 10.1038/srep27846 JIF: 5.5

D. Khaghani et al.,"Laser-driven proton acceleration enhancement using micro-pillar arrays", Sci. Rep. 7, 11366 (2017). JIF: 5.5

General relativistic simulations of binary neutron star mergers

Binary neutron star systems are among the most violent astrophysical events, giving off gravitational waves and electromagnetic signals. **Professor Bruno Giacomazzo** of the University of Trento has been investigating this using the magnetohydrodynamic code WhiskyMHD, while at the same time looking into whether these systems could also be the source of so-called short gamma-ray bursts.


Bruno Giacomazzo

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

Sometimes, two neutron stars can bind together in a binary system, orbiting around each other for millions of years as they get closer and closer. Eventually, the two merge together in an instant (just a few milliseconds), resulting either in a black hole or in a rapidly rotating super-heavy neutron star, which can still collapse to a black hole later on.

Binary neutron star mergers are among the most violent astrophysical events and can be observed in several ways. First, they emit gravitational waves – perturbations in the geometry of spacetime that propagate at the speed of light and can be observed as tiny periodic changes in the distance between free-floating objects. Gravitational wave emission is responsible for the decrease in orbital separation in binary neutron star systems and becomes particularly strong towards the end of the long inspiral phase and the time of merger. Modern laser interferometers such as LIGO and Virgo are searching for those signals, and have reached a sensitivity that in August 2017 allowed them for the first time to detect a binary neutron star merger that occurred more than 100 million light-years away. Both LIGO and Virgo will soon start looking again for such signals in 2019.

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

The GRSimStar project, led by Professor Bruno Giacomazzo of the University of Trento in Italy, studied the dynamics of different binary neutron star systems, considering neutron stars with different masses, different mass ratios, different equations of state (describing the internal composition of neutron stars), and different magnetic field configurations. The aim was to compute the gravitational wave signals and electromagnetic emissions that these systems can produce.

"The two neutron stars eventually merge together in the space of a few milliseconds"

The simulations used the WhiskyMHD code and were run on the SuperMUC supercomputer at LRZ. A typical simulation used 2048 cores for around 20 days. More than 100 terabytes of data were generated and analysed by the researchers. Some of the data, including the gravitational wave signals, were also made publicly available. Data of this kind can be compared with future detections of gravitational waves in order to distinguish between different neutron star physical models.

Simulations of a typical merger resulted in a long-lived heavy neutron star surrounded by a dense and thick accretion disk. This was an interesting case for gravitational wave astronomy, since the signal also carried an imprint of the strongly oscillating remnant.

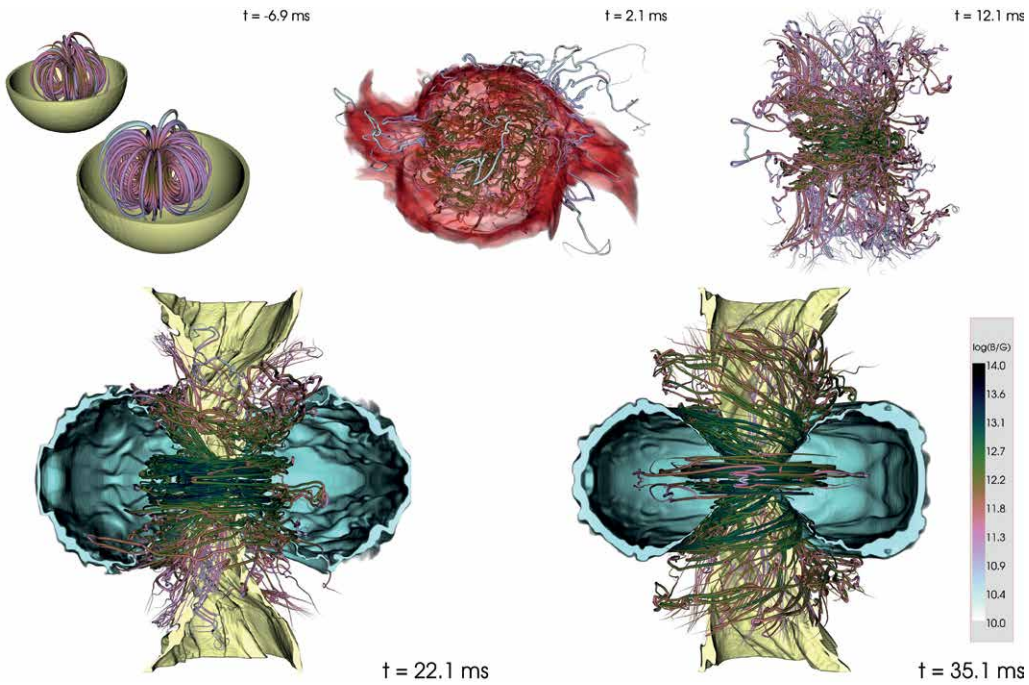

Simulation of a binary neutron star merger. The different panels show different stages of the evolution highlighting in particular the magnetic field structure. The last two panels show the final stages of the simulation where a black hole surrounded by an accretion disk is formed. In this case a magnetised low-density funnel is also formed, which could launch a relativistic jet and possibly produce a short gamma-ray burst at a later time. The color of the field lines gives a rough indication of the field strength. Source: Kawamura et al 2016, Phys. Rev. D 94, 064012.

The researchers also investigated in detail the structure of long-lived remnants, their mass distribution and their rotation profile, which represents key information for the interpretation of future observations. For larger masses, the merger remnant quickly collapsed into a black hole. In this case, the researchers investigated the evolution of the magnetic field, in order to understand if the conditions could support the formation of a relativistic jet, which can lead to short gamma-ray bursts. Their simulations showed the development of a strong toroidal field in the orbital plane followed by the emergence of a strong conical field at the edges of the accretion disk, with a structure reminiscent of a tornado.

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

Currently, the team are running additional simulations to investigate additional models for the properties of dense matter. Furthermore, they used remote visualisation resources provided by LRZ to produce movies showing 3D visualisations of their simulations, which are available now on their website.

One of the main challenges of the simulations was the fact that some important effects leading to magnetic field amplification happen on small length scales. This made it very difficult to resolve them numerically. In order to further improve the accuracy, the team are planning a follow-up study in which they will evolve one or more models with very high resolution and then use the results to calibrate a so-called sub-grid model, which is designed to capture the field amplification on scales not resolved with the lower, more affordable resolutions. Once calibrated, the sub-grid approach will allow the team to investigate a large number of models without the need for very high resolutions.

For more information

www.brunogiacomazzo.org
www.nature.com/news/colliding-stars-spark-rush-to-solve-cosmic-mysteries-1.22829

Resources awarded by PRACE

This project was awarded 15.2 million core hours on SuperMUC hosted by GCS at LRZ, Germany

Publications

Kawamura T., Giacomazzo B., Kastaun W., Ciolfi R., Endrizzi A., Baiotti L., Perna R. 2016, Binary Neutron Star Mergers and Short Gamma-Ray Bursts: Effects of Magnetic Field Orientation, Equation of State, and Mass Ratio, Phys. Rev. D, 94, 064012

Ciolfi R., Kastaun W., Giacomazzo B., Endrizzi A., Siegel D., Perna R. 2017, General Relativistic Magnetohydrodynamic Simulations of Binary Neutron Star Mergers Forming a Long-Lived Neutron Star, Phys. Rev. D, 95, 063016


Jens Harting

Colloids in multiphase flow

Particle-stabilised emulsions have long been studied for their unique properties, which have a number of different industrial applications. **Professor Jens Harting** of the Eindhoven University of Technology has been using simulations to investigate these systems, the results of which are now being picked up by experimental groups and realised in practice.

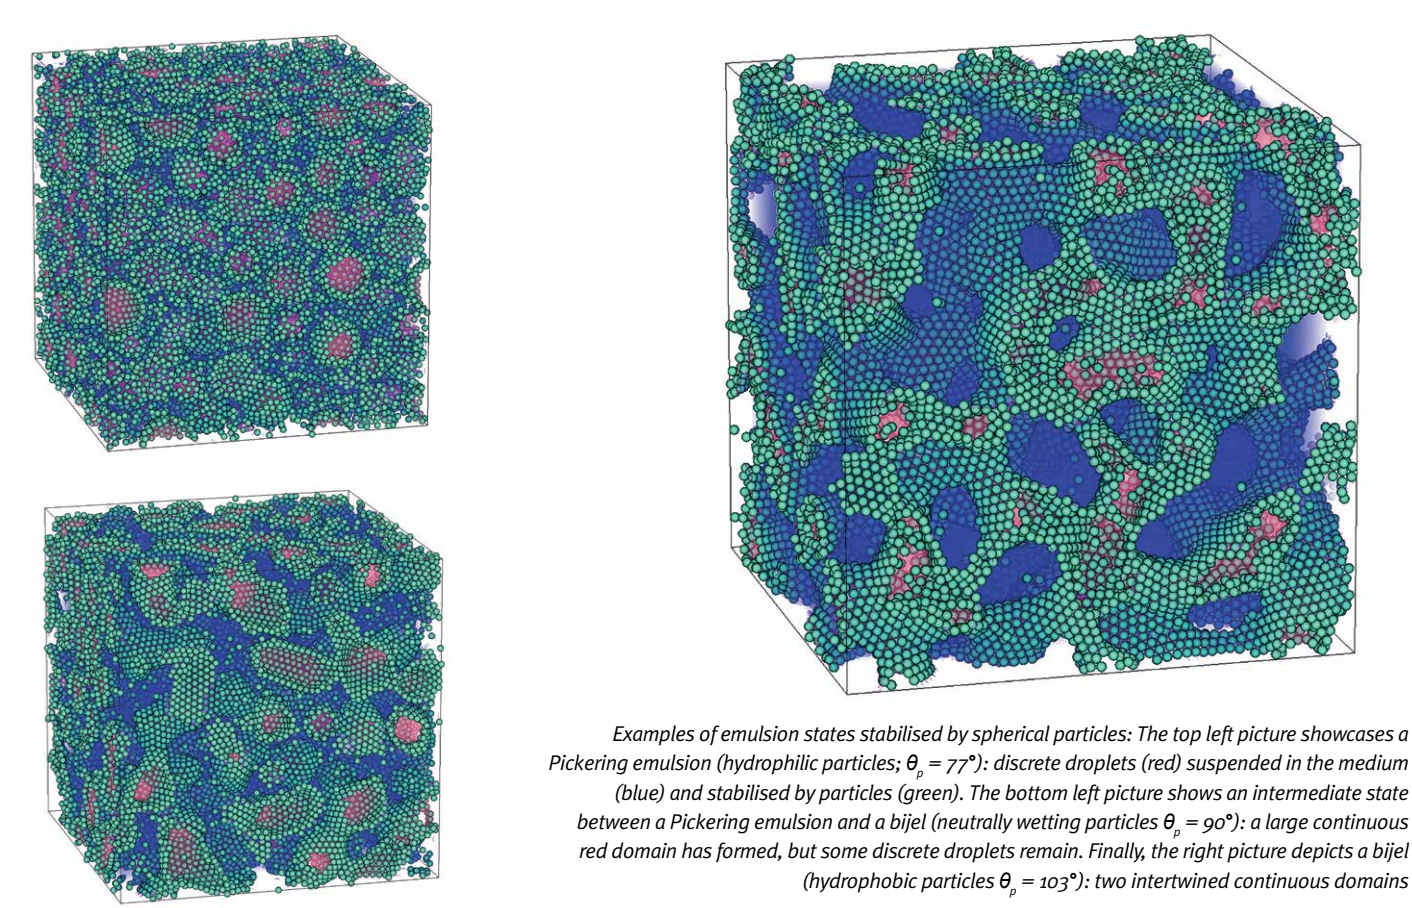

Examples of emulsion states stabilised by spherical particles: The top left picture showcases a Pickering emulsion (hydrophilic particles; $\theta_p = 77°$): discrete droplets (red) suspended in the medium (blue) and stabilised by particles (green). The bottom left picture shows an intermediate state between a Pickering emulsion and a bijel (neutrally wetting particles $\theta_p = 90°$): a large continuous red domain has formed, but some discrete droplets remain. Finally, the right picture depicts a bijel (hydrophobic particles $\theta_p = 103°$): two intertwined continuous domains

Particles as stabilisers for emulsions are used in the food and cosmetics industries, for crude oil recovery, and waste water treatment.However, the properties of such systems are only poorly understood.

challenges in this field are understanding the complex particle-particle and particle-fluid interactions. A team of researchers led by Professor Jens Harting of the Eindhoven University of Technology has been investigating the behaviour of individual particle-covered droplets, as well as the influence of the particle properties on the stability of emulsions.

Such systems are too complicated to be treated with theoretical models, while experimental trial-and-error searches for the best stabilisation method are both expensive and time consuming. “Our large scale simulations, which use combined lattice Boltzmann and molecular dynamics simulations, are able to cover the relevant parameters of such systems,” says Harting. “They have not only led to better fundamental knowledge about the systems, but may also directly lead to improved production or spark the interest of companies in using particle stabilised multiphase systems for their own products.”

Harting and his colleagues originally developed a computational method based on a combination of the lattice Boltzmann method to simulate fluid flows and the discrete element method to simulate suspended particles. “We combined the two and, from the beginning, it was clear that it would be an HPC application because of the different scales involved in the systems,” says Harting.

Harting’s group has always made sure to write their codes with parallel architectures in mind, and received support in this from a number of centres including the Edinburgh Parallel Computing Centre, the Jülich Supercomputing Centre and SURFsara in

Amsterdam. This put them in a unique position to be able to simulate these complex systems on large scales. "The world of HPC is always moving on to new architectures, and we have always had one eye on this," he explains. "This has meant that we have only had to apply minimum effort to move on to the next set of machines, such as the new modular architectures with booster modules."

Particle-stabilised emulsions are commonly used in various industrial applications. These emulsions can present in different forms, such as Pickering emulsions or bijels, which can be distinguished by their different topologies and rheology. Harting's PRACE-supported project, called "COMFLOW – Colloids in multiphase flow", covered a number of sub-topics. The first of these looked at how the distribution of droplet size and structure sizes in particle-stabilised emulsions is dependent on fluid parameters, such as viscosity, surface tension, particle shape, and particle concentration.

"The team's parallel codes put them in the unique position of being able to simulate these complex systems"

The team numerically investigated the effect of the volume fraction and the shape or wettability of the stabilising particles in mixtures of two fluids. For this, they used the well-established three-dimensional lattice Boltzmann method, extended to allow for the added colloidal particles with non-neutral wetting properties. "We obtained data on the domain sizes in the emulsions by using both structure functions and our own Hoshen-Kopelman cluster detection algorithm," explains Harting. "We confirmed an inverse dependence between the concentration of particles and the average radius of the stabilised droplets. Furthermore, we were able to demonstrate the effect of particles detaching from interfaces on the emulsion properties and domain-size measurements, and identified various timescales in the emulsion formation in the case of anisotropic particles."

Harting's team also investigated how to dynamically change the properties of such emulsions from one state to the other, i.e. from a Pickering emulsion to a bijel. "We wanted to be able to switch back and forth between the two just by changing the contact angle, so that we would have a dynamic system whereby changing some property on the outside can quickly change the rheology and other properties," explains Harting. "This could be useful for industrial purposes where, for instance, you may want to switch between low and high viscosity depending on the application."

Following on from this work, the researchers developed a new mechanism for creating self-assembled porous media with highly tunable geometrical properties and permeabilities. This mechanism involves first allowing a particle-stabilised emulsion to form from a mixture of two fluids and colloidal particles. Then, either one fluid phase or the particle layer is solidified, which can be achieved by techniques such as polymerisation or freezing. "Based on computer simulations, we demonstrated that modifying only the particle wettability or concentration results in porous structures with a wide range of pore sizes and a permeability that can be varied by up to three orders of magnitude," says Harting. "These could be perfect for use in chemical reactors or filters."

The final topic that Harting and his colleagues looked at involved similar systems, but with particles at the fluid-fluid interface that have magnetic cores. With these systems, it is then possible to use a magnetic field to apply torque to the particles, and thus deform the interface and create very fine structures. Although these systems have not been studied extensively yet, there are a number of possible uses for them, including the printing of electronics and photovoltaics, as well as acting as an alternative electronic ink in e-readers.

The results obtained in the PRACE-supported project have led to a number of publications and are a substantial part of the PhD theses of Stefan Frijters, Florian Günther, Qingguang Xie, and a collaborator from University College London, Gary Davies.

The results of this project were recently picked up by an experimental group, who have been able to create the particles that Harting's group have theorised. "They have shown experimentally that what we proposed via simulation actually works in practice," says Harting. "Since then, a number of other groups have been working in a similar direction."

For more information
www.hi-ern.de

Resources awarded by PRACE
This project was awarded 26 million core hours on JUQUEEN hosted by JSC

Publications
Self-assembled porous media from particle-stabilized emulsions. S. Frijters, J. Harting
Submitted for publication (2014), arXiv:1408.2974

Controlled capillary assembly of magnetic Janus particles at fluid-fluid Interfaces. Q. Xie, G. B. Davies, J. Harting
Soft Matter 12, 6566 - 6574 (2016)

Soft particles at a fluid interface. H. Mehrabian, J. Harting, J. H. Snoeijer
Soft Matter 12, 1062-1073 (2016)

3D atmospheric simulations of astrophysical objects

PHOENIX is a general-purpose state-of-the-art stellar and planetary atmosphere code, used to calculate atmospheres and spectra of astrophysical objects. **Professor Peter Hauschildt** and his colleagues from the University of Hamburg are currently working on the first 3D version of the code, in preparation for the computing power which will become available over the next 5-10 years.

Over the last decade, a number of detailed 3D simulations of the gas dynamics in the outer layers of stars have become feasible. The most recent of these simulations also include the effect of magnetic fields on the gas (and vice versa). Such calculations allow for the investigation of the time dependent behaviour of the outer layers of active cool stars (M dwarfs), called the photosphere and the chromosphere. In order to 'connect' such gas dynamic simulations with the observable light from the star, a very complex non-local thermodynamic equilibrium radiative transfer simulation is necessary.

The PHOENIX code is a general-purpose model atmosphere simulation package, designed to model the structures and spectra of objects ranging from extrasolar planets (both terrestrial and gas giants) to brown dwarfs and all classes of stars, extending to novae and supernovae. The main results from the calculations are synthetic images and spectra (and derived quantities, such as colours), which can be directly compared to observed spectra and, in 3D simulations, images. By adjusting the simulation parameters and comparing the results to the observed data, the physical parameters of the stars or planets can be determined.

"We knew from the beginning that it would be 25 years before we could start properly using the code"

Scientists using state-of-the-art telescopes such as Hubble and the ESO are constantly gathering observations from a wide range of astrophysical objects. "The data that our observational colleagues are gathering now is really amazing," says Professor Peter Hauschildt of Hamburger Sternwarte. "It is this data that is the driving force behind the PRACE-supported project PHOENIX, because if you have excellent data, you need to put a lot of time, effort and resources into analysing it.

Peter Hauschildt

"Our group has been modelling stellar and planetary atmospheres for around 25 years now," he continues. "When we began, we were doing this in one dimensional mode, meaning that everything was perfectly spherical and symmetrical. Now, working in three dimensions, the whole problem is a lot more complex and harder to calculate."

The 3D version of the code has been in development for around 10 years now and, compared to the original 1D version, is somewhat of a resource monster, requiring both MPI and OpenMP to get the maximum performance out of the systems currently available. "Currently the largest version of the model runs on about 2.5 million MPI x OpenMP threads, which is the largest we can get the machine to run for more than a few hours before the whole thing goes down," says Hauschildt.

The computing power and time requirements for a realistic PHOENIX/3D model are staggering. This is due to the fundamentally 7D nature of the radiative transfer modelling (for a 4D space-time). For the simulations at HLRS on Hazel Hen, a typical full iteration with 50 000 MPI processes and 2-4 OpenMP threads per process requires approximately three hours, and about 100-200 iterations are needed for convergence. Therefore, about 30 million core hours are needed to compute a single PHOENIX/3D model for a high-resolution gas dynamics snapshot. A minimum of 50 GB of output are generated per iteration, and the final output

The figures below are from the dissertation of Ivan de Gennaro-Aquino (Hamburg 2017).They show the z (vertical) component of the radiation flux at the surface of the simulation box (this is approximately what can be observed) for 3 wavelengths. x-y are the horizontal coordinates, the color tables give the 'brightness' (linear scale).

Figure 1: 2796.352 Angstrom: resonance line of singly ionized Magnesium in the UV

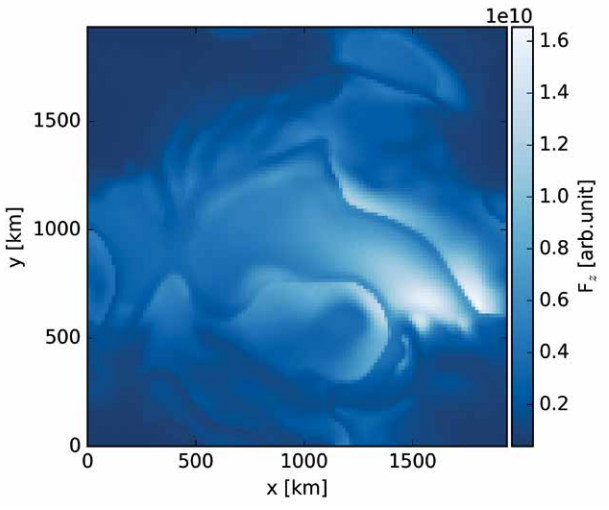

Figure 2: 6300.000 Angstrom: optical (red) wavelength

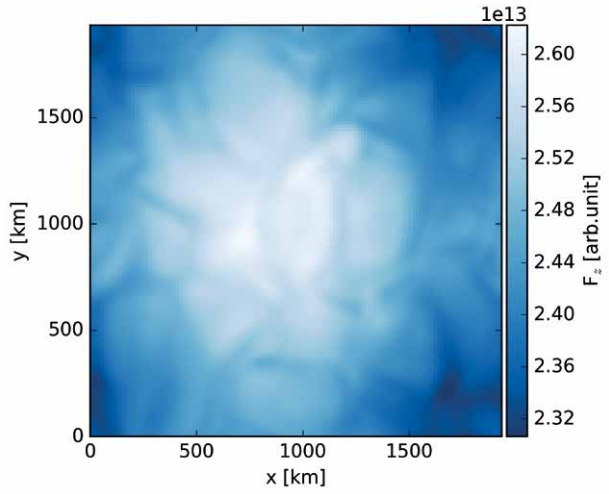

Figure3: 16000.000 Angstrom (1.6 micro-meter): Near infrared wavelength.

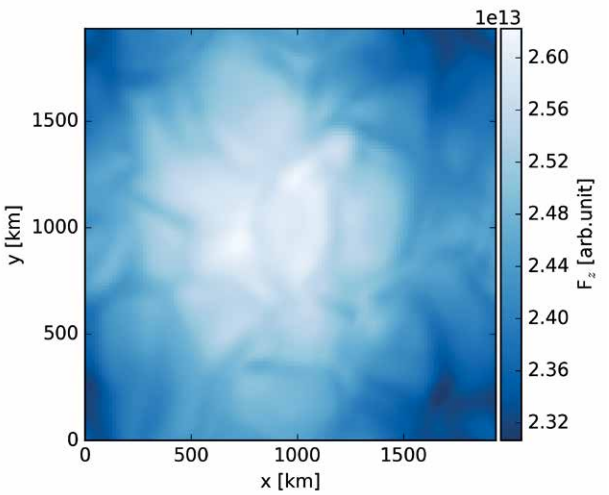

required for the generation of images and observables is well in excess of 1 TB.

The basic physics modules of the code are now ready, but current computers are too small to run full-scale productions of it without restricting it in some way. "We knew from the beginning when we started working on this code 10 years ago that it would be in the realm of 20-25 years before we could start properly using the code," explains Hauschildt. "In principle we can now run the model, but it would take around a year and a half for a single run, which is clearly impractical. So it is now a matter of waiting for the computers to get faster while running simulations at the bleeding edge of HPC."

The code can be used for a variety of purposes. For instance, it could be used to simulate the spectrum of light reflected and transmitted on the secondary star of a binary star system, or any other similar system. As a general radiation transport code, it could also be used for terrestrial applications.

At the moment, one of Hauschildt's PhD students is simulating the spectrum of a star system which already has a lot of observational data and therefore can be used as a blueprint for test runs. This will provide them with invaluable experience for the future when they eventually move on to exoplanet spectrum analysis.

Other continuing work involves scaling up and optimising the programme for specific architectures such as KNL (which has been done with great success) and NEC Aurora, and also to involve the use of GPUs, following the trend of many of the new supercomputing devices that will be available in the coming years. This will involve re-writing the critical pieces of code that take most of the time. "By optimising just a few thousand lines of code, we can account for around 80-85% of the total running time," says Hauschildt. "This is where a lot of our effort will be focused in the coming years."

For more information

http://hobbes.hs.uni-hamburg.de/
https://www.hs.uni-hamburg.de/index.php?option=com_content&view=article&id=14&Itemid=294&lang=en

Resources awarded by PRACE

This project was awarded 28.8 million core hours on HORNET hosted by GCS at HLRS, Germany

Publications

"Preparing NERSC users for Cori, a Cray XC40 system with Intel many integrated cores", Concurrency and Computation Practice and Experience 30(4), August 2017

Simulating accretion at the centre of the Milky Way

Many accretion disks — most notably the one feeding the black hole at the centre of our galaxy — are too hot and diffuse to be treated as fluids in simulations. **Professor Matthew Kunz** of Princeton University has been carrying out ab initio calculations of these bodies of plasma to try and more rigorously explain their behaviour.

Ninety-nine per cent of the universe is made up of plasma, explains Princeton University's Matthew Kunz as he begins describing his research. "That's the standard line scientists in this field give in interviews – it's how we legitimise ourselves!".

Plasmas – the fourth state of matter, along with solids, liquids and gases – are, like Kunz says, extremely common on an astronomical scale. All stars are made of plasma, as is most of the interstellar medium from which stars are born. Plasmas differ from gases in that most of the electrons in them have been stripped away from their atomic nuclei, leaving a mixture of positively and negatively charged particles. This means that, unlike gases, they can be strongly influenced by electric and magnetic fields.

The behaviour of many of the most energetic phenomena in the universe is dictated by plasma. At the centre of our galaxy there is a supermassive black hole with the mass of about four million suns. It is possible to infer from electromagnetic emissions that there is a large, donut-shaped reservoir of plasma — an accretion disk — orbiting and falling into this black hole.

In order for matter to spiral inwards towards a black hole, it must lose a certain amount of its angular momentum. As it does so, the released gravitational energy is often emitted in the form of radiation. However, researchers have noticed that a number of accretions disks — including the one surrounding the black hole at the centre of our galaxy, known as Sagittarius A* — do not give off as much radiation as one would expect. Therefore, the angular momentum must be being dissipated via another channel.

As it turns out, this can happen through what is known as the magneto-rotational instability. The phenomenon allows angular momentum to be transferred from one area of plasma to another via a shared magnetic field, which acts like a tether between the disparate elements. In the context of an accretion disk, this causes the areas of plasma with a lot of angular momentum to gain even more, while the areas with less lose it. "This means that some of the plasma falls out and some of it falls in, and overall we see a net inward accretion of material," explains Kunz.

Since the early 1990s, physicists have been using computers to try and answer questions about how plasma behaves in such a

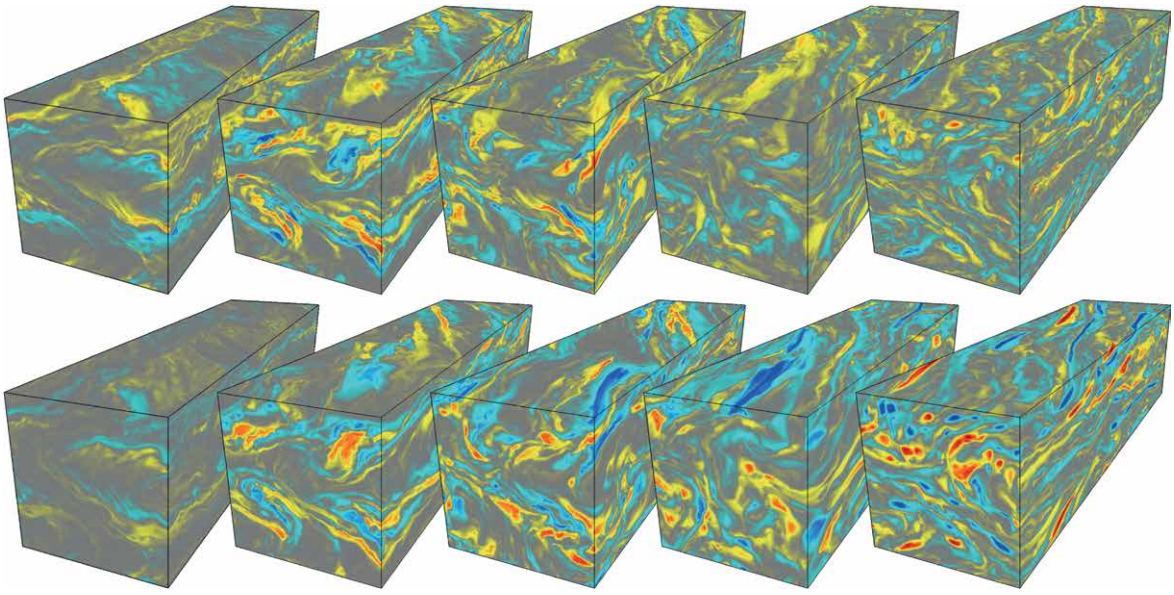

Pseudo-colour images of magnetic-field fluctuations from a simulation of magneto-rotational instability in a collisionless plasma. Time flows from left to right, highlighting the development of the instability and its breakdown into fully developed turbulence.

system. However, because of the limitations of computing power, this is often done using equations normally used to describe fluids (liquids and gases) under the influence of a magnetic field. This approach describes some black hole systems remarkably well, but falls short when used on plasmas which are extremely hot and diffuse, such as in the case of the Sagittarius A* hole.

The failure of these fluid equations to adequately describe the behaviour of plasma in such systems comes down to fundamental differences between the states of matter. "In liquids and gases, the molecules are constantly rubbing against each other," explains Kunz. "In brighter accretion disks, the plasma behaves similarly and therefore can be treated as such in simulations. But extremely hot and diffuse plasmas are better described as collections of charged particles which, in reality, rarely interact." Because of this, fluid equations are not guaranteed to accurately describe the process that causes collisionless disks to grow unstable and spiral down.

Kunz and his team from Princeton University have developed a rigorous new ab initio method for modelling these low-luminosity accretion disks, which systematically traces the paths of individual collisionless particles. This complex approach, conducted using the Pegasus computer code developed at Princeton by Kunz, Stone and Xuening Bai, produced a set of equations better able to model the behaviour of the disk that orbits the Sagittarius A* hole.

"Plasma is better described as a collection of charged particles than as a fluid"

The new method was applied in a PRACE-supported project using 26 million core hours on the Curie supercomputer at GENCI. The project marks the first instance of an ab initio 3D simulation of magneto-rotational turbulence in a hot accretion flow. Kunz's team has established that such a calculation, on the frontier of plasma physics and computational research, is feasible, and since the project finished, a number of other groups have begun carrying out similar investigations with particle-in-cell codes.

One of the main outcomes of the project has been to show that the results obtained in the past using simplified fluid equations are actually fairly accurate. "Proving that fluid equations give a good idea about how plasma in these systems behaves is probably a great comfort to many of the people working in this field," says Kunz. "We have shown definitively that two decades of work in the field has not been based on false assumptions!"

A lot of what is known about plasma in an astrophysical context is based upon measurements of solar winds. Solar winds — the


Matthew Kunz

phenomena which cause the spectacular auroras seen at the poles of our planet — are caused by ejections of hot diffuse plasma from the sun. Kunz's simulations have shown that many of the assumptions made about other hot diffuse plasmas based on measurements of solar wind are correct and can be applied to highly energetic systems observed throughout the universe. "What we have shown is that our extrapolations of knowledge about plasma are, by and large, correct," says Kunz. "This reinforces the notion that we can take knowledge about one system that we can measure and extrapolate it to something that we can't measure."

For many groups, simulations such as the ones carried out by the Princeton researchers are still not feasible due to the computational power needed. However, the results of this project are already being fed back into simplified models that are much less expensive to run. This means that the microphysical details learned from the more rigorous calculations can now be more accurately predicted by groups with less computational resources at hand.

For more information
https://www.astro.princeton.edu/~kunz/Site/Welcome.html

Resources awarded by PRACE
This project was awarded 26 million core hours on Curie hosted by GENCI at CEA, France

Publications
Kunz, Matthew W.; Stone, James M.; Quataert, Eliot "Magnetorotational Turbulence and Dynamo in a Collisionless Plasma", Physical Review Letters, Volume 117, Issue 23, id.235101

Accelerating dehydration to probe water persistency


Walter Rocchia

Water molecules can play a significant role in the binding and unbinding of pharmaceutical drugs to their target proteins. **Dr Walter Rocchia** of the IIT Central Research Lab Genova has been investigating this using new computational methods that could offer a platform for the technological transfer of molecular dynamics-based methods into the fast-paced field of drug discovery.

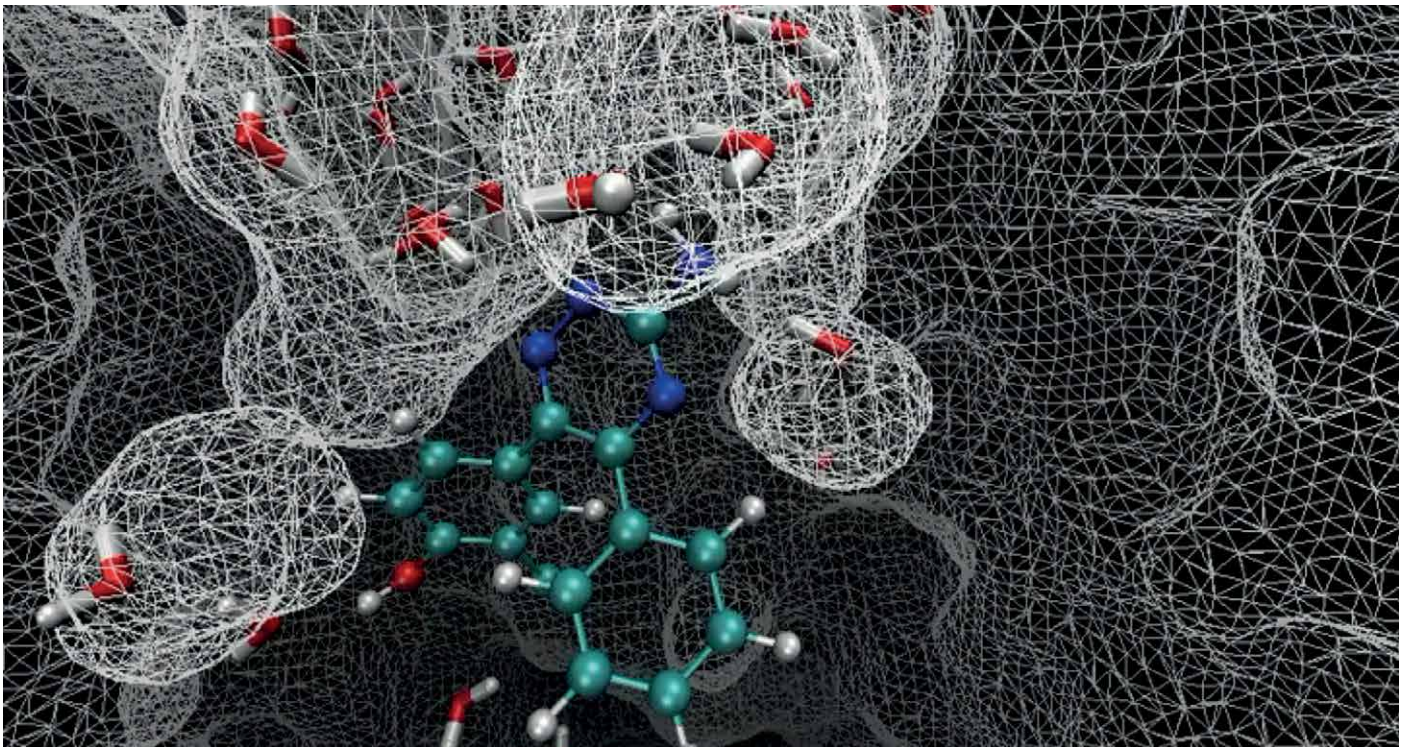

Example of application of the method for comparing two different inhibitors of the A2A G- Protein Coupled Receptor.

Massively parallel computing has brought forth a new age of discovery in the scientific world, allowing for seemingly insurmountable problems to be broken down into smaller ones which can be calculated simultaneously to reduce the time it takes to solve. Much of the work done on modern codes for supercomputers now revolves around optimising them for use on parallel computers.

In spite of this, there are certain problems which are not so easily addressed by this method. Molecular dynamics simulations, which deal with the evolution of systems containing hundreds of thousands of atoms, are restricted to using very small time-steps in the region of femtoseconds. It is, of course, impossible to break down a continuous timeline into smaller parts that can be calculated simultaneously, because the initial state of each step is dependent on the previous one. This constraint restricts the amount of time that can be recreated to the range of microseconds.

This becomes an issue when dealing with phenomena that only happen in rare circumstances at the biomolecular level, such as the binding or, even more markedly, the unbinding of a drug to a target protein. Under typical experimental conditions, it may take in the range of seconds or minutes for these types of uncommon events to occur, making it impossible to significantly sample them or obtain thermodynamic data about them using ordinary simulation methods.

In order to get around this issue, researchers have developed ways to bias their simulations in order to force specific events to occur. These methods are known as enhanced sampling techniques, and

it is necessary to take steps afterwards to remove the effects of the bias in order to gain realistic data.

Walter Rocchia and colleagues from the IIT Central Research Lab Genova have been leading a PRACE-supported project which has been testing a new tool they have developed for analysing hydration patterns in biomolecular systems. Water is an important variable when considering how, for example, a drug will bind to the binding site of a target protein. This is because the interaction between any two biological molecules must compete with their interaction with surrounding water molecules.

Increasingly often, researchers are finding water molecules that are involved in maintaining protein structure and dynamics, in forming protein–ligand complexes, in allosteric regulation of receptors, in protein–protein interactions, and in other biological mechanisms, many of which are crucial for drug discovery. This is because displacing these water molecules can indirectly affect the strength of the molecular complex one would like to modulate. However only a small subset of the water molecules can be termed "structural" in that they are crucial for a protein's structure (and thus its activity), for instance, by forming a bridge between different protein regions, or being relevant for binding.

"Water is an important variable when considering how a drug will bind to a target protein"

Most enhanced sampling techniques require some knowledge of the variables and the degrees of freedom required to make the specific event occur. However, Rocchia is aiming for approaches that do not require such information. "I wanted to see whether it was possible to find biases that are general enough to help with the initial exploration of such phenomena, especially in the field of molecular recognition," he explains.

Rocchia's lab previously developed a piece of software called NanoShaper for building the surface of molecules, according to the definition of a spherical nanoscale probe (representing the solvent) that rolls over the surface to create the so-called Solvent Excluded Surface of, for example, a protein. Using molecular dynamics, the researchers have been carrying out simulations of the adenosine A2A receptor, a target of significant pharmaceutical relevance. In these simulations, a contrived electrostatic field is applied between water molecules and the binding site to see whether, under the effect of a repulsive bias, it was possible to distinguish between water molecules that have a prominent interaction with the binding site and those that do not.

The team then post-processed the trajectories obtained from the simulations using NanoShaper, and were able to identify regions of the surface of the protein which were easily accessible for water molecules and those that were not. Using this information, they then differentiated between the water molecules which were persistent in the hard-to-access regions and those which were persistent in more stably accessible regions, assuming that the latter were the most likely to play an important role in drug binding. The results were not only interesting on their own, but have allowed the researchers to validate their tool.

Now that the tool, which can be used to apply a contrived electric field anywhere in a simulation, has been validated, the researchers are looking to apply it elsewhere. For instance, in a paper published in *The Journal of Chemical Theory and Computation*, the team has moved beyond applying the bias between the protein and water molecules, and has started to explore applying it between the binding site and the drug itself. "We have created what is essentially a highway of electric field lines which dynamically adapt to the possible deformation of the binding site," says Rocchia. "Our simulation results have been shown to be extremely similar to comparable experimental results showing the crystallographic structures of the complex."

Rocchia is now collaborating with the PhD student who ran the original project, Dr Rehana Syeda Zia, using the same tool to see whether they can accelerate the binding of drug molecules by changing the effective hydrophilicity of both the protein and the drug. "By applying this approach, we can again see whether we can discover something about the binding process without any previous knowledge of which atoms are involved," says Rocchia. The fact that hydration is such a general factor in these interactions means that there is plenty of room for further exploration of the new tool.

For more information
https://concept.iit.it

Resources awarded by PRACE
This project was awarded 67.9 million core hours on FERMI hosted by CINECA, Italy

Publications
A Spitaleri, S Decherchi, A Cavalli, W Rocchia, "Fast dynamic docking guided by adaptive electrostatic bias: the MD-Binding approach", J. Chem. Theory and Comp., 14 (3), pp 1727–1736, 2018

S. Decherchi, G. Bottegoni, A. Spitaleri, W. Rocchia, A. Cavalli, "BiKi Life Sciences: a New Suite for Molecular Dynamics and Related Methods in Drug Discovery", J. Chem. Inf. Mod. 58, 2, 219-224, 2018


Jost von Hardenberg

Climate SPHINX reloaded

Climate modelling is a hugely computationally demanding task, and sometimes a model's limitations can stop it from adequately describing certain phenomena. **Dr Jost von Hardenberg** of ISAC-CNR and collaborators have been exploring the effects of resolution and stochastic parameterisations using Tier-0 supercomputers

Climate prediction is currently one of the most computationally challenging problems in science and yet also one of the most urgent problems for the future of society. It is well known that a typical climate model (with a typical resolution of ~120-km in the atmosphere and ~100-km in the ocean) is unable to represent many important climate features, such as the Euro-Atlantic variability.

Indeed, recent studies have shown that climate models at much higher resolutions (up to ~16km) simulate the Earth's climate more realistically. Whilst few would doubt the desirability of being able to integrate climate models at such a high resolution, there are numerous other areas of climate model development which compete for the given computing resources (for example, the need to incorporate additional Earth system complexity).

Instead of explicitly resolving small scale processes by increasing the resolution of climate models, a computationally cheaper alternative is to use stochastic parameterisation schemes. The main motivation for including stochastic approaches in climate models is related to the observed upscale propagation of errors, whereby errors at very small scales (only resolved in high horizontal resolution models) can grow and ultimately contaminate the accuracy of larger scales in a finite time.

The Climate SPHINX (Stochastic Physics HIgh resolutioN eXperiments) project has investigated the sensitivity of climate simulations to both model horizontal resolution and stochastic parameterisations and exploring extremely high resolutions. Led by Dr Jost von Hardenberg at ISAC-CNR, the experiments included both historical and scenario projection in order to unveil whether resolution or stochastic schemes impact the way models represent climate change.

It was shown that both resolution and stochastic perturbations lead to improvements in the representation of the climate variability

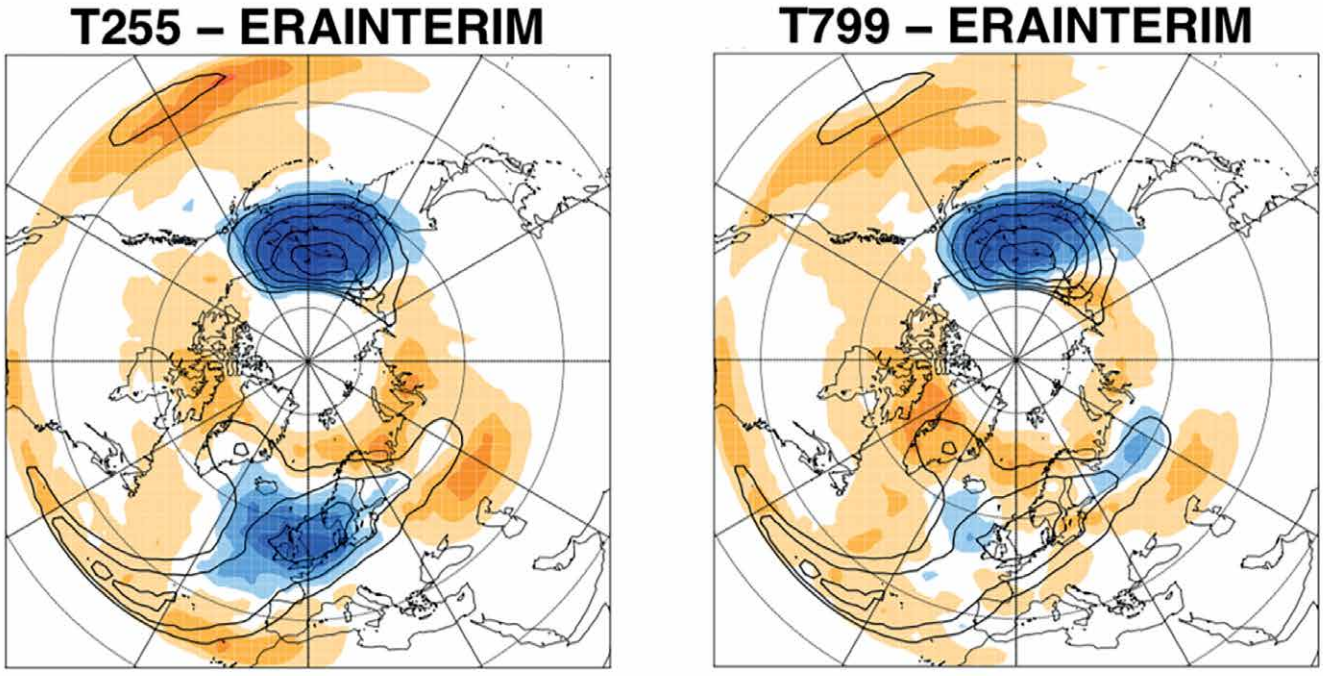

Atmospheric blocking bias for the present-day climate SPHINX simulations compared to a reanalysis product (ERA-Interim). The model bias is reduced considerably with increasing resolution (see the blue section over central Europe disappearing at T799, i.e. at about 40km resolution).

rather than of the mean climate. However – unsurprisingly – different phenomena showed different sensitivities.

For instance, tropical rainfall variability seemed to benefit from both increased horizontal resolution and stochastic parameterisation, whereas a tropical convective phenomenon known as Madden–Julian oscillation only saw improvements when stochastic perturbations were added. Another group has been using the data produced in these simulations to study the frequency and the intensity of tropical cyclones. It appears that stochastic physics give a much better representation of their frequency.

In general in the tropics, applying stochastic schemes at low resolution led to interesting improvements; this was also observed in coupled simulations when looking at the impact of stochastic perturbations on the El Niño–Southern Oscillation. On the other hand, increasing resolution beyond a certain point did not seem to further improve the tropical variability.

"Typical climate models are unable to represent many important climate features"

Conversely, in the mid-latitudes, where atmospheric blocking frequencies were analysed, no statistical difference was found between stochastic and deterministic runs. Here, resolution played a key role, especially over the Euro-Atlantic sector.

In addition, a phenomenon known as elevation dependent warming was explored, which shows that some mountainous regions of the Earth are warming up faster than others. Increasing resolution showed that certain mechanisms which were previously thought to be less important to this were actually the main drivers behind it.

A natural complementary follow up of the experiments carried out under the Climate SPHINX umbrella is to include other components of the climate system in the experimental framework, namely the ocean and land components. This was done in a follow up project called Climate SPHINX reloaded. This has explored the impact of stochastic physics (applied to different components of the climate system) in long climate coupled integrations as a function of model resolution. The systematic comparison between uncoupled simulations from Climate SPHINX and coupled simulations from Climate SPHINX reloaded will help produce new process-oriented studies tackling the origins of well-known biases affecting climate models, such as the double-ITCZ tropical bias.

During the Climate SPHINX projects, the University of Oxford developed a stochastic parameterisation package for the ocean and sea ice models in EC-Earth. In the ocean model, the scheme represents variability in ocean mixing, firstly by perturbing eddy mixing in the Gent McWilliams scheme, secondly by perturbing vertical mixing in the case of unstable stratification, and thirdly by perturbing the shear and buoyancy tendencies required for the calculation of turbulent vertical mixing. In the sea ice model, the sea ice strength parameter is stochastically perturbed to represent sub-grid scale variability in the ice properties. It is also important to represent model uncertainty in other components of the global climate system, and the University of Oxford has helped in this endeavour through the development of a stochastic land surface scheme to represent uncertainties in this component.

The projects have now produced over 220 terabyte of data, and that is after having made many compromises about which climate variables to save at different time frequencies. So far, analysis of the experimental results from the project has led to the publication of four papers and the submission of another three. Von Hardenberg's team members are continuing to analyse their results and are also looking to add more simulations to their repertoire to investigate specific scientific issues arising from the analysis.

For more information

www.to.isac.cnr.it/sphinx

Resources awarded by PRACE

This project was awarded 10.1 million core hours on MARCONI hosted by CINECA, Italy

Publications

Palazzi, E., Mortarini, L., Terzago, S., von Hardenberg, J., Elevation-dependent warming in global climate model simulations at high spatial resolution, Climate Dynamics (2018). **https://doi.org/10.1007/s00382-018-4287-z**

Watson, P. A. G., Berner, J., Corti, S., Davini, P., von Hardenberg, J., Sanchez, C., Weisheimer, A., & Palmer, T. N. (2017). The impact of stochastic physics on tropical rainfall variability in global climate models on daily to weekly timescales. Journal of Geophysical Research: Atmospheres, 122(11), 5738-5762, **https://doi.org/10.1002/2016JD026386**

Davini, P., von Hardenberg, J., Corti, S., Christensen, H. M., Juricke, S., Subramanian, A., Watson, P. A. G., Weisheimer, A., & Palmer, T. N. (2017). Climate SPHINX: evaluating the impact of resolution and stochastic physics parameterisations in the EC-Earth global climate model. Geoscientific Model Development, 10(3), 1383–1402. **https://doi.org/10.5194/gmd-10-1383-2017**

Measuring the transition temperature of quark-gluon plasma

In the beginning — shortly after the Big Bang — there was quark-gluon plasma. **Dr Christian Schmidt** of the University of Bielefeld has been using simulations to examine this embryonic state of matter and compare findings with experimentalists using particle colliders.

Hadrons, such as the protons and neutrons in atomic nuclei, are made up of elementary particles called quarks. Quarks generally only exist in these bound states, held together by other particles called gluons which transfer the fundamental force known as the strong interaction.

However, under extreme conditions, it is possible to break this strong interaction and release the quarks and gluons from their bound states. The Large Hadron Collider (LHC) at CERN and the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory in the US routinely smash atomic nuclei together at close to the speed of light. The extremely high temperatures created by these collisions cause the nuclei to melt into tiny droplets of quark-gluon plasma — otherwise known as quark soup — in which the bound states are broken and the quarks and gluons move freely.

The free movement of these quarks and gluons is only short-lived. As the quark-gluon plasma expands and then cools, the quarks and gluons return back to their bound states as hadrons. This causes thousands of these newly-formed particles to explode in every direction, which are then measured by detectors in the collider. The temperature at which the freely moving quarks and gluons return back to bound states is known as the freeze-out point.

Christian Schmidt

"Under extreme conditions, it is possible to break apart hadrons and release the quarks and gluons from their bound states"

Dr Christian Schmidt of the University of Bielefeld is a particle physicist who uses high performance computers to power simulations of this transition. His interest in quark-gluon plasma stems from the fact that it is likely the state at which matter existed in the early universe shortly after the Big Bang. He uses his simulations to try and explain the experimental data that is gathered from colliders such as the LHC.

The points of interest in these explosions of particles at the freeze-out are known as event-by-event fluctuations. These are subtle changes in parameters such as electric charge, baryon number charge, and strangeness (the occurrence of 'strange' quarks). These are measured in collider experiments but can also be examined using what are known as lattice simulations. Schmidt and his colleagues compare data from these simulations and experiments to try and deduce quantities such as the freeze-out temperature and the density at which the freeze-out occurs.

Schmidt's simulations consist of huge four-dimensional lattices in which each point describes a point in space-time. A simple description of the lattice would assign a quark to each point of this lattice, with the links between them representing

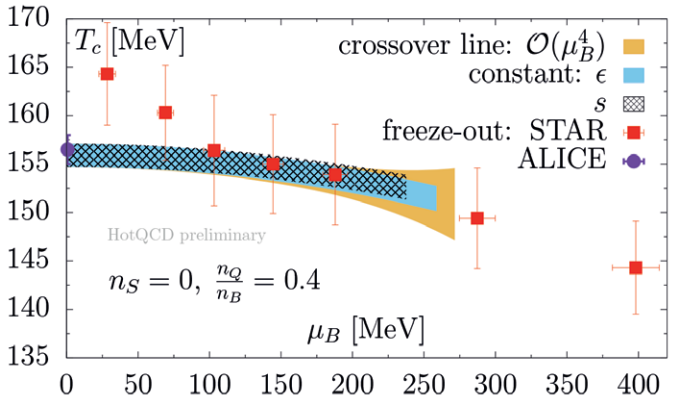

The transition temperature in QCD as a function of the baryon chemical potential, which is the parameter that controls the baryon-number density. Also plotted is the freeze-out temperature as measured in different experiments (STAR at RHIC and ALICE at LHC), based on the various measured particle yields that have been fitted to a Hadron Resonance Gas (HRG) model.

gluons. However, as a representation of a quantum field theory, the lattice is actually an approximation of a continuous field — a wave function — of quarks, with the gluons acting as the interaction between points in the continuum. Each point in the lattice can interact with its nearest neighbours and their neighbours.

Many different versions of these matrices have to be generated, due to the nature of the hybrid Monte Carlo algorithm that is used. The fluctuations of different parameters throughout the lattice such as strangeness and charge are measured at different temperatures. Each lattice is then examined to see whether the energy has decreased or increased, and is accepted or rejected on this basis. Once the lattice has been fully resolved it has to be inverted a number of times to make sense of the data. It is for these reasons that high-performance computers are essential for this kind of investigation.

One of the main results to come out of Schmidt's research has been the most precise measurement of the transition temperature ever recorded, clocked at 156.5 +/-1.5 mega electron volts, which equates to a temperature about 100 000 times hotter than the core of the sun. The points at which these collisions occur in particle collider experiments are ephemerally some of the hottest places in the universe.

"The simulations consist of huge four-dimensional lattices in which each point describes a point in space-time"

Another result relates to how this transition temperature can change when baryon density is varied. By looking at a number of different densities in their simulations, Schmidt and his team have shown that the transition temperature decreases in a parabola-like curve as baryon density increases. These results fit well with the changes in freeze-out temperature that have been measured from experimental data.

The idea that the temperature of freeze-out decreases as baryon density increases was originally postulated through the extrapolation of experimental data using relatively simple models. Now, Schmidt and his team have offered an alternative to these models by providing ab initio calculations from their lattice simulations which, as it turns out, agree very well with these models.

At the current level of discretisation, Schmidt and his team have not been able to obtain good continuum extrapolation or statistical certainty on some of the quantities they wished to measure. In the future, they intend to carry out simulations with even finer lattices to address this.

The knowledge gained from this research will allow for a deeper understanding of the evolution that occurred in the early universe and that finally led — luckily for mankind — to the evolution of life. Furthermore, places may exist in the universe where quark-gluon plasma can be found even today. Apart from the man-made tiny droplets found in colliders, it could be possible that quark-gluon plasma still exists in the core of some super heavy neutron stars.

For more information
www.physik.uni-bielefeld.de

Resources awarded by PRACE
This project was awarded 3 million core hours on MARCONI hosted by CINECA, Italy

Publications
A. Bazavov et al. [HotQCD Collaboration], Skewness and kurtosis of net baryon-number distributions at small values of the baryon chemical potential, Phys. Rev. D 96 (2017) no.7, 074510, [arXiv:1708.04897 [hep-lat]].

C. Schmidt and S. Sharma, The phase structure of QCD, J. Phys. G 44 (2017) no.10, 104002, [arXiv:1701.04707 [hep-lat]].

Mechanisms of drug binding to the influenza A virus

New strains of influenza are resistant to the standard drug used to treat it, amantadine. **Professor F. Javier Luque** of the University of Barcelona has been investigating how amantadine binds to the virus, and has been looking for potential new binding sites with the idea of creating more powerful "multi-target" therapies

Influenza A is a viral infection which can be highly contagious among animals and humans. Outbreaks of the infection provoke symptoms such as acute respiratory infections, temperature and muscle pain, but with a different morbidity and death pattern than common influenza.

Professor F. Javier Luque of the University of Barcelona has has spent much of his career helping with the development of new drugs that can be used to fight influenza A. Up until recently, a drug known as amantadine was the main compound used for this task, but recent years have seen the appearance of mutant viruses which are resistant to it.

During an influenza A infection, a protein on the virus known as the M2 proton channel enables the entrance of protons inside the virus. This facilitates the release of the genetic material, and eventually the replication of the viral genome in the infected cell, leading to the formation of new virions. Amantadine targets the M2 channel of the virus, blocking the flow of protons and stopping the infection process and influenza A virus reproduction. In amantadine-resistant versions of the virus, it is this proton channel which has mutated.

Now, Luque and his team are investigating new ways of targeting the virus, based on his previous research into drugs aimed at inhibiting neurodegenerative diseases. One of the recent developments in this field is the use of "multi-target" drugs, which are able to simultaneously interact with two or more targets, and are therefore more therapeutically effective. Luque now wants to apply this approach to infectious diseases like influenza A.

Several groups around the world have already begun developing drugs that target proteins other than the M2 proton channel, including neuroaminidase, which enables the virus to be released from the host cell, and haemagglutinin, which is responsible for the fusion of the viral envelope to the endosome membrane of the host cell. Luque and his team are working alongside other groups,

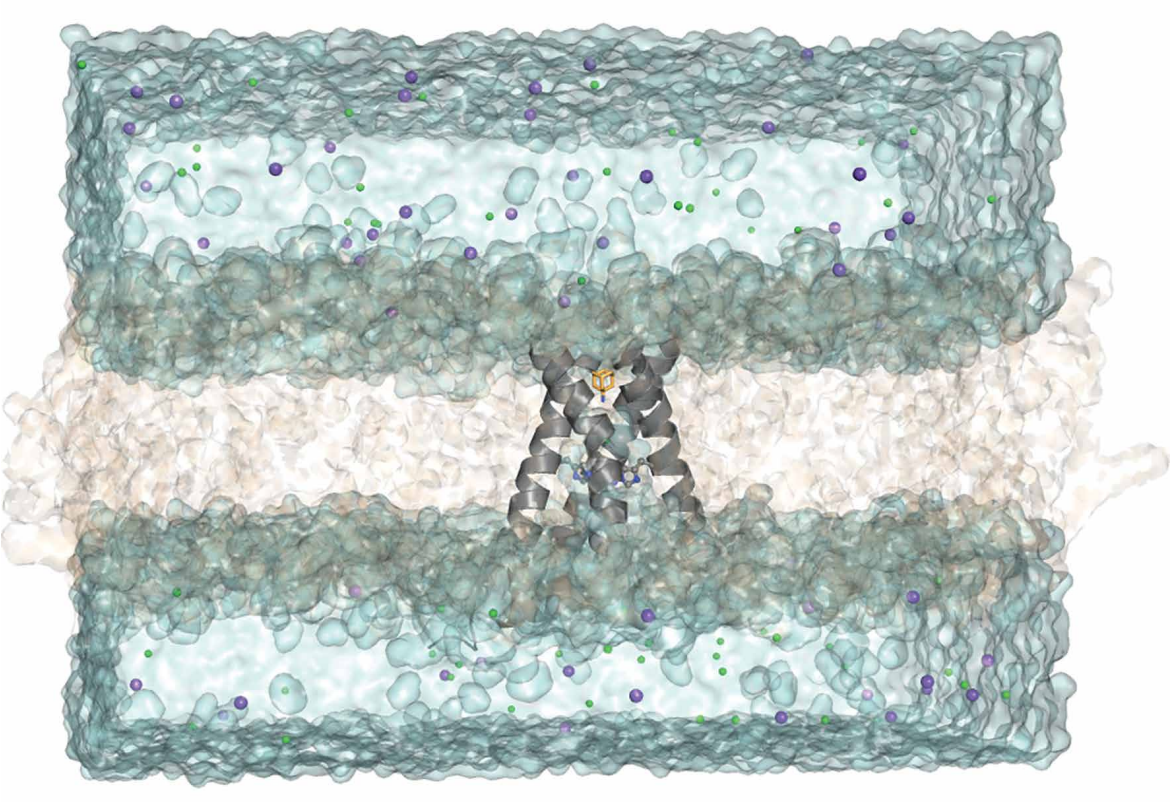

Representation of M2 proton channel, displayed as grey helices embedded into the membrane (yellow surface). The amantadine inhibitor is depicted in orange in the "down" orientation. The histidine tetrad is shown in ball-and-sticks. Chloride and potassium ions and the water surrounding the membrane are displayed as spheres and blue surface, respectively.

From left to right: Carolina Estarellas, Tiziana Ginex and F. Javier Luque.

to develop compounds that simultaneously target the M2 proton channel and haemagglutinin.

Luque's group began investigating this area of research around eight years ago, collaborating with the medicinal chemistry group headed by Professor Santiago Vázquez at the University of Barcelona, who were synthesising compounds to target the M2 proton channel. Luque and his colleagues use molecular modelling and computational studies to try and understand the compounds' inhibitory activity.

"Multi-target drugs which simultaneously interact with two or more targets are more therapeutically effective"

The type of simulations the group use involve atomistic molecular dynamics simulations, in conjunction with enhanced sampling techniques, such as constant pH simulations. Typically, in a molecular dynamics simulation, the protonation states of ionisable groups of a protein are set in the beginning of the simulation. This is not always an obvious task, in particular for histidine, as its pKa is close to the physiological pH. In contrast, in a constant pH molecular dynamics simulation, the protonation state of an ionisable group of a protein is allowed to change during the simulation according to the local electrostatic environment and the pH of the solution. The pKa values of the ionisable groups can then be obtained from the distributions of the protonation states.

"Our techniques allow us to facilitate the sampling of events that normally would take a long time to occur in a molecular dynamics simulation," explains Carolina Estarellas, a junior researcher in the group. "As well as this, they also allow us to quantify the energetics associated with different chemical processes like conformational changes or the binding of a drug."

Recently, the team were granted a PRACE allocation to continue their work. One of the main outcomes of the project has been the publication of a paper which describes a hypothesis for a potential binding mechanism of amantadine to both the normal version of the influenza A virus and the amantadine-resistant version with the altered M2 proton channel. The process with the normal type involves amantadine changing its orientation inside the channel into a binding mode that prevents the M2 channel from transporting protons to the inside of the virus.

However, with the amantadine-resistant version of influenza A, amantadine binds differently, decreasing the drug's affinity and causing a loss of its inhibiting capacity.

More recently, the team also started looking at haemagglutinin. "We believe that this target is more interesting from a medicinal chemistry point of view as it should offer a larger number of suitable binding targets," explains Tiziana Ginex, a postdoc in the research team. Working with their colleague Santiago Vázquez, the team again used molecular dynamics simulations in order to identify potential binding sites on the protein.

In the long term, the group are now working alongside others to develop compounds that target both the M2 proton channel and haemagglutinin simultaneously. Such a compound would be able to prevent the proliferation of the virus even if the proton channel were to allow protons through, creating a two-pronged mechanism for fighting the virus. "On one side we want to prevent the transport of protons, but on the other we want to prevent the release of the fusion peptide even if those protons do get transported," says Luque. "In this way, such a drug would be far more effective."

For more information
www.ub.edu/cbdd

Resources awarded by PRACE
This project was awarded 22 million core hours on MareNostrum hosted by BSC, Spain

Publications
Mechanism of the Pseudoirreversible Binding of Amantadine to the M2 Proton Channel
Salomé Llabrés, Jordi Juárez-Jiménez, Matteo Masetti, Rosana Leiva, Santiago Vázquez, Sabrina Gazzarrini, Anna Moroni, Andrea Cavalli, and F. Javier Luque

Journal of the American Chemical Society 2016, 138, 15345-15358

Simulating turbulence in a virtual wind tunnel

Many simulations of turbulent flows rely on the use of simplified turbulence models, which approximate the behaviour of fluids at smaller scales. However, the availability of more powerful machines is now making it possible to simulate these flows fully. **Professor Philipp Schlatter** has been doing this to investigate the flow of air over a section of an aeroplane wing.

Automotive, aeronautic, and maritime transport of people and goods play important roles in the globalised world, but also use up about five billion barrels of oil per year. Roughly half of the energy being spent worldwide in such transport activities is dissipated by undesired turbulent motion in the interface between moving objects and surrounding fluid. Knowledge of the behaviour of turbulence close to these surfaces can help improve design and reduce drag.

For many years, the shapes of aeroplane wings and cars were designed by testing them in wind tunnels and measuring the flow around them. With the rise of more powerful computers, much of this can now done via simulations rather than using physical experiments. This concept of a "virtual wind tunnel" actually provides a much larger wealth of data relevant for design purposes.

"Our simulation was easier, cheaper and more accurate than an equivalent experiment in a wind tunnel"

Simulating the entirety of the flow around a structure such as an aeroplane is a computationally huge task. Most of the time, approximate models are used by studying the smallest structures of turbulent flow in just one area and then extrapolating this to the rest of the flow. However, computers are now reaching the point at which it has become possible to simulate everything without resorting to these turbulence models.

Professor Philipp Schlatter of the Linné FLOW Centre at the Royal Institute of Technology (KTH) has spent the last 15 years carrying out simulations of flows in which all of the turbulence is resolved. Having started off with simple geometries such as pipe flows and channel flows, these simulations have been increasing in complexity to include features such as surface curvature and other important aspects of fluid dynamics.

In a recent PRACE-supported project, Schlatter and his colleagues did a complete simulation of turbulent flow around a section of an aeroplane wing. This was, at the date of commencement, the first time such a simulation at high speeds had been attempted. "Just the act of carrying out these simulations is a success in itself, as it proves to a lot of people, including engineers and experimentalists, that we are capable of simulating these kinds of complex flows," says Schlatter. "As it turns out, our simulation of a wing was both easier, cheaper and more accurate than an equivalent experiment in a wind tunnel."

Many studies have been made of aeroplane wings and boundary layers in the past, but these have all been based on turbulence models. These models used to take their coefficients and other "ingredients" from experiments, but this approach has always

Philipp Schlatter © Tobias Ohls

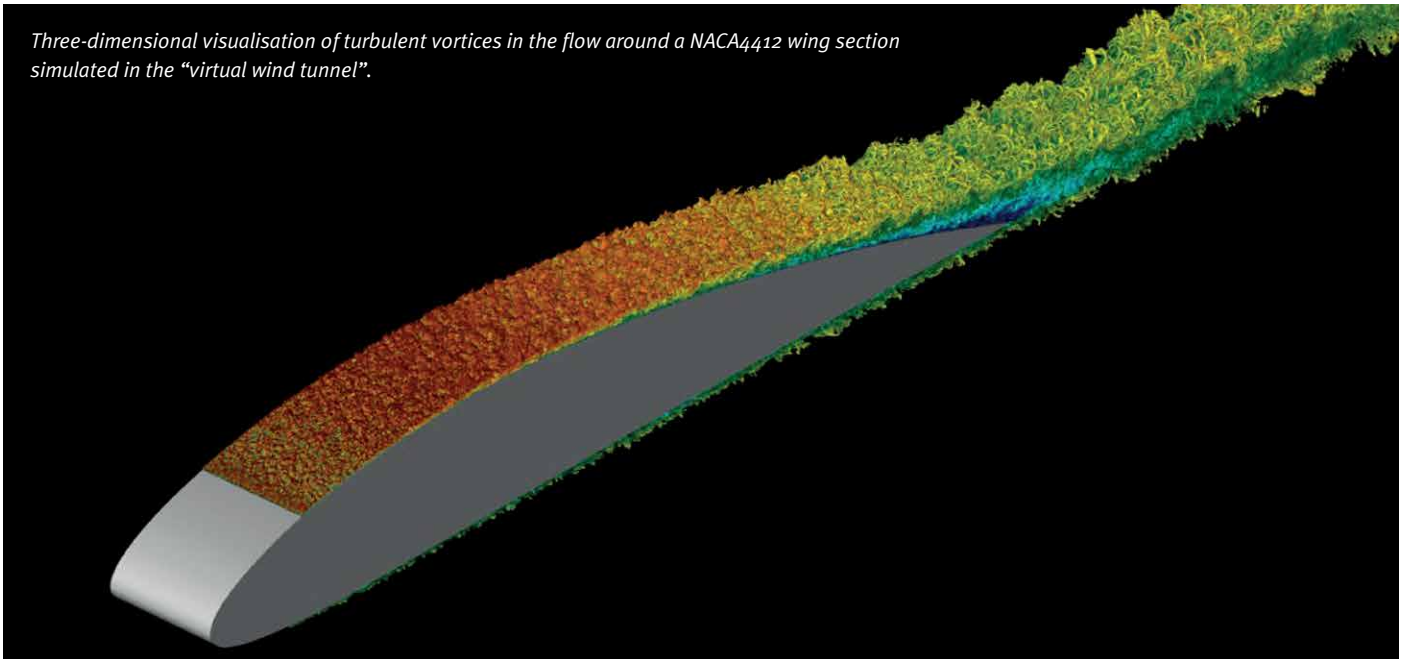

Three-dimensional visualisation of turbulent vortices in the flow around a NACA4412 wing section simulated in the "virtual wind tunnel".

been limited by the accuracy of the data or unavailability of it for certain regions of flow. Schlatter's simulations have provided a vast new dataset that can be used to tune turbulence models by providing more detail and greater accuracy.

This new database of turbulent flow around curved surfaces is extremely valuable to those studying the nature of turbulence. "My work in the past began with studying flows in flat plates, channels, pipes, and other simple geometries," says Schlatter. "These are interesting, but in order to really understand turbulent flow and to see the dynamics and energy exchange processes that occur in them, you need to look at more complex geometries than these."

The database gained from this project is extremely valuable to Schlatter and his colleagues, who are using it to assess the influence of curvature, pressure gradient, and other factors on the behaviour of turbulence and how it can change. It spans one order of magnitude of Reynolds numbers — an important dimensionless quantity in fluid mechanics that signifies how turbulent a flow is — and they have published a number of papers based on the analysis of this data.

The simulations were carried out on the MareNostrum computer at the Barcelona Supercomputing Centre. "I have been working with supercomputers for a number of years now, and so I am quite used to working with them," says Schlatter. "However, there is always a human aspect to these things, and I have to say that the staff in Barcelona were extremely helpful in supporting us throughout the process, from setting our codes up to retrieving the data."

One of the most important outcomes of the project was that the team were able to simulate these types of flow at a level of accuracy never achieved before. They are still in the process of analysing their data and are trying to deduce some physical laws and correlations from it in order to explain the turbulent flows more precisely.

Aside from this, Schlatter and his colleagues came to the conclusion during the project that they had reached a limitation with the strategy they were using to generate their computational meshes. "We rely on research codes to carry out our simulations," explains Schlatter. "In this particular case, we have reached the limit of what we can do using the current code we have for preparing our meshes. As a consequence, we have redirected some of our efforts into improving the way that we do this. We now have a number of new results based on our new methods, which allow us to automate the process of mesh generation and ensure that the accuracy of our simulations is consistent throughout our computational domain."

For more information

https://www.mech.kth.se/mech/info_staff.xhtml?ID=216

Resources awarded by PRACE

This project was awarded 31 million core hours on MareNostrum hosted by BSC, Spain

Publications

S. Rezaeiravesh et al., "Assessment of uncertainties in hot-wire anemometry and oil-film interferometry measurements for wall-bounded turbulent flows," European journal of mechanics. B, Fluids, vol. 72, pp. 57-73, 2018.

High-performance computing using Feel++

Feel++ is an open-source software project designed to create improved methods for solving partial differential equations on high-performance computers. **Professor Christophe Prud'homme** of the University of Strasbourg has been leading the project and has been using a PRACE allocation to test to out some of their new methods.

A partial differential equation is a relation between a function of several variables and its (partial) derivatives. It is used in physics, mathematics and engineering to model most natural phenomena such as propagation of sound, electrodynamics, fluid flow, elasticity or more complex problems.

The Feel++ project has been developing an open-source library that provides a large range of numerical methods for solving partial differential equations. Led by Professor Christophe Prud'homme of the University of Strasbourg, it has been a large collaborative effort, working alongside other mathematicians as well as engineers, doctors, industry experts and others to develop scalable solutions for HPC and apply them to applications that are useful for researchers.

For instance, one of the main collaborations has been with the Laboratoire des Champs Magnétiques Intenses (LNCMI), part of the European Magnetic Field Lab, which is dedicated to generating extremely powerful magnetic fields that can be used for scientific research and making them available to the scientific community. "Many winners of the Nobel Prize have worked with LNCMI on the way to their accolades," says Prud'homme. "We have helped LNCMI by providing them with the software to help them design their magnets."

The team has been using their PRACE allocation to test out the methods they are developing on Tier-0 computers. "The design of our library allows us to quickly put our methods into applications," Prud'homme explains. "There are a few which are still in development and will require further adaptations and improvements, but some researchers are already using them."

The library has two types of interface. One of them is designed for mathematicians, using C++ to describe things in a way that is very

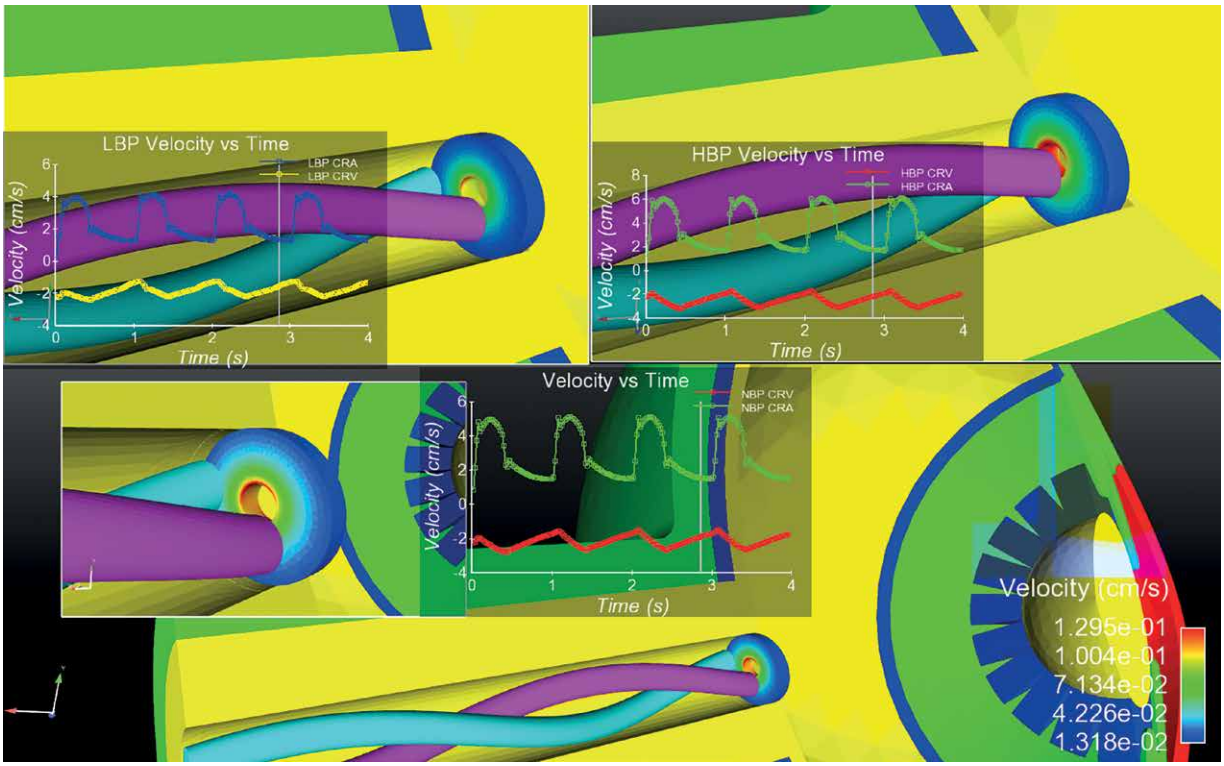

The Eye2Brain system requires mathematical models parameterised with patient-specific data that can both quantitatively describe the fluid-dynamical and metabolic connections between the eye and the brain, and identify the main factors that influence these connections.

close to the mathematical language of the methodologies. This is useful for those developing the methodologies and for deploying applications on to supercomputers. The other type of interface uses multi-physics toolboxes, providing a more user-friendly, application oriented way for researchers from a wide array of disciplines to access and work with the library.

Over the course of the PRACE allocation, the Feel++ project has developed two different kinds of method. The first of these is based on domain decomposition and is called a mortar method. "For this one we can reach extremely good scalability and excellent effectiveness, meaning that there is both minimal communication between processing elements as well as good performance when larger numbers of cores are used," says Prud'homme. The design of this numerical method and the results of scaling it has been published in M2AN.

"We want to provide useful HPC applications to people like doctors, physicists and engineers"

The other type of method developed by Feel++ is a so-called "fictitious domain" method. One real-life application that this has been used to explore has been tracking the rheology and movement of red blood cells in blood flows. There are already a number of ways of doing this, but this particular method disconnects the discrete geometrical model of the fluid from the geometrical model of the cells. From a mathematical point of view, the method developed by Feel++ has been shown to work well and is able to simulate extremely large numbers of cells at once to provide a realistic picture.

Prud'homme and his colleagues used part of their allocation to compare their fictitious domain method with other methods such as level-set methods, and showed that there are advantages and disadvantages for both. "It is not quite clear what the best compromise here is yet, but what we did show is that scaling the fictitious domain method is quite difficult compared to our mortar method," he says. "However, this was expected because that method was designed from the ground up to scale well, whereas this method was more a case of trying to handle a specific physical problem in a flexible way."

The mortar method has now been integrated into the Feel++ framework and is being used to help the team tackle large-scale problems. This has required a number of design changes and adaptations, but the rewards have been significant in terms of speedup, effectivity and scaling.

The PRACE allocation has also allowed the team to strengthen their core framework and start moving into new territories. They have just finished a European project called MSO4SC, in which they have been developing an HPC cloud framework that connects their software directly to supercomputers via the cloud. "As far as we know this is the first type of framework like this," says Prud'homme. "This allows users to access supercomputers in a seamless way, using our applications that have already been tuned for the specific machines being used. All we do is provide some guidelines about how to select the proper resources for the task at hand."

One pilot application being tested out is Eye2Brain, which uses a code to simulate flows and biomechanical behaviour inside the eye and part of the brain. The application provides insight into some of the mechanisms behind the development of certain conditions such as glaucoma, and will work with more neurodegenerative diseases in the future. This provides an in silico measurement tool alongside the physical equipment in clinics, and the team worked alongside medical professionals to develop the models and parts of the software. They hope to provide doctors and medical technicians access to this software soon so that it can be used in a clinical research setting.

This PRACE allocation has provided Feel++ with the platform to work on their codes and bring them up to a level of maturity so that they can scale seamlessly to larger machines. "What we want is to be able to provide really useful applications to a wide variety of people like doctors, physicists and engineers," says Prud'homme. "That is what we are now achieving thanks to PRACE."

For more information
http://docs.feelpp.org

Resources awarded by PRACE
This project was awarded 6 million core hours on Curie hosted by GENCI at CEA, France

Publications
Thibaut Metivet, Vincent Chabannes, Mourad Ismail, Christophe Prud 'Homme. High-order finite-element framework for the efficient simulation of multifluid flows. 2018.

Cécile Daversin, Christophe Prudhomme, Christophe Trophime. Full 3D MultiPhysics Model of High Field PolyHelices Magnets. IEEE Transactions on Applied Superconductivity, Institute of Electrical and Electronics Engineers, 2016, 26 (4), pp.1-4.

Silvia Bertoluzza, Micol Pennacchio, Christophe Prud'Homme, Abdoulaye Samake. Substructuring Preconditioners for h-p Mortar FEM
ESAIM: Mathematical Modelling and Numerical Analysis, EDP Sciences, 2016, 50 (4), pp.1057-1082.

PRACE

PRACE, the Partnership for Advanced Computing in Europe, provides access to Europe's world-class High Performance Computing Research Infrastructure for scientists and researchers from academia and industry.

HPC ACCESS

CALL ANNOUNCEMENTS

Stay tuned for news on the 19th Call planned to open on 5 March 2019: **www.prace-ri.eu/call-announcements**

TRAINING

PRACE TRAINING PROGRAMMES

These programmes offer free, top-level training to an ever-growing number of academic and industrial researchers, and actively train the HPC researchers of tomorrow. **www.training.prace-ri.eu**

INDUSTRY

PRACE INDUSTRY ACCESS

97% of industrial companies that use HPC consider it essential to their ability to innovate, compete and grow.
www.prace-ri.eu/industry

SMEs

PRACE SHAPE SME HPC Adoption Programme in Europe

Apply to 8th Call until 1 December 2018.
www.prace-ri.eu/hpc-access/shape-programme

STAY IN TOUCH

SUBSCRIBE TO PRACE

www.prace-ri.eu/subscribe-to-prace-gdpr-compliance

www.prace-ri.eu

The Partnership for Advanced Computing in Europe (PRACE) is an international non-profit association with its seat in Brussels. The PRACE Research Infrastructure provides a persistent world-class high performance computing service for scientists and researchers from academia and industry in Europe. The computer systems and their operations accessible through PRACE are provided by 5 PRACE members (BSC representing Spain, CINECA representing Italy, ETH Zurich/CSCS representing Switzerland, GCS representing Germany and GENCI representing France). The Implementation Phase of PRACE receives funding from the EU's Horizon 2020 Research and Innovation Programme (2014-2020) under grant agreement 730913. For more information, see www.prace-ri.eu.

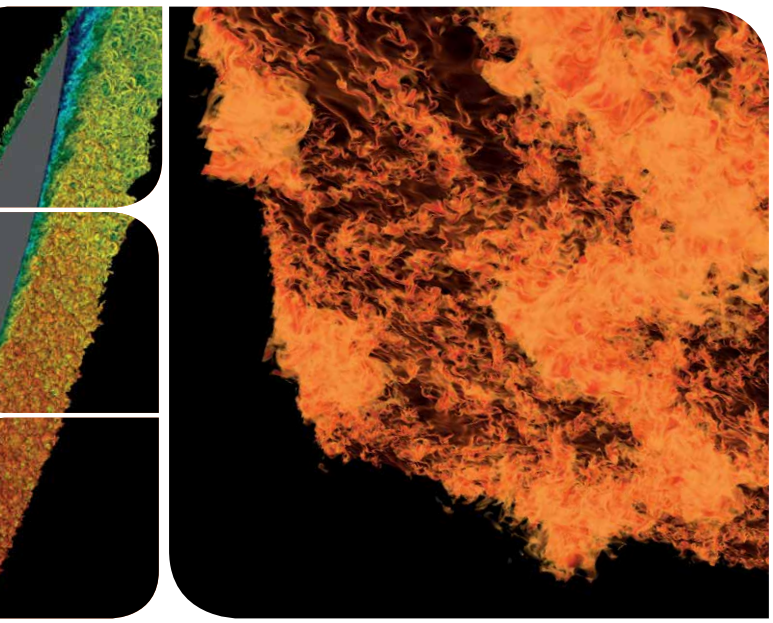
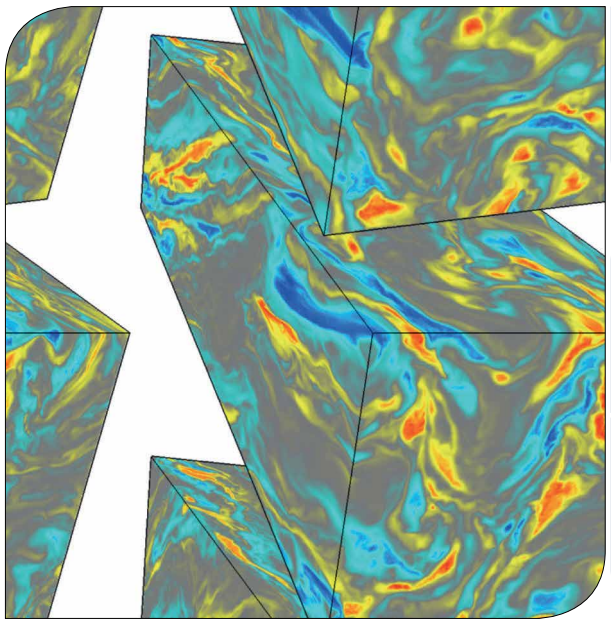
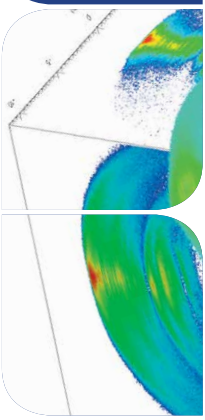
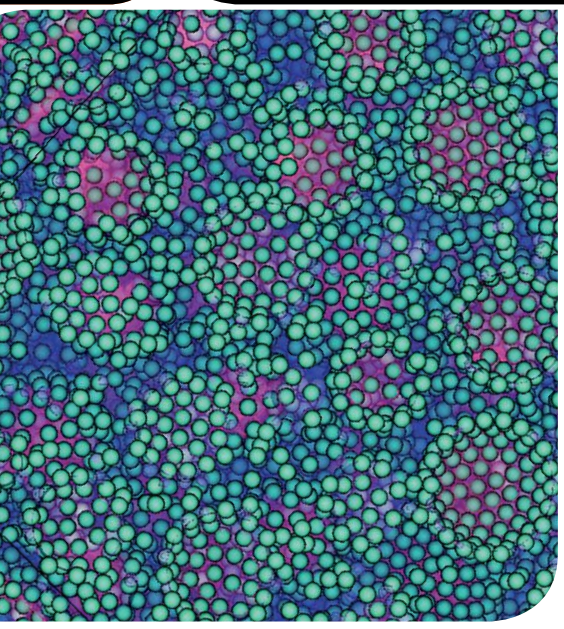
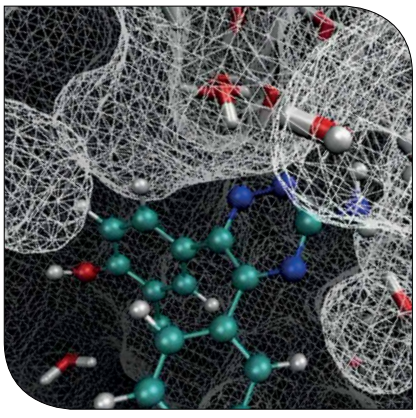

The Implementation Phase of PRACE receives funding from the EU's Horizon 2020 Research and Innovation Programme (2014-2020) under grant agreement 730913