PRACE

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

PRACE DIGEST 2019

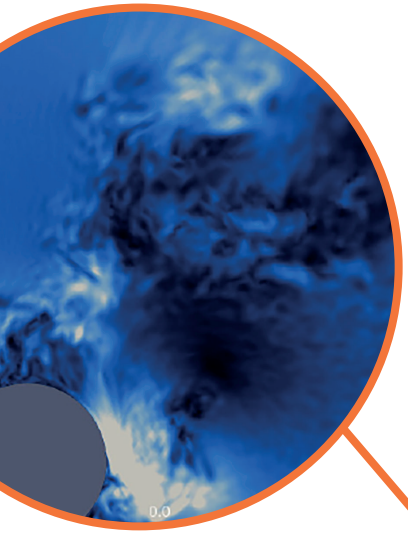

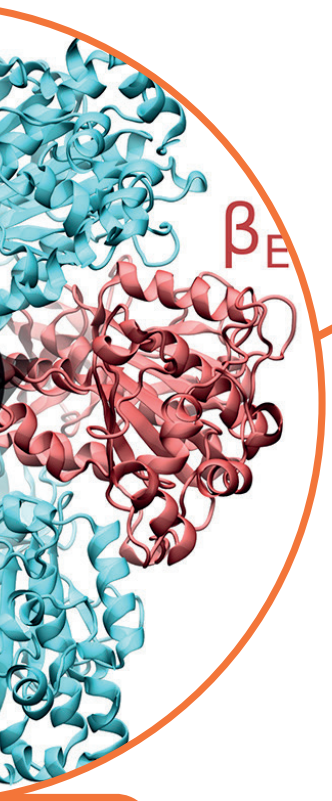

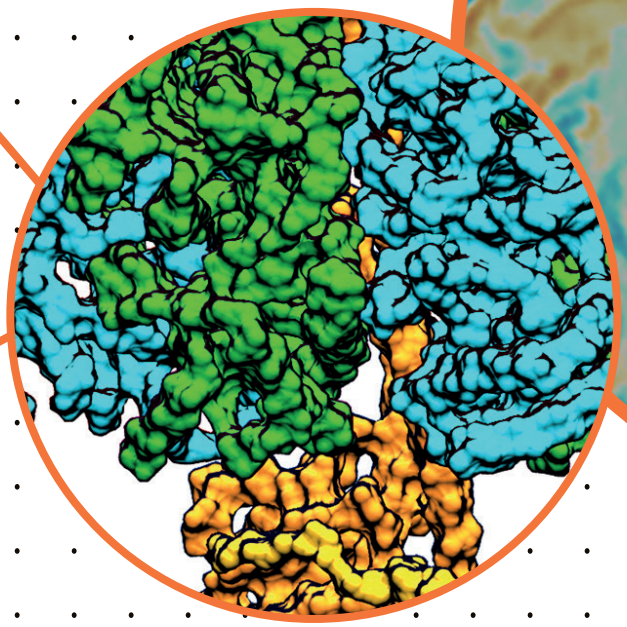

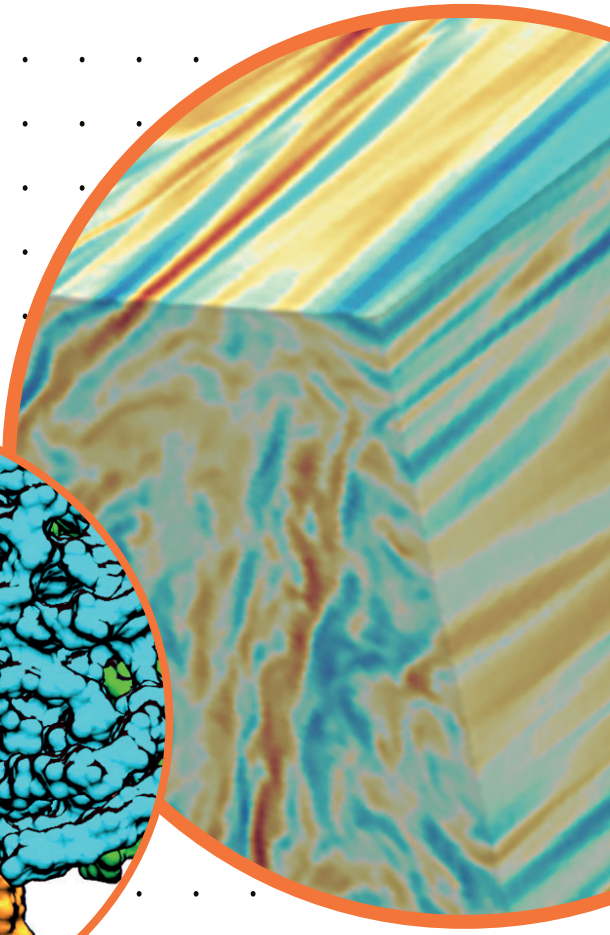

www.prace-ri.eu

PRACE Digest 2019

PRACE Digest 2019
www.prace-ri.eu

Editor-in-Chief
Marjolein Oorsprong – PRACE Communications Officer

Editors
Veronica Teodor, FZ Jülich
Silke Lang, PRACE

Article Editor
Harry O'Neill, Insight Media

Communications Assistant
Jeet Landa, Insight Media

Graphic Design
Crystal Mather, William Davis, Insight Media

Published by Insight Media Group Ltd, 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, 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 823767.

Published: November 2019
ISBN: 9789082942248

Foreword

Welcome to the ninth edition of the PRACE Digest. I hope to meet you at the launch of this publication at SC19 in Denver, Colorado, where we will be reconnecting with people from across the world of High-Performance Computing.

Each year, PRACE funds an incredible diversity of projects, each led by pioneering researchers looking to push the boundaries of their respective fields. The PRACE Digest is our way of sharing this with you, giving you a taste of the range of discovery and invention that we enable by providing world-class researchers with world-class computing power.

The origin of life on Earth will forever be a topic that piques people's curiosity. Marie-Pierre Gaigeot's simulations have shed some light on how the first peptides – one of the fundamental building blocks of life – may have arisen from the prebiotic soup of the oceans. Jacek Czub's work addresses life further along its development, using molecular dynamics simulations to gain a better understanding of ATP synthase – an enzyme shared by lifeforms ranging from primitive bacteria to humans.

Even more fundamental questions about the universe can also be addressed using HPC. The behaviour of plasma in space was the focus of Francesco Califano's recent project, and the AI tools that he is developing to analyse the huge datasets produced through his simulations are now being used in a similar fashion to make sense of the data sent back from some of our most advanced satellite missions.

Engineers have for many decades now relied upon the calculations provided by supercomputers to remodel and refine the designs of some of our most complex pieces of machinery. Marianna Braza's work on a new generation of bio-inspired wings that can morph and vibrate in order to improve the efficiency of aircraft as they move through the air would not have been possible without HPC. Nor would the research of Neil Sandham, who similarly has been investigating how to reduce the effects of shock-related buffeting on aircraft wings.

Of course, all of this is just a small taster of what can be found within this publication. We are very proud to have worked with everyone who appears in the coming pages, and have enjoyed the process of finding out more about their outstanding work so that we could tell you all about it.

All there is left to say is that I hope you enjoy reading this as much as I did!

Serge Bogaerts
Managing Director of PRACE aisbl

Contents

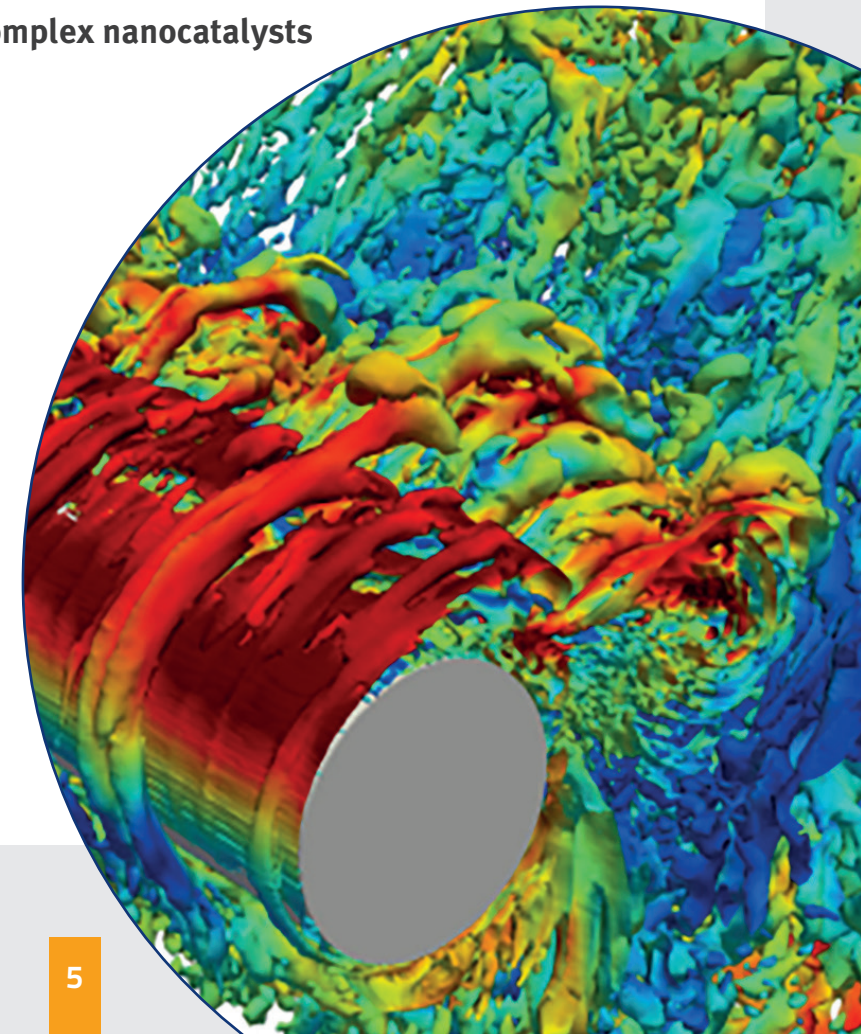

Unsteady aerodynamics of oscillating airfoils

New natural laminar flow wing designs have the potential to reduce the energy costs of aircraft, but show instabilities that cannot be modelled by classical methods. **Dr Prabal Negi** of the KTH Royal Institute of Technology has been carrying out large-eddy simulations of these wings to try and gain a better understanding of their aerodynamic behaviour.

Prabal Negi

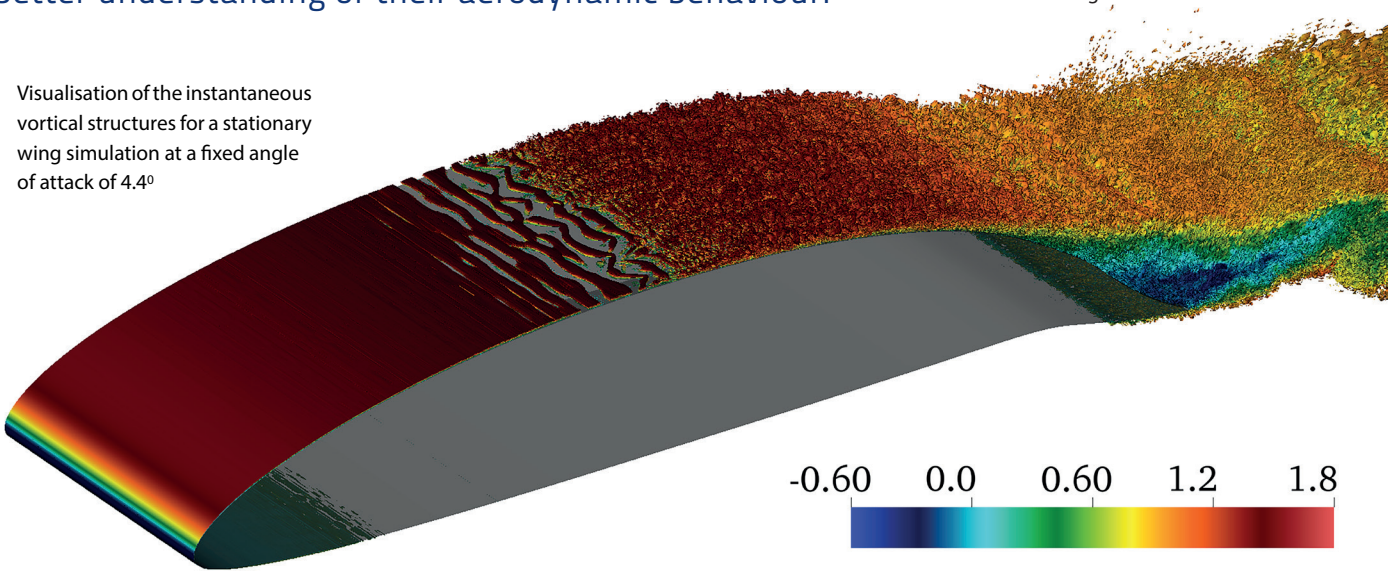

Visualisation of the instantaneous vortical structures for a stationary wing simulation at a fixed angle of attack of 4.4^0

Laminar flow control and natural laminar flow (NLF) wing design have been proposed as one of the main options for reducing the drag and consequently the fuel consumption of commercial aircrafts, which could lower the carbon footprint of the industry.

Dynamic stability is a concern for aircraft designers and the occurrence of aeroelastic instability in mid-flight can have catastrophic consequences. Simplified mathematical models which were laid down in the early 1930s are normally used to examine these kinds of instability, but recent developments have found these models to be somewhat unreliable for NLF airfoils.

Working out how NLF airfoils behave is important for working out the stability characteristics of aircraft that use them. Recent experimental studies on NLF airfoils have shown that their behavioural dynamics differ from those of turbulent airfoils. Most notably, NLF airfoils exhibit non-linear aerodynamic responses to small-amplitude pitch oscillations.

These experimental results have shown that the classical models used for the prediction of dynamic stability are no longer valid for NLF airfoils. This failure of the classical models has been traced to the point of transition on the airfoil where flow changes state from laminar to turbulent. The flow state affects the point at which the flow separates from the airfoil, which in turn affects the point of transition itself. Thus, these two boundary layer characteristics – the point of transition and point of separation – are coupled and lead to new dynamic behaviour in unsteady cases, which is not shown by the classical theories.

NASA identified the prediction of transition and separation in general unsteady flows as one of the fundamental challenges of computational fluid dynamics

Prediction of this transition and flow separation is crucial for adequate design and control of aircraft systems. This requires an in-depth understanding of boundary layer characteristics in unsteady applications, where both transition and separation may dramatically alter flow characteristics. Such scenarios are

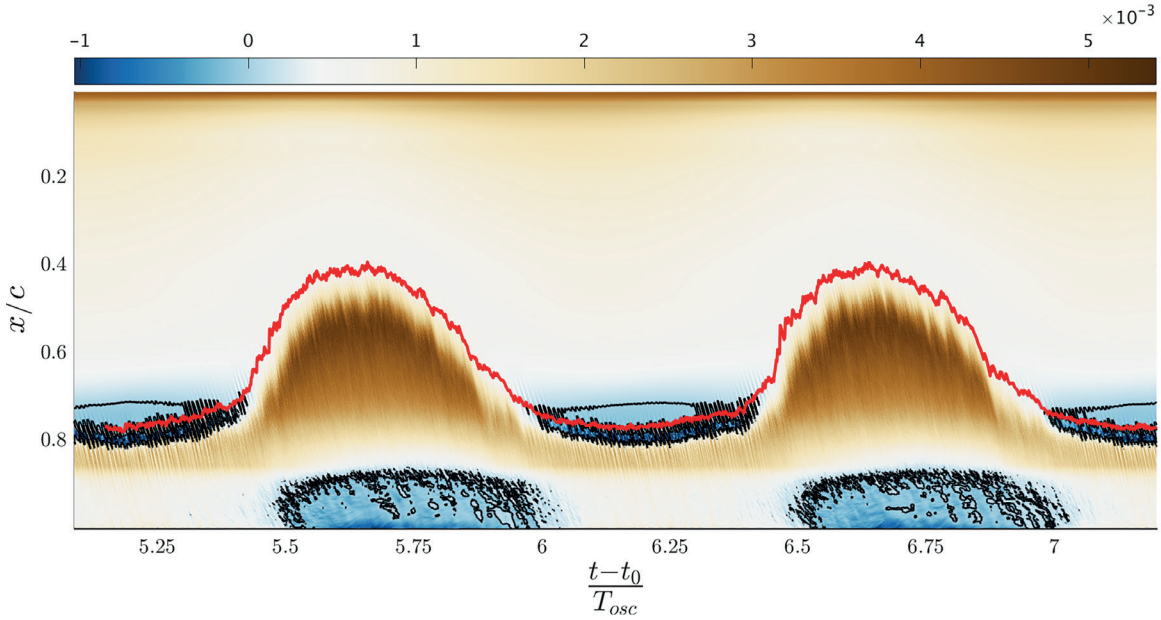

Spatio-temporal variation of the locally tangent wall-shear stress. The y-axis represents the chord-wise location while the x-axis represents the normalised simulation time. Black contours denote the points of zero tangential shear stress. Empirically determined transition location (red curve) is superposed on the space-time plot.

especially relevant for NLF airfoils, which can exhibit a sensitive dependence on the angle of attack, leading to large changes in aerodynamic forces for even small changes in the angle of attack. Such small changes are inevitable in normal flight conditions and can arise from small structural deformations of the wing structure or incoming gusts. In such cases, current simple models of unsteady flow prediction are not reliable and a more detailed understanding is required for better control strategies and aircraft wing designs.

The non-linearities observed in NLF airfoils were the focus of a recent PRACE-supported project called "Large-Eddy-Simulations of the Unsteady Aerodynamics of Oscillating Airfoils at Moderately High Reynolds Numbers", led by Dan S Henningson of the KTH Royal Institute of Technology, Stockholm. In this work, the dynamic coupling between the boundary layer characteristics of unsteady airfoils has been simulated, with the aim to of creating an improved mathematical model to predict the aeroelastic stability characteristics of airfoils. A recent report by NASA identified the prediction of transition and separation in general unsteady flows as one of the fundamental challenges of computational fluid dynamics for the coming decade. Such studies necessarily need to be undertaken in the relevant Reynolds number – the ratio of the largest to the smallest spatial scales in the flow – range of the intended application since transition and separation characteristics on wings can change substantially with Reynolds numbers.

"In our work, we numerically simulate a Reynolds number of 750 000 which is close to the Reynolds number range of high-altitude long-endurance aircraft," says Prabal Negi, a researcher from Henningson's lab at KTH. "A high-order spectral element method code Nek5000 was used for the numerical simulation and 1.4 x 10^9 grid points were necessary to discretise the problem."

The project investigated the unsteady flow characteristics for a specific case of small-amplitude pitch oscillations in an NLF airfoil. Such small-amplitude pitch oscillations mimic the changes in operating conditions due to structural deformations. The NLF airfoil exhibits the sensitive dependence on angle of attack, with large changes in transition and separation characteristics within a small angle of attack change.

The simulation results have shown that the non-linear aerodynamic response can be modelled using a simple three-parameter empirical model. A good agreement has been found between experimental data and the model for a wide range of frequencies. The results obtained so far suggest that it may be possible to model and predict the non-linear aerodynamic response using computational tools.

In terms of further work, several subsequent analyses are in progress with the data generated from the simulations. The simplified model for the aerodynamic forces can hopefully be turned into a predictive model. In addition, linear growth of small disturbances in the flow in such unsteady wing cases has not been investigated. These changes in the growth of small disturbances dictate the unsteady transition point on the airfoil.

For more information

arc.aiaa.org/doi/10.2514/6.2018-3824

kth.diva-portal.org/smash/record.jsf?pid=diva2%3A1365798 (PhD Thesis)

Resources awarded by PRACE

This project was awarded 40 million core hours on Hazel Hen hosted by GCS at HLRS, Germany

Publications

Negi, P., Hanifi, A., & Henningson, D. (2018). LES of the Unsteady Response of a Natural Laminar Flow Airfoil. In 2018 Applied Aerodynamics Conference (p. 3824).

P. S. Negi (2019), Stability and transition in pitching wings. Doctoral thesis, KTH Royal Institute of Technology

Genome wide association studies for type 2 diabetes


David Torrents

The risk of an individual getting a complex disease such as type 2 diabetes and cardiovascular diseases is wrapped up in many genetic factors. **Professor David Torrents** of the Barcelona Supercomputing Centre has been carrying out huge genome-wide association studies to try and find out what they are.

Complex diseases are a major healthcare problem worldwide. These diseases, which include type 2 diabetes and cardiovascular disease, have multiple factors involved in their risk. The work of David Torrents of the Barcelona Supercomputing Centre involves looking for genetic variants that predispose people to these diseases.

"With some diseases, it is a case of one particular mutation – if you have it, then you have the disease," explains Torrents. "With complex diseases, the picture is more complicated. There can be a number of genetic determinants involved, and even those only give you a predisposition to the disease. For example, you can be genetically predisposed to cardiovascular disease, but if you live and eat healthily, you won't necessarily get it."

Because of the multifactorial nature of these diseases, analysing their genetic origin is a complex task. For each of the several million genetic differences that are seen in the population, the researchers test whether or not that particular variant is associated with a disease, using extremely large cohorts of data. Torrents looks only at the genetic factors behind these diseases, as getting good quality data about the lifestyle of patients is difficult – people often lie about their health or misunderstand questions when surveyed.

"The goal of our research is to find biomarkers – genes that can be used to identify the presence of a disease. We live in a time where big data is starting to drive a lot of research, and this is especially true in genomics. New generations of sequencing and genotyping are arising, and our community is analysing larger datasets of individuals than ever before. This is giving us amazing results, but as a consequence we burn through CPUs like crazy!"

Having evaluated a number of previous studies done by other groups in the field, Torrents realised that many of the genome-

wide association studies done in the past had failed to perform comprehensive analyses of the data due to lack of computing resources. "Groups have to prioritise what they look at, often excluding analysis of the x chromosome, due to the methodological complications that it brings.

"Their studies also tend to only consider one inheritance model out of the five that exist. What we have done in our analyses is to include both the x chromosome and all of the inheritance models. When analysing the genomes of around one million individuals, this requires a huge amount of computational power."

As researchers at the Barcelona Supercomputing Centre (BSC), Torrents and his colleagues are lucky enough to be situated next to one of the most powerful supercomputers in Europe – MareNostrum. One recent PRACE allocation of 16.5 million hours of CPU time on this machine has allowed them to carry out a genome-wide association studies of a huge cohort of individuals.

The study focused on type 2 diabetes – a well-studied disease that crucially has a lot of public data available about it. The paper from the study has not been published yet and so the full results cannot

"We have shown that by not doing a full analysis of a dataset, you can miss up to 20-30% of the variants involved in a disease."

be disclosed, but Torrents is enthusiastic about their findings. "Essentially we have shown that by not doing a full analysis of a dataset, you can miss up to 20-30% of the variants involved in a disease. By including the X chromosome and all of the inheritance models, we have found a number of new genetic factors involved in type 2 diabetes."

The researchers use a programming model created at BSC called COMPS, which parallelises the code and optimises the use of the MareNostrum computer. It ensures that all of the nodes are used at the same time, and balances the loads across the nodes.

The team are now working on a follow-up study, again using a PRACE allocation. Postdoctoral researcher Ignasi Morán explains the aims of the new study: "The ongoing project is working within the same framework of analysis, looking for the genetic components of complex diseases. But in our previous studies, we have always looked at a single genetic variant at a time. We analyse whether the variant is associated with the disease, then when we have an answer we move on to the next one.

"With this new study, we are delving a bit further and examining the effects of combinations of variants. Instead of looking at each variant in isolation, we look at the effects of pairs of variants. One variant might have a risk associated with it, but if there are two of these together then the risk associated with both of them might be higher than the sum of each individual variant. In this way we can get a better picture of the network of interactions between all of these genetic factors."

Analysing combinations of variants like this is even more computationally expensive than studies of single variants, and so access to HPC has again been essential for the research. "As far as we are aware, this is the largest genome-wide association study to date," says Torrents.

Torrents and Morán are now continuing with this analysis of pairs of genetic variants, and are hoping to show that certain pairs are associated with increased risk. As well as this, they are now employing an entirely different methodology – machine learning – to approach their work from a different angle.

"Using machine learning for our work is like turning the problem on its head," explains Morán. "The system looks at the entire set of genomic information as a whole, and then works out which parts of it are likely to be associated with the disease. It is interesting to compare the two methods, and we think they can complement each other in providing us with a better understanding of the network of genetic determinants."

For more information
www.bsc.es

Resources awarded by PRACE
This project was awarded 16.5 million core hours on MareNostrum hosted by BSC, Spain

Publications
Publications (most relevant three):

Re-analysis of public genetic data reveals a rare X-chromosomal variant associated with type 2 diabetes
S Bonàs-Guarch, et al.
Nature Communications 9 (1), 321 19 2018

A Loss-Of-Function Splice Acceptor Variant in IGF2 is Protective for Type 2 Diabetes
JM Mercader, et al.
Diabetes, db170187

Genome-wide associations for birth weight and correlations with adult disease
M Horikoshi et al.
Nature 538 (7624), 248-252

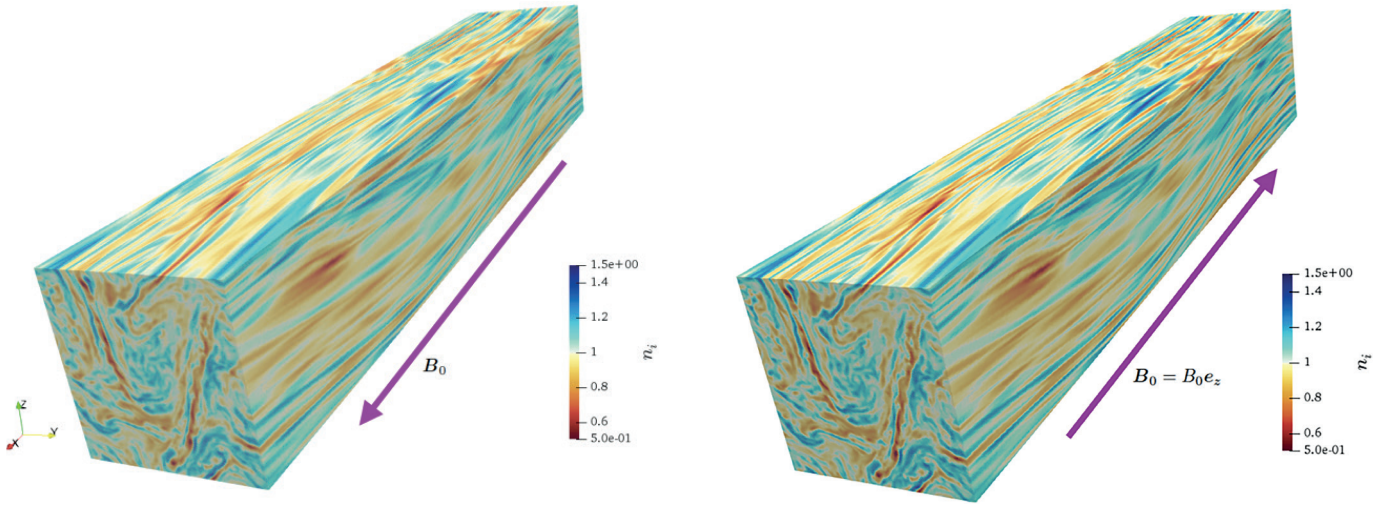

3D-3V Hybrid-Vlasov simulation of strong plasma turbulence simulations with beta ~ 1. The shaded iso-contours represent the total current and its projection on a plane at 45 degrees. We see many magnetic structures generated by the turbulence, most of them showing the development of magnetic reconnection which in turn feeds the turbulence at sub-ion scale.

Plasma turbulence in solar winds

Solar wind interacting with the Earth's magnetosphere represents an excellent laboratory for studying the behaviour of space plasma because of the high quality of in-situ satellite measurements. **Professor Francesco Califano** has been using HPC to simulate this phenomenon, the results of which could aid future satellite missions.


Francesco Califano

Professor Francesco Califano works in the field of plasma physics, studying the solar wind that occurs in between the surface of the sun and the Earth's magnetosphere. "The magnetosphere can roughly be imagined as a magnetic screen that protects us from the particles that are ejected from the sun," he explains. "Without it, we would be in trouble!"

Califano's work involves modelling what happens in this region to try and understand the main physical processes that occur there. The medium that supports the processes in this region is called plasma, which is made up of an immense number of charged particles. One of the main characteristics that defines plasma is that its dynamics is mainly driven by electromagnetic forces, with collisions between individual particles almost negligible. To examine such a system, one has to look at how each particle interacts with all of the others simultaneously. To do this by brute force is impossible even with today's most powerful computers, and so instead the problem is approached in a statistical way using what is known as a distribution function.

"The way we look at plasma is to consider it as a system in which the current and the electromagnetic field self-consistently generated by the global system is much more important than single interactions between individual particles. The main force acting on each particle is due to the global system," explains Califano. To model the system, a Vlasov equation is used to solve the position and the velocity of the particles through time in phase space. This equation is then coupled with the Maxwell equations that solves the electromagnetic field. Once the density and the current produced by the system of charged particles has been solved, the corresponding electromagnetic field can then be calculated. This field in turn plays a role in the dynamics of the global system of charged particles.

The partial differential equations that have to be solved to model the dynamics can then show how turbulence develops in the solar wind. "Turbulence is one of the most important phenomena involved in solar wind, as it is responsible for the transport of heat and particles," says Califano. "The equations we use to examine turbulence are in some ways very similar to the Navier-Stokes equations used to study fluid dynamics in fields such as meteorology. However, plasmas are not fluids – they are categorised as a separate state of matter. From a mathematical point of view, investigating them is mathematically much more complicated. Fluids are composed of neutral particles, with local collisions dominating the way that they behave. Plasmas are made up of many species of particles, including electrons, protons, ions, heavy ions and more. As a consequence, the physics acting at the different scales and frequencies is very different.

"With plasma, you have to look at the behaviour occurring at different scales, from fluid to kinetic ions to electron scales. Local fluctuations due to electron dynamics occur at very small scales,

while the dynamics at larger ion scales are more and more similar to fluid dynamics but including ion kinetic effects. So when solving plasma physics problems, you not only have to deal with partial differential equations, but also have to look at scales and frequencies which differ by many orders of magnitude depending on the system of interest."

Examining the energy exchange between solar wind and the Earth's magnetosphere is important for a number of reasons. Although most of the charged particles emitted from the sun towards the Earth are stopped by the magnetosphere, conditions can occur allowing particles to break through. The fundamental process allowing such abrupt entry of particles is known as magnetic reconnection, one of the most important processes in plasma physics. The auroras seen at the Earth's poles are a visible effect of this occurrence but, although beautiful, the charged particles involved can cause huge damage to electrical systems, to people crossing the polar regions, and can affect mobile phone networks.

High-Performance Computing is essential for solving the equations needed to study the behaviour of solar wind, and this has been made possible through a PRACE allocation of 60 million core hours on MARCONI-KNL at CINECA. A typical simulation solving the distribution function in the 3D-3V phase space uses up around 15-20 million CPU hours, and these simulations produce terabytes of data that have to be stored. "Dealing with all of this data is difficult – it is a big data problem," says Califano. "We are looking at ways of using artificial intelligence tools like machine learning to filter through this data and find the signatures of the physical processes

"The magnetosphere can roughly be imagined as a magnetic screen that protects us from the particles that are ejected from the sun"

that we are interested in. Without tools like this, you are left with the almost impossible task of trying to sift through these huge amounts of data manually."

Space plasmas are not only interesting from a fundamental science perspective – they represent one of the best laboratories for gaining measurements of plasma physics that would be impossible to gather on Earth. "As well as this, our work is useful for those working on satellite missions to space, such as NASA's Magnetospheric Multiscale Mission that is studying the Earth's magnetosphere," says Califano. "The instruments on these satellites can gather huge amounts of data, for example about electron physics. The people working on the missions are then left

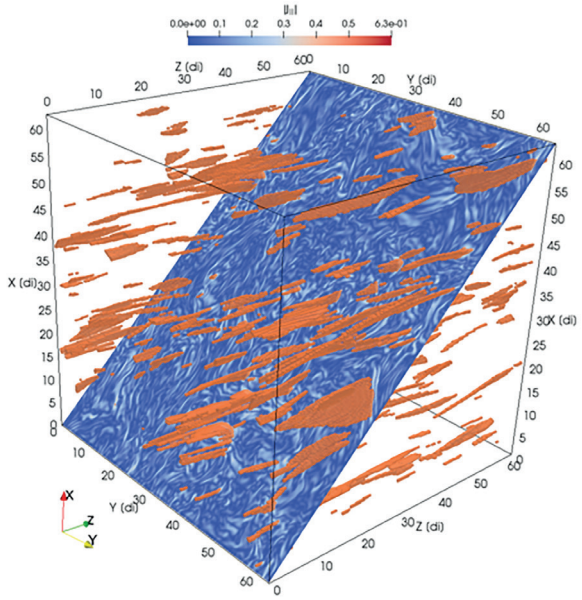

3D PIC simulation of critically balanced driven turbulence in the low-beta regime beta ~ 1. The picture shows the density fluctuations . The character of the fluctuations are clearly anisotropic with respect to the mean magnetic-field direction showing also structures associated with magnetic reconnection.

with a similar problem to us – how to extract the information about events that they want to study from the huge amounts of data they have gathered. The artificial intelligence tools that we are developing in the framework of the "AIDA" EU project to look at our simulation data can be transferred to look at real data from space, as we are trying to do."

The "AIDA" Horizon 2020 project that Califano's work is a part of has now been running for about one year, and the researchers are now gearing up to start using the AI tools they are developing on the numerical simulations.

For more information

osiris.df.unipi.it/~califano/Francesco/Welcome.html
aida2020.web.cern.ch/

Resources awarded by PRACE

This project was awarded 60 million core hours on MARCONI-KNL hosted by CINECA, Italy

Publications

M. Sisti, M. Faganello, F. Califano, B. Lavraud, Satellite data-based 3D simulation of Kelvin-Helmholtz instability and induced magnetic reconnection at the 3 Earth's magnetopause, Geophysical Research Letters, 46 (2019) https://doi.org/10.1029/2019GL083282

J. Dargent, F. Lavorenti, F. Califano, P. Henri, F. Pucci, S.S. Cerri, Interplay between Kelvin-Helmholtz and Lower-Hybrid Drift instabilities , J. Plasma Physics, in press.

S.S. Cerri, M. Kunz, F. Califano, Dual Phase-space Cascades in 3D Hybrid-Vlasov–Maxwell Turbulence, Astrophysical J. Lett. 856:L13 (2018)

Investigating ATP synthase – the splendid molecular machine

ATP synthase is an enzyme found in organisms ranging from primitive bacteria to some of the most complex lifeforms, such as humans. Its energetic efficiency is unrivalled, but not well understood. **Professor Jacek Czub** of Gdansk University of Technology has been using HPC to study this remarkable enzyme at a level of detail never seen before.


Jacek Czub

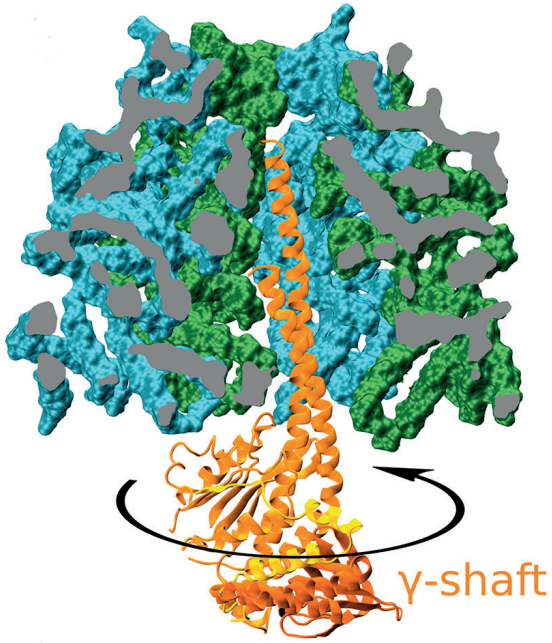

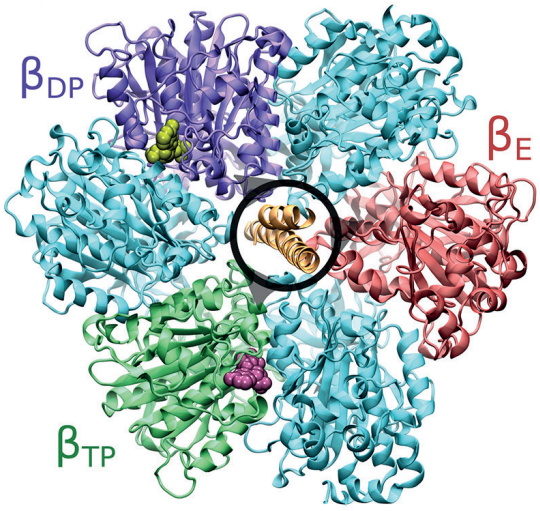

Left - Longitudinal section cut through the catalytic domain shows the structure of the rotary γ subunit inside the α3β3 hexamer. The arrow shows the direction of rotation during ATP synthesis.

Right - Cross section of the catalytic domain of F1 -ATPase showing the three coupled catalytic sites at the β subunits.

The enzyme ATP synthase has been described by one of its foremost experts as "a splendid molecular machine". It is capable of using energy stored in the form of protein gradients across the mitochondrial membrane to synthesise ATP – the ubiquitous universal energy carrier which provides energy for nearly all of the processes in the cell, including protein synthesis.

One of the most remarkable characteristics of ATP synthase is its efficiency. In terms of energy conversion, it beats anything that has been made by artificially by humans, with almost no energy being lost up to the limits of the second law of thermodynamics.

Jacek Czub of Gdansk University of Technology has been using HPC to try and understand the mechanisms behind this high level of efficiency. "We want to find out the design choices that evolution has taken to produce such an efficient enzyme," says Czub. "ATP synthase is extremely well conserved evolutionarily – the version we see in humans is almost exactly the same as the version we see in primitive bacteria, so it is a molecule that is seen across most of the lifeforms that are present today on Earth."

To be able to study the mechanisms of ATP synthase at an atomistic level, computer simulations are essential. X-ray crystallography and, more recently, cryo-electron microscopy, have provided a plethora of extremely valuable structural information about the enzyme at work, but these are limited to what are essentially still images – snapshots – of the mechanism, from which it is difficult to elucidate the entire dynamic cycle.

Czub and his colleagues have taken these still images from the field of structural biology as a starting point for their research, and are now using molecular dynamics simulations to try and fill in the gaps of the entire dynamic cycle of the enzyme, which involves a number conformational changes and movements of the active sites that are responsible for the synthetic chemical part of the mechanism. These changes are coupled to the motions of a rotary element that connects the ion transportation portion of the molecule to the chemical portion.

"The entire complex of the protein is huge, and along with solution that we simulate it in, consists of around one million atoms," says Czub. "This is why we needed the HPC resources that we have been

awarded by PRACE, in order to study the enzyme at the level of precision that we are interested in."

The two functional parts of the enzyme – known as F1 and F0 – are, due to computational limitations, usually studied separately, but with the PRACE access received by Czub and his colleagues they have been able to model the entire protein as one.

Although it is at a stage where it is difficult to give any definite answers about the findings, Czub is optimistic about the project's outlook. "We have produced somewhere between 10 and 20 terabytes of numerical data describing the dynamic evolution of all the atoms in the system, and we are now in the process of analysing that data in detail.

"We were very pleased with the transparent nature of PRACE's evaluation procedure, and with the interesting modifications to our project that they suggested."

"We are now fairly confident that we should be able to solve one of the central puzzles of the mechanism of ATP synthase, which is in what way the ATP release – which is not spontaneous – is coupled to the binding of ADP, the substrate ATP is made of. This question of how the enzyme harnesses the energy of ADP binding to drive the release of the product has been unknown for many years, and so we will be very pleased if we can shed some light on it."

The team is also hoping to explain the energy transmission mechanisms of the enzyme from the rotary element to its active sites. This stage of the processing of the data collected will no longer require the use of HPC, so it is now a case of carefully making sense of the findings.

In terms of the code used in the project, most of the work was done using the standard GROMACS code, which is well optimised for parallel computing. However, Czub has also been collaborating with colleagues from Germany and Sweden who have provided a computational tool for calculating electron densities.

Enhanced sampling has also played a large part in the project to allow the team to study how the enzyme operates at crucial stages of its mechanism. The technique – which uses the application of artificial external forcing to push the system in a certain way – allows the researchers to observe rare events in their simulations that would be unlikely to occur otherwise. The effects of the artificial forcing can then be subtracted to give a genuine picture

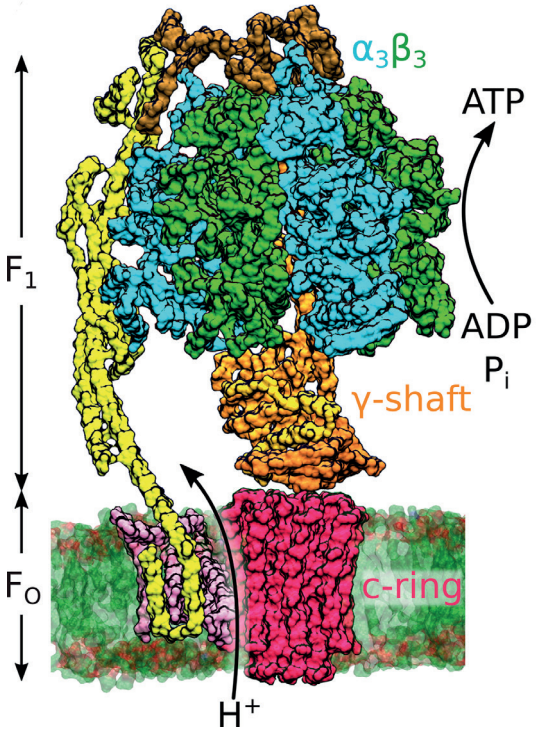

Molecular structure of ATP synthasereveals a remarkable architecture with two coupled oligomeric rotary motors: the transmembrane FO responsible for H+ translocation across the membrane and the soluble F1 which is involved in ATP synthesis.

of what happens, at a much lower cost of just using brute force and waiting for the rare events to occur. Czub is keen to highlight the positive experience of working with PRACE. "The whole project has gone very smoothly," he says. "We were very pleased with the transparent nature of PRACE's evaluation procedure, and we were also pleasantly surprised with the interesting modifications to our project that they suggested. This reassured us that we were being evaluated by real experts in the field, which is not always the case in this kind of situation. Although the architecture was fairly familiar to us, the people at the Leibniz Supercomputing Centre, where the SuperMUC system is hosted, were very helpful when we needed them, and we would definitely consider applying for resources with them in the future."

For more information
chem.pg.edu.pl/kbm/main-page

Resources awarded by PRACE
This project was awarded 40 million core hours on SuperMUC hosted by GCS at LRZ, Germany

Publications
J. Czub, M. Wieczór, B. Prokopowicz, H. Grubmüller. Mechanochemical Energy Transduction during the Main Rotary Step in the Synthesis Cycle of F1-ATPase. J. Am. Chem. Soc. 2017, 139, 4025–4034.

J. Czub, H. Grubmüller. Rotation Triggers Nucleotide-Independent Conformational Transition of the Empty βSubunit of F1-ATPase. J. Am. Chem. Soc. 2014, 136, 6960–6968.

J. Czub, H. Grubmüller. Torsional elasticity and energetics of F1-ATPase. Proc. Natl. Acad. Sci. USA 2011, 108:7408–7414.

Wings of the future

The high speed flight of aircraft creates turbulence that leads to increased drag and energy requirements. **Dr Marianna Braza**, Director of Research CNRS at the Institute of Fluid Mechanics of Toulouse has been using HPC to test out advanced technologies that morph the shape of the wing and produce vibrations in order to attenuate the effects of turbulence.

Marianna Braza of the Institute of Fluid Mechanics of Toulouse coordinates a Horizon 2020 European research project which is designing the aeroplane wings of the future. Partially inspired by bird wings, their designs incorporate a flexible wing structure with electroactive smart actuators (placed under the "skin" of the lifting surface) that provide optimal deformations and vibrations for reducing drag, increasing lift and reducing the aerodynamic noise. The combination of multiple kinds of actuations at different time and length scales is known as hybrid electroactive morphing. This is studied in collaboration with Prof. J.F. Rouchon's team at the IMFT Laboratoire Plasma and Energy Conversion in Toulouse.

The PRACE-supported project FWING - "Future smart wing design" has provided Braza and her colleagues with 15 million core hours on the Joliot Curie computer to carry out simulations to support their work. The FWING project is linked to the larger European project SMS – Smart Morphing & Sensing – which also includes experimental work using small prototypes as well as full scale models of the wings. Two types of electrodynamic actuators are used in the new wing design. The first take the form of wires made of shape memory alloys that are embedded inside the wing. When excited by electricity, these wires can deform the shape of the wing – known as the camber – in order to provide more lift. The second type of actuators involved are a new generation of piezo-actuators – strips added to the edge of the wing that can create higher frequency vibrations. These vibrations produce small vortices that effectively cancel out larger destabilising vortices created by the aerodynamics of the wing. In addition, these small vortices create extra lift. The overall effect is to make the virtual shape of the wing – the physical wing plus the wake it produces – thinner, reducing drag and therefore reducing the energy needed to propel the plane through the air. The simulations on the Joliot Curie machine have been used to find the optimal way for using these actuators in real time during flight.

"We have demonstrated, using a combination of numerical simulations and experimental data, that the piezo-actuators alone can increase lift by around 3%,"says Braza. "If we camber the wings as much as we can with the shape memory alloys then we can increase lift by around 10-15%, but as you do this you also increase drag, so we have used our simulations to find the optimal

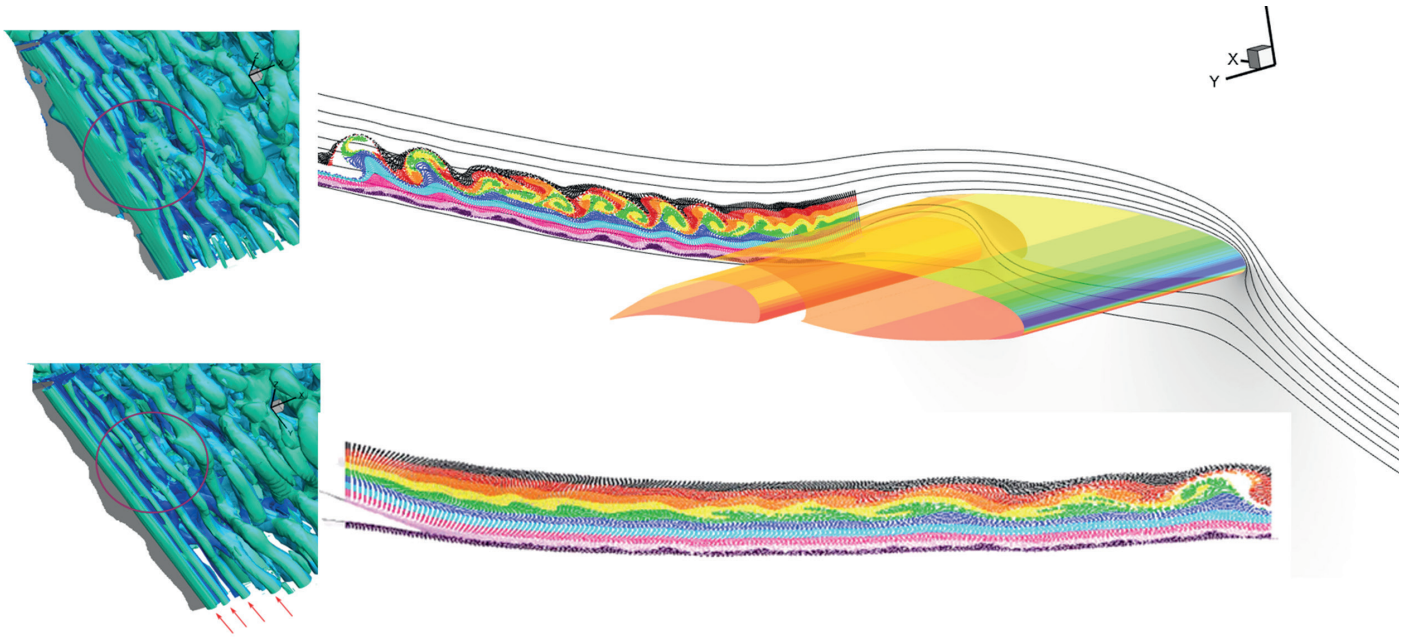

High fidelity simulations around the large-scale morphing wing with high-lift flap near full-scale in take-off configuration. Top: Non-morphing: formation of coherent vortices around the flap and in the wake. Bottom: morphing of the rear part by piezoactuators producing slight deformation and vibration of order 300 Hz: attenuation of these vortices and suppression of three-dimensional effects reducing the drag and aerodynamic noise and increasing the lift.

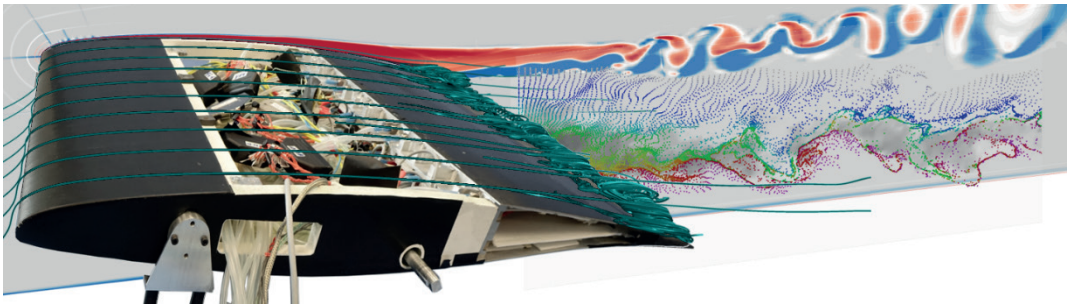

Large scale A320 Morphing Prototype. Simulations in the PRACE – FWING Project of IMFT by means of the NSMB code (Navier Stokes MultiBlock)

compromise for the ratio of lift to drag." Although the design of the wing is partially inspired by the wings of birds, at cruise speeds of 850-960 kmph the aerodynamic phenomena around a plane's wing start to look very different to those experienced by birds. "At these speeds you start to see the effects of air being compressed, which produces shock waves. We have demonstrated that our actuators can almost completely suppress the negative effects of these shockwaves and retain a good lift to drag ratio."

Having proved that they can provide increased performance, the members of the SMS project are now closely collaborating with Airbus in Toulouse. The aviation company plans to construct wings with these technologies incorporated for their A3xx commercial passenger planes, after first testing on large unmanned aerial vehicles in real flights within the next year, in order to prepare future testing on commercial airplanes.The code used for the simulations is called NSMB, which stands for Navier Stokes Multi-Block. It was created by a European consortium that includes a number of companies and institutions from across the continent, and is continuously upgraded and refined in terms of the numerical methods for defining the shape of the wing, and the methods for calculating the flow of air that surrounds the wing.

"Our actuators can almost completely suppress the negative effects of shockwaves."

To simulate a wing, a representative grid must be created, made up of millions of three-dimensional elements. Around the surface of the wing, the elements become smaller and smaller in order to get the most accurate depiction of the flow around it. The code then calculates the pressure on the wing and the velocity of the air around it. "The pressure and velocity are directly linked with the lift and drag," explains Braza. "If the code is well modelled, we can then use integration to calculate the lift and drag forces at a given speed."

Even using a grid with millions of points, it is not possible to calculate all of the details of the turbulence that surrounds the wing at high speeds. To get around this, the researchers use what is known as turbulence modelling to take into account the effect of these smaller vortices on the larger vortices that have been directly calculated. "Imagine our grid like a fisherman's net. The fish that are bigger than the holes in a net will be caught, while smaller ones will swim through," says Braza. "Our grid can capture and calculate the larger vortices, but the smaller ones cannot be seen. To get around this, the NSMB code includes within it a very advanced turbulence modelling component that takes into account the effect of these smaller vortices on the aerodynamic forces."

The code has excellent scalability, meaning that the time taken for information to pass between the processors does not increase a lot when run on higher numbers of processors at once. As such, the code was able to carry out the calculations while being run on 3000 processors simultaneously. With the FWING project now complete, the group is preparing a continuation of its work with PRACE. This will allow them to simulate the performance of its wing design at higher speeds and full-scale conditions that are closer to those found in flight, bringing them one step closer to incorporating their technology into real airplanes, increasing safety and reducing the energy needs by reducing the harmful instabilities in real time as well as the CO_2 and NO_x emissions.

For more information

www.smartwing.org/SMS
lejournal.cnrs.fr/videos/the-wings-of-the-future

Resources awarded by PRACE

This project was awarded 15 million core hours on Joliot-Curie hosted by GENCI at CEA, France

Publications

"Morphing of a supercritical wing by means of trailing edge deformation and vibration at high Reynolds numbers: experimental and numerical investigation", Journal of Fluids & Structures, https://doi.org/10.1016/j.jfluidstructs.2019.06.016

Effects of vibrating and deformed trailing edge of a morphing supercritical airfoil in transonic regime by numerical simulation at high Reynolds number », Journal of Fluids & Structures, 2019, https://doi.org/10.1016/j.jfluidstructs.2019.02.011

"Optimized design of real-scale A320 morphing high-lift flap with shape memory alloys and innovative skin", Smart Materials and Structures, 27, N° 11, 2018

Peptide bonds: simulating the origins of life

It is not known exactly how life on Earth first arose, but it is likely that it all began with the formation of a peptide. **Professor Marie-Pierre Gaigeot** of the Université d'Évry Val-d'Essonne has been investigating how this first step towards a biotic planet may have occurred at the interface between water and air.

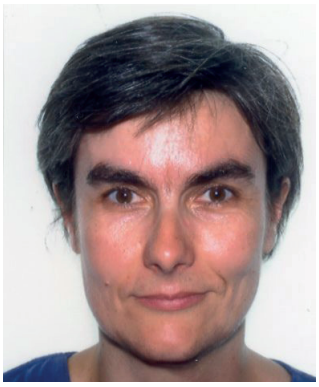

Marie-Pierre Gaigeot

Peptides are one of the fundamental building blocks of life. Finding out how they might have first formed from the primitive precursors of amino acids could provide a huge clue to how life first formed on Earth.

There have been many hypotheses about how the so-called peptide bond condensation reaction may have first occurred in the prebiotic soup that existed billions of years ago. Recent experiments by Professor Veronica Vaida of the University of Boulder in the USA showed that it is possible for this reaction to happen at interfaces between air and water, which would have been common in prebiotic times when the Earth was dominated by oceans.

Marie-Pierre Gaigeot of the Université d'Évry Val-d'Essonne is a researcher who specialises in investigating the structure and dynamics of water at interfaces, and how this structure impacts chemical reactivity. One of her recent papers showed the existence of a special two-dimensional network of hydrogen bonds that exists at interfaces between air and water, and she is now looking at whether this network might provide the right conditions for peptide bond condensation to happen. "Our investigations of this topic use density functional theory molecular dynamics simulations (DFT-MD)," says Gaigeot. "These are extremely computationally expensive, but it is necessary for us to model all of the electrons involved in this way as they play a fundamental role in whether a chemical reaction will occur or not in specific conditions."

The chemical reaction in question would be extremely unlikely to occur in a simulation without it being "forced" in some way. The researchers therefore bias their simulations so that they can not only observe the reaction happening, but also measure how much energy it costs to make it happen and decide whether this amount of energy would likely have been available in prebiotic conditions.

"To create the bias which leads you along the path of the chemical reaction, we had to decide coordinates of the reaction that made sense," explains Gaigeot. "We did this by carefully examining literature about peptide bond condensation reactions, as well as using basic common sense from a chemistry perspective. Although this does introduce a level of guesswork on our behalf, there is no other way of observing this reaction in a simulation." Gaigeot and her colleagues carried out DFT-MD simulations of the reaction in a number of different situations. One of the main parts of their hypothesis was that the 2D-network they discovered in their previous work helps the reaction to occur at the interface. To test this, they did one set of simulations at the air-water interface, and

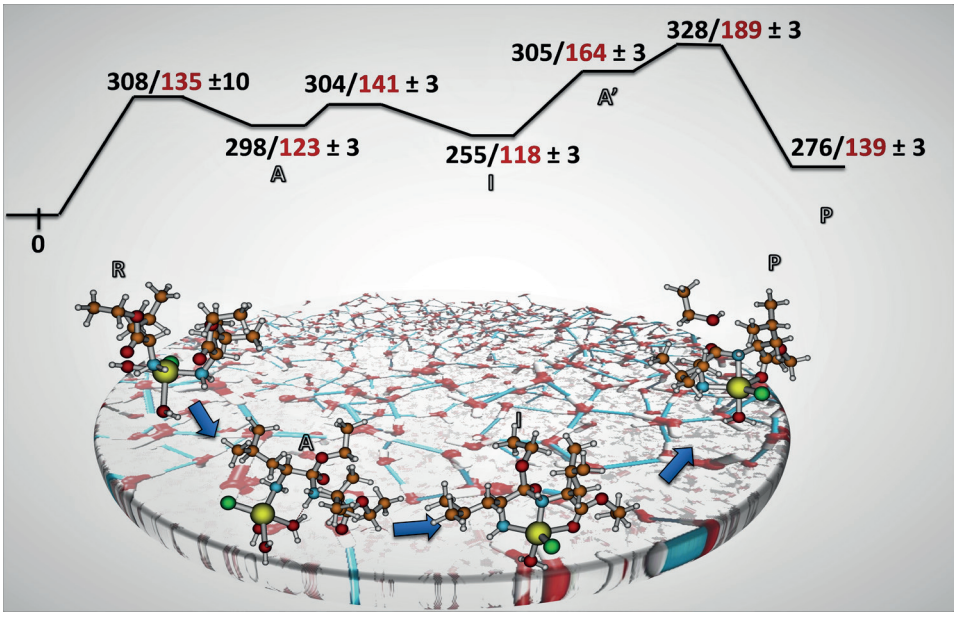

General scheme obtained by biased DFT-MD metadynamics of the peptide bond formation at the air-liquid water interface. Bottom: 3D-structures of reactants/products and intermediates; Top: Free energy landscape (energies in kJ/mol, black/red: simulations without/with added salt).

another just within water. "We guessed that in the liquid – without the interface – the energy barrier for the reaction to occur would be much higher, and that the structure found at the interface reduces the energy barrier and allows the reaction to happen."

As well as this, the researchers wanted to investigate whether the presence of salt affected the likelihood of the reaction occurring. In the experiments gathered by Veronica Vaida, copper chloride salt was diluted in the water. This salt is known to have probably been present in the oceans of primitive Earth, but it was not clear from Vaida's experiments whether its presence was important for the reaction to occur. The results of the simulations have confirmed many of Gaigeot's suspicions about the peptide bond condensation reaction. "We have found out the exact mechanism for the chemical reaction occurring at the interface and have quantified the different energy barriers along the pathway. We can say that it is more than probable that this reaction could have occurred in prebiotic times.

"We are now working on a paper that we believe will be of high research impact, and we could not have done this without PRACE."

"On top of that, we have shown that the presence of copper chloride salt is absolutely essential for the reaction to occur, as it acts as a catalyst, reducing the energy barriers and allowing certain portions of the pathway of the reaction to happen."

The researchers have also found out why the reaction does not happen when there is no interface. "We saw that the amino acids need to be in a particular conformation for the reaction to occur, and while this conformation is unlikely to occur in bulk water, it is the most likely conformation at the air-water interface. As well as that, the 2D-network that we discovered previously plays an essential role in creating this conformation."

Gaigeot plans to continue this line of research by looking more closely at the catalysis by the copper chloride salt. The salt is known to act as a catalyst in other situations, but the demonstration of this happening at the air-water interface is a first. "Overall, I would like to stress the importance of the allocation we received from PRACE. We were awarded 29.3 million CPU hours on the Curie supercomputer, which sounds enormous on paper, but was absolutely necessary for our investigation. My PhD student who oversaw the simulations told me that in total we managed to simulate one nanosecond of trajectory of our chemical system, which in terms of DFT-MD simulations is huge. We are now working on a paper that we believe will be of high research impact, and we could not have done this without PRACE."

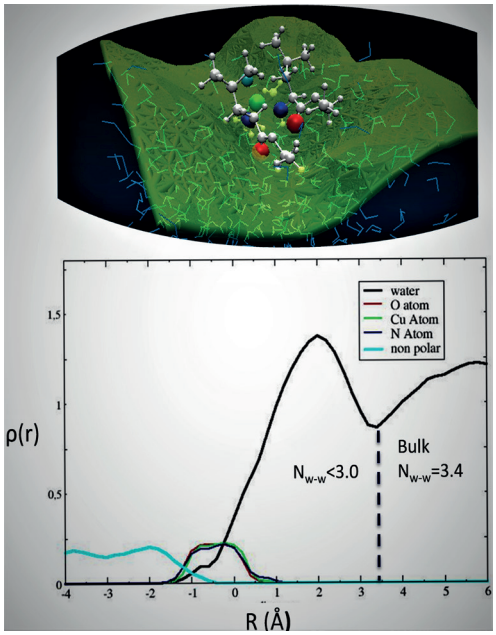

Top: image of the amino acid reactants at the air-liquid water interface. The interface is defined and materialised by the instantaneous Willard and Chandler's surface (in green), the water is located below and schematised with the blue links, the reactants are shown "surfing" above the surface. Bottom: density profile of the water (black), the limit of the BIL (Binding Interfacial Layer) interface is materialised with the vertical dashed line, beyond r>3Ang the centrosymmetric bulk liquid water is recovered. R is the vertical distance from the instantaneous Willard and Chandler's surface. In colors, density profiles of the different atoms of the reactants, showing that the reactants are surfing above the BIL interface. See ref [J.Phys.Chem.Letters, 8:2133 (2017)] for a detailed description of the special BIL 2D-HBond-Network existing at the air-liquid water interface.

For more information

mpgaigeot-research.fr

www.univ-evry.fr/recherche/unites-de-recherche/sciences-du-vivant/laboratoire-analyse-et-modelisation-pour-la-biologie-lenvironnement-lambe.html

Resources awarded by PRACE

This project was awarded 29.3 million core hours on Curie hosted by GENCI at CEA, France

Publications

J.D. Cyran, M.A. Donovan, D. Vollmer, F. Siro-Brigiano, S. Pezzotti, D.R. Galimberti, M.-P. Gaigeot, M. Bonn, E.H.G. Backus - Molecular Hydrophobicity at a Hydrophilic Surface. PNAS 116:1520-1525 (2019)

F. Creazzo, D.R. Galimberti, S. Pezzotti, M.-P. Gaigeot DFT-MD of the (110)-Co3O4 cobalt oxide semiconductor in contact with liquid water, preliminary chemical and physical insights into the electrochemical environment. J. Chem. Phys. 150:041721-18 (2019)

S. Pezzotti, D.R. Galimberti, M.-P. Gaigeot 2D H-Bond network as topmost skin to the air-water interface. J. Phys. Chem. Letters, 8:3133 (2017)

This research was carried out by PhD students Flavio Siro-Brigiano and Simone Pezzotti, and postdoctoral researcher Daria Ruth Galimberti.

Shock-related buffeting in aeroplanes

Shock-related buffeting is a phenomenon that occurs when air passes over the wing of an aeroplane under extreme conditions and can have profound consequences for how wings are engineered and their durability. **Professor Neil Sandham** of the University of Southampton has been investigating this using direct numerical simulations.

Professor Neil Sandham

The flow of air over aircraft wings is an important application for scale-resolved numerical simulations that is only just becoming feasible with improved algorithms and computer hardware. An improved understanding of the complex physical behaviour of turbulence, including transition to turbulence and the interaction of turbulence with shock waves, will pave the way to improved engineering design. This includes reducing air resistance (and thus lowering fuel consumption), increasing the robustness of wings, and improving wing characteristics.

For passenger transport aircraft, the flow over the wing becomes locally supersonic, terminating in a shock wave. The changed flow physics relative to low speed flight leads to an additional source of drag, known as wave drag, which must be included in the design optimisation of complete wings. Additionally, these shock waves can strongly interact with the boundary layer, leading to local flow separation and eventually to an undesirable phenomenon involving low-frequency oscillations known as transonic buffet, which can affect the structural life of wings.

Neil Sandham and his colleagues from the University of Southampton have been working on a PRACE-supported project that has been investigating shock-related buffet. Their simulations have gone beyond previous research on the topic, increasing the complexity of the flow and moving from low speed flow over the aerofoil to higher speeds where shock waves are seen. "Previous studies of buffet have been limited to the averaged equations of motion where what you have to do is add empirical turbulence models, which are unreliable when you have flow separation as you see in buffeting," explains Sandham.

Using the Hazel Hen system hosted by GCS at HLRS, the researchers have run simulations where all of the scales of the turbulent flow are resolved (excluding the internal structure of the shock waves). Although they have not been able to simulate at the Reynolds numbers (the governing parameter of turbulent flow) seen in full flight mode, they have been able to reach figures that provide similar physics seen in flight. This provides useful insight into the physics of buffets.

The simulations carried out ran on around 30 000 cores, and were able to successfully capture the buffet process. "We saw variations in the boundary layer separation and the interactions with shock waves," says Sandham. "One novelty we found is that these shock wave motions do not occur at the same frequency as the buffet. In the simulations we used large computational domains, so we are pretty confident that this is not the cause."

The team was also able to run cases with modified geometries to investigate how this might affect the buffet process. In particular, the team was able to add a splitter plate to the trailing edge of the wing to see how the buffet responded. As well as this, the researchers investigated the use of simpler methods called large eddy simulations, where the large and medium scales of flow are resolved but the smaller scales are modelled. "The idea is to create a similarly accurate method that uses less computational resources," explains PhD student Markus Zauner. "We were successful in creating a large eddy method that reproduced the buffet and used a factor of 16 times less computational resources than direct numerical simulations, although we did find that these were sensitive to resolution." One difficulty of the project was that it required a lot of effort to

"You do not know at the beginning which parts of the simulations you are going to be interested in"

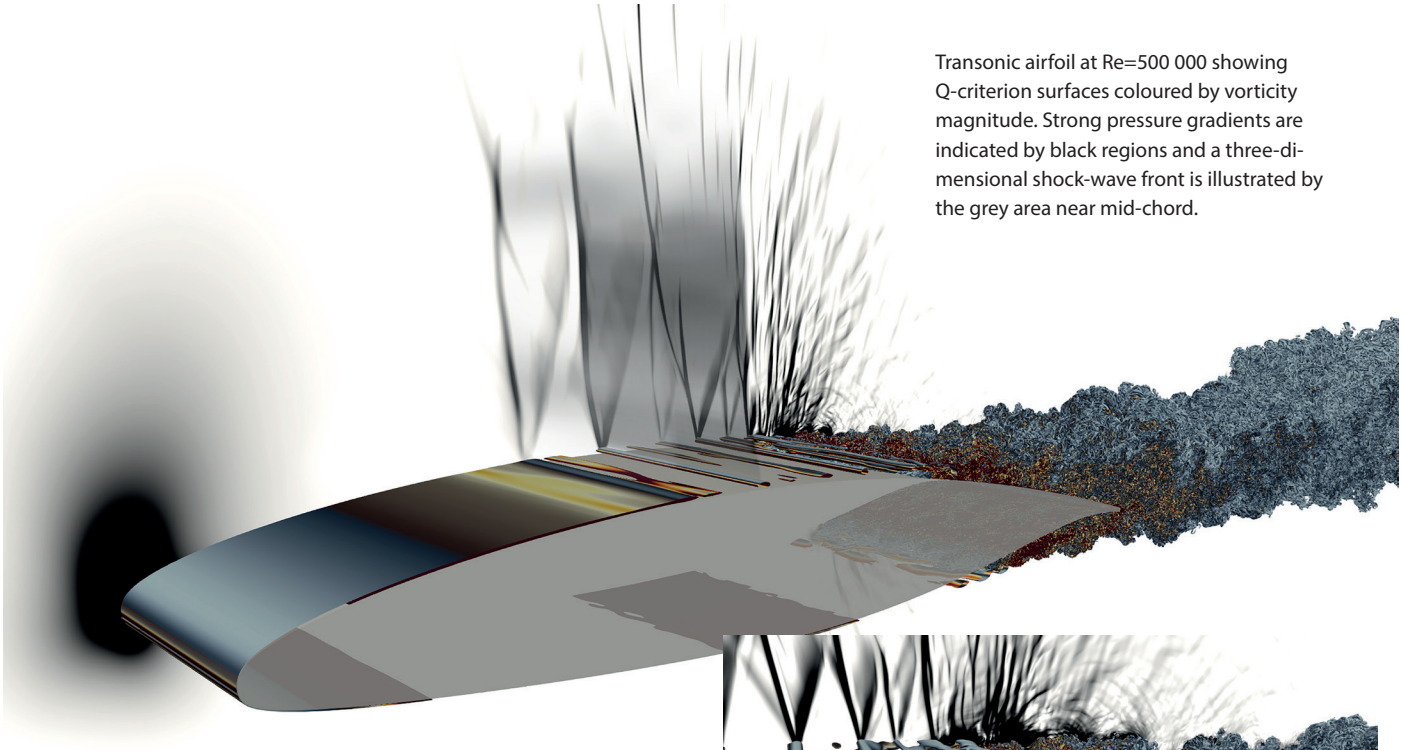

Transonic airfoil at Re=500 000 showing Q-criterion surfaces coloured by vorticity magnitude. Strong pressure gradients are indicated by black regions and a three-dimensional shock-wave front is illustrated by the grey area near mid-chord.

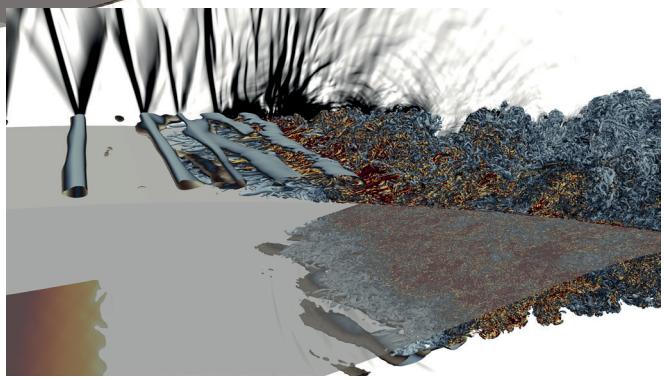

Close-up of the transition region showing turbulent break-down of 2D roll-ups. The full simulation consists of more than five billion grid points.

generate the grids upon which the calculations were performed. This labour-intensive process was part of Markus Zauner's sub-project, in which he was able to apply some error indicators and set up a grid generation system that allowed them to optimise the grid much more effectively than in the past. This software has now been released as open source and so is available to other researchers working within the same field.

"My work involved setting up the simulations and ensuring that we had a suitable grid to work with," Zauner explains. "Then we had to set up the production runs. This sounds straightforward, but because we were producing quite a lot of data, we had to be very careful to establish a process to avoid loss of data."

The amount of data produced during the project has been astounding, even compared to other High-Performance Computing projects. "In our simulations, we produced huge amounts of data, but of course you do not actually need all of it to reconstruct the three-dimensional flow dynamics," explains Zauner. "You do not know at the beginning which parts of the simulations you are going to be interested in, so after you have done them you go through post-processing and find interesting examples that show particular flow behaviour. Even after all of this, though, we were still left with around 30 terabytes of selected data."

Sandham and his team will continue to exploit the data produced by the simulations long after the project has ended, and smaller versions of their database will be released to the public in association with their publications. This will ensure that the results of the project get used as much as possible and will hopefully yield some more findings in the near future.

For more information

www.southampton.ac.uk/engineering/about/staff/nds9.page

Resources awarded by PRACE

This project was awarded 35 million core hours on Hazel Hen hosted by GCS at HLRS, Germany

Publications

Zauner, M, Sandham, N.D. 2019 Modal analysis of a laminar-flow airfoil under buffet conditions at Re=500,000. Accepted for publication in Flow Turbulence and Combustion.

Zauner, M., De Tullio, N, Sandham, N.D. 2019 Direct Numerical Simulations of Transonic Flow Around an Airfoil at Moderate Reynolds Numbers. AIAA Journal 57(2), 597-607.

Jacobs, C.T., Zauner, M., De Tullio, N., Jammy, S., Lusher, D and Sandham, N.D. 2018 An error indicator for finite difference methods using spectral techniques with application to aerofoil simulation. Computers & Fluids 168, 67-72

Vortex-induced vibrations for energy harvesting

New wind harnessing generators that gather energy through a phenomenon known as vortex-induced vibrations could represent a new frontier for renewable energy. **Dr Oriol Lehmkuhl** of the Barcelona Supercomputing Centre has been using supercomputers to help advance this technology.


Oriol Lehmkuhl

Over the last decade, the renewable energy sector has been growing at a much faster pace than the rest of the economy in Europe. Wind power especially has seen substantial growth, and a lot of research is currently being done to work out the best way to extract this bountiful energy source.

Nearly all wind energy is currently harvested using what are known as horizontal axis turbines – the large windmill-liked turbines that are already common across Europe. However, there are other ways of harnessing wind energy that do not use this type of structure.

A relatively new type of wind generator that does not use rotational movement to capture energy is now being studied extensively across the world. The basis of this technology is a phenomenon called vorticity. Bladeless structures, resembling little more than tall poles, convert wind power into electricity through the oscillation that happens when the structure of the device reaches the same frequency of resonance as the wind vortexes created behind. This is known in fluid dynamics as vortex-induced vibration (VIV).

Oriol Lehmkuhl, a postdoctoral researcher from the Barcelona Supercomputing Centre, recently won the PRACE Best Industrial Presentation during the PRACEdays19, for his work on the topic. The research project, titled "VIVALdI - HPC of Vortex Induced VibrAtions for flow controL and energy harvestIng", was awarded 27 million core hours on the German supercomputer SuperMUC.

"The impact of the results will help industry to build more competitive wind-harnessing structures."

"When you have a fluid interacting with a solid, vibrations are generated and the solid begins to move," explains Lehmkuhl. "This phenomenon is well known in bridges. If the wind is the right speed and direction, it can cause the supporting structures of the bridge to vibrate. In some cases, this has even caused bridges to collapse." Because of the inherent danger of these types of vibrations, most research into them has revolved around safety. But Lehmkuhl's research has taken a different view of the

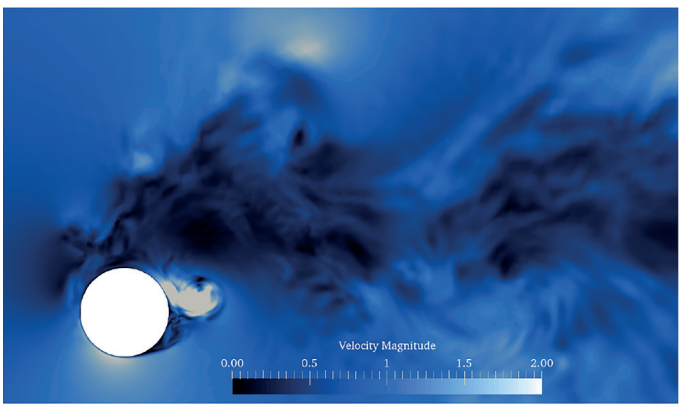

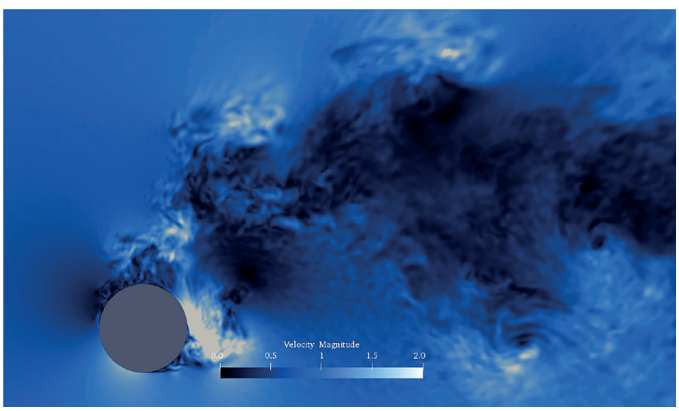

Instantaneous flow in the subcritical regime in the super-upper branch at U* = 5.5. Re = 5300 (left) and Re = 105 (right).

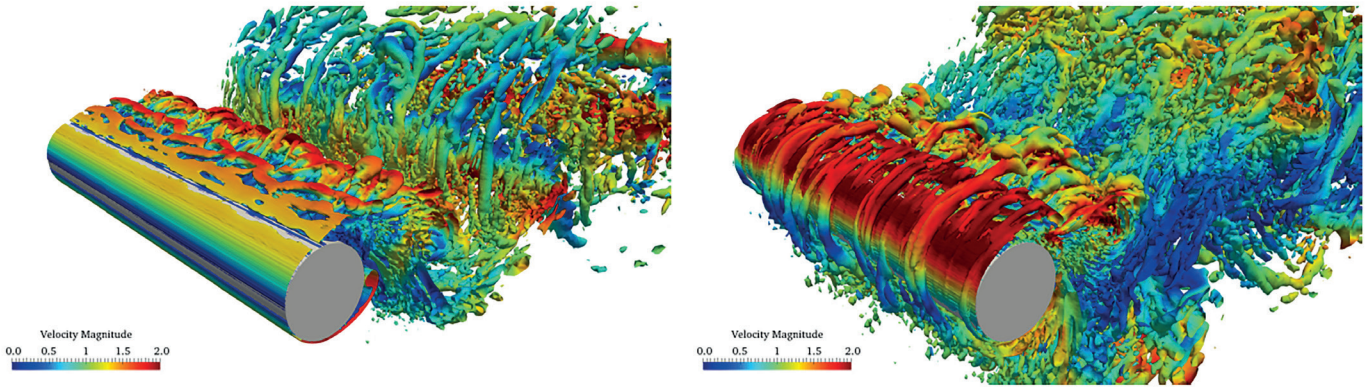

Instantaneous vortical structures close to the cylinder surface. Initial branch (left), superupper (right)

phenomenon – one of opportunity. He and his colleagues are interested in seeing whether it is possible to harness the energy of these vibrations, using tall, flexible structures that move in the wind.

This topic has been investigated to a certain extent, but using small cylinder structures and with low wind speeds. Lehmkuhl's research has extended this to more realistic wind speeds with cylinders closer to the size of standard wind turbines. "Our hope is that in the future this type of technology could be complementary to standard wind energy turbines."

In the VIVALdI PRACE project, high fidelity simulations of a cylindrical body oscillating in a free-stream from sub-critical to super-critical Reynolds number have been carried out by means of wall-resolved large-eddy simulations using thousands of CPUs and meshes of hundreds of million of elements. This is the first time these kind of simulations have been performed at this level of modelisation, providing a step forward in the understanding of the physics of fluid-structure interaction in the range of industrial applications.

"Thanks to the PRACE Tier-0 resources, we have been able to carry out large scale simulations that will help build bigger and more efficient VIV wind turbines. The impact of the results will cover very basic research aspects of the turbulent fluid-structure physics, but will also help industry to build more competitive wind-harnessing structures like these", says Lehmkuhl.

"In our work we simulate the fluid and we simulate the structure, and we couple them together in the same simulation domain," says Lehmkuhl. "For the fluids, we use large eddy simulation models that directly simulate the larger scales of the flow and then use turbulence models for the smaller structures of the flow."

The simulations are challenging computationally. The model of the fluid in itself is extremely computationally intensive and, on top of this, because the solid is moving all of the time, the mesh needs to be changed throughout the simulation to make the calculation. As such, High-Performance Computing is needed, with many CPUs working at the same time to solve the system. The team is still finishing some of the calculations, but overall Lehmkuhl is pleased with what they have achieved so far in the project. "We have identified how these vortex-induced vibration mechanisms happen at very high Reynolds numbers, and have gained more insight into other mechanisms involved with these technologies. Hopefully this research will help others to push this technology further towards becoming a reality."

For more information

scholar.google.com/citations?hl=es&user=xM19OIAAAAAJ&view_op=list_works&sortby=pubdate#

Resources awarded by PRACE

This project was awarded 27 million core hours on SuperMUC hosted by GCS at LRZ, Germany

Publications

D. Pastrana, I. Rodriguez, J. C. Cajas, O. Lehmkuhl, and G. Houzeaux. On the formation of Taylor-Görtler structures in the vortex induced vibration phenomenon, submitted to International Journal of Heat and Fluid Flow (2019).

I Rodriguez, O. Lehmkuhl, D. Pastrana, J.C. Cajas and G. Houzeaux. Wakes and instabilities of static and freely vibrating cylinder. Direct and Large Eddy Simulations XII, ERCOFTAC Series, Springer, Cham.

O. Lehmkuhl, I Rodriguez, D. Pastrana, J.C. Cajas and G. Houzeaux. High fidelity simulation of vortex induced vibrations for flow control and energy harvesting, EuroHPC Summit Week 2019, PRACE Days, Poznan, Poland, 2019.

Oriol won the PRACE Best Industrial Presentation Award at PRACEdays19 during the EuroHPC Summit Week 2019 www.prace-ri.eu/prace-awards-eurohpc-summit-week-2019/

Advanced control of wind farms

Turbines in wind farms are programmed to act selfishly – they optimise their own power output and structural loading while ignoring the impact of their wake on other turbines. **Dr Paolo Schito** of the Politecnico di Milano has been using HPC to develop methods for making turbines work together in harmony to increase power production.

Current practice in wind farm operation involves each turbine being separately controlled to optimise its own performance in terms of energy capture. This way of operating means that each turbine operates independently, based on the information it gathers about its surrounding environment.

However, turbines within a wind farm do not exist in isolation from each other. Due to their relative proximity, the wake of one turbine can lead to loss of power generation for turbines downwind of it. Wind turbines in the wake of another turbine experience a wind field with a lower average wind speed and higher turbulence intensity, causing a decrease in the available energy while increasing mechanical loads. One study of the Lillgrund Wind Farm off the coast of southern Sweden showed power generation losses of up to 23% due to these wake effects.

"Wind turbines in the wake of other turbines have less energy available to them."

The major reasons for this non-optimal approach to wind farm operation is the fact that turbines were initially developed as single machines, and only subsequently were placed together in wind farms. There is also a lack of knowledge and tools which can model the dynamics of the flow inside the wind farm, how wind turbines modify this flow, and how other wind turbines are affected by the perturbed flow.

Paolo Schito of the Politecnico di Milano is part of the Horizon 2020 project CL-Windcon, which is looking at ways to optimise energy harvesting in wind farms. It will develop new innovative advanced open and closed loop wind farm control algorithms, allowing entire wind farms to be treated as an integrated optimisation problem.

Paolo Schito

"These algorithms will be validated using data from three sources," explains Schito. "Firstly, we are doing wind tunnel testing at our institute in Milan. As well as that, we are gathering information in the field from a wind farm in Sardinia. Finally, we are doing high-fidelity numerical simulations to investigate the interactions between turbines."

The numerical simulations are being carried out as part of the PRACE-supported project "ALISIOS – Advanced controL of wInd farm uSing computatIOnal fluid dynamicS", using 30 million core hours on the MARCONI-KNL machine hosted at CINECA. The numerical model being used was validated against the experimental data gathered from wind tunnel and full-scale farms. Testing was then carried out using open-source software managed by the USA National Renewable Energy Laboratory (NREL).

"The software we use, called the Simulator for Onshore/ Offshore Wind Farm Applications (SOWFA), is the only open-

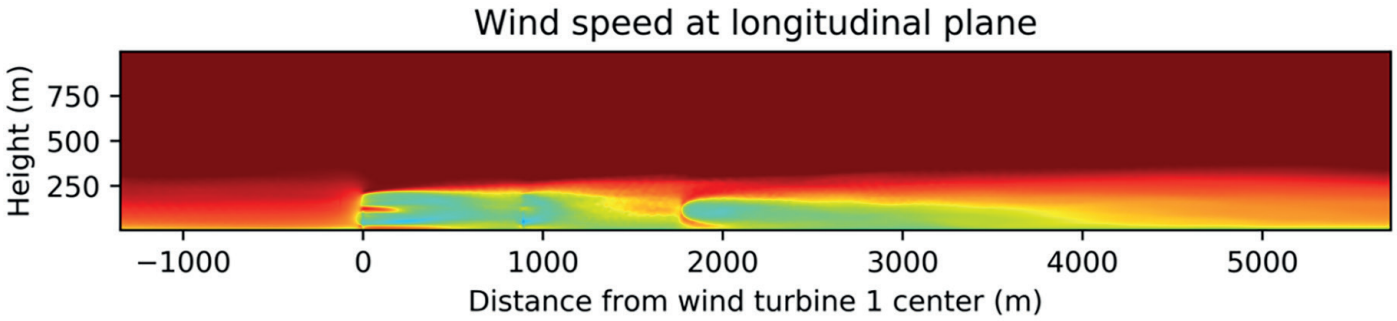

Lateral view of the wake interaction between three aligned wind turbines

source software that is capable of coupling the aerodynamics of the rotor with the structural dynamics of the wind turbine and the control strategy," explains Schito. "Developed and supported by NREL, this software has had many contributions from experts around the world to make it so functional."

Initially, the simulations were used to understand how the wake of a single wind turbine develops, in order to have a model of a single turbine wake. "After that, we were looking at what we call the superposition of turbine wakes, seeing how they mix and how they overlay on top of each other," says Schito. "This was first done with 3x1 wind farms and then with 3x3 wind farms, totalling nine turbines. We then applied various control techniques to see how this affected overall power generation."

Schito and his colleagues have investigated two different control techniques for more optimised energy harvesting in wind farms. One of them is to reduce the power extraction of the windward turbines, so that more energy remains in the wind that comes through to the downwind turbines. The other control technique is to steer the windward turbines so that their wake is directed away from the downwind turbines.

The setups of the wind farms being tested in the simulations are idealised, with very narrow spacing between the turbines and uniform, flat terrain, but this does not mean that the data gathered is not useful. For groups using lower resolution methods to develop and further explore wind farm control techniques, this data will be invaluable for tuning their models. The data has been made publicly available as one of the main outcomes of the project.

Overall, the simulations have shown that an increase of power extraction of between 8-10% can be achieved in good conditions. Looking beyond this particular PRACE-supported project, the team is thinking about applying for more access to investigate the effects of more complex terrain. With wind energy offering one of the best opportunities for Europe to reduce its CO_2 emissions and increase its share of renewable energy, this kind of invaluable research will be essential for the future health of our planet.

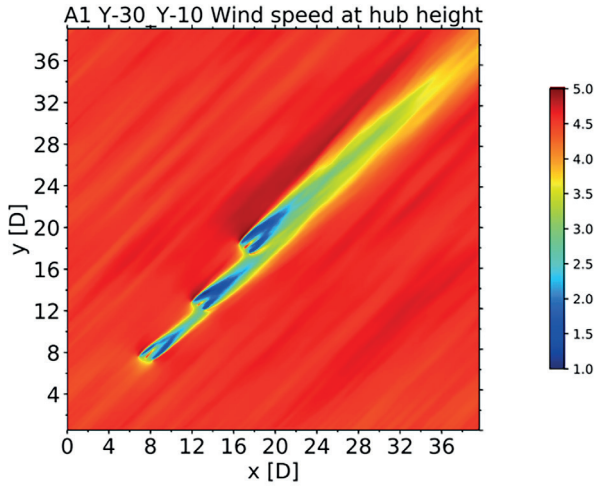

View of the average wind field for three wind turbines with yaw control for wake redirection and maximum power production

For more information

www4.ceda.polimi.it/manifesti/manifesti/controller/ricerche/RicercaPerDocentiPublic.do?EVN_PRODOTTI=evento&lang=EN&k_doc=151473&aa=2018&n_docente=schito&tab_ricerca=1&jaf_currentWFID=main

Resources awarded by PRACE

This project was awarded 30 million core hours on MARCONI-KNL hosted by CINECA, Italy

Publications

Melani P.F., Schito P., Persico G., "Experimental Assessment of an Actuator-Line Simulation Tool for VAWTs", Wind Energy Exploitation in Urban Environment. TUrbWind 2018 (ISBN: 978-3-030-13530-0; 978-3-030-13531-7)

Schito P., Bayati I., Belloli M., Bernini L., Dossena V., Zasso A., "Numerical wind tunnel tests of an open data IPC-VAWT", Wind Energy Exploitation in Urban Environment. TUrbWind 2017 (ISBN: 978-3-319-74943-3; 978-3-319-74944-0)

Ferrari G., Federici D., Schito P., Inzoli F., Mereu R., "CFD study of Savonius wind turbine: 3D model validation and parametric analysis", RENEWABLE ENERGY 2017 (ISSN: 0960-1481)

Sand and dust forecasting in Europe

Sand and dust storms affect many areas of our lives. They change the quality of the air we breathe and can have negative consequences for industries such as aviation and solar energy production. **Dr Sara Basart** of the Barcelona Supercomputing Centre has been using High-Performance Computing to create unrivalled predictions of these events.

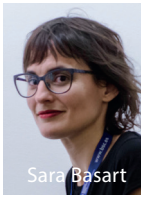
Sara Basart

Over the past decade, there has been a growing recognition of the crucial role of sand and dust storms on weather, climate and ecosystems, along with their important adverse impacts upon life, health, property and the economy.

Sara Basart is a researcher at the Earth Sciences Department in the Barcelona Supercomputing Centre (BSC). A core activity of the department is sand and dust storm modelling and forecasting from regional to global scales. The expertise of the department is the result of the group's participation in several European and international projects. Basart's work examines desert dust and how High-Performance Computing can be used to predict sand and dust storms at global and regional scales. Applied topics being explored by her group include looking at the impact of dust on weather and creating predictions for potential sand and dust storms in the region of North Africa, the Middle East and Europe.

So what exactly are dust forecasts used for? During the strong heatwave earlier this year, high levels of dust were brought over to Europe from Africa. High levels of atmospheric dust affects the quality of air that we breathe, which is significant when considering the European legislation on air quality.

For more information
dust.aemet.es/Barcelona Dust Forecast Center

Resources awarded by PRACE
This project was awarded 21 million core hours on MareNostrum hosted by BSC, Spain

Publications
Status and future of numerical atmospheric aerosol prediction with a focus on data requirements, Atmos. Chem. Phys., 18, 10615–10643, https://doi.org/10.5194/acp-18-10615-2018, 2018.

High-resolution dust modelling over complex terrains in West Asia. Aeolian research, 23, 37-50.

Current state of the global operational aerosol multi-model ensemble: An update from the International Cooperative for Aerosol Prediction (ICAP). Quarterly Journal of the Royal Meteorological Society.

"When there are higher levels of particles in the atmosphere, it is important to ascertain whether the source of the peak is natural or anthropogenic," says Basart. "If you can show that the peak is due to naturally occurring events such as dust and sand being whipped up from the Sahara, then there is less blame on societal sources. And if this affects large parts of the population, then it is important to make these people aware of the potentially negative impacts on their health."

The south of Europe tends to be more affected by dust and sand storms than the north, although sometimes signs of these events are seen as far up as Scandinavia. However, it is countries such as Cyprus, Greece, Italy, Spain and Portugal that tend to bear the brunt of atmospheric dust.

The success of the predictions made by Basart's group has led to them being recognised by the World Meteorological Organisation (WMO). The BSC hosts the WMO Regional Specialised Meteorological Centre which is the reference centre for providing operational sand and dust products for Northern Africa, the Middle East and Europe. In the context of this short-term forecasting they began working on a PRACE-supported project titled "High-resolution regional dust reanalysis based on ensemble data assimilation (eDUST)". "We wanted to improve the accuracy and performance of our prediction using data assimilation techniques," she explains. "This involved the inclusion of observations available into our dust predictions. Of course, this kind of project requires a lot of computing resources, which is the reason we came to PRACE."

It is in fact possible to run predictive simulations without HPC, but the coarse resolution and the limited area of simulation restricts the potential of research like this. For example, certain microscale sand and desert storms called haboobs are only possible to reproduce at high resolutions (below 4km). To do these kinds of simulations, supercomputers become essential. "We are using parallel codes that scale well with many processors, which can give us a final result within a few hours or even minutes," says Basart. "This is impossible with a normal computer."

The project involves creating a model ensemble that is taken to derive the best solution. Additionally, observational data from satellites is also included to correct errors in the model's

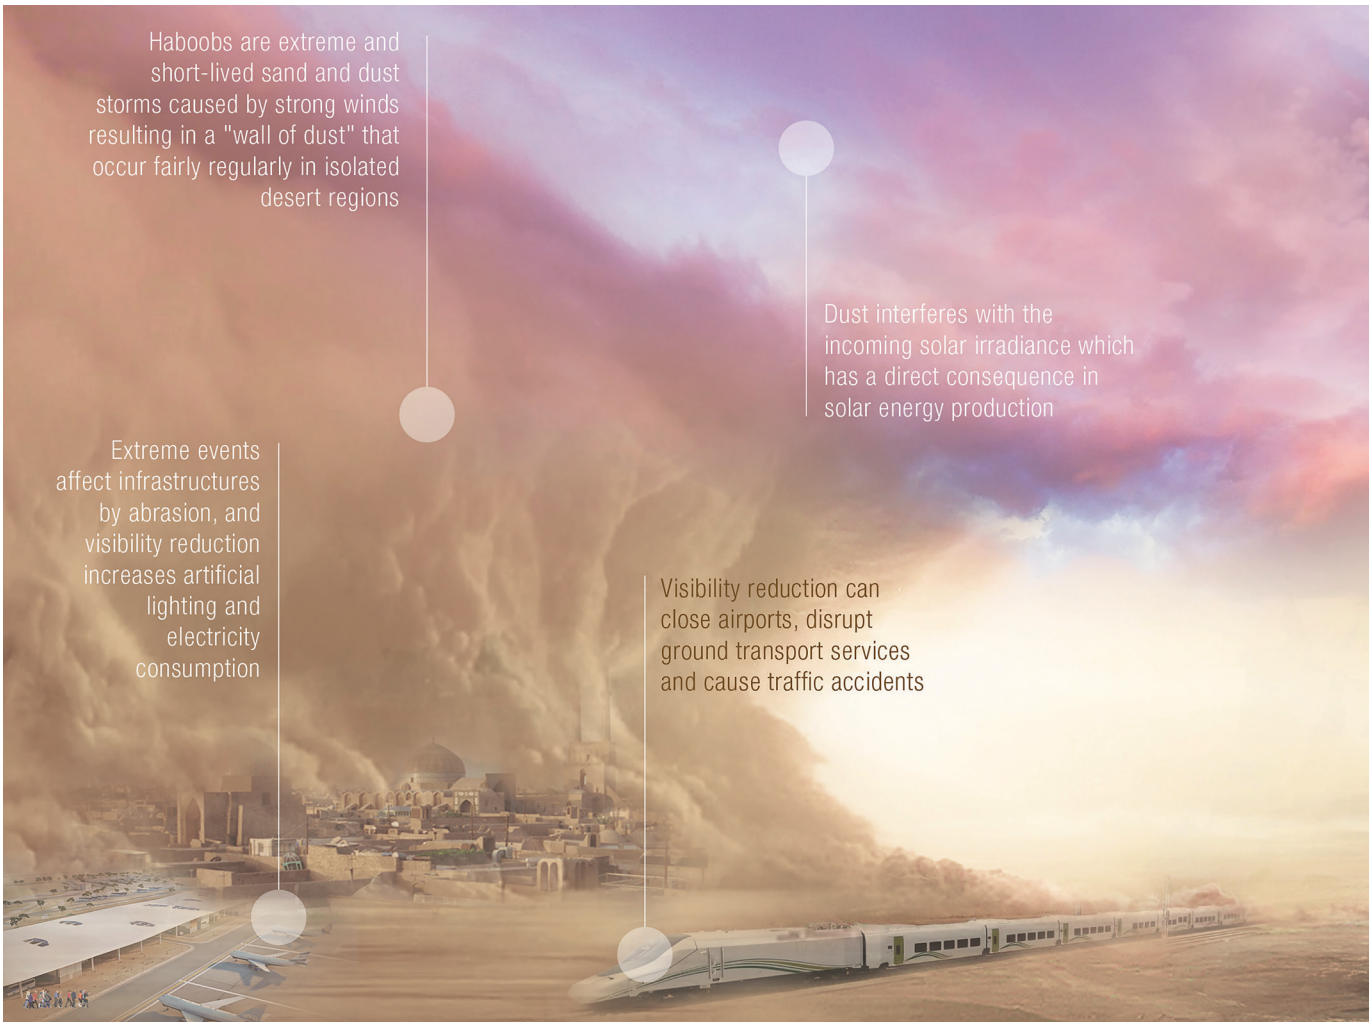

predictions. "We converge on to our best solution by using multiple and independent simulations based on the same model to account for model uncertainty in the calculation of the data assimilation corrections. This allows for the uncertainty of forecasting to be included in our final prediction," says Basart. "This is an advanced way of improving predictions and, as you can imagine, requires a lot of computing time."

Initially, Basart and her colleagues ran a reanalysis of past data, carrying out simulations of the period between 2000 to 2015 using observations that had already been collected. Such ensemble runs make the optimisation of the model efficiency even more crucial. Then, the PRACE-supported project was used to fix the implementation and developments needed for running such simulations, from the design and set-up of a specific workflow created for this purpose, looking at how the data from satellites impacted the results and how they should be used in the future. Thanks to the PRACE-supported project, the configuration for doing these predictions has now been finalised.

Air quality has been shown to have a marked effect on health, but it can also have impacts in other socioeconomic sectors, such as aviation. "When there is a strong sand or desert storm in Egypt, for example, it is likely that the airports over there will be closed shortly after, which can have a knock-on effect on air traffic in Europe," Basart explains. "This is an area of research that we are just beginning to look into now – connecting these kinds of phenomena with specific activities."

Basart and her team have been in contact with meteorologists who work with airports to try and help them with these issues, and they hope that their operational dust forecasts can go on to be applied in many other sectors. The harvesting of solar energy, for instance, is greatly influence by the presence of dust in the atmosphere, and so accurate forecasts could potentially help to improve predictions of how much energy will be produced by solar plants in the coming days.

Looking to the future, Basart hopes to improve the accuracy of their simulations in order to show up more and more of the phenomena that are known to occur in the real world. But she also foresees a potential bottleneck for the future – data storage. The effective use of current gigaflop-scale and future teraflop-scale computing systems requires a balance between compute power and I/O capability. "The problem we have now is that we are producing so much data. This is a real problem which I think many scientists working with HPC are facing nowadays."

New approaches in radiotherapy

Treating cancer by firing high doses of radiation at the tumour is often the most effective method, but it is difficult to not damage the surrounding tissue in the process. **Dr Yolanda Prezado** of CNRS has been using HPC to investigate a new method for radiotherapy that reduces the damage caused to nearby tissues, allowing for greater doses to be delivered to the tumour.

Radiotherapy is one of the most frequently used methods for treating cancer, with over 50% of all cancer patients receiving the treatment during their illness. Although advances in the technologies used to deliver radiation therapy have enabled improved accuracy towards tumours and reduced dose to nearby tissues, the dose tolerance of normal tissues remains a limiting factor for the effectiveness of radiotherapy. This leads to difficulties in treating radioresistant tumours, such as gliomas, or cancers near highly sensitive structures, such as paediatric brain tumours.

Finding new approaches that increase the resistance of normal tissues would make it possible to escalate tumour dose, resulting in an improvement in cure rate. Yolanda Prezado, a researcher from CNRS, has been investigating a new approach called proton and heavy ions minibeam radiation therapy. It combines the advantages of using protons and heavier ions – superior dose distribution and better precision – with the use of submillimetric field sizes and spatial fractionation of the dose, as is done in minibeam radiation therapy.

Minibeams have a width of between 500 and 1000 microns Recently, Prezado and her team carried out the first technical implementation of this combined approach at the Curie Institut-Orsay Proton Therapy Centre, confirming its technical feasibility.

"This way of irradiating seems to activate biological mechanisms that are not well understood yet, but lead to a remarkable increase in normal tissue resistance, opening the door for a dose escalation in the tumour," explains Prezado.

The use of a pencil beam scanning (PBS) system could further improve the technique and bring it to patients, since it allows a higher degree of precision and minimises the overall exposure and radiation to healthy tissue. By employing Monte Carlo simulations and thanks to access awarded by PRACE to the MareNostrum system hosted by the Barcelona Supercomputing Centre, Prezado and her team have been studying how to adapt this system for proton minibeam radiation therapy.

Monte Carlo simulations have aided the researchers in many ways, from the conception of the idea, its refinement and optimisation, and as a guide for its technical implementation

Yolanda Prezado

and for further biological experiments. "These simulations are helping us to design and prepare our future clinical trials and allow us to calculate what a treatment plan would look like for a patient in terms of dosage."

The work with the MareNostrum machine is now over, and the team have gained enough biological data about proton minibeam radiation therapy to allow them to start thinking about clinical trials. From the more physics-based side of the research, they are using their PRACE access to optimise beam generation for the treatment of patients in the future, and to calculate doses for potential future treatment plans.

Having recently received an ERC consolidator grant, Prezado and her colleagues plan to continue their work on this exciting new therapeutic approach for several more years and hopefully up to the point where it can begin to be used to help people with cancer. "If everything goes well, we could get clinical trials up and running in three to five years' time," she says. "We have had some very positive initial discussions with medical doctors but, as is always the case with these things, it is difficult to tell exactly how things will play out. What we do know is that the access to HPC awarded by PRACE will help us to get the physical data we need to proceed with this."

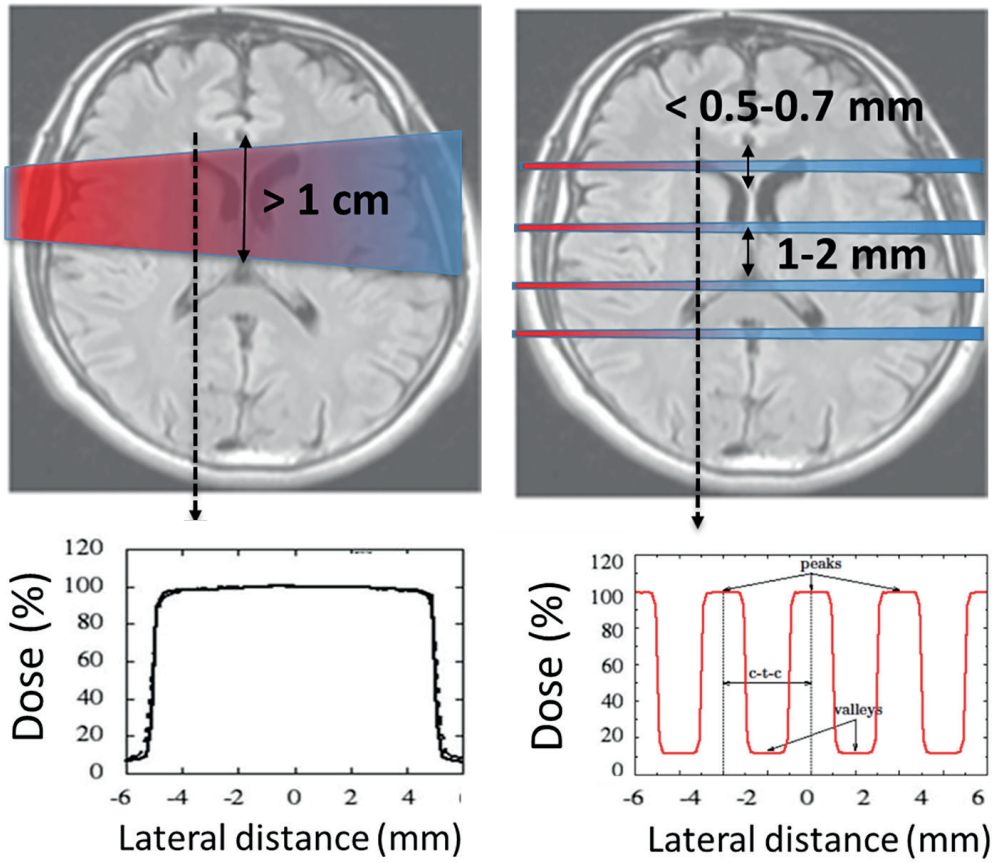

Upper row: one port irradiations with conventional RT (left) and x-rays MBRT (right). Lower row: corresponding lateral dose profiles. The dose profiles in MBRT consist of peaks and valleys in contrast to the flat profiles in standard RT

"These simulations are helping us to design and prepare our future clinical trials and allow us to calculate what a treatment plan would look like for a patient"

Another potential approach for the future could be to employ minibeams of heavy ions. Ion beams of 500 µm in width can be generated to give an extremely high peak-to-valley dose ratio, meaning that they can deliver a lot of radiation to the target tumour without damaging the surrounding tissue.

Prezado describes working with the MareNostrum computer as an "excellent" experience. "The technical support we received from the people at BSC was amazing. They were very reactive to our needs and very supportive. When you work with HPC, there are always going to be some technical difficulties, but everything that did go wrong was rapidly solved, which we were very grateful for."

Although the team have been using the MareNostrum machine at the Barcelona Supercomputing Centre up until now, they have got access to continue their work on the Joliot-Curie machine hosted by GENCI at CEA in France. It is thanks to access to HPC that Prezado and her colleagues have been able to conceive and design these new approaches in radiotherapy. The continued exploration and investigation of proton minibeam radiation therapy could soon lead to real treatments that can cure cancer more effectively, so it is hard to think of a better use of the most powerful computers in Europe.

For more information
www.researchgate.net/profile/Yolanda_Prezado

Resources awarded by PRACE
This project was awarded 16 million core hours on MareNostrum hosted by BSC, Spain

Publications

Y. Prezado, G. Jouvion, A. Patriarca, C. Nauraye, C. Guardiola, M. Juchaux, C. Lamirault, D. Labiod, L. Jourdain, C. Sebrie, R. Dendale, W. Gonzalez and F. Pouzoulet. Proton minibeam radiation therapy widens the therapeutic index for high-grade gliomas. Scientific Reports 8, article number 16479 (2018).

L. De Marzi, A. Patriarca, C. Nauraye, R. Dendale, E. Hierso, C. Guardiola and Y. Prezado. Implementation of planar proton minibeam radiation therapy using a pencil beam scanning system: A proof of concept study. Med. Phys. 2018

T. Schneider, A. Patriarca and Y. Prezado. Improving the dose distributions in minibeam radiation therapy: Helium versus proton. Med. Phys. 46, 3640-3648 (2019).

Modelling the mechanisms of complex nanocatalysts

Ziegler-Natta catalysts are important for industry, but determining exactly how they work is difficult due to their complex nature which involves a number of different active compounds on nano-sized structures. **Dr Maddalena D'Amore** of the University of Turin has been using Density Functional Theory (DFT) to try to find out more about these types of systems.

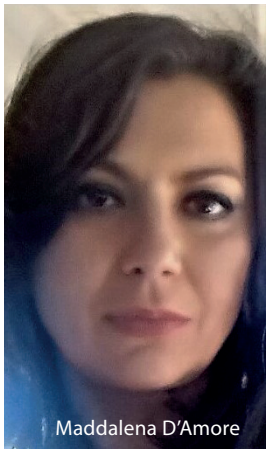
Maddalena D'Amore

Heterogeneous catalysts have a complex nature that makes it difficult to determine atomic-level knowledge of their active sites, especially in reaction conditions. Traditional experimental methods are only able to provide a partial picture of how these catalysts function.

Now, due to a combination of developments in DFT and the availability of powerful computational facilities, it is starting to become possible to design efficient heterogeneous catalysts from first principles. The NANOCAT-ZN project, led by Dr Maddalena D'Amore at the Department of Chemistry of the University of Turin, is using an allocation on the SuperMUC machine hosted by GCS at LRZ, Germany, to model the mechanism of reactions and active sites in Ziegler-Natta nanocatalysts.

Ziegler-Natta catalysts are an important class of mixtures of chemical compounds that were developed in the 1950s by German chemist Karl Ziegler, and are well known for their ability to polymerise "olefins" (hydrocarbons containing a double carbon–carbon bond) into polymers of high molecular weights and highly ordered structures. These complex nanometre-sized materials, composed of $MgCl_2$, $TiCl_4$, an aluminium alkyl and organic Lewis bases, are computationally demanding systems to investigate, as D' Amore explains. "To understand how Ziegler-Natta catalysts work, you must take into account a number of things. Firstly, there are multiple active sites on different $MgCl_2$ surfaces and on defective positions originated by both disorder and mostly nanostructuration phenomena. Secondly, the chemical compounds of catalytic interest – involving $TiCl_4$, aluminium alkyls and very bulky aromatic donors – have a very low degree of coverage. This is why it is necessary to model very large systems that are very far from perfection to achieve a reliable description of catalysts."

"Finally, after the reduction of titanium species, unrestricted DFT calculations are necessary. For these reasons, an accurate ab initio prediction of the thermodynamics and kinetics of these systems requires highly expensive calculations and huge computational resources."

The team led by D'Amore used an efficient highly parallelised code called MPPCRYSTAL to carry out ab initio simulations, correlating the structure of the Ziegler-Natta catalysts at the nanoscale with their activity and selectivity to understand the role of each component in the reaction mixture. The code was specifically tailored to achieve an excellent scalability in terms of both speed-up and memory usage, allowing the researchers to treat very large unit cell systems with large memory requirements and few or no symmetry operators. The project has been truly multidisciplinary, representing a collaboration between theoreticians, spectroscopists and industrial partners. Their joint aim was to find out more about how the structure of the active sites in these catalysts contribute to their function and the role of "each player" and synthetic procedure in the game. "As well as this, we

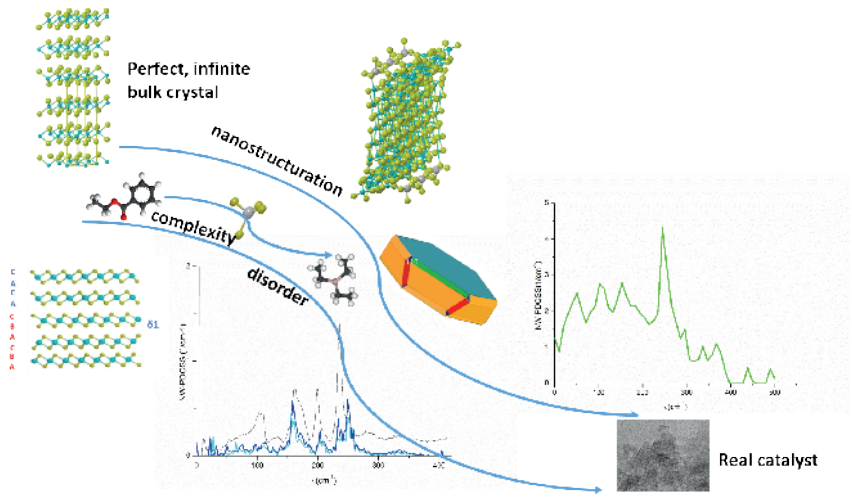

Figure 1: Scheme of computational approach adopted to model "real" $MgCl_2$ supported Ziegler-Natta catalysts for olefin polymerisation. Simulations of Inelastic Neutron Scattering (INS) and G_s Wulff polyhedron are also reported

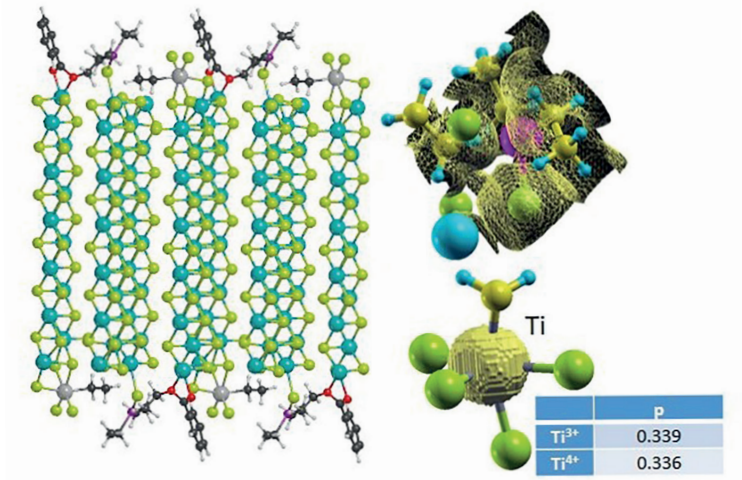

Figure 2: Possible product of transalkylation and active site in 'activated catalyst': the binding of aluminium alkyl co-catalyst (Al2R6) to catalytic (110) surfaces with adsorbed $TiCl_4EB$ adduct (EB=Ethyl Benzoate) to form quaternary $MgCl2/TixCl_{4x}/EB/Al_2R_6$ adducts lead to a minimum with the donor grafted to the surface and a trans-alkylation towards Ti center (B3LYP-D*/customised TZVP level of computation). Some results of ELF and MPD analyses are also sketched (preliminary MPD estimates of probability of different oxidation states for Ti active species are also reported).

wanted to both improve the industrial performance of these catalysts and the performance of our code when running on thousands of processors," says D' Amore.

The simulations of the catalysts used models that involved hundreds of atoms – a huge amount in the world of accurate DFT. Due to the large size of the investigated systems and the high level of accuracy of calculations, that employs both hybrid exchange-correlation functionals and extended basis sets, the feasibility of simulations was possible only thanks to the HPC resources granted by PRACE.

The realistic description of these catalysts required advanced DFT methods accounting for dispersion and employing hybrid functionals or highly parameterised m-GGA exchange-correlation functionals to describe metal catalytic sites. Thanks to the use of thousands of cores simultaneously, it was possible to simulate nano-crystals of up to one thousand atoms, analysing plausible reactivity indicators such as electron density transfers in order to evaluate the possible occurrence of specific reaction mechanisms, on the basis of ELF and MPD (Maximum Probability Domains) theories.

The investigation modelled low-dimension systems: surfaces, nano-rods and energetically stable primary particles, adopting (when periodicity is still kept) a supercell approach to build huge unit cells to simulate adducts with poor loadings (Figure 1). Thanks to vibrational analysis, the entropic contribution to surface formation energy was estimated in reaction conditions. "To our surprise, the (110) surfaces, identified by recent theoretical models as the place for stereo-selective active sites of catalytic processes involving Ti_xCl_y species, arose on our simulated Wulff polyhedron," says D'Amore. "This is the first time that the active catalytic surface has emerged in a quantum mechanical simulation in reaction conditions."

These results of the project have helped to contribute to the understanding of the relation between the structure, reactivity and selectivity of the catalysts' active sites towards olefins. Starting from the identified precatalytic sites and assuming that the adjacent donors affect the electron density of metal active sites, D' Amore and her team plan to continue investigating the mechanism of titanium reduction and olefin insertion starting from identified quaternary adducts (i.e. Figure 2). In this way, it could well be possible to work out how to improve the productivity and stereospecificity of these catalysts once a better understanding of the process by which they function is achieved, making the "trial and error" method used by industry obsolete.

At the same time, this will help to reduce the health and environmental impacts that Ziegler-Natta catalysts can have. The team used poorly validated theoretical methodologies as MPD to surface catalysis; (Figure 2) the application of these methodologies in other fields of heterogeneous catalysis would be a valuable and long-term contribution of the project out of the specific scientific field of polymer catalysis.

For more information

www.crystal.unito.it

Resources awarded by PRACE

This project was awarded 30 million core hours on SuperMUC hosted by GCS at LRZ, Germany

Publications

"A Periodic Hybrid DFT Approach (Including Dispersion) to MgCl2-Supported Ziegler-Natta Catalysts, 1. TiCl4 Adsorption on MgCl2 Crystal Surfaces", Journal of Catalysis, 2012.

" Surface Investigation And Nanoscale Morphological Analysis Of Coordinatively Unsaturated Surfaces In Structurally Disordered MgCl2 And MgCl2/TiCl4 Ziegler-Natta Catalysts", ACS Catalysis, 2016, 6, 5786-5796.

"The Bond Analysis Techniques (ELF and Maximum Probability Domains) Application to a Family of Models Relevant to Bio-Inorganic Chemistry", in "Applications of Density Functional Theory to Biological and Bioinorganic Chemistry". Series: Structure and Bonding, Vol. 150.

"Periodic and High-Temperature Disordered Conformation of Polytetrafluoroethylene Chains: an Ab-Initio Modelling", Journal of American Chemical Society 2006,128(4); 1099-1108.

The Partnership for Advanced Computing in Europe

The mission of PRACE is to enable high-impact scientific discovery and engineering research and development across all disciplines to enhance European competitiveness for the benefit of society. The PRACE Research Infrastructure seeks to realise this mission by offering world class computing and data management resources and services through a peer review process. PRACE also seeks to strengthen the European users of HPC in industry through various initiatives.

HPC ACCESS

CALL ANNOUNCEMENTS

Stay tuned! **www.prace-ri.eu/call-announcements**

TRAINING

PRACE TRAINING PROGRAMMES

These programmes offer free-of-charge, 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/prace-supports-industry**

SMEs

PRACE SHAPE SME HPC Adoption Programme in Europe

SHAPE offers European SMEs the opportunity free-of-charge to use HPC for business innovation and increase competitiveness. Stay tuned! **www.prace-ri.eu/shape**

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 823767. For more information, see www.prace-ri.eu.

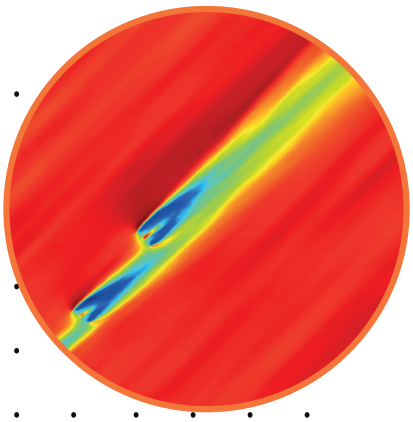

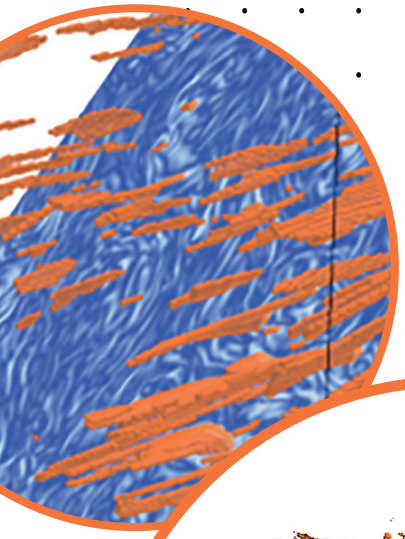

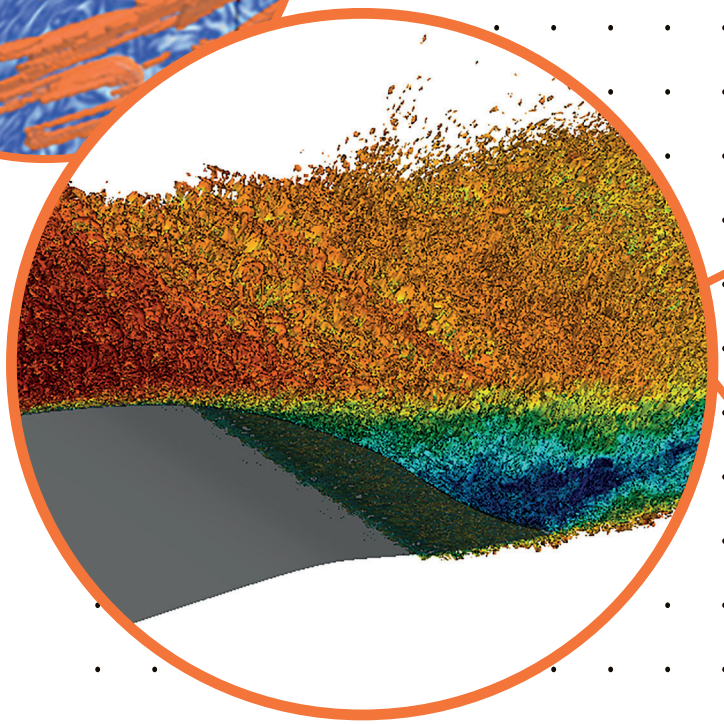

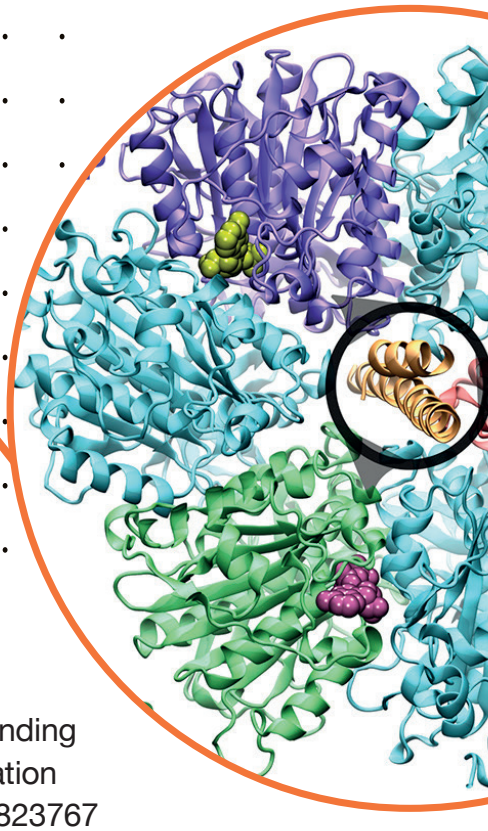

The Implementation Phase of PRACE receives funding from the EU's Horizon 2020 Research and Innovation Programme (2014-2020) under grant agreement 823767