PDC Center for High Performance Computing

Scientific Discovery through HPC

An assortment of research projects from PDC

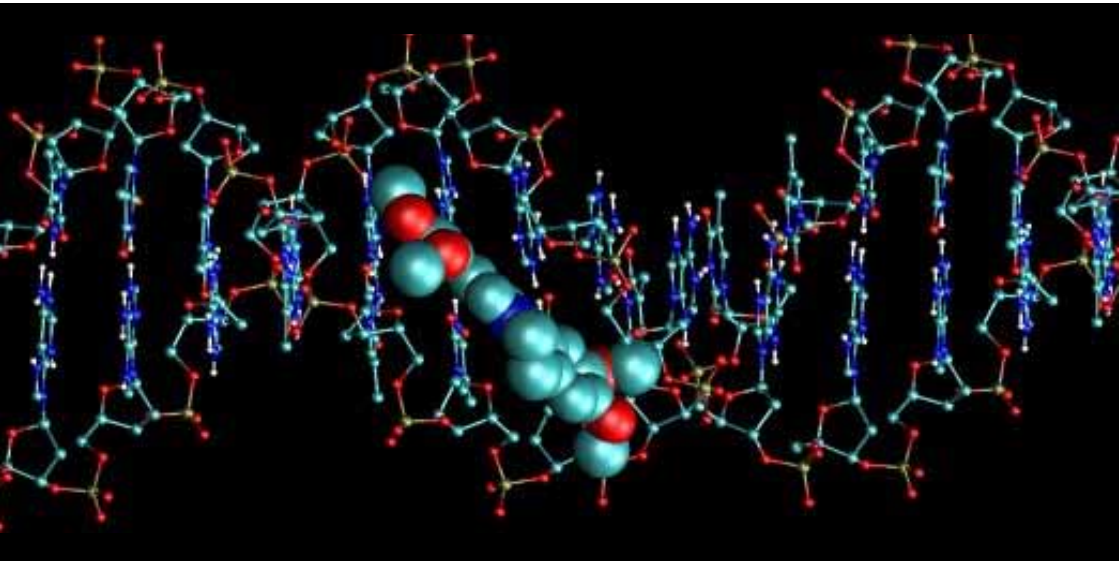

Contents

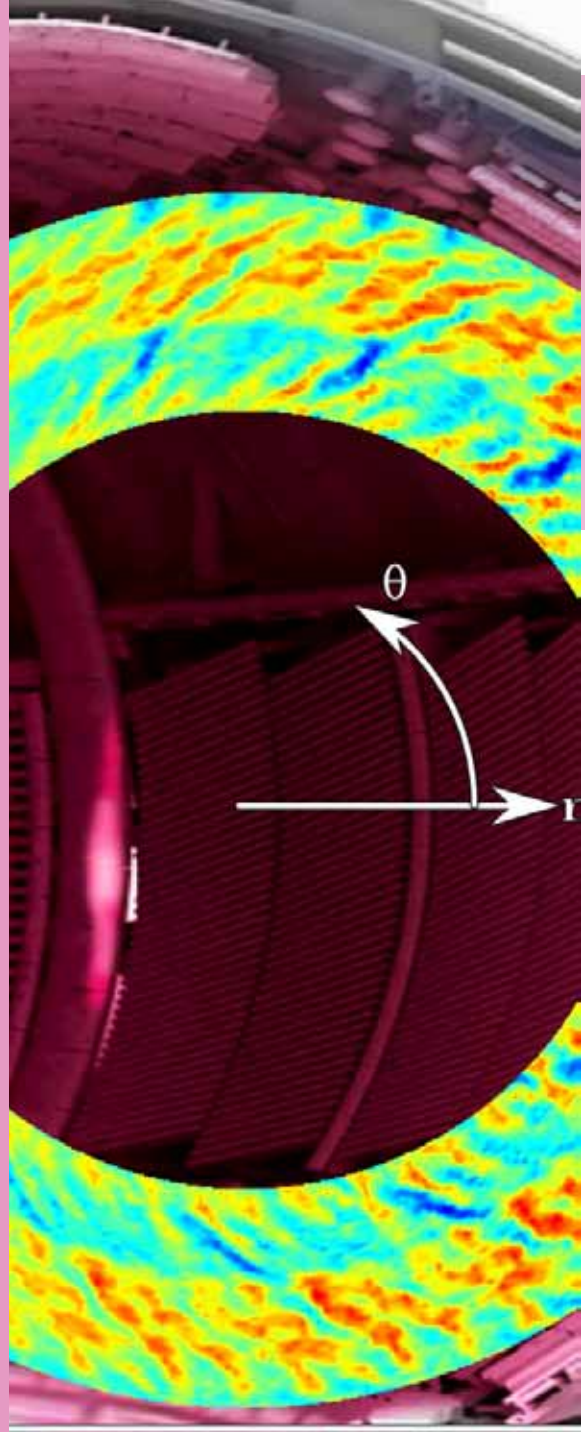

Foreword

Large-scale computational resources, commonly known as “supercomputers”, have become an indispensable tool for many scientific disciplines. The PDC Center for High Performance Computing at the KTH Royal Institute of Technology in Stockholm has been providing these kinds of resources to Swedish academia and industry for twenty-five years with a constant drive to stay at the forefront of technology. PDC is a member of the Swedish National Infrastructure for Computing (SNIC), which coordinates the national supercomputing resources in Sweden, and of the Swedish e-Science Research Centre (SeRC). Through SNIC, PDC is also a member of the Partnership for Advanced Computing in Europe (PRACE), which enables Swedish researchers to access the even larger European supercomputers.

Even though we are generally not aware of it, the impact of supercomputers on day-to-day life is tremendous! Supercomputers provide essential contributions to many of the challenges that our global society faces, such as helping in the quest for new and greener forms of energy, discovering new and better drugs and treatments for a whole gamut of problems from the purely physical through to neuropsychiatric disorders, developing new and more effective materials for the production of items such as electronic circuits and solar cells, as well as monitoring and understanding climate change. And these are just a few of the many areas where supercomputers are being used! (Did you know that yacht-building companies use supercomputers to design hulls for new boats nowadays?) Here we present an assortment of the research activities that are supported by SNIC’s and PDC’s supercomputing resources. Although this collection is far from being comprehensive, we hope it gives a good overview of the kind of research supported by PDC and its societal impact.

Erwin Laure, Director, PDC

Improving wind farm design

Stefan Ivanell

Dan Henningson

Philipp Schlatter

Nowadays responsible industrialised nations are placing more and more focus on developing clean sources of power that neither destroy our natural resources nor leave a legacy of poisonous waste for our descendants. Over the last two decades, Sweden has seen a steady increase in the use of wind power, with yet more wind farms under construction, and, during the last decade, a drastic increase in the number of planned wind farm sites has been proposed worldwide.

For wind farms to be as effective as possible, they need good designs for the turbines that "harvest" the wind. If one considers the enormous size of typical wind turbines, it becomes obvious why it is economically impractical to test a range of different designs by building them all physically. Consequently we rely on mathematical modelling and simulations to predict suitable designs for wind turbines and for the layout of turbines in a wind farm. As you can imagine, this involves both complex modelling of the airflow in and around individual turbines and also between multiple turbines. However, many issues related to the analysis of the airflow in these situations still need to be understood in much more detail.

For example, when air flows past an object, a swirling flow of air (known as a vortex) can be created. With wind turbines, vortices arise as a result of the air flowing past the roots and tips of the turbine blades. The tip vortices or spirals then continue moving in the direction of the prevailing wind, which means that, in a wind farm with multiple turbines, the rotating air mass of a tip vortex can encounter one or more turbines downwind of the initial turbine. Thus the flow of air arriving at the downstream turbines is likely to be significantly different to the direct flow of wind that the initial turbine experiences. This can result in damage to the blades of the downwind turbines and can also reduce the amount of power that is generated by the turbines.

The swirls above show a tip vortex formed behind the wing of a plane. The vortex has been made visible by having coloured smoke rising from the ground as the plane passes above it, so the smoke shows the air movement caused by the plane.

Source: NASA Langley Research Center

Not surprisingly, there are many factors to be taken into account when trying to model these processes, including the detailed flow physics associated with the instability of the tip vortices that are produced or shed from the turbine blades, the interactions of these tip vortices with downstream turbines (which affect the overall amount of power that is produced), and the influence of atmospheric boundary-layer density stratification and associated turbulence properties. Understanding the dynamics of tip vortices and their interactions when multiple turbines are present is absolutely crucial to improving our current understanding of wind turbines and the associated modelling.

All of these phenomena need to be understood and subsequently incorporated into the mathematical models used to design wind turbines and wind farms. This vital area has given rise to a large amount of research dealing with the modelling of wind farms. In Sweden, research on wind power is being coordinated by STandUp for Wind (which is a collaboration between researchers at

The researchers are working on optimising the numerical methods that are used to simulate individual wind turbines (and also complete wind farms) using advanced methods. Their implementation in the highly accurate and parallel code Nek5000 (potentially running on millions of processes using MPI — a system for passing messages on parallel computers) reproduces the well-known tip and root vortices. Due to the rotation of the turbine blades, these vortices form a helical pattern downwind of the turbine as shown below. The figure shows the spiral vortices from a rotating three-bladed turbine, with the colour indicating the distance downwind from the blade (with blue being closest to the blade). The tip vortices are well-structured close to the blades, but start to interact and break down into turbulence further downwind (not shown here).

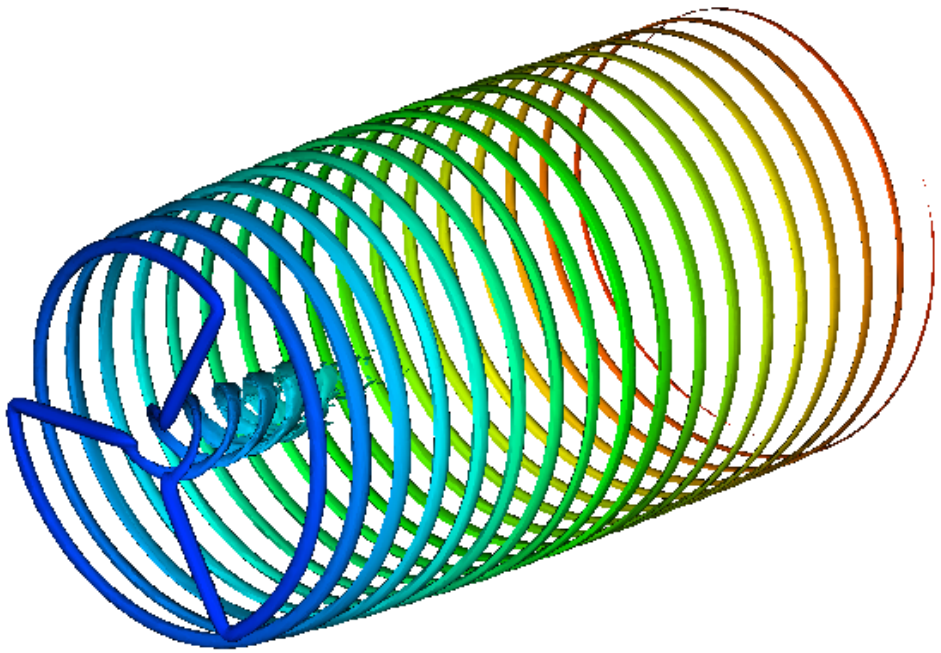

the KTH Royal Institute of Technology and Uppsala University that is led by Stefan Ivanell), and through a relatively new programme which is funded by the Swedish Research Council (VR) and led by Dan Henningson from KTH. Close collaboration between research groups within a Nordic consortium from Sweden (KTH and the Gotland campus of Uppsala University) and Denmark (Technical University of Denmark) has led to interesting studies on the topic of tip vortices. For instance, a relationship between ambient flow properties and the length of the tip spiral vortices has been established. The figure above illustrates one of the outcomes of a new implementation of a wind turbine, showing the tip and root vortices, and how they slowly decay further downstream. Previous work by the group from KTH has established new knowledge about how the tip spirals interact under different ambient conditions, and proposed a model of how far downstream these vortices persist.

Having established accurate ways of simulating individual turbines, the next step, both within the Nordic consortium that was mentioned earlier and in the recent VR-funded framework programme, is to consider what happens in a wind farm where we have multiple turbines. Wind turbines placed downstream of other wind turbines have a reduced power output due to the greater turbulence in the approaching airflow. To optimise the layout of our wind farms, we must develop a better understanding of the airflow surrounding a single wind turbine, and then extend that to understand what happens in a wind farm where the air flows around and between multiple turbines. This will enable us to realistically model complete wind farms and thus give better estimates of their total potential power production.

Within these projects, the airflow for a wind turbine is analysed using higher-order numerical methods implemented in two sets of computer code, Nek5000 and SIMSON, which run on massively parallel computer systems such as those at PDC. Often simulations like this divide the relevant area (such as the blade of the wind turbine) into a mesh of points, and then performing calculations to work out what happens at each of the points. However using this approach for the wind turbine blades would be too costly in computational terms when it comes to achieving the required level of detail. Instead the researchers are modelling the wind turbines as body forces (that is, forces that act throughout the volume of a body – like gravity – rather than just at the point of contact – like friction) using a particular technique known as the actuator line method (ACL). With this method, the body forces mimic the action of the blade on the incoming flow of air. Thus the bulk of the supercomputer computational resources can be allocated to resolving what happens when the air has flowed past a turbine, that is, in the turbine wake flow. Using the geometric flexibility of the Nek5000 code, single turbines (in cylindrical grids) as well as turbines close to surfaces and entire wind farms can be implemented. Stefan Ivanell and Philipp Schlatter, along with other researchers from KTH, intend to use the Nek5000 code to perform detailed comparisons of experimental setups to improve the engineering predictions for large wind farms.

Modelling proteins and protein-drug interactions

Cancer? Parkinson's disease? Diabetes? Heart disease? None of us want to suffer from diseases like these. Researchers are working hard to develop drugs to treat such diseases effectively. The study of G-protein-coupled-receptors (GPCRs) is one area that offers hope for developing new drugs to treat various diseases.

GPCRs are a type of protein, known as membrane proteins, that occur all over our bodies - they exist in the membranes of our cells and change the way that each cell responds to stimuli from outside the cell. GPCRs are closely related to the occurrence, development, and treatment of many illnesses, such as immune disorders, infectious diseases, neurological conditions and psychiatric disorders, as well as diabetes, heart disease and cancer. Drugs that alter the activity of GPCRs in the body are already being used in the treatment of various diseases. (In the biomedical world they say that a drug that affects a GPCR in this way is "targeting" the GPCR.) In fact, the world's top twenty best-selling drugs are ones that target GPCRs, and those twenty drugs have an annual total sales value of around 200 billion US dollars!

Yaoquan Tu and his doctoral students at the KTH Department of Theoretical Chemistry have been using the supercomputers at PDC to model protein structures, as well as interactions between various proteins, and between proteins and drug molecules. Modelling protein structures is valuable as it helps us to predict protein structures that can

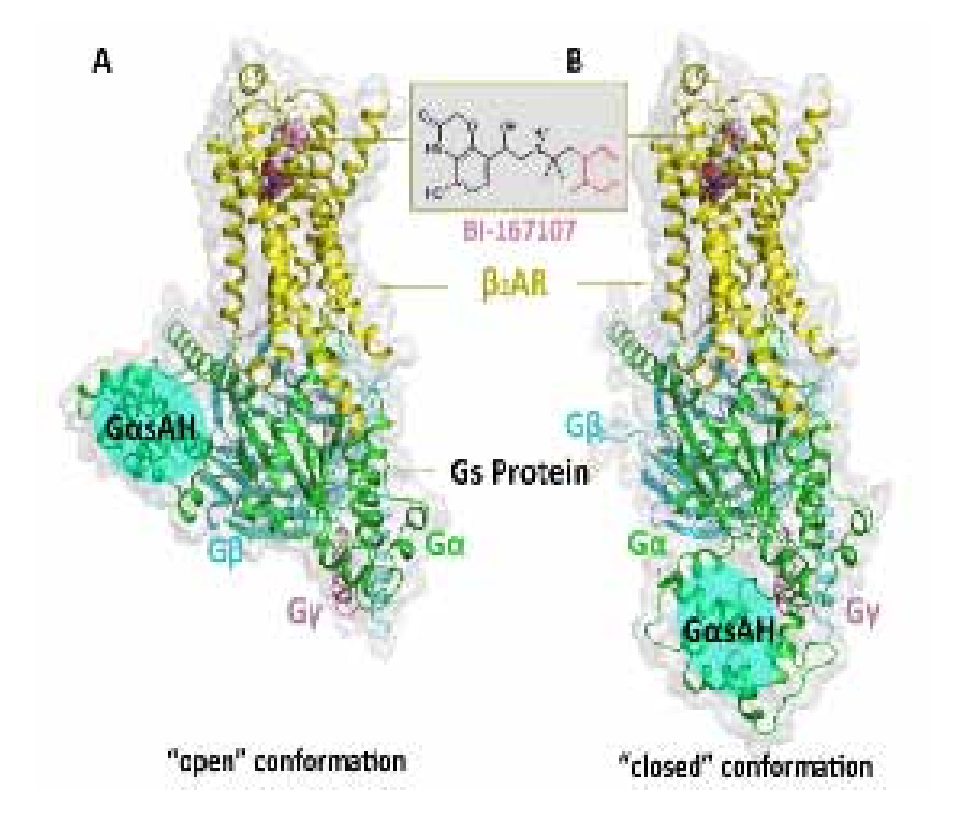

In the "open" conformation above (A), the Gs protein prefers to bind to the B2AR, while in the "closed" conformation (B) uncoupling of the Gs protein from the B2AR occurs.

Microsecond molecular dynamics simulations (below and opposite) provide insight into the uncoupling of the Gs protein from the B2AR.

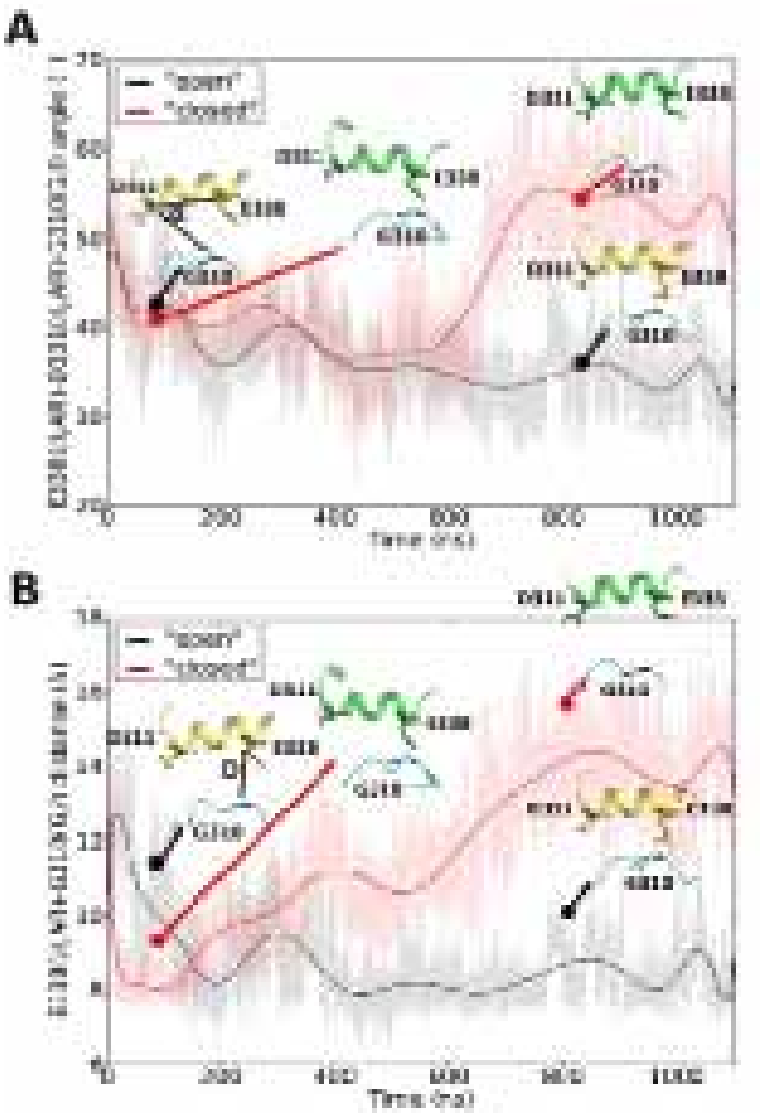

be used for designing drugs, while the results from studies of protein-drug and protein-protein interactions are of great benefit in helping us to understand many biological processes in the body, and thus the

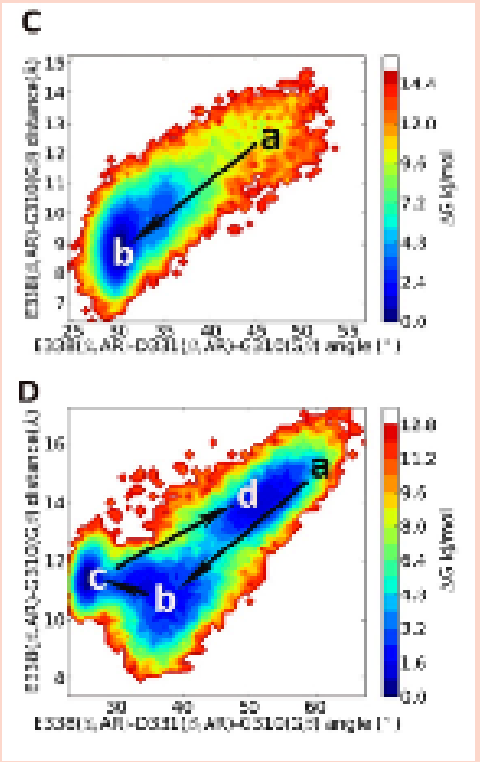

resulting information can be used to develop new drugs to address serious illnesses.

One protein that Yaoquan's group has been investigating is the beta-2 adrenergen receptor (also known as β2AR or B2AR). The B2AR protein is a GPCR that reacts with adrenaline to affect muscles or organs. Drugs that target B2AR can be used in the treatment of respiratory diseases, such as asthma. Recent studies have also shown that long-acting agonists targeting B2AR can have therapeutic potential in the treatment of neurodegenerative disorders, such as Parkinson's disease. (Agonists are chemicals that bind to receptors and make them respond in a particular way.)

The diagrams illustrate a hypothesis that Yaoquan's group is testing. They show a possible mechanism for the way in which a particular protein, known as the Gs protein, decouples (or disconnects) from the B2AR. Finding out how this process works is important for the design of drugs targeting GPCRs, as it is related to how signals propagate through the body at a cellular level, which in turn influences the effect that a drug will have on the human body.

Although this is, in a sense, one very small part of how the whole complex human body works, it is through the dedicated efforts of many biochemistry researchers around the world that we are gradually building up a better and better understanding of how our bodies work, and how various substances cause the effects that they have on us. Thus this kind of research is gradually leading us towards being able to treat illnesses in efficient ways with compounds that specifically address the problem, and influence as little else as possible. While we may not yet have complete cures for diseases like cancer, each little building block in our understanding that comes from research like this is bringing us closer and closer to such cures.

Clean power through fusion


Hans Nordman


Pär Strand

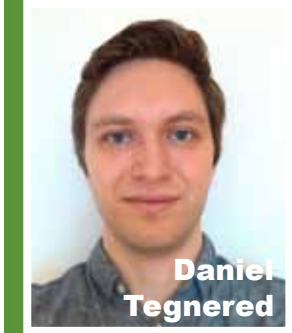
Daniel Tegnered

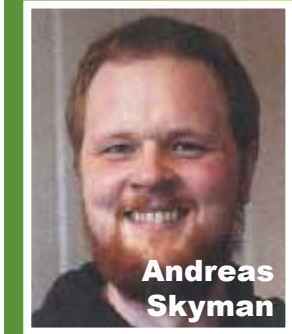
Andreas Skyman

Hans Nordman and Pär Strand are leading researchers at the Plasma Physics and Fusion Energy Group in the Department of Earth and Space Sciences at Chalmers University of Technology in Gothenburg where they are working on ways to provide clean energy through fusion. Their work involves modelling fusion processes in devices known as tokamaks, which are used to control the fusion reaction and to contain the plasma that provides the fuel for the reaction. Plasma is the fourth state of matter and occurs where temperatures are so high that the atoms are stripped of their electrons. The particles in a plasma are therefore charged and will gyrate around and follow magnetic field lines. This makes it possible to confine the hot plasma, and prevent it from touching the walls of the tokamak, using magnetic fields. In tokamaks, the magnetic field resembles a doughnut, often with a D-shaped cross-section as shown in the model of the ITER tokamak on the next page.

Why fusion?

Daniel Tegnered (a postgraduate student in the group) and Andreas Skyman (a recent postgraduate student in the group) explain that the advantages of fusion are that the reactions release vastly higher amounts of energy per nucleon than fission reactions, and that fuel suitable for fusion is available globally anywhere on Earth. Fusion thus has an enormous potential as a clean, safe, and environmentally friendly energy source and, due to the large amounts of energy that it is releases, fusion is one of the few energy sources that could potentially replace all the fossil-fuel-based energy sources. At present, researchers around the world are working on developing safe and effective fusion reactors, and this requires performing experiments to determine what

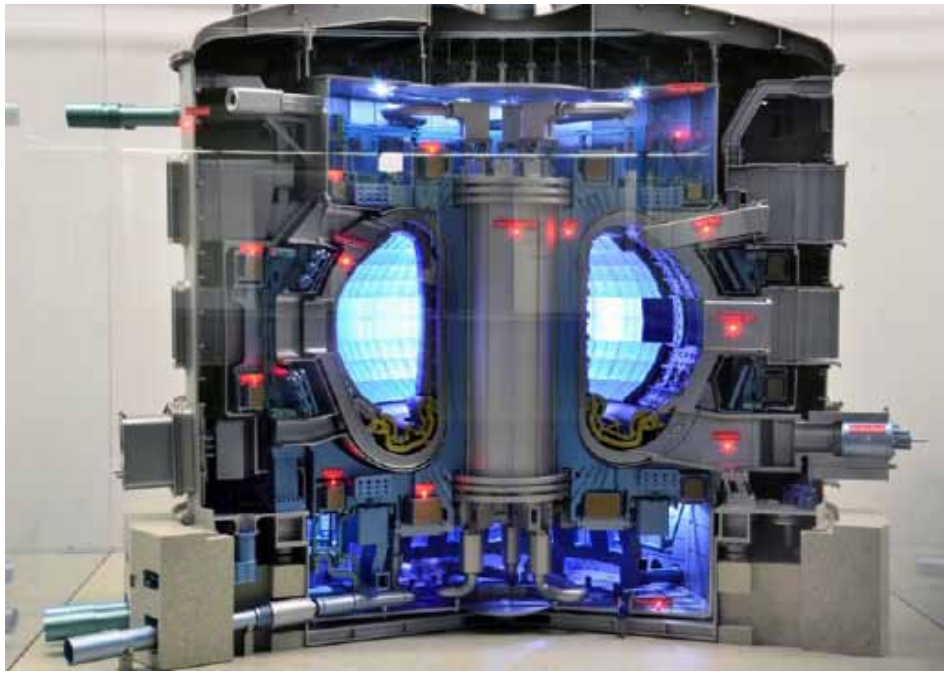

This model of ITER, the world's largest tokamak, shows the doughnut-shaped vacuum vessel where the plasma is confined. The fuel — which is a mixture of deuterium and tritium, two isotopes of hydrogen — is heated to temperatures in excess of 150 million degrees Celsius, forming a hot plasma. Strong magnetic fields are used to keep the plasma away from the walls; these are produced by superconducting coils surrounding the vessel, and by an electrical current driven through the plasma.

Source: ITER Organization

type of reactor design and fuel are most suitable.

Having to build reactors of different shapes and using different materials, and then test them all in various scenarios would be prohibitively expensive, and hence the research group is concentrating on modelling fusion reactions mathematically. Being able to model and predict the behaviour of future fusion experiments, like the experimental reactor ITER that is currently being built in the south of France, is important both in order to understand the physics of fusion reactions and to dictate what experimental setups are actually considered. The economic and time constraints mean it is likely that all the fusion experiments that are performed physically will be preceded by extensive predictive modelling.

The simulations that the Chalmers group have performed on PDC's supercomputer resources, amongst others, have resulted in a better understanding of the scaling of transport properties (such as the way that heat or particles flux inwards or outwards within the plasma) with reactor-relevant parameters. In particular, the shape of the impurity density profile in the fusion plasma has a strong impact on the performance of fusion reactors and these effects have been studied in detail.

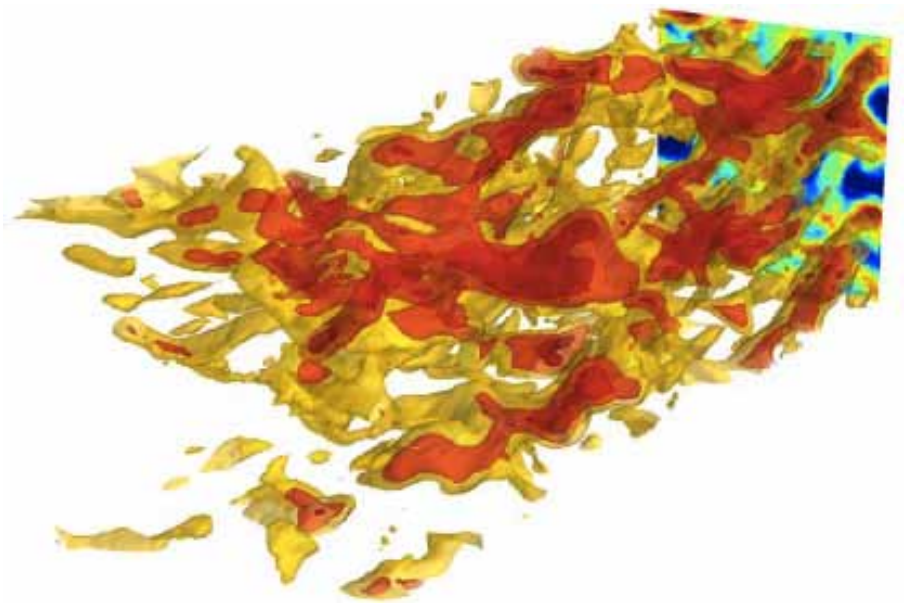

This diagram shows the change in the electrostatic potential in a cross-section of a flux-tube over time. A flux tube is a doughnut-shaped area in space bounded by a flux surface, which is a surface where the magnetic field is parallel at all points. A tokamak is essentially a plasma-filled flux tube. If there were no turbulence or collisions in the plasma, the charged particles would follow the magnetic field lines and stay on these surfaces. The square area in the background is a part of a cross-section of a flux tube.

How fusion works

Fusion relies fundamentally on the same physical principle as fission: transmuting elements in such a way that the reactants (that is, the elements that are the fuel) lose mass per nucleon, which yields a net excess of energy. In fusion processes, lighter and normally stable elements are fused to form heavier elements, for example, fusing hydrogen to produce helium and energy.

Fusion requires exotic conditions: intense pressure or fantastically high temperatures. In the Sun, pressures far beyond those that have been achieved on Earth facilitate the fusion of protons, but the most favourable route to fusion for power production relies on the latter alternative, namely extreme temperatures. In the largest currently operational fusion experiment, the Joint European Torus (JET), temperatures of 100,000,000 K are routinely achieved, which is what is necessary to efficiently fuse the hydrogen isotopes deuterium (D) and tritium (T). JET and ITER are both examples of tokamaks, which are the most widely researched class of device for controlling fusion-plasma using magnetic confinement.

Turbulence and designing an efficient tokamak

To design a tokamak that is suitable for producing energy via fusion, we have to consider how the fusion plasma will behave inside the tokamak. To this end, we need to look at the dynamics of fusion plasmas, that is, how the plasma will move and behave. The dynamics are generally divided into two categories: stability and transport. Here "stability" refers to global-scale dynamics, affecting the whole bulk of the fuel plasma, whereas "transport" encompasses smaller-scale phenomena occurring within the plasma. The stability of fusion plasmas can often be understood and studied using Magnetohydrodynamics (MHD), whereas studying the transport phenomena requires more advanced physical models (which are closer to the first principles of plasma dynamics as described by the kinetic Boltzmann–Vlasov equation). When it comes to determining feasible operating scenarios for future fusion power plants, both stability and transport are important, although modelling the transport phenomena is the harder problem.

To produce cheaper, optimised fusion devices with a good level of power output, we need to control the transport. This is because, to achieve a power breakeven, where we get as much (or more) energy from the fusion reactions as we

put into the plasma, the so-called fusion triple product (which is the product of density, temperature and confinement time – a measure of how quickly the energy would be lost from the plasma if power was not supplied) must exceed a critical value. However, the temperature and density are constrained respectively by the need to maximise the fusion cross-section and to maintain global stability. Consequently the only way to improve the energy return from the tokamak is to increase the energy confinement time.

Unfortunately turbulent transport, which leads to a shorter confinement time, is unavoidable as small-scale instabilities in the plasma are caused by the necessarily large variations in temperature and pressure across the plasma, and these small-scale perturbations turn turbulent. Being able to model the properties of these processes is crucial for improving our theoretical understanding of the way fusion plasmas behave and thus predicting suitable tokamak configurations. The mathematical computations for this kind of modelling work are extremely time consuming and complex, requiring hundreds of hours on the world's fastest supercomputers. Researchers are seeking ways to make the computations for modelling this kind of turbulence less expensive; the gyrokinetic method is one such approach – it uses the fact that the gyrating motion of the charged particles around the magnetic field lines is small and fast compared to other phenomena to reduce the overall complexity of the problem.

Fusion plasmas contain different impurities, in addition to the isotopes of hydrogen that form the fusion fuel, and such impurities are also subject to this type of turbulent transport. These impurities can come from the wall of the tokamak reactor chamber where materials such as tungsten and beryllium are used. These additional impurities can release energy in the form of radiation and also effectively dilute the plasma, thus lowering the power output of the reactor. There are also other sources of impurities such as gases that are injected for control purposes and the helium ash that results from the fusion reaction.

To reduce the effects of these types of impurities on the final power output, it is important to determine whether the impurities and fuel ions move outwards towards the wall or inwards towards the centre of the plasma. The GENE program code (http://gene.rzg.mpg.de) has been used to investigate various experimental parameters relevant to the JET tokamak, such as how the shape of the cross-section of the plasma — the so-called magnetic geometry — affects the impurity transport. The dependence

The image below shows a contour plot of the fluctuating part of the electrostatic potential superimposed on a cross-sectional image of the JET tokamak. The electrostatic potential is related to the potential energy of the charged particles in the electric field. The fluctuations in this electrostatic potential give a measure of how strong the turbulence is.

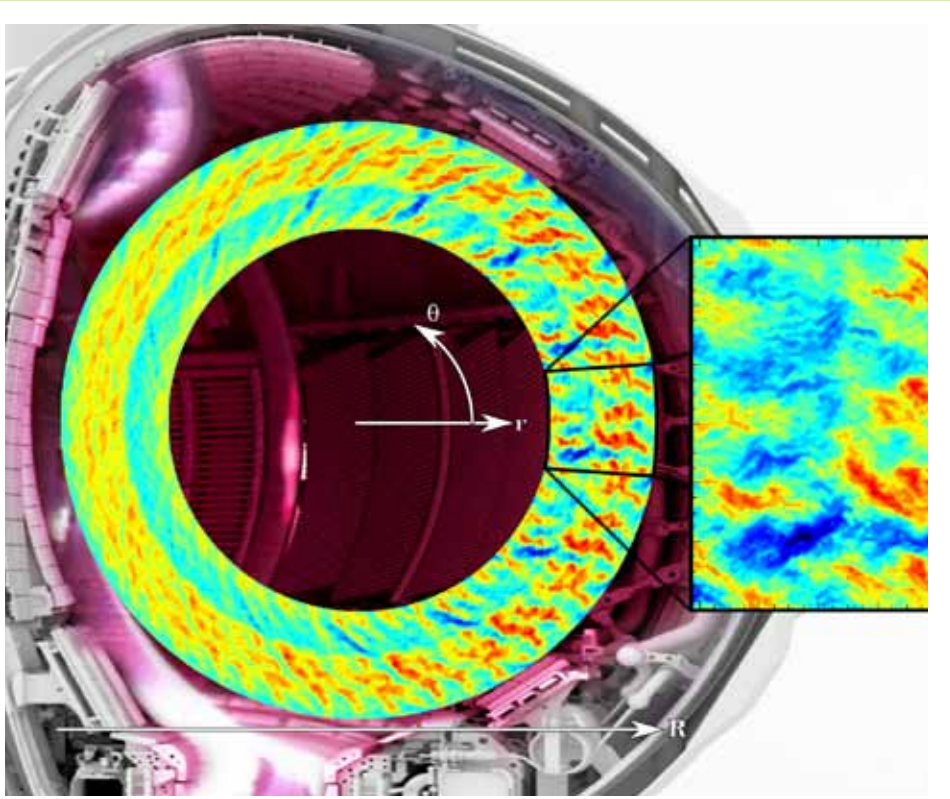

of the scaling of relevant parameters on the main ion transport (that is the transport of the most common types of ions in the plasma, which are usually deuterium in most current experiments) and how this, in turn, affects the impurity transport have also been explored. Currently the researchers are investigating the differences in the transport of heat and particles resulting from changing the materials in the reactor walls at JET from carbon-based to metal-based (beryllium and tungsten). These simulations have been run on PDC's Lindgren supercomputer, which is currently the only Swedish HPC resource capable of meeting the computational demands of the non-linear GENE simulations required, and the group is looking forward to being able to extend their simulations with Beskow's increased computing capacity.

Tracking climate change

With scientists currently arguing over whether recent changes in the Antarctic ice presage a rise of nearly five metres in the sea level over the next two centuries (which would submerge large parts of the inhabited world), we are all well aware of the importance of predicting changes in the climate in time to take remedial steps.

Researchers Gunilla Svensson, Rodrigo Caballero and Laurent Brodeau, from the Department of Meteorology at Stockholm University (MISU), along with their colleagues, Torben Königk and Klaus Wyser from the Swedish Meteorological and Hydrological Institute (SMHI), are working to improve our understanding of climate change and thus enable us to make better-informed decisions on future policies (for example, in relation to limiting carbon emissions). They have performed a large ensemble of historical and future climate simulations using the state-of-the-art climate model EC-Earth as part of the Coupled Model Intercomparison Project Stage 5 (CMIP5, http://cmip-pcmdi.llnl.gov/cmip5), which is being coordinated by the World Climate Research Programme. In this exercise, current-generation climate models from institutions around the world are run under a standardised protocol to permit intercomparison of the results under uniform conditions. The model output is stored and made publicly available through the Earth System Grid, a distributed data storage system. They have also been active in the scientific analysis and interpretation of the CMIP5 climate simulations. Much of their climate model simulation work has been performed on PDC's supercomputer resources.

This important climate modelling work has been taken into consideration in the Fifth Assessment Report (AR5) of the Intergovernmental Panel on Climate Change (IPCC), which is a summary of current knowledge in climate science that will be used by global decision makers. This report is a key vehicle for the transmission of climate information to society.

The work on CMIP5 climate simulations has fed into the most recent IPCC climate assessment report (shown above).

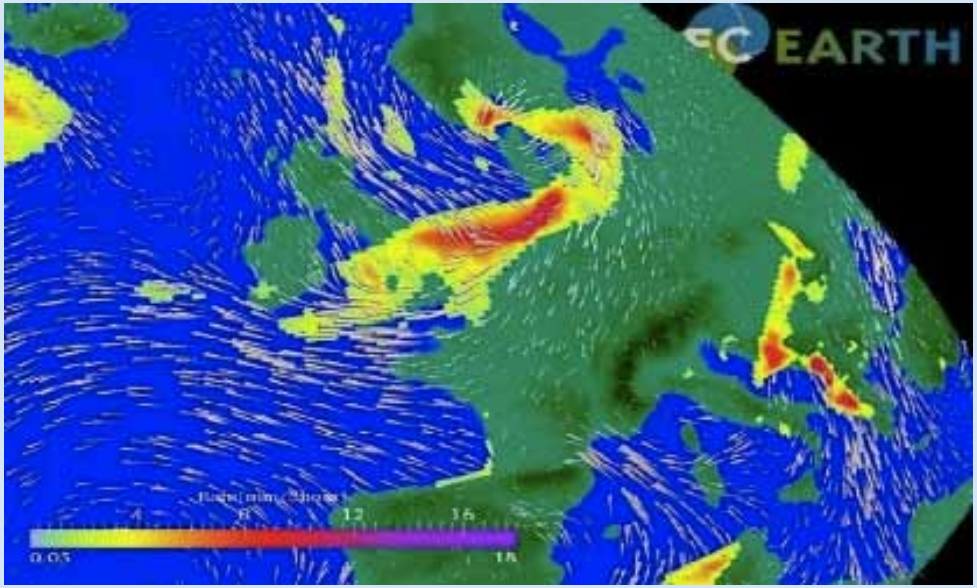

This work has involved development tuning and validation of the latest model release of the EC-Earth climate model. Source: Royal Netherlands Meteorological Institute (KNMI)

Paving the way to better health and medical care


Erik Lindahl


Mark Abraham


Magnus Andersson

Erik Lindahl is the head of the Biophysics group in the Department of Theoretical Physics at the KTH Royal Institute of Technology and is also a professor of Biophysics in the Department of Biochemistry and Biophysics, Faculty of Science, at Stockholm University. Erik's group is a large team whose research and development work spans the full gamut from fundamental methodology and software development for biophysics to the development of actual biological applications. The main people involved in the method and software development side of the research are Berk Hess (Associate professor), Mark Abraham (researcher), and Szilárd Páll (PhD student), all of whom are affiliated with the Theoretical Biophysics group within the Department of Theoretical Physics at KTH. On the application side, their software is used for large-scale simulations by a number of people, such as Magnus Andersson (researcher), who is also from Theoretical Biophysics. The whole team is closely affiliated with the Swedish e-Science Research Centre (SeRC) and the Science for Life Laboratory (SciLifeLab), both of which come under the Strategic research areas financed by the Swedish Research Council (VR) in the Stockholm region.


Erik Lindahl, Szilárd Páll and Berk Hess

The team's overall goal in advancing the state of the art in biomolecular modelling and simulation is to use these methods as a computational microscope for complex biological molecules. They are primarily working with membrane proteins – the small molecules that are embedded in our cellular walls and that act as the doors and windows of each cell. Every nerve signal and heartbeat in your body is due to these membrane proteins operating as microscopic machines and altering their conformations as a result of changes in the electric potential or ion concentration across the cell membrane. This behaviour is almost impossible to study experimentally since it only occurs transiently at the scale of microseconds. However, by using molecular simulations, it is

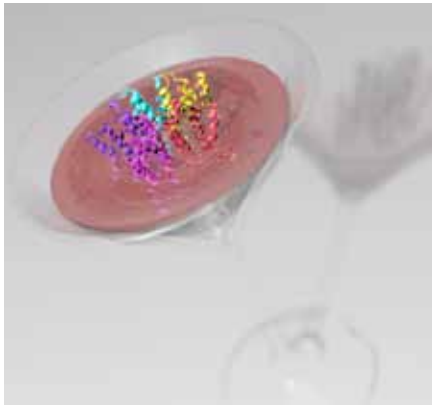

Ligand-gated ion channels such as the glycine receptor (shown in a martini glass above) are responsible for our nerve signals, but also highly sensitive to other molecules. This is how the function of our brain is influenced by alcohol and other drugs.

To simulate how an entire ion channel works it is necessary to include both a cell membrane and lots of water and ions, which can amount to hundreds of thousands of atoms - this type of research is simply not possible without supercomputers such as Beskow.

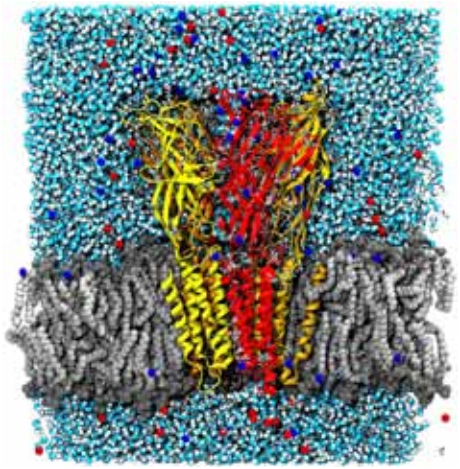

possible to literally track the atomic-level motions to understand both how the proteins work and how we can change their behaviour.

While that all sounds very technical and complicated, one of the amazing things with this research is the way these techniques have evolved – in just a decade – from relatively inaccessible hard-core theoretical physics into widely available program code for simulating these processes (which has become a cornerstone in most structural biology studies today). The code is called GROMACS and it is a molecular dynamics package that is primarily designed for simulating proteins, lipids and nucleic acids.

You might wonder what use that is in day-to-day terms! In actual fact, the entire early stage of the pharmaceutical drug design pipeline is carried out on computers these days, and GROMACS is used by thousands of academic researchers around the world, as well as by all the large pharmaceutical companies in the world. In just the last two to three years, Lindahl's group has been able to use simulations to predict how the receptors in our nerve cells are regulated by external drugs such as anaesthetics and alcohol. This is a really complex process since

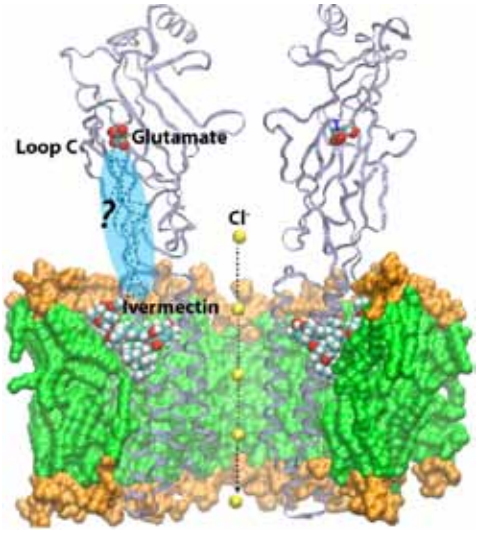

A neurotransmitter (glutamate) couples to the outside part of the ion channel, while drugs or anaesthetics bind to regions in the membrane. These two molecules interact to either increase or decrease signals. Understanding the mechanisms behind this will enable researchers to design molecules that alter their behaviour and help explain the molecular basis of dependency disorders.

some molecules up-regulate (or amplify) the response while others down-regulate (that is, dampen) the nerve signals. The drugs themselves have no effect whatsoever, but they act roughly like transistors: if they are present, the membrane protein channel will respond differently to the normal nerve signals. The research group were able to show that this odd-sounding effect is due to multiple different binding sites in the membrane part of the protein, which have opposite effects. This might sound technical and advanced, but the cool result is that it may make it possible to design new pairs of drugs which would, for example, enable us to fine-tune anaesthesia to safely sedate elderly or ill patients (where anaesthetics can sometimes have unexpected effects compared to with patients in a normal state of health). This discovery also has the potential to improve the treatment of many addiction disorders, including alcoholism as ethanol hits these ligand-gated channels!

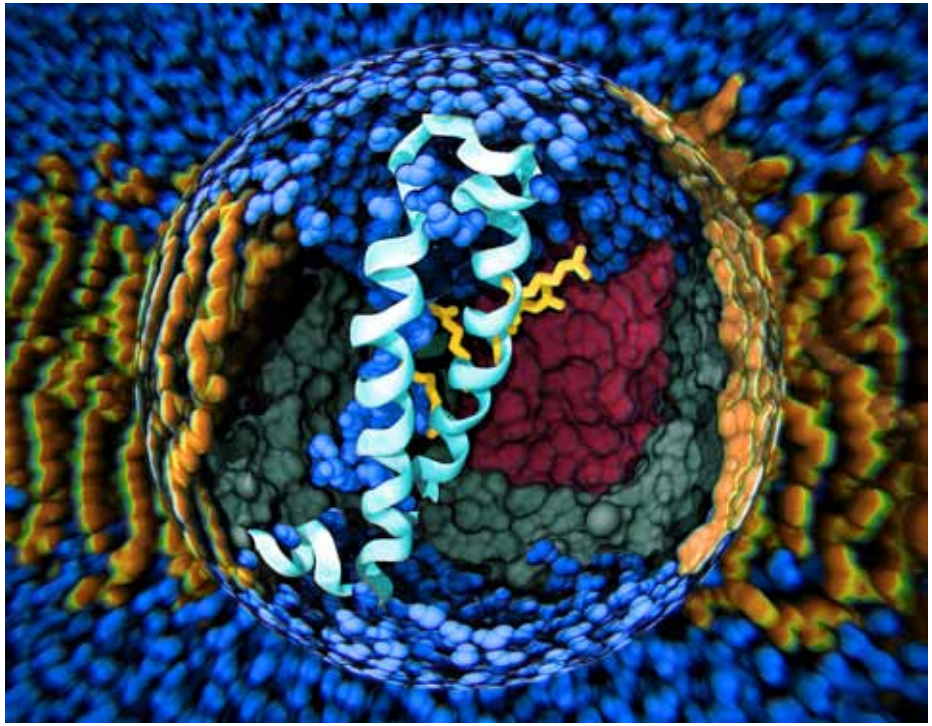

The proteins in our cells are not merely molecules, but amazing microscopic machines whose motions we can simulate with supercomputers. Every heartbeat in our body is the result of nerve signals that are generated by new voltage-gated channels opening in response to changes in the electrical potential of the nerve cells.

Yi Luo and his research group

Producing nanocircuits with graphene fluoride

Imagine a sheet of carbon that is just one molecule thick! That is graphene - the thinnest and strongest form of carbon. It is a hundred times stronger than steel, conducts heat better than any other known material, and conducts electricity just as well as copper does. Now add a fluorine atom to each of the carbon atoms and you get graphene fluoride or fluorographene, which is stronger and more chemically stable than graphene. Fluorographene however has lower electrical conductivity, and so it is classed as an insulator.

Thanks to these properties, fluorographene has been investigated for use in very, very small electronic circuits, known as nanocircuits – tiny electrical circuits with measurements in the millionths of a millimetre. As you can imagine, we cannot build circuits on such a small scale by hand, so special techniques are necessary to construct these incredibly small circuits. As fluorographene is such a thin substance, it is potentially useful in nanocircuits where it could be used as an extremely thin insulator to separate conducting graphene channels from each other and thus form a minuscule electronic circuit.

Fluorographene nanocircuits are potentially vital in supercomputers, as the speed of our computers is determined by the number of field-effect transistors on each of the computational chips in the computer. To put more transistors into a computer chip, we need to decrease the size of them, and that is where these tiny graphene-based nanoscale electronic circuits could be used.

The development of the first graphene-based field-effect transistor was reported in 2007, however that transistor could not be switched off because of its large conductive

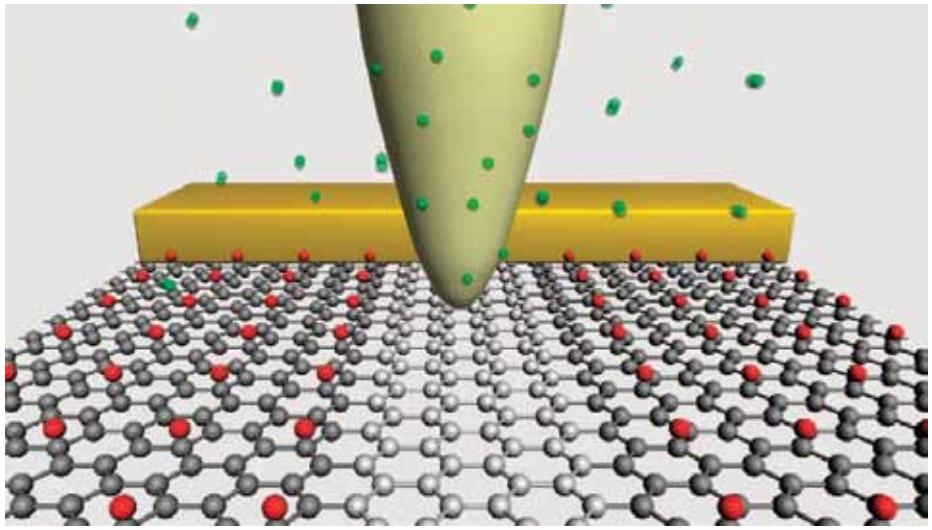

Fluorographene is being turned into an electronic nanocircuit using the cSPL technique. A microscopic stylus with a chemical catalyst is moved across the surface of the fluorographene and reduces it chemically, so the fluorine is removed. This leaves a conducting graphene channel, which you can see in the centre. The channel connects to a metal electrode which is shown in yellow at the back.

graphene area, which made it unsuitable for realistic applications. In 2012, Yi Luo's research group at the Department of Theoretical Chemistry, KTH Royal Institute of Technology, together with their experimental collaborator, Prof. Xiaoping Wang from the University of Science and Technology of China (USTC), proposed using catalytic scanning probe lithography (cSPL) to directly write electronic devices on graphene oxide (which is the single layer form of graphene with oxygen, rather than fluorine, added to it).

In the cSPL technique, a microscopic or nanoscopic stylus bearing a catalyst is moved mechanically across a surface to form a pattern, which results from the interaction of the catalyst with the surface. Using these probes to directly catalyse a variety of localised chemical reactions makes it possible to generate nanoscale features with a high degree of chemical complexity in a "direct-write" manner, almost as if one were printing the resulting chemical substance onto the surface.

Graphene oxide has been considered as a possible means of producing graphene nanocircuits, however there have been problems

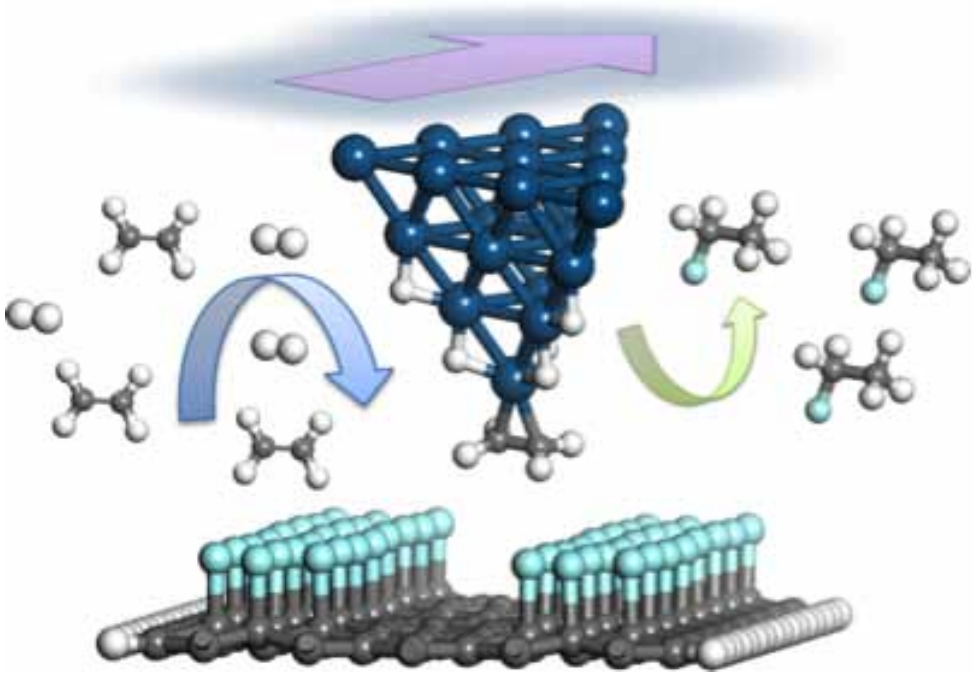

Schematic diagram of the local reduction reaction of graphene fluoride via the cSPL technique with the assistance of a mixture of hydrogen and ethylene.

with producing graphene oxide of a suitable quality. In comparison, fluorographene can be readily obtained by exposing graphene under xenon difluoride gas, making it a better material for practical applications. Therefore, it is worth considering using the new cSPL method to partially reduce fluorographene (rather than graphene oxide) in order to build nanocircuits. Xinrui Cao, Yongfei Ji, Wei Hu and Sai Duan from Yi Luo's research group at the KTH Department of Theoretical Chemistry have been using Lindgren to investigate ways to do just that. With the help of high accuracy first principle simulations, they have proposed a new strategy for fabricating nanoscale electronic circuits on graphene fluoride using the cSPL technique.

The scheme that the KTH researchers have proposed uses the cSPL "direct writing" technique to "write" electronic circuits onto fluorographene at a low temperature compared to that which would be needed for graphene oxide – this is a promising method for practically producing nanocircuits. Their research is mainly focused on simulating these processes using first principles and using the resulting information to design suitable molecular electronic devices for production. It is very important to simulate the production method first because the results obtained from the theoretical simulations make it possible to avoid damage to the equipment and instruments, as well as to operators, when it comes to performing the experiments. For example, simulation results indicate that the tip-induced reduction of fluorographene in a hydrogen atmosphere would release the hazardous gas, hydrogen fluoride. Consequently, when practical experiments are started, the researchers would definitely not perform the experiments in an atmosphere of hydrogen.

In one particularly interesting simulation project that was run at PDC, Yi Luo's group studied the possibility of directly writing nanoelectronic devices on a single layer of graphene fluoride with the tip of an atomic force microscope that is coated in platinum. The diagram on the opposite page shows how the graphene fluoride is chemically reduced (that is, how it gains additional electrons) with the cSPL technique using a mixture of hydrogen and ethylene. The reduction reaction removes some parts of the fluorine when scanning the fluorographene surface with the platinum-coated tip. This changes the fluorographene in those areas to graphene and thus a conductive graphene channel is created on the surface. Using this method nanocircuits can be built without using or producing any hazardous products.

Greening the skies


Dan Henningson


Ardeshir Hanifi


Philipp Schlatter

Global warming and concerns about greenhouse gases are major issues these days. While carbon offsets are one way to balance carbon emissions in the airline industry, Professor Dan Henningson and his team in the Department of Mechanics at the KTH Royal Institute of Technology and the Linné FLOW Centre – which includes Ardeshir Hanifi, Philipp Schlatter and a number of postdoctoral researchers and graduate students – are looking at ways to reduce fuel consumption (and hence the amount of carbon dioxide in aviation exhaust gases) by reducing the air resistance (or drag) on aeroplane wings.

To this end they have been conducting large-scale computer simulations of the airflow over aircraft wings to help understand what effects different wing shapes and different wing surfaces have on air resistance.

Philipp Schlatter explains that a significant portion of the drag on an aircraft is caused by what is referred to as the turbulent flow of air over the surface of the aircraft. To understand the difference between this turbulent flow (which causes a lot of the air resistance) and what we ideally want (which is known as laminar flow), think of water flowing in a twisting river bed where there are many stones and tree roots – then compare that to water flowing in a smooth straight canal. In the river, there will be lots of eddies and the water will rush back on itself when it flows towards a rock – in other words, the flow will be turbulent.

In the canal however, the water flows relatively smoothly with no interruptions, which is more like laminar flow (as shown in the diagram below). If you look at the figure on the next page, which shows the results from one of their simulations of turbulence around a wing, you will see swirls in the colours next to the narrower part of the wing shape. These swirls indicate air turbulence.

These results can be used to understand how varying the wing profile and surface can cause the airflow around the wing to change from a smooth laminar flow to a turbulent flow. (The more turbulent the flow is, the greater the air resistance is, and hence the harder the engines of the aircraft have to work and the more fuel they consume.) The resulting knowledge is vital for

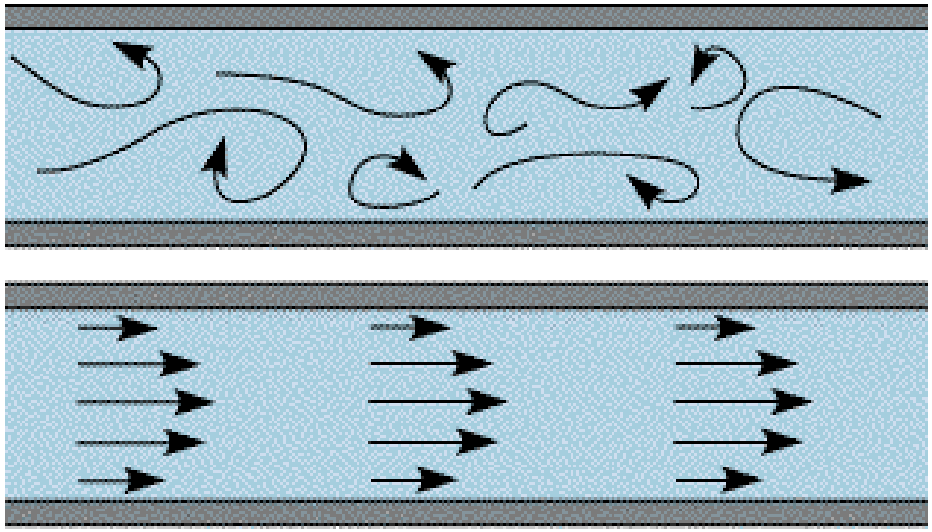

The sketches show the difference between turbulent flow (top) and laminar flow (bottom) in a parallel channel. The arrows indicate the trajectory of particles in the fluid. As the particles in the turbulent case follow a more chaotic route, greater force is needed to keep the flow moving.

The image below shows an instantaneous snapshot of the turbulent flow of air around a section of an aeroplane wing. The colour code indicates the velocity of fluid particles, ranging from blue (slow) to red (fast). One can clearly see the formation of a very thin boundary layer adjacent to the wing surface, which is characterised by intense turbulent motions.

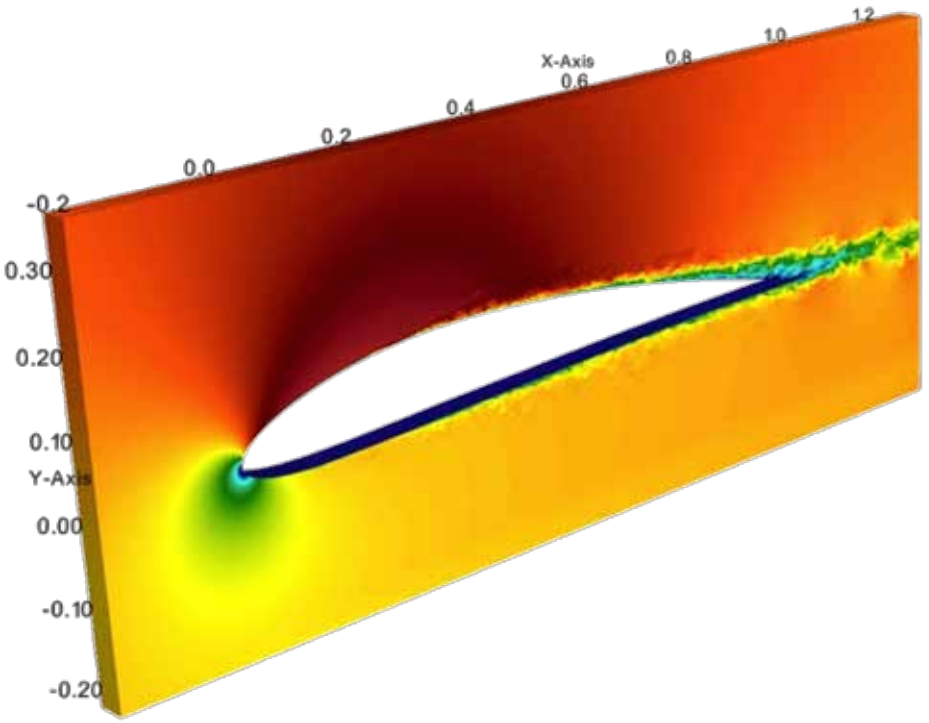

developing low-drag wing designs that reduce the area of turbulent airflow and thus decrease the plane's fuel consumption and its impact on the environment.

As you can imagine, it takes a lot of equations and computations to make an accurate model of air resistance on even a small area of a wing. Using a total of about 10 billion points to represent the surface of the wing and a significant portion of the air around it, the group at KTH Mechanics has been able to simulate air turbulence on 50 cm of an aircraft wing at flight conditions. These results, which are shown in figure on the right, have been used to clarify some scaling laws (which can be used for improving models of turbulence which in turn allow cheaper development times). The results have also helped other researchers to calibrate some of the measurement probes that they use when building physical models and testing them in a wind tunnel. In

the past, that was the only way to investigate the effects of changing the shape of a wing, or the material that it was made of. However it was, and still is, both time-consuming and costly as a means of testing various options for wing designs. Consequently, researchers are ultimately aiming to model the air currents around a whole aeroplane (although it is predicted that it will take another 20 or 30 years to reach the point where that can be done). As a step in that direction, the group at KTH Mechanics is now working on extending their simulations to cover a whole wing, which is something that no one has yet achieved. These computationally very intensive projects are supported both by the Swedish e-Science Research Centre (SeRC) and the Wallenberg Foundation.

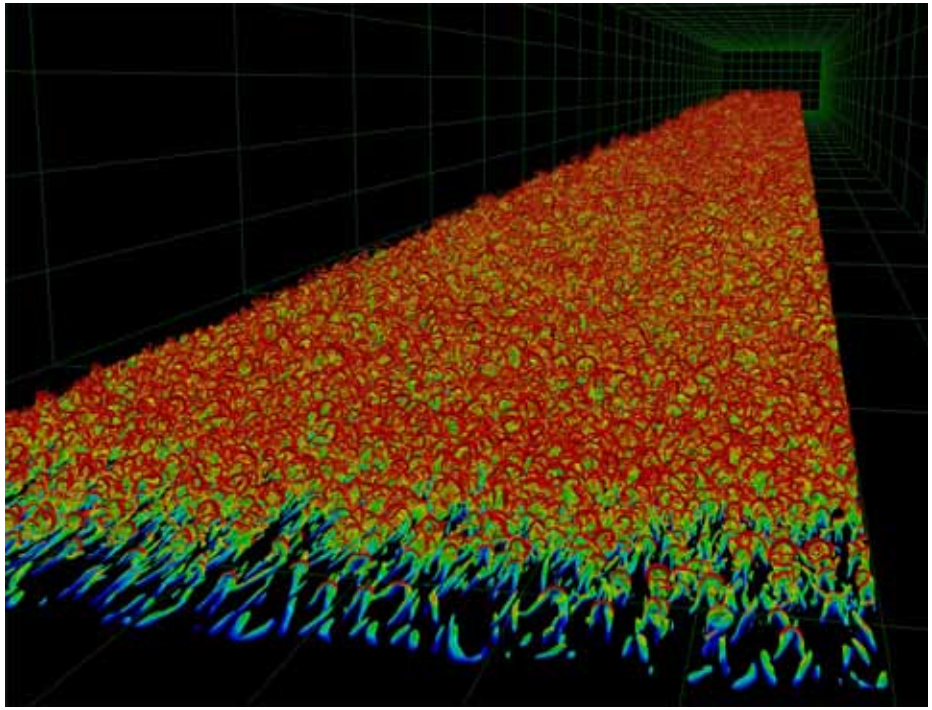

The canonical version of an aeroplane wing is the so-called flat-plate boundary layer. The figure shows a snapshot of the complex turbulent flow structures that appear around the surface of a wing, with the colour indicating the local velocity of the air. Just imagine all these turbulent worms of air being continuously generated while flying in an aeroplane! To make accurate predictions of the lift and drag forces on a wing, the researchers need to take all of these phenomena into account in their calculations.

Computational brain dynamics and function

Researchers from the Department of Computational Biology at the KTH Royal Institute of Technology and the Department of Numerical Analysis and Computer Science at Stockholm University are studying how our brains work. Anders Lansner, Jeanette Hellgren Kotaleski and Pawel Herman are working on modelling the neural systems of the brain to understand how it functions, both at a biologically detailed level and from a less detailed computational perspective. Along with postdoctoral researchers, Omar Gutierrez Arenas and Olivia Eriksson, and other colleagues and postgraduate students, they develop and simulate multi-scale and large-scale brain networks in an attempt to integrate the huge amounts of available biological data into coherent mathematical models that give insights into the mechanisms underlying the brain's impressive information processing capabilities.

This research is being conducted in close collaboration with experimental neuroscientists at the Karolinska Institute in Stockholm and with European researchers via EU projects (such as the Human Brain Project) and together with the Swedish e-Science Research Centre (SeRC). The brain systems that are being studied include the cortex (the largest part of the brain), the hippocampus (a small part of the cortex that is important for learning and memory) and the basal ganglia (which is critical for decision-making and is primarily affected in conditions like Parkinson's disease).

Biologically detailed models of the brain each address different levels of brain systems, such as the cascades of biochemical signals within individual brain cells that are triggered by neurotransmitters, or the electrical properties of individual neurons in the brain, or the networks of neurons that form brain microcircuits and connections between areas of the brain. These models vary in the amount of detail that they contain and in the extent to which they take into account information about how other levels of the brain work.

To make a more comprehensive model of the brain that builds on biological data at different levels of brain organization, the researchers are using the multi-scale computational approach. This will help to bridge the gaps in our understanding of how human behaviour and brain phenomena observed with the use of popular neuroimaging techniques such as functional magnetic imaging (fMRI) depend on the cascades of electrical signals between the brain cells in normal healthy brains.

Understanding this may prove crucial for the diagnosis and treat-

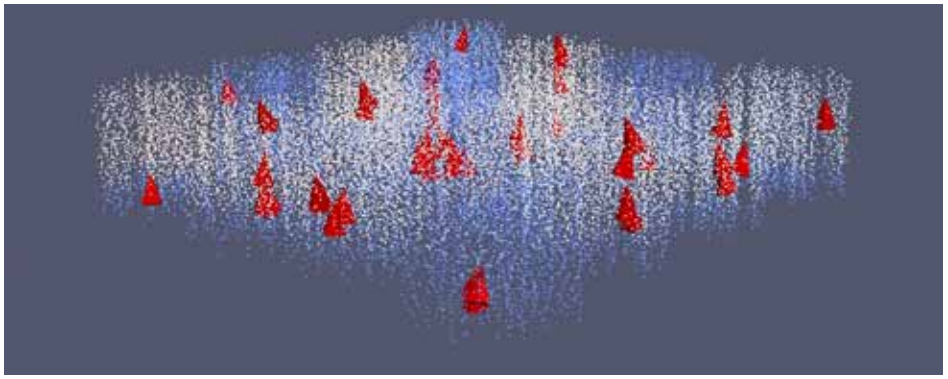

The image above is from a simulation of the neural activity of an area of the neocortex. The individual dots represent neurons, which become larger when firing. The red pyramids indicate active pyramidal neurons (which are a particular type of neuron).

ment of a variety of illnesses. For example, it could help us to develop new drugs, since drugs typically have an effect on proteins inside brain cells, while the purpose of the drug is to bring relief at the level of the brain's neural network or even at a behavioural level. Thus the multi-scale computational brain models can be used for experiments to understand brain disease processes and to explore the effects of drug therapies.

The researchers have recently modelled the cascade of biochemical signals inside a cell that is triggered by the activation of receptors involved in reward-dependent learning in part of the basal ganglia. Modelling such systems is necessary for understanding the mechanisms involved in memory and learning in the brain. However this is very complex: multiple sources of information are required when building models of this kind of subcellular signalling. To simulate such systems is computationally intensive and requires the use of high-speed supercomputers at PDC and elsewhere.

So far we have considered using computers to model brains, but it is also possible to look at how the brain processes information and to use that to develop new computing technologies. This leads us to another type of brain simulation that the researchers have been running on the supercomputer facilities at PDC. The long-term goal of this work is to identify the most significant aspects of the detailed brain models and to incorporate them into a framework based on more abstract brain theories. This could then be used to produce a system that can analyse large amounts of complex data in the same way that the brain processes multiple streams of sensory information. To date the researchers have run some of these types of simulations on the Cray supercomputers at PDC. As a result they have developed a "cortex-inspired" computational framework for performing advanced analysis of brain imaging recordings from various types of medical scanning, such as fMRI and positron emission tomography (PET).

To diagnose and cure human brain diseases, we need to be able to identify abnormal patterns in the ongoing brain activity that correspond to the actual illness. The vast amounts of data that are being generated by modern medical scanning equipment are too large for clinicians to analyse by eye and too complex to be assessed with the current conventional methods. Thus society at large has a problem where vital information (which might, for example, lead to being able to detect and diagnose illnesses early enough to effect a cure before the problem becomes serious) is not being fully utilised.

To address this problem, the KTH researchers are currently working towards enhancing their existing analysis method to make it possible to identify and study relevant patterns in electroencephalography (EEG) and magnetoencephalography (MEG) neuroimaging data. They want to be able to identify patterns both in spatial terms (relating to different areas in the brain) and in temporal terms (looking at what happens over time in specific parts of the brain). At present they are focusing on developing the large-scale computational framework to adequately model the brain's information processing machinery and to analyse neuroimaging data – the main diagnostic aspects of the work are yet to come.

Interestingly, the work may also be beneficial in a broader perspective. As the methods for analysing brain imaging data that are based on models of the human brain are general enough to be used in many different areas, this line of research appears to be attractive for the emerging field of big data (which looks at ways to handle huge amounts of data that cannot be dealt with by normal data processing methods).

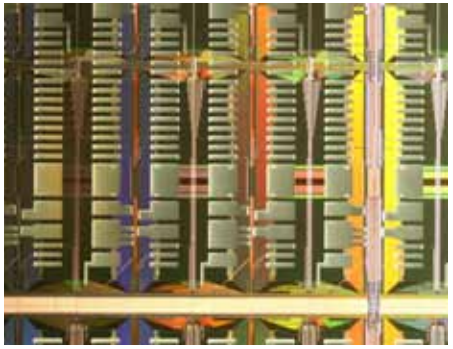

Neuromorphic chips are VLSI chips that implement models of the neural system.

Source: Kirchhoff Institute for Physics, University of Heidelberg

To effectively handle such vast quantities of data, the researchers' neural-network-based brain models need to be very large and hence require so many calculations that the simulations have to be run on very fast supercomputers. However, the configurations of different supercomputer systems and their components, along with the way the code takes advantage of the parallel processing in each supercomputer, mean that there can be variations in how fast and how efficiently code runs on different systems. Therefore the researchers are also investigating ways to implement their simplified large-scale brain network models efficiently on novel types of massively parallel hardware like the SpiNNaker system (a computer architecture specifically designed to simulate the human brain) and the Epiphany chip from Adapteva, as well as on a custom-designed very-large-scale integration (VLSI) chip, which is being developed in collaboration with the KTH School of Information and Communication Technology.

Would you like to use PDC's supercomputers for your research?

Beskow, Cray XC40, and Milner, Cray XC30

Contact PDC for more information:

support@pdc.kth.se

PDC Center for High Performance Computing
KTH Royal Institute of Technology
Teknikringen 14, fourth floor
100 44 Stockholm
Sweden
Phone: +46 8 7906000

www.pdc.kth.se