#2 | 2015

A magazine published by UNINETT Sigma2

meta

Looking inside the earth

using e-infrastructure resources

Exit interview with Hans A. Eide

Neuronal modelling for schizophrenia research

Renewal of the national e-infrastructure resources

#2 | 2015

A magazine published by UNINETT Sigma2

EDITORIAL DIRECTOR:
Gunnar Bøe
CONTRIBUTING EDITOR:
Vigdis Guldseth

Subscription: An electronic version is available on www.sigma2.no

E-mail: sigma2@uninett.no
Phone: +47 73 55 79 00
Layout: HK reklamebyrå
Print: Skipnes kommunikasjon AS

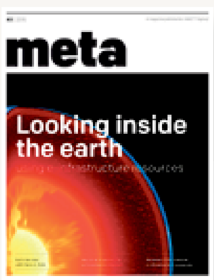

Cover picture: Shutterstock

CONTENTS

One of the objectives of Sigma2 is to take care of the procurement of new e-infrastructure for the research community in Norway. Activities relating to a new procurement started in the spring of 2015 with a pre-study involving many people from the universities of NTNU, UiB, UiO and UiT.

Gunnar Bøe
Managing Director
UNINETT Sigma 2 AS

Procuring e-infrastructure is in itself a very daunting task, getting all the user requirements to match with the right services and technology. On top of this, we have a new consortium where Sigma2 is trying to coordinate the interests of four universities into the same procurement. It takes a lot of goodwill on many parties to move such a complex process forward. We are very grateful the good spirit of cooperation shown so far and how many people from the universities that have joined enthusiastically into this work. Sigma2 is trying very hard to make this a joint effort with the universities and we are very dependent on the input from the experts working at the metacentre.

The timing of the grant from the Norwegian Research Council for new e-infrastructure was very good. At the end of November, we were informed that we have been granted 75.7 MNOK for procuring new storage facilities and HPC systems. This was one of two important milestones remaining in the pre-study. The other important milestone was for the board to approve our strategy and plan for the new procurement. The board approved this in their meeting on December 4th. With the completion of the requirement specifications, planned for the beginning of January, we can now start the procurement process early next year.

Dissemination of research is a very important task. It is also challenging task for the researcher to prioritize time for it and to find the right technical level for the targeted audience. This might be one reason for the challenge we have to get contributions from researchers presenting their research in Meta. We will encourage all of our researchers to think about how you can present your excellent research in Meta (and other channels). It is important to keep our stakeholders and the public in general, up to speed on the scientific progress. A lot of money is spent on research infrastructures and we have a joint responsibility to disseminate the results and usefulness of the e-infrastructure.

In this issue, we can present some of the excellent research in two very different fields. Research in geosciences has traditionally consumed large amounts of CPU-hours. Some of these CPU-hours have now been spent on the creation of high-resolution, detailed dynamical models of planets. Centre for Earth Evolution and Dynamics (CEED) is doing very interesting research into the earth. One of their aims is to get a better understanding of the workings of our own planet.

To many it might be a surprise that our e-infrastructure also is used for research in psychiatry. Mental illness is causing many people at lot of grievance and it is very satisfying to see our e-infrastructure being used for research that in due time hopefully can improve the lives of these people. It is neuroscience used in psychiatry that makes a bridge into the e-infrastructure. Neuroscience has actually quite a long history in applying computational methods on simpler models running on PCs. However, when modelling network of neurons, a lot more computational resources are needed, hence the use of Stallo and Abel for this research.

We foresee a shortage of HPC resources in the next period. This quite common when systems come to the end of their lifetime and new systems are procured. We hope you will bear with us until we again can provide you with generous awards for computing resources.

This edition of Meta is very special, since it will be the last printed edition. With Sigma2s new communication plan we will have a stronger focus on using electronic media. Meta will continue to provide you with excellent articles on advanced research performed on the national e-infrastructure using a more dynamic format in the future. We hope this will further increase the value of Meta.

FEATURES

FOCUS

At CEED, strong numerical skills combined with cross-disciplinary geological and geophysical expertise, and the accessibility of high-performance computing systems, such as Abel and Stallo, enables the creation of high-resolution, detailed dynamical models of our planet, and other bodies.

AUTHORS

Dr. Abigail Bull, Dr. Trond Torsvik, Ms. Eva Hauk Fritzell, Dr. Chris Mohn, Mr. Nils Prieur, Dr. Tobias Rolf, Dr. Grace Shephard and Dr. Reidar Trønnes
Centre for Earth Evolution and Dynamics, University of Oslo

Photos: Shutterstock

Looking inside the Earth

using e-infrastructure resources

The Earth is a complex and dynamic system. For the most part inaccessible and observable only through-remote-sensing techniques, the inner workings of our planet have serious catastrophic potential for humans in the form of earthquakes and volcanic eruptions. Using the computing power of Notur, a research group at the University of Oslo is investigating the link between processes which occur deep in Earth's interior and those at the surface, aiming to better understand the workings of our planet and to unravel the clues as to how Earth has evolved to the remarkable planet that it is today.

While the interior of the Earth cannot be directly observed, much information about its structure has been derived from global seismic tomography studies. Based on mathematical similarities with the medical 'CAT' scan x-rays, seismic tomography provides key observational constraints on the present state of heterogeneity within the Earth. Such studies have revealed a radially stratified Earth, separated into the crust, mantle and outer and inner cores (Fig. 1). Quantitatively, the mantle represents Earth's most significant layer; comprising nearly 84% of the planet's total volume and 67% of its mass. Understanding the complex structure and dynamics of the mantle is a fundamental goal for researchers at the Centre for Earth Evolution and Dynamics (CEED), as mantle behaviour has important implications for numerous aspects of Earth and space sciences, including earthquake mechanisms,

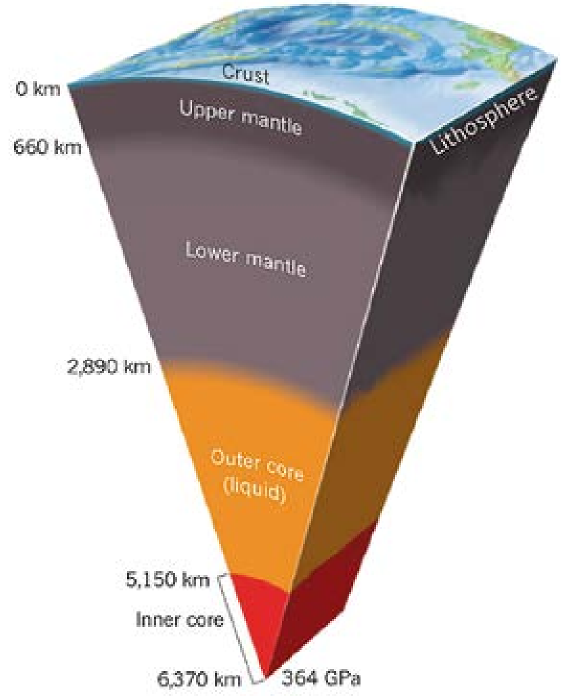

Figure 1.
Earth's interior (Duffy, 2011).

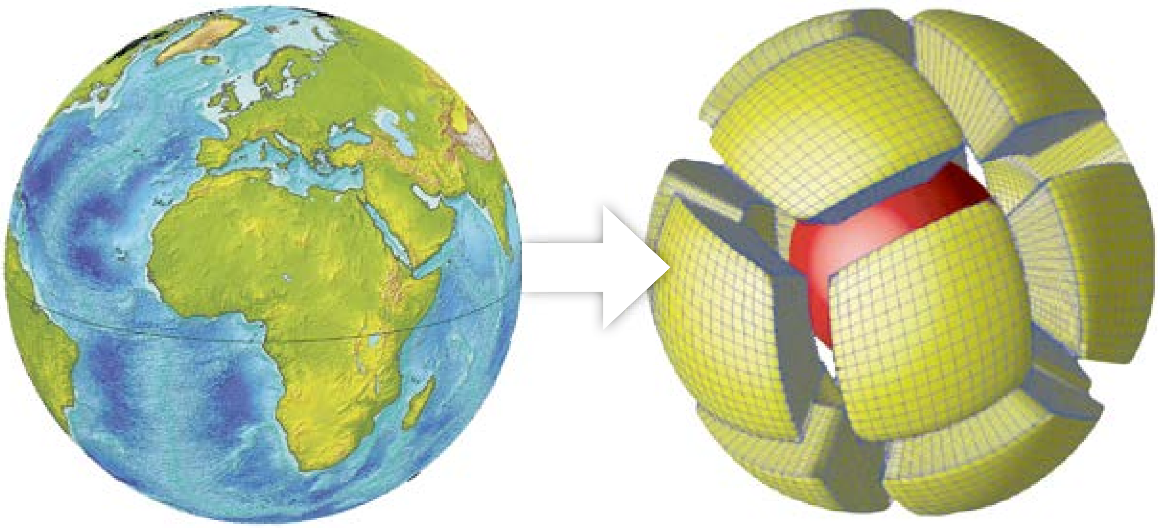

Figure 2.

To simulate convection, the entire Earth is represented as a 3D grid consisting of several millions of points. Mathematical equations, which describe the physics of mantle convection, are repeatedly solved at each discrete point, allowing researchers to track several variables—including the temperature of the convecting material, its velocity and its composition. It takes several days to weeks to run a simulation, depending on the complexity of the model.

volcanic systems, planetary evolution and climate change. On human timescales, the mantle seems rigid, solid. However, over much longer periods - hundreds of thousands to millions of years - the rocks of the mantle flow.

This flow, termed 'mantle convection', is the primary way in which heat is transferred from the interior of the Earth to its surface, as hot, less dense, rock rises up at regions termed 'mid-ocean ridges' and cold material is recycled into the mantle at regions termed 'subduction zones'. Although extensive knowledge on the dynamics of mantle convection can be derived from geological, geophysical and geochemical observations and experiments, many aspects of convection remain ambiguous. However our present knowledge of the deep Earth, combined with an understanding of the fundamental physics and chemistry, provides a solid basis from which to investigate the dynamics of this enigmatic region of our planet.

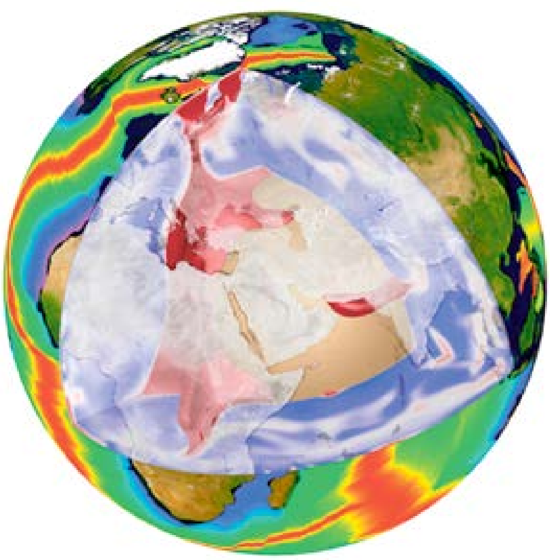

The nature of mantle convection can be described by mathematical models - a set of partial differential equations and boundary conditions which govern the movement of material in Earth's mantle. In the field of geodynamics, such models predict quantitatively what will happen when the mantle deforms over time. Because the process of convection is very complex, the corresponding mathematical models are too complicated to be solved analytically. To overcome this, the equations are transformed into discrete equations to be solved numerically using computers. This branch of Earth science dealing with numerical modelling has greatly evolved in the past decade. As a result, detailed computer models of the mantle now exist and the dynamics of the deep Earth can be explored through highresolution computer simulations (Fig. 2). Utilising detailed numerical codes that describe mantle convection and other dynamic Earth processes, a versatile team of researchers at CEED use the

Figure 3.

Example of a computer model (Bull et al., 2014) of mantle convection within the Earth. The red colour represents material moving upwards towards the surface: note the large plume rising beneath Iceland. The age of the oceanic crust is shown as bright colours: orange-red=youngest, purple=oldest.

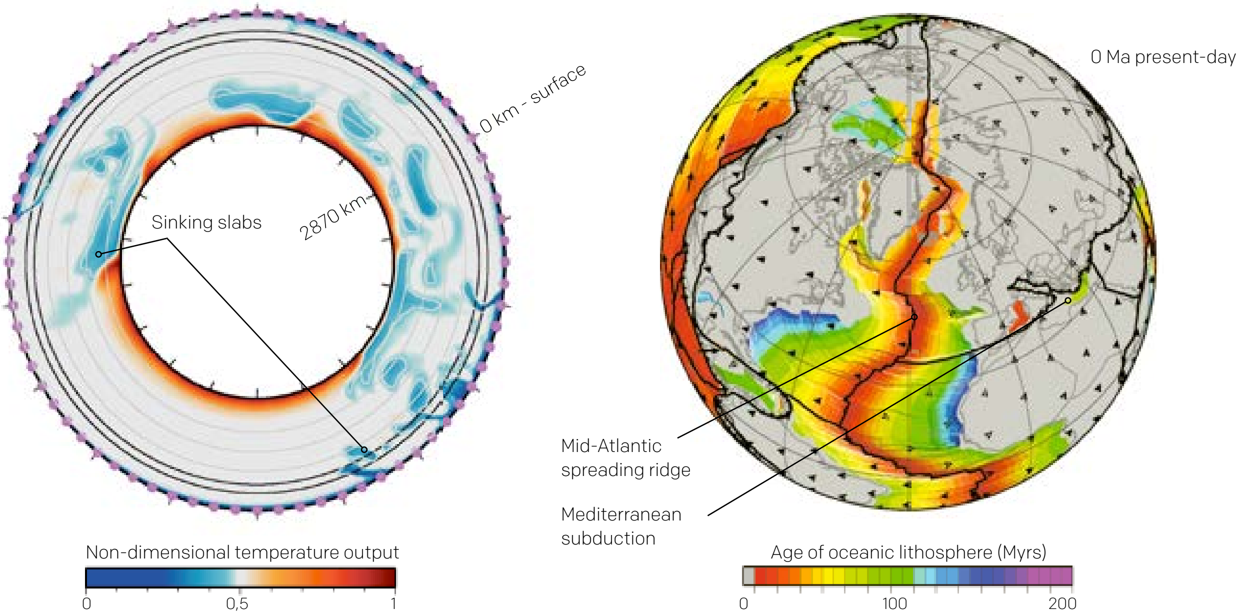

Figure 4.

Example of input (right) and output (left) of a 4D (space plus time) model of mantle convection (Shephard et al., 2014). The surface perspective (right), shows the present-day age of oceanic seafloor (colours; Müller et al., 2008) and continents (grey). The deep mantle view (left) shows the predicted location of subducted slabs in the mantle in a vertical slice around the equator. Blue anomalies are cold subducted slabs of past oceans and red are hotter than average regions near the core-mantle boundary.

supercomputers Abel (UiO) and Stallo (UiT) to create intricate computer models of Earth's interior (Fig. 3). Such numerical models can be used to study the evolution of a planet and its role on current surface observations. It is this evolution of the mantle that CEED researchers are interested in.

Exactly how did Earth's mantle come to have the structure that it has today? Has the mantle always looked the same, or has it evolved through time due to convection? Indeed, a fundamental mission of the research centre is to develop an Earth model that explains how mantle processes drive plate tectonics. With this goal in mind, a method to integrate data from several geo-disciplines into high-resolution numerical models of the mantle is now being employed at CEED, by combining the complementary techniques of seismology, geochemistry, and palaeomagnetism. Due to the inaccessibility of the deep Earth, such a multidisciplinary approach is crucial.

One aspect of mantle evolution being investigated at CEED is the fate of crustal material when it is recycled into the mantle at regions termed 'subduction zones'. These boundaries mark the collision between two tectonic plates, with one plate sliding beneath the other, curving down into the mantle. Subduction continues to be one of the most powerful and dynamic processes on planet Earth and as numerical modelling improves we can come to understand more about this process. By integrating data detailing the movements of the tectonic plates over the last 400 million years into the numerical convection models, it is possible to shed light on the fate of subducted material in the mantle (Figs. 4, 5 and 6).

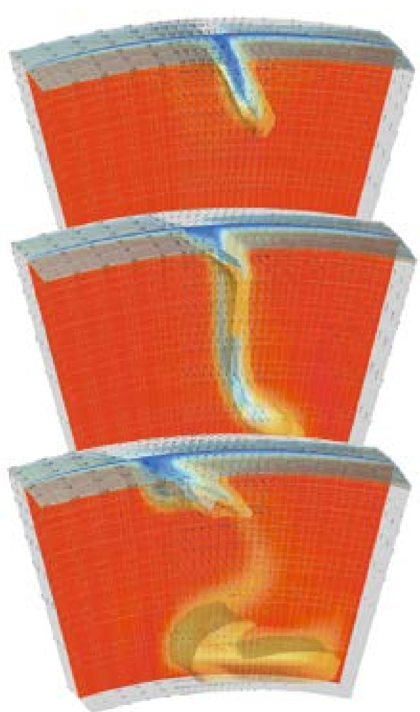

Figure 5.

Excerpt from a computer model of mantle convection (Bull et al., 2014) showing a slab of cold, dense crustal material (blue) descending into the mantle. As the model progresses through time (top to bottom), the slab breaks and part of it sinks to the base of the mantle. The fate of slabs such as this is a focal point of research at CEED.

In addition to the fate of subducting material within the mantle, researchers at CEED strive to understand the relation between mantle convection and the drift of continents at the surface, and to elucidate the conditions under which continents assemble to form a supercontinent, as was the case on Earth around 260 Myr ago (Fig. 7).

Complementary to the numerical simulations of large-scale mantle convection, CEED researchers are also undertaking numerical simulations on the atomistic scale, endeavouring to discover how minerals behave at the high-pressures and temperatures of the deep mantle. The Earth's most important mineral, bridgmanite, constitutes 70-75% of the lower mantle and 35-40% of the entire Earth. Bridgmanite (bm) transforms to a high-pressure and low-temperature crystal structure, referred to as post-bridgmanite (pbm), in the coldest parts of the deep mantle (Fig. 8). Together with the mineral ferropericlase, bridgmanite comprises the main rock type, peridotite. Using the Abel computer cluster, Density Functional Theory (DFT) computations are performed which enable researchers to determine the strength, stiffness and length of each of the cation-anion bonds within each mineral, and thereby the atomic positions and elastic constants. The computations also derive volume, enthalpy and entropy and the relative stability of coexisting minerals at various pressures and temperatures. Such studies have important implications for certain areas of the deep mantle.

While most interest is focused on our home planet, CEED researchers also study mantle evolution for other planets, e.g., Venus, using the supercomputer Stallo. Contrasting with Earth, plate tectonics and continental drift are not observed on Venus. To understand this dichotomy, Venus' gravity field, measured by satellites, is compared to results from numerical model predictions, enabling researchers to constrain the structure and evolution of Venus' silicate mantle. It is important to investigate our neighbouring planets to gain insight into the whole range of dynamic processes that occur throughout planetary evolution.

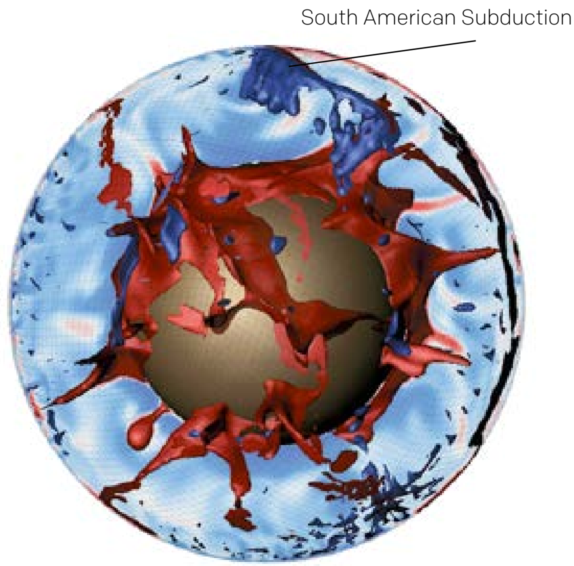

Figure 6.
A computer model of mantle convection (Bull et al., 2014) showing a slab of cold, dense crustal material (labelled, blue) descending into the mantle beneath South America. Hot material, which rises up towards the surface from the base of the mantle is shown in red.

Figure 7.
Left: The supercontinent Pangea, formed on Earth ca. 260 million years ago (Domeier and Torsvik, 2014) . Right: Supercontinent formation in a mantle convection model. The 3D temperature field is shown as bright colours, with subducting slabs indicated in blue. Grey contours mark the position of continents at the surface (Rolf et al., 2012).

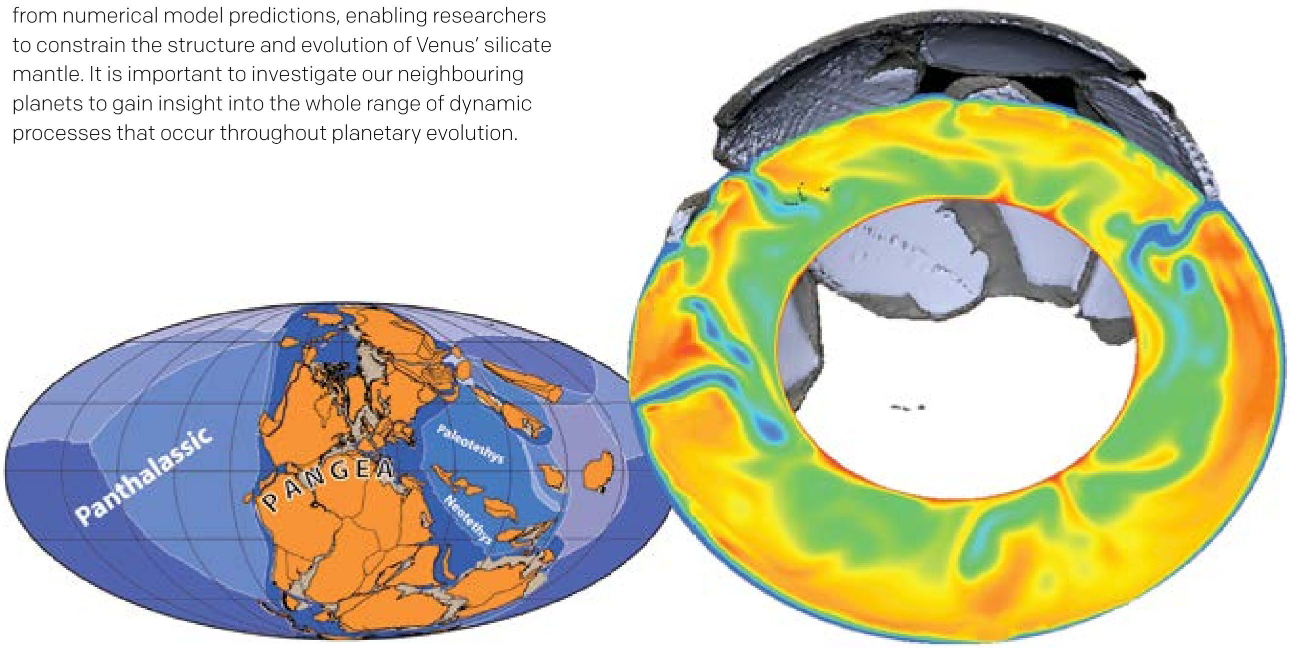

Figure 8.

Crystal structures of MgSiO3-based minerals. The yellow SiO6-octahedra have O-corners with Si (and minor Al) in the middle. Red spheres are Mg, Fe and Al in larger positions. The octahedra are corner-linked in all three directions in bridgmanite but edge-linked along the aaxis in post-bridgmanite, which has rigid ac-planes in a layered structure with a long and compressible b-axis. The resulting anisotropy in strength and diffusion rate causes weak rheology and low viscosity.

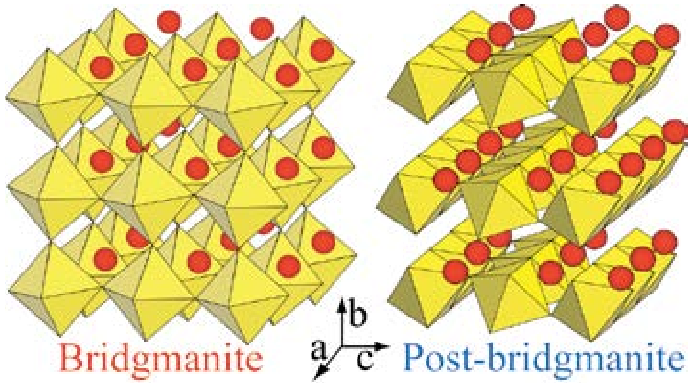

Figure 9.

Crater profile 0.8 seconds after the impact of a spherical (1600 m diameter), granitic projectile on a granitic target. The shock wave generated by the collision expands radially from the point of impact. Shown is the pressure in the range from 0.1 GPa (blue) to 5 GPa (red).

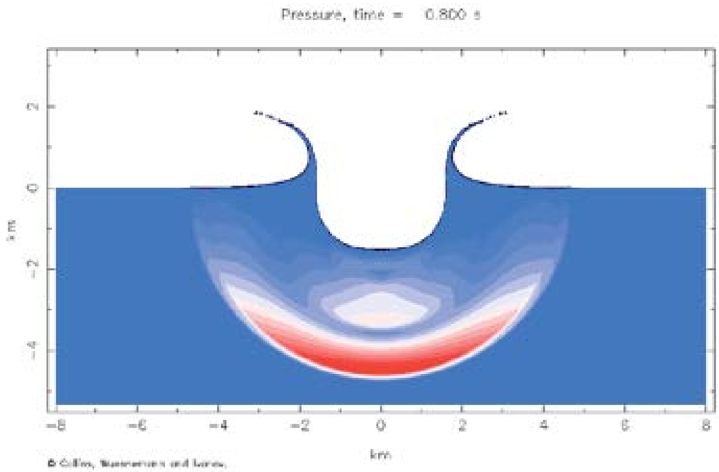

One such process is crater formation caused by impacts. On Earth, presently less than 200 crater structures are known due to the continuous action of resurfacing processes such as erosion and plate tectonics. Contrastingly, the surface of the Moon is heavily cratered and provides evidence for a much older surface than that of Earth, a phenomenon usually explained by the absence of modification processes.

The process of impact cratering itself is rarely observed, thus, numerical simulations provide a powerful alternative to investigate crater formation. The use of such numerical codes on the supercomputer infrastructure Stallo allows insight into the processes behind crater formation and how they are influenced by various target properties such as porosity, friction and material strength (Fig. 9). Both the frequency and the size of craters on a homogeneous surface can be used to derive crater age. Target properties influence the final size of craters substantially, and therefore the age derived for the target surface. Information retrieved from numerical models is therefore crucial to reconstruct the surface histories of planetary bodies accurately and to improve the constraints on the evolution of planets, moons, comets and asteroids during the last 4.5 billion years.

Numerical modelling, such as that described in this article is a prime tool for studying Earth evolution and the dynamics of the deep Earth. At CEED, strong numerical skills combined with cross-disciplinary geological and geophysical expertise, and the accessibility of high-performance computing systems, such as Abel and Stallo, enables the creation of high-resolution, detailed dynamical models of our planet, and other bodies. Through these models, CEED aims to better understand how the Earth has evolved to be the unique planet that we live on today.

REFERENCES

1. Bull, A.L, Domeier, M, Torsvik, T.H.
The effect of plate motion history of the longevity of deep mantle heterogeneities, Earth Planet. Sci. Lett. (2014).

2. Domeier, M., Torsvik, T.H.
Plate tectonics in the Late Paleozoic, Geoscience Frontiers (2014).

3. Duffy, T.S.
Earth Science: Probing the core's light elements, Nature (2011).

4. Müller, R.D., Sdrolias,M., Gaina, C., Roest,W.R.
Age, spreading rates and spreading asymmetry of the world's ocean crust. Geochemistry, Geophysics, Geosystems v9. (2008).

5. Rolf, T., Coltice, N., Tackley, P.J.
Linking continental drift, plate tectonics and the thermal state of the Earth's mantle, Earth Planet. Sci. Lett. (2012).

6. Shephard, G., Flament, N., Williams, S., Seton, M., Gurnis, M., Müller, R. D.
Circum-Arctic mantle structure and long-wavelength topography since the Jurassic. J. Geophys. Res. (2014).

***also: Natural History Museum, University of Oslo.*

Photos: Shutterstock

Neuronal modelling for schizophrenia research

AUTHORS

Tuomo Mäki-Marttunen
Ole A. Andreassen
CoE NORMENT, KG Jebsen Centre for Psychosis Research, UiO, Oslo University Hospital
Gaute T. Einevoll
IMT, NMBU and Department of Physics, UiO

Schizophrenia is a complex mental disorder characterized by impaired perception of reality, and often abnormal social behavior. Scientific studies have found that schizophrenia is highly heritable, meaning that a large proportion of the risk is due to genetic factors. This has inspired decades of research trying to find the "schizophrenia gene" that would explain the mechanism of the disorder.

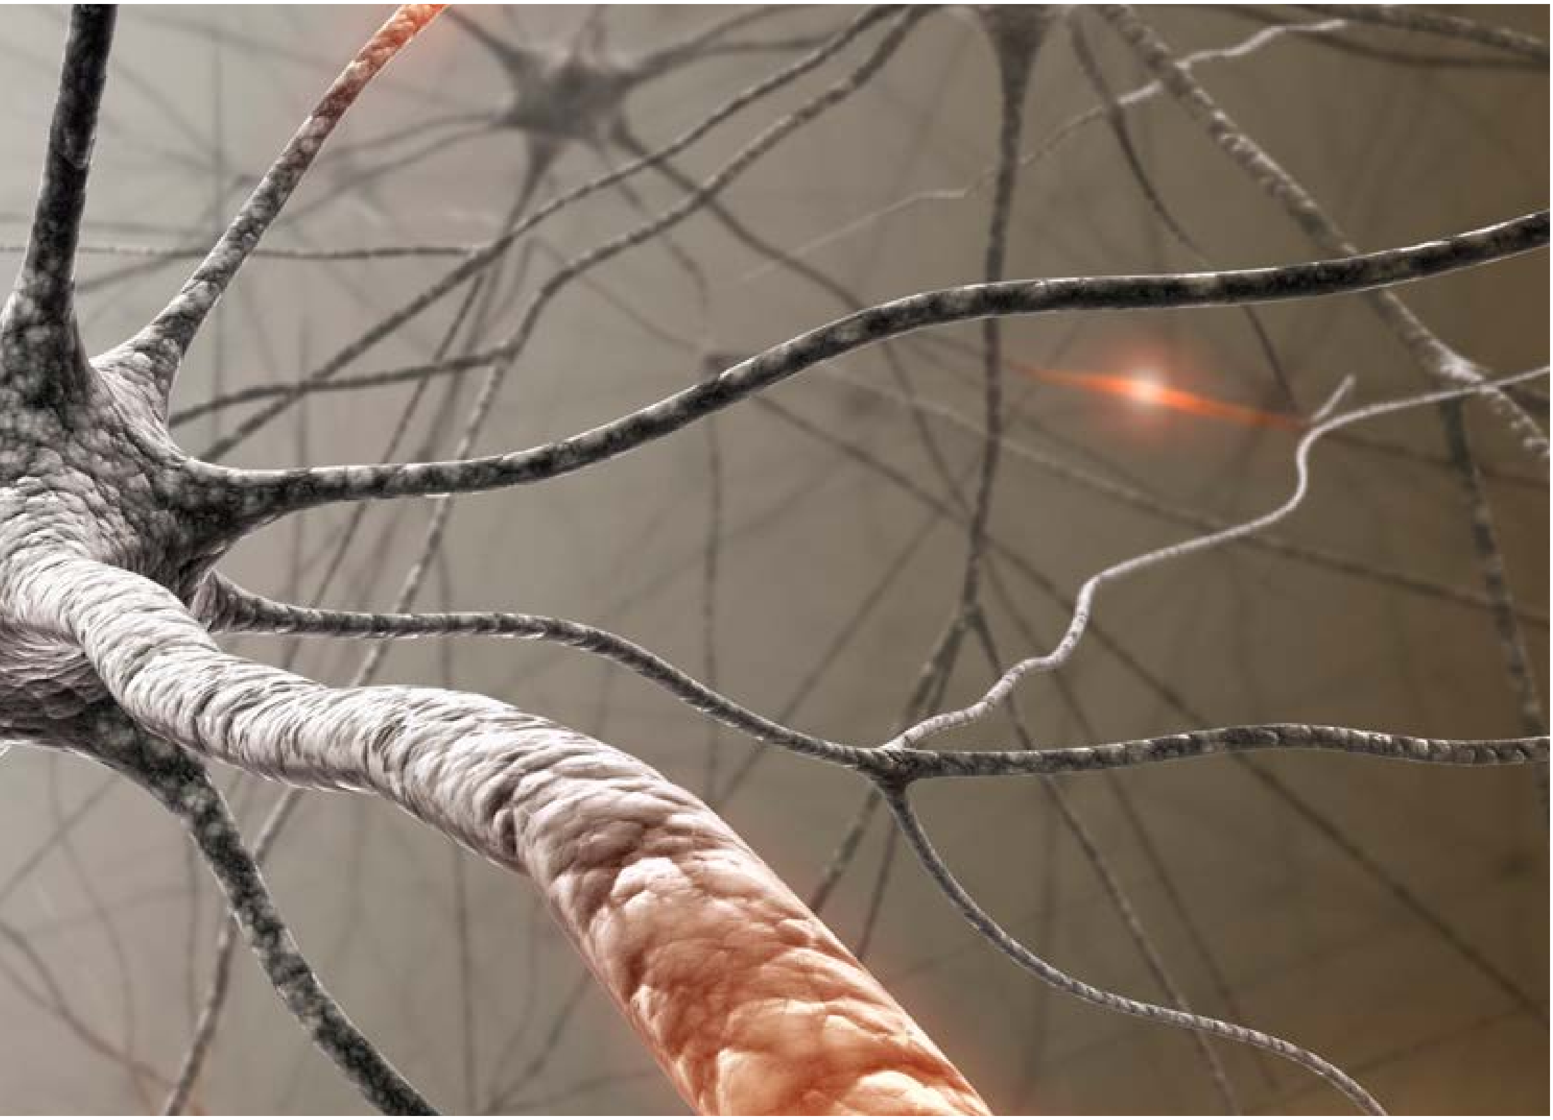

Despite the past efforts to discover the schizophrenia gene, the genetic background of the mental disorder has long remained obscure. However, increasing amounts of genetic data have become available for researchers over the past ten years, meaning that larger populations of people have been genotyped in more detail than before. Since 2007, wide international collaborations to join genetic data from different countries have given more insight into the genetic background of schizophrenia. A recent genome-wide association study from 2014 showed that hundreds of genes are involved in the risk of schizophrenia, which confirmed the view that the disorder cannot be explained by just one or a few genetic mechanisms.

COMMON GENETIC VARIANTS CONTRIBUTE TO SCHIZOPHRENIA

All humans have the same set of genes, but there are variations in the genetic code (amino acid sequences) of these genes. There are millions of locations along the human genome (DNA), where the genetic information varies across people, making us different from each other. Simplest of these variations are single nucleotide polymorphisms (SNP or "snip"), where a single amino acid (A, C, G or T) is changed to one of the three alternatives without affecting the rest of the sequence. As an example, the aforementioned study found 128 SNPs that were statistically more frequently present in schizophrenic population than in healthy controls. Some of these SNPs may not affect the translation of any known gene, whereas others may contribute to many different genes.

The changes in the genome are conserved when the gene is translated into a protein, and the function of the protein might accordingly be changed. A great challenge in genetics is to understand how the gene variants in certain genetic locations lead to the development of the disorder and to the expression of the symptoms. This is especially true for schizophrenia, which involves interplay of many genes and is characterized by complex symptoms, such as hallucinations and delutions. In other heritable diseases, where the outcome may depend on a single rarely occurring gene variant, there are numerous experimental studies describing the effect of this "mutation" both on the function of a single cell and further on the system function. Such studies form the field of functional genomics.

BRAIN GENES ARE HIGHLIGHTED IN SCHIZOPHRENIA

An intriguing discovery from the recent genetic studies of schizophrenia is that multiple risk genes encode for ion channels. The ion channels are trans-membrane proteins that are important in controlling the ionic concentrations inside a cell. Especially, they are important for signal transmission between neurons in the brain. It is therefore of importance to understand how changes in the identified genes could underlie altered brain function in the disorder.

FOCUS

Our effort is an early attempt toward understanding the disease mechanisms of polygenic psychiatric disorders by computational means. It binds together genetic and single neuron levels of abstraction in the big picture of modelling psychiatric illnesses.

Traditionally, the genetic mechanisms underlying a disease have first been studied in transgenic animals, i.e. a mouse with an engineered gene mutation. This, however, requires large resources from the experimenting laboratory. Computational modelling has recently grown to a popular method of creating hypotheses that can first be tested in animal or cell experiments and then refined to serve drug development. Once a computational model has been built and tested, it can be effectively used to produce detailed information of the function of the studied system, both in normal and altered conditions. Therefore, although collecting enough data to build a computational model of a biological system demands a lot of resources, it is a low-cost way of performing the following early steps of the research.

COMPUTER SIMULATIONS IN THE STUDY OF THE BRAIN

Neuroscience has a relatively long history in applying computational methods. Traditionally, computational neuroscience has studied the interplay of neurons in coding and transmission of information using relatively simple models, e.g. integrate-and-fire models. In these models, the state of the neuron is described by a single value of the membrane potential, the difference in electrical potential between the inside and outside of the cell. If the membrane potential crosses a threshold, the neuron fires an action potential that is transmitted to the connected neurons. On the other hand, the pioneering work by Hodgkin and Huxley in the 1950's to describe the biological details of single-cell action potential generation has been applied and further developed by neuroscientists to bring in more biophysical details into the neuron modelling. Furthermore, the cable theory developed by Wilfrid Rall in the 1950's and 1960's to describe the propagation of electric signals in neurons has been effectively used in neuroscience, leading to standardized ways of reconstructing whole-neuron morphologies and using them in computer simulations.

RISE OF MEGA-SCALE COMPUTATIONAL NEUROSCIENCE PROJECTS

To date, there are many computational models of single neurons involving complex morphologies and detailed description of different types of ion channels, which allow the study of different aspects of neuron biophysics and their impact on neuronal network behaviour. This has occurred in parallel to several large-scale efforts, such as the Blue Brain Project (bluebrain.epfl.ch), the Human Brain Project (www.humanbrainproject.eu), the B.R.A.I.N initiative (braininitiative.nih.gov), and Project MindScope at the Allen Brain Institute (alleninstitute.org/our-research/research-science/mindscope) which contribute to both building these biophysically detailed models and applying them to increase the know-ledge on how the brain works.

BIOPHYSICAL NEURON MODELLING CAN BRIDGE THE GAP BETWEEN GENETICS AND BRAIN DYNAMICS

The core of our effort is to apply the available biophysically detailed neuron models to the study of schizophrenia. We aim at combining data from genome-wide association studies and studies of functional genomics, where changes in ion channel-encoding genes have been shown to cause an effect in the properties of the ion channel or its subunits. These changes are implemented in our neuron model to study the behaviour of the altered neuron in different situations. The model neurons can further be connected to each other to form a model of a neuronal network, which will help in translating the results found in single neurons to network behaviour.

LARGE-SCALE NEURON MODELLING

While the simulations of single neurons can usually be performed on a regular PC, the simulations of networks of neurons can demand large amount of computation time and memory. Nevertheless, these network simulations can often be parallelized to different computers to facilitate the computations.

The simulations consist of calculating numerical solutions to sets of thousands or millions of ordinary differential equations, where each equation describes the dynamics (for quantities such as membrane potential, ion concentration, or degree of activation or inactivation of a specific ion channel) in a small section of a neuron. Moreover, when building or modifying these models, they have to be simulated numerous times with different parameters in order to find a parameter set that brings up optimal model behaviour. Typically, software specifically designed for neuron simulations is used to gain maximum simulation performance.

KEYS TO BETTER UNDERSTANDING OF SCHIZOPHRENIA

Our effort is an early attempt toward understanding the disease mechanisms of polygenic psychiatric disorders by computational means. It binds together genetic and single neuron levels of abstraction in the big picture of modelling psychiatric illnesses. In order to lead to breakthroughs in the treatment of mental disorders, it should be complemented with appropriate animal and human studies that allow the design of novel medications and therapies. Although the field advances in small steps, the future looks bright in this dawn of biophysical psychiatry – an era of mechanistic modelling of mental disorders.

NORMENT:

Norwegian Centre for Mental Disorders Research

- NORMENT KG Jebsen Center for psychosis research is a Centre of Excellence established by the Norwegian Research Council of Norway in 2013.
- NORMENT is integrated with the K. G. Jebsen Centre for Psychosis Research, funded by the K.G. Jebsen Foundation.
- The center has the University of Oslo as host institution, with the University of Bergen and Oslo University Hospital as partners.
- The centre has about 85 employees and is headed by Professor Ole A. Andreassen.
- The research is organized into eight interdisciplinary research groups:
  - Translational Psychiatry
  - Clinical Psychosis Research
  - Psychiatric Molecular Genetics
  - Brain Imaging, Cognition
  - Neurocognition
  - Structural MRI
  - Convergent functional genomics
  - Animal Models

"At the Norwegian Centre for Mental Disorders Research we strive to find answers to why some people develop perceptual disturbances, delusions, deep depressions or elated, manic phases".
www.med.uio.no/norment

LITERATURE

Hodgkin, A. L.; Huxley, A. F.
A quantitative description of membrane current and its application to conduction and excitation in nerve.
J Physiol. 117: 500–544 (1952).

Markram, H.
The human brain project. Scientific American, 306, 50-55 (2012).

Rall, W.
Branching dendritic trees and motoneuron membrane resistivity. Experimental neurology, 1, 491-527 (1959).

International Schizophrenia Consortium.
Common polygenic variation contributes to risk of schizophrenia and bipolar disorder. Nature 460, 748–752 (2009).

Schizophrenia Working Group of the Psychiatric Genomics Consortium.
Biological insights from 108 schizophrenia-associated genetic loci.
Nature 511, 421-427 (2014).

Photo: Yngve Vogt

Exit interview with Hans A. Eide

- previously head of the Department for Research Computing at USIT

Having been engaged at USIT's Department for Research Computing at the University of Oslo for several years, Dr. Hans A. Eide has been a central figure in the development of the e-infrastructure in Norway. After ten years, he has decided to step down from his position as department head and and relocate to Stavanger. Hans has valuable insight into the e-infrastructure in Norway, IT service needs of user communities and Norwegian participation in international collaborations in this field.

AUTHOR

Vigdis Guldseth
UNINETT Sigma2

Q: What was the reason you decided to leave the position as department head and move to Stavanger?

It was purely a family decision. It has been the plan for a long time, and some important pieces suddenly fell in place. Time goes fast when you have fun, and if you don't act, you may one day find out it is too late.

Q: Tell us about your background, what did you do before you started at USIT, how did you end up in HPC?

I have a background as a researcher, as many of us in the e-infrastructure business do. I did my master at the physics department, University of Bergen. From there I moved to the University of Alaska, Fairbanks, and obtained my PhD in atmospheric sciences at the Geophysical institute. Then I moved, you could say, from one extreme to the other; to the greater New York City area and Stevens Institute of Technology in Hoboken, where I was at first a post-doc and later research associate professor.

Already in Bergen I got acquainted with computational physics, large UNIX systems and early Linux. I was modelling electromagnetic waves, light scattering and radiative transfer. When in Alaska, I was the typical go-to person for all of my colleagues in the lab when it came to computational problems. I was basically the system administrator and one-man research support operation. I enjoyed, and still do, solving other people's computer problems and making things work.

This continued during my period as a researcher, but when the time came to move back to Norway, which was a family decision as well – there were no openings in my field. Nothing relevant turned up for months, and I started to look outside research and at alternative options. I sent an application to USIT and the rest is history. Today I'm really happy that this was how it turned out, and I have enjoyed every minute of this work since I started. I still try to stay on-top of research developments, and I tinker with electronics, home automation and Linux system programming, in my spare time – which haven't been much in later years, unfortunately.

Q: You have experience as a researcher, how has that affected your work in HPC and provider of IT services for researchers?

I think it is crucial to be able to talk to researchers with their language, to understand the research process and to have seen how HPC, or IT in general, can be leveraged in a research project. I also believe researchers listen to our advice and bring us into their projects more easily when they know that the IT support person has research experience. At the Department for Research Computing we have always been striving to recruit people with their own research background, and

I'm confident that my successor, Gard Thomassen, and all my colleagues at the Department of Research Computing at USIT, will keep up the good work and bring new and even better e-infrastructure services to our researchers.

we try to diversify the fields that we cover with our in-house experts in this way. At present we have 14 colleagues with PhD in the department. Also, my own research background has made me quite experienced in writing proposals for projects or for funding. That is always handy.

I also like to add that having a research background gives me an extra motivation for what we are doing. I feel more like a collaborator than an IT-person. What we do for researchers directly supports important research that make our lives healthier, safer and better.

Q: What was the greatest achievement of the previous e-infrastructure company?

I don't think there is a single great achievement that I would put above everything else, but looking back there are series of important steps taken towards where we are today. Establishing NorStore and utilizing the Advanced User Support mechanism to develop computational portals, for example. We also had some memorable and engaging Notur conferences, in some ways we were closer to the users and the user communities when there were fewer of them.

Q: In your opinion, what should be the major goals for the further development of the national e-infrastructure?

I think we need to address more broadly the needs of researchers in the research process. I often say tongue-in-cheek that HPC is solved, and by that I mean that we have reached a level of proficiency in this field, and need to address new areas and address more of user's needs besides HPC, be ahead of the curve. There is a huge amount of researchers in no need for what we can call classical HPC, that have never used the command-line or even know how to program. However, they still need to collect data, analyze data, visualize data. Collaboration is ever more important, researchers need to share data and work on them together, at the same time and regardless of where they are located. Researchers need to discover existing data and reuse them, they need to be able to reproduce each other's results, and so on. These are some of the grand new challenges.

Therefore, in my view the major goals for the further development of the national e-infrastructure should be to broaden the scope of what we do. That is, we should listen to the user's needs and make available easy-to-use services to support the research process. It is necessary to build our own competence and be proactive in new areas, for example in data-driven research. We need to tell our researchers about new methods and technologies, and then help them start using these. We need to help building the competence and skills of our researchers.

Today we are used to on-demand services, that things are immediately available when we need them. We can't have services and resources that one might need to wait a long time after an application process to get access to. This is especially important for busy researchers on a tight schedule, and in my opinion cloud-based services and portals are a way to go. We should therefore start thinking about what we do as RSaaS Reseach Services as a Service. Finally, we must improve the marketing effort and show that we make a difference, by highlighting the success stories. This will lead to in e-infrastructure being considered a strategic resource rather than a passive resource, which I think is crucial. Only this way can we increase our user-base and funding, and bring the national e-infrastructure to the level it should be.

Q: The core activities of the e-infrastructure services in Norway are organized through a national Metacenter. What do you consider to be the most important strengths of the Metacenter?

Norway is a small country, and there are not many with the competences and skills needed to support all types of research and operate the national e-infrastructure. Sharing this competence, sharing different types of resources and working together is a must, and can be characterized as the strength of the Metacenter. Even though we live in a wealthy country, we are not able to support all kinds of research and work-loads locally at each university. The Metacenter makes all this possible by enabling the sharing of competence and resources. In addition, the Metacenter facilitates our participation in international e-infrastructure projects and the collaboration across borders, which we get a lot in return for and something I find very important.

Q: Where do you see the weakest spot?

As I mentioned earlier, researchers need more or less immediate access to both support and resources. The way we tend to do things, and many of the services we offer, are introducing delays and the offerings are sometimes hard to navigate and understand. Researchers often need other

types of support beyond classical HPC and data storage. I think the Metacenter has a weak spot in being a bit "remote" for many users. By that I mean that many researchers perceive a gap between their needs and what we offer. This is a recurrent complaint from researchers who struggle to fill out a Notur application form the first time.
I hope these things can be addressed as we go forward with the new organization. Our shortage of resources and limited scope may of course be considered additional weaknesses. In the Metacenter we do not have much in support for humanities and social sciences, for example statistics, qualitative methods, data harvesting from social networks to mention a few. There would also be little in support for research involving sensitive data if it were not for extra effort from the University of Oslo, Oslo University Hospital and several research groups.

Q: We are in the middle of a large procurement process, the first in the new organization where we are all in it together. How do you see this process so far?

The current procurement process is extremely important and will in many ways decide how several things will look for many years to come. We are doing this at the same time that we are redesigning how we do operations and how we do user support. It is unfortunate, then, that we do not have enough time and resources to do all of this as thoroughly as we, in my view, should. We might end up learning some very expensive lessons.

The idea of doing strategic investments in a national and coordinated perspective, where the intention is to reduce operational costs, to maximize research production, operating the infrastructure together and putting the user experience in the center, are all very good. The thing is, in order to simplify the complex problem of realizing these goals we are not doing the right thing, which would be to look for solutions that optimize *all* of these these parameters at same time. Instead we are reducing the complexity of the problem by considering separate sub-problems, like housing, and end up optimizing in compartments. As an example, consider the task of minimizing operational costs. It is easy on paper, we need as few sites as possible, but this solution is not optimum with respect to other costs and factors, such as need for initial investments in data-center capacity, but most importantly, it might not be optimum with respect to research production.

Different computational needs are most cost-efficiently tackled on different HPC resources. It is possible that it could be better, give more research over total costs, to have more than two HPC systems, and let the systems be technically different and more tuned to the applications in important aspects. By doing this we would probably neither have to do expensive initial investments in data-centers to house large systems, and could instead utilize already existing data-center capacity. Research is very dynamic and needs are changing fast. To be able to change and scale quickly I think we ultimately must take advantage of the large-scale external data-centers that are out there, where we would be more free of such constraints.

Q: How do you consider the balance between national and international e-infrastructure resources?

In my opinion, the current balance is good in theory, but maybe not so good in practice. The purpose of PRACE, for example, was to cover the needs of our largest users (in terms of scale), but we have had only A few projects covered by PRACE. With EUDAT the problem is rather that we struggle with the relationship to our own national research data, for example the role of NorStore versus NSD. Due to this, I believe that we are not all in, nor are we all out of EUDAT. This makes it difficult to get enough resources and motivation to participate and to really get our users in on it. One success story of a sufficient balance between national and international e-infrastructure resources is CERN and the WLCG. The balance is also good when considering how we, the Metacenter, benefit from participating in international e-infrastructures. We gain competence, learn from other's experiences, learn about new services and so on. This is true for both PRACE and EUDAT, and I'm sad to see EGI go now that we should be looking into more cloud-based services.

Some have said that our participation in international e-infrastructure projects in a sense steal away resources from the Metacenter and our national e-infrastructure. I couldn't disagree more. There is a great amount of synergies from this activity, there is external funding and this activity provides access to resources, for our researchers as well as to us as an e-infrastructure provider. There is a great amount of synergies from this activity, a whole lot of resources for our researchers to access, and these activities have made it possible to recruit more skilled people, including from abroad, and to build a critical mass of diversified competence and human resources.

Q: What are your plans for the future?

I think it is reasonably safe to assume that I will continue to work with HPC, research data and research services. Exactly how this will work out remains to be seen, but I will stay on at USIT as an adviser and be a part of the Metacenter for some time to come. I am now in a position to contribute in different and maybe more important ways than before, and I'm confident that my successor, Gard Thomassen, and all my colleagues at the Department of Research Computing at USIT, will keep up the good work and bring new and even better e-infrastructure services to our researchers.

FOCUS

We need to tell our researchers about new methods and technologies, and then help them start using these. We need to help build the competence and skills of our researchers.

The Stallo compute cluster at UiT The Arctic University of Norway will be replaced by a new HPC-machine in 2018. Photo: Thilo Bubek.

Renewal of the national e-infrastructure for science

AUTHOR

Arild Halsetrønning UNINETT Sigma2

One of the important objectives for the new UNINETT Sigma2 company is to have a strong focus on procurement of new HPC and storage equipment. All of the current Notur HPC facilities were procured during 2012 and are expected to reach the end of their lifetime by 2017/2018. The NorStore data storage resources will also need to be replaced in the following years.

System	Type	Capability (cpu-hours)	Performance (Tflop)	Production start
Abel	Capacity	76413480	181,6	October 2012
Hexagon	Capability	102807360	107,9	April 2012
Vilje	Capability	113012760	268,3	October 2012
Stallo	Capacity	78804960	195,7	October 2012

UNINETT Sigma2 therefore started a pre-study for a procurement project in April for new HPC and data storage equipment.

The pre-study had amongst others, the following objectives:

- collect user requirements
- propose suitable technologies
- consider pros and cons of having one or several HPC and storage facilities
- evaluate if external (of universities) housing facilities can be beneficial
- propose future operation model(s) for e-infrastructure.

In the process of collecting requirements a survey was sent to all current Notur and NorStore projects, and potential new user communities, requesting input concerning existing and future service requirements related to the national e-infrastructure. In addition, meetings were arranged with some research communities to inform about the plans and elicit more detailed information from the communities.

DATA-CENTRIC INFRASTRUCTURE

Several of the communities that were interviewed expressed a need for improving the environments for the development and testing of new software. One consequence of this is a decision to establish a small low entry system in addition to the production systems to facilitate this request. A better integration of services and an expansion of services offered was also requested by many user communities, resulting in a proposal to establish a data-centric infrastructure. The goal is to achieve a single infrastructure platform where all research data can be maintained while offering an extended portfolio of relevant services, including data analytics, visualisation, community specific services and the traditional HPC provisioning. The infrastructure should also enable the integration of external data sources such as research labs and facilities. Figure 1 illustrates this concept in the future national infrastructure.

PLAN FOR INVESTMENTS

During the autumn of 2015, the UNINETT Sigma2 board has made decisions relating to an investment and development strategy. It has been decided to reduce the number of HPC systems from four to two and to establish a data-centric infrastructure. The initial data infrastructure storage system, the first HPC production system and HPC entry level system will be placed at UiT The Arctic University of Norway. The new HPC production system will replace Hexagon and Vilje.

The HPC tender will be announced in January 2016, with a pre-announcement in December 2015. The goal is to have the first HPC system in production in Tromsø early 2017. The Storage tender will be announced in February 2016 and new storage equipment ready for operational service in early 2017.

Early 2017 UNINETT Sigma2 plans to start another procurement project, aiming at replacing Abel and Stallo. The goal is to have this new HPC system in production early 2018. After this the strategy will be to replace each of the two HPC machines every fourth year.

UNINETT Sigma2 will be the owner the new equipment and be accountable for support and operation. Personnel from the universities will still be responsible for the day-to-day support and operation.

The housing of the second HPC system has not yet been decided. The alternatives are UiB, NTNU or somewhere external, i.e. hiring facilities at a commercial housing center. This decision is expected during spring/early summer next year.

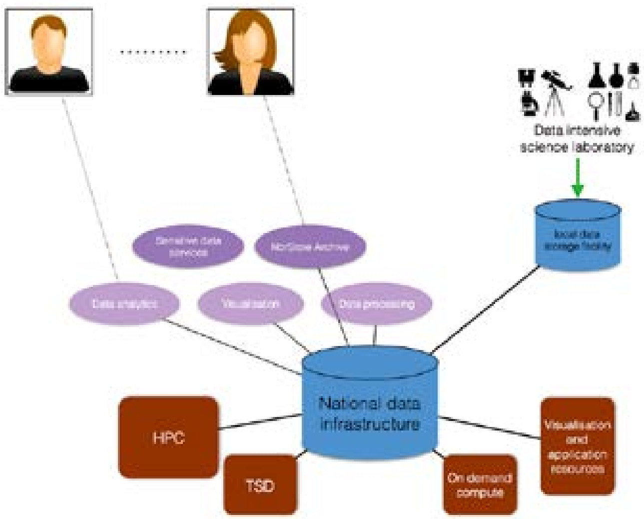

Figure 1:
Data centric infrastructure

Call for e-infrastructure resources

The deadline for applying for applications for period 2016.1 is 26 January 2016. Earliest project start is 1 April 2016

The current allocations by UNINETT Sigma2 of the High-Performance and Storage facilities of the Notur and NorStore projects expire 31 March 2016.

- ***Projects who want to continue using the facilities after this date, must apply to this call.***
- ***Applicants that already applied for access for the period 2016.1 in a previous call do not need to apply.***

- ***For more information, please visit www.sigma2.no***

ONLINE APPLICATIONS

An on-line application procedure is used. There is a new application tool service for researchers using the national e-infrastructure resources in Notur and NorStore. Old applications submitted in previous periods will be available directly from the new application forms. The online application form is available on www.metacenter.no/mas/application/project

WHO CAN APPLY

Individuals and groups from the following Norwegian organizations can apply for access to the resources in the national e-infrastructure:

- Norwegian universities and university colleges
- Meteorological Institute
- Research centers
- Projects or institutions funded by the Norwegian Ministry of Education and Research or the Research Council of Norway (RCN)
- Industrial parties who contribute to the funding of the e-infrastructure

The applicant must hold a permanent position at his/her organization. Master students and PhD students cannot apply directly, but must do this via a colleague (e.g., supervisor) that holds a permanent position at the organization.

The resources in the national e-infrastructure are available for publicly funded research and education. Up to 20% of the infrastructure can be offered to the commercial sector or other parties, in which case actual costs will be charged.

APPLICATION EVALUATION

The Resource Allocation Committee assigns allocations to projects a few weeks before the start of each allocation period. It is therefore important to submit an application during the regular calls. Applications that are received outside the regular calls, can usually only be granted a limited amount of low-priority processor core hours or may have to be allocated on facilities that can be considered suboptimal. Request for storage resources are not guaranteed and only limited amount of resources can be offered between calls.

The decisions of the Resource Allocation Committee are guided by the Guidelines for Resource Allocation on the National e-infrastructure.

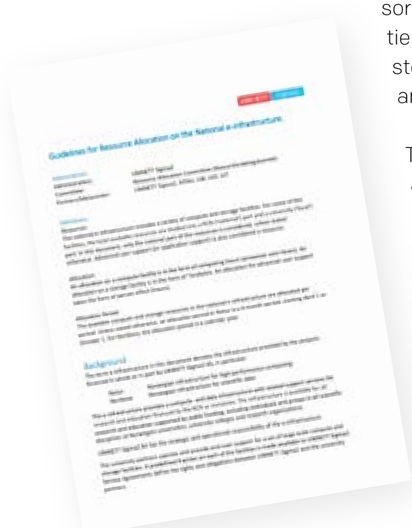

Guidelines for Resource Allocation of the National e-infrastructure.

Advanced user support for HPC and data storage

Scientists can apply for advanced support by sending an application to UNINETT Sigma2.

- ***Download the project description at www.sigma2.no/content/how-apply-aus.***
- ***The deadline ordinarily coincides with the deadline for application for computer time, biannually at end of January and end of August.***
- ***Science communities interested in discipline specific AUS are encouraged to contact UNINETT Sigma2 for more information.***
- ***For more information, please visit www.sigma2.no***

Advanced user support is an important part of a HPC environment and for the more recent area of data archives and open access. Both these areas are highly specialized and in order to utilize the resources in the best way, advanced user support is often needed.

WHAT IS ADVANCED USER SUPPORT?

Advanced user support is special support given to scientists using the HPC services in Notur or storage services in NorStore or a combination of these. The intention is to help the scientists to make better use of the services and hence create better science.

For the HPC services, advanced user support aims at helping scientists to improve or extend the performance and capabilities of their applications. This can be done in a number of ways, including code parallelization, porting, optimization, benchmarking and improving user-interfaces. AUS is also for scientists in need of high-end resources, but are unfamiliar with parallel computing architectures and parallel programming languages. For the storage services, advanced user support aims at helping scientists to create data plans, classify data as sensitive or not, define meta-data, finding good storage formats and so on.

DIFFERENT TYPES OF AUS

Advanced user support is classified into Project-based AUS and Discipline specific AUS.

Project-based AUS can be the sole initiative of a researcher, a researcher group or a scientific community. Project- based AUS will be granted through the Research Allocation Committee with 2-3 PMs spent over a maximum of 6 months. The expected results should have value and benefits beyond the project itself. An exception is for scientific disciplines new to e-infrastructure services in need of help to start using the e-infrastructure.

Discipline specific advanced user support can be initiated, in cooperation between Sigma2 or the universities, with a scientific community, focusing on the need for more long- term support. The Resource Allocation Committee does not govern the allocation because it is considered a strategic area where dedicated long term support is needed and has board approval. Discipline specific AUS can have allocations of more than 12 personnel months spent over a maximum for 2 years.

Photo: Alexandr Oltu

Expect limited HPC resources

With the expansion of the RFK share of Notur with 4800 cores prior to the computing period 2015.2, starting 1 October 2015, 208 million core hours is available for allocation during this 6 month allocation period. The Gardar system will be phased out by 1 January, and until new hardware is acquired, the RFK capacity will be 204-205 million core hours per allocation period. A new compute resource is not expected to be in place before beginning 2017. This system will be hosted by the University of Tromsø, and it will also replace the current Vilje and Hexagon systems.

AUTHOR

Jørn Amundsen UNINETT Sigma2

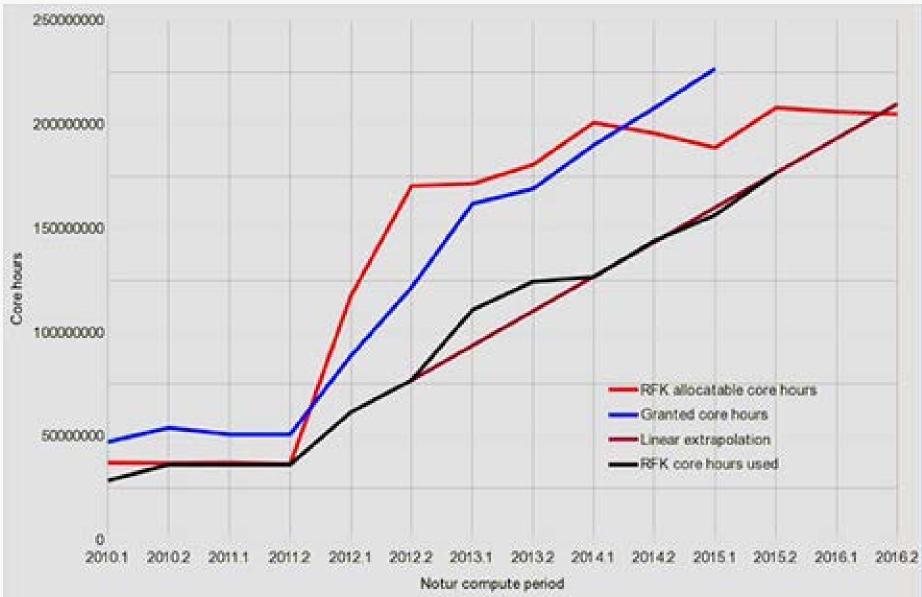

With an average demand growth of approximately 9%, the total available capacity will be exceeded end 2016.1, and the 85% capacity line will be exceeded in 2015.2.

To the extent possible, UNINETT Sigma2 will try to avoid expanding the capacity of current systems before 2017. The current hardware infrastructure was acquired in 2012, and expanding any of those will require old systems to be operated after their planned end of life, resulting in very high maintenance costs compared to a new system. Also, expanding an old system gives much less core hours per currency unit invested compared to a new system.

With an average demand growth of approximately 9%, the total available capacity will be exceeded end 2016.1, and the 85% capacity line will be exceeded in 2015.2. On a system with checkpoint-restart queues (and C-R jobs), 90% of the capacity is considered being the maximum attainable utilization, while 85% is considered the maximum attainable without a checkpoint-restart service.

UNINETT Sigma2 will consider the following measures for a better utilization of installed capacity, before resorting to expansions:

- introduction of prioritized and non-prioritized grants
- provision of checkpoint-restart queues for applications with internal checkpointing capabilities
- moving allocations to other systems for better overall utilization
- closer monitoring and follow-ups of inefficient usage

By the end of November, 230 million core hours is granted for 2015.2, an overbooking of 10%.

The amount of allocations requested for both Notur and NorStore indicates that the demand for resources continues to increase.

Notur and NorStore in statistics

At the end of 2015 there are more than 1100 active e-infrastructure users altogether making use of the Notur supercomputing and NorStore data storage.

There are respectively 208 computing projects making use of the Notur facilities and 76 storage projects on NorStore. The amount of allocations requested for both Notur and NorStore indicates that the demand for resources continues to increase.

In particular, researchers in climate modelling, life sciences, materials technology, physics, chemistry, geosciences and language technology have a wide-ranging need for reliable computing power and data storage.

GEOSCIENCES AND PHYSICS CONSUMES THE MOST COMPUTER TIME

Figure 1 shows the overall number of projects by discipline on the HPC-facilities in Notur in 2015. The HPC-customers from the scientific fields of geosciences and biosciences are the largest group of researchers in need for computing resources, representing respectively 33 and 27 of the allocated projects. These are followed by chemistry (25), physics (25) and computational fluid dynamics (16). There is a growing trend that non- traditional natural disciplines, such as medical sciences and social sciences, are slightly increasing their presence.

The size of computing projects in terms of consumed CPU-hours has grown year after year. In the ongoing allocation period, the largest individual computation project is Solar Atmospheric Modelling, led by University of Oslo Professor Mats Carlsson. The astrophysics project will consume 35 million core hours in the ongoing allocation period. All of Notur's top 25 projects are in the millions-range in CPU-hours of computing projects.

For Norstore, the number of storage facilities by scientific field are dominated by geosciences (21) and biosciences (10). These disciplines are followed by computational fluid dynamics. In terms of storage customers, researchers in terms of climate modelling are by far the largest group in need for storage resources. Furthermore, the national service for archiving and sharing large-scale research data in February 2014 has resulted in an increased number of users. Similarly, services for sensitive data has also expanded the NorStore-community.

AUTHOR

Vigdis Guldseth
UNINETT Sigma2

NOTUR'S TOP 10 PROJECTS

Project	Project leader
Solar Atmospheric Modelling	Mats Carlsson
Periodic boundary calculations of metal-organic frameworks	Trygve Helgaker
NorClim - Climate of Norway and the Arctic in the 21st Century	Christoph Heinze
Mechanical properties of nano-scale polymer particles by large-scale molecular dynamics	Zhiliang Zhang
Hydrogen production in a high temperature membrane reactor: advanced catalysis from atoms to process	De Chen
Seasonal-to-decadal prediction of climate in the Atlantic sector	Noel Keenlyside
Impacts of Sea-Ice and Snow-Cover Changes on Climate, Green Growth and Society	Noel Keenlyside
Quantum chemical modelling of catalysis and spectroscopy	Knut J. Børve
3D forward modelling of lithosphere extension	Ritske Huismans
Simulations of microscopic systems	Ladislav Kocbach

Figure 1

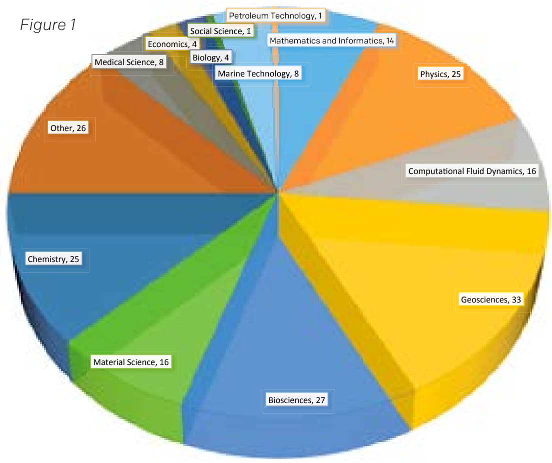

Hazel Hen is Europe's fastest supercomputer on the High Performance Conjugate Gradient (HPCG) benchmark list. This was announced at the Supercomputing Conference 2015 in Austin, Texas. © Boris Lehner for HLRS

Accelerator node investments in the next generation system

AUTHOR

Jørn Amundsen, UNINETT Sigma2

It is hard to accurately predict the need for accelerator nodes in future HPC systems. In particular, there is a risk of acquiring too many nodes compared to what can be exploited with the current software stack and usage pattern. The current software stack is very diverse, and specific staffing for accelerating and maintaining accelerated applications is not available. In this regard, we are highly dependent on international efforts within open source and 3rd party vendors.

Looking to other European HPC installations, it appears that international sites also are cautious on installing a larger accelerator share initially. The most recent top European HPC installation, the Hazel Hen system ranked number eight on TOP500 (second in Europe), appears to be a system with CPU nodes only.

When installing the current Notur hardware generation in 2012, four GPU nodes were added to the Abel system in Oslo and integrated with the queueing system. During the first years, the GPUs experienced very little use, and it is only during the last year we have seen significant constant load, and even queues. During the last two months, the average utilization has been 70%. The GPU software stack is very limited, however, and the load is dominated by the Amber and NAMD applications.

Notur GPU enabled applications

- ADF
- Amber
- ANSYS Fluent
- BarraCUDA
- Beagle
- Blast
- CP2K
- GAMESS-US
- Gaussian (in development)
- GROMACS
- LAMMPS
- LSDalton
- LS-DYNA (beta version)
- MATLAB
- Molcas
- MrBayes
- NAMD
- NWChem
- OpenFOAM (GPU library plug-in)
- Quantum ESPRESSO
- S3D
- VASP
- WRF

By comparing the Notur software stack with available GPU enabled applications, the following applications emerge.

The potential usage footprint of this collection is quite significant, although it depends on the functionality provided in the GPU enabled vs the CPU versions of the applications.

Notur MIC enabled applications

- Amber
- Blast
- CESM (in progress)
- GAMESS-US
- GROMACS
- Harmonie (in progress)
- LAMMPS
- MATLAB
- NAMD
- NWChem
- OpenFOAM (in progress)
- Quantum ESPRESSO
- R
- VASP (in progress)
- WRF

The MIC architecture is not that mature as a platform, and the list of MIC enabled applications present in Notur is hence shorter.

BIFROST

In order to possibly expand on future accelerator footprint, UNINETT Sigma2 has launched a strategic project to investigate on the acceleration of the stellar atmosphere simulation code Bifrost. Bifrost is the main application in the largest Notur project, NN2834K, «Solar Atmospheric Modelling». This project has applied for 85 million CPU-hours during the last year, and has been able to consume 54 million CPU-hours. Furthermore, with the advent of a new compute resource late 2016 or early 2017, it is important to serve the needs of this project with the best possible cost efficiency.

Bifrost is a MPI application scaling very well, and has been executed on 64000 cores on a US HPC system. This proves that the application possess a large degree of parallelism, making it a candidate for both future MIC and GPU architecture execution. The goal of the strategic project is to try to assess cost efficiency of Bifrost on accelerators on our next generation system, as compared to CPU execution. This might justify investing in a significant initial share of accelerator technology, contrary to investing in accelerators for development only.

NN2834K - Solar Atmospheric Modelling

Source: www.sigma2.no/user-projects

Abstract

Solar magnetism lies at the root of most solar and heliospheric physics. The intricate structure of the solar field, the activity cycle and the influence of the field on the heliosphere represent major quests of (astro-) physics which bear directly on the human environment. The Sun's magnetic field is generated by enigmatic dynamo processes in the solar interior, is organized into the highly complex patterns of solar activity observed in the solar photosphere, dominates the structure of the outer solar atmosphere (chromosphere, transition region, corona), regulates the solar wind, and affects the whole extended heliosphere into the Earth's upper atmosphere. Solar activity modulation affects satellite orbits, influences jet stream patterns and contributes to the causes of minor, and possibly major, ice ages. Our group is world-leading in modelling the solar atmosphere as one system; from the convection zone where the motions feed energy into the magnetic field and all the way to the corona where the release of magnetic energy is more or less violent. The computational challenge is both in simplifying the complex physics without loosing the main properties and in treating a large enough volume to encompass the large coronal structures with enough resolution to capture the dynamics of the system. We have developed a massively parallel code, called Bifrost, to tackle this challenge. The resulting simulations are very time-consuming but crucial for the understanding of the magnetic outer atmosphere of the Sun.

Space Research

Source: E-INFRA 2014 A national e-infrastructure for science

Space Research in Norway is concentrated around two areas: Solar-terrestrial physics, with an emphasis on knowledge about the fundamental processes of the sun and the solar atmosphere and on how solar wind and solar activity affect the global environment. The development of the universe, with an emphasis on knowledge about fundamental astrophysical processes.

In both areas the use of high-performance computing is vital. One example is cutting edge models of the solar atmosphere where Norway is world leading and one of the biggest users of computing resources in Notur. High resolution is needed to catch the essential driving mechanisms that take place in the solar convection zone at small spatial scales. At the same time it is important to have a large enough computational domain to describe the outer atmosphere phenomena under study. This leads to very large computational volumes: modelling one hour of solar time takes several million CPU hours and this is only possible through parallelisation over several thousand cores (the codes have been run on up to 64000 cores). There are several other examples in space research where Norway is in the absolute research front internationally (8 papers in Science and Nature since 2011). The national and international collaborations are extensive with international infrastructures partly hosted in Norway (EISCAT) and in space (ESA and NASA satellites) and a number of international projects including two ERC Advanced grants and one ERC starting grant.
From E-INFRA 2014

Graphical user interface for project leaders

Photo: Thinkstock

Project leaders in Notur and NorStore may now administrate their national e-infrastructure projects from UNINETT Sigma2s project leader - GUI: www.metacenter.no

As a Notur or Norstore project leader you are the responsible for administrating your Notur and/or Norstore-projects. This involves:

- being the overall responsible for the project
- applying for (extra) resources
- signing user applications
- extend user access
- add users and administrate user lists.

The project leader graphical user interface allows you to administrate all these tasks. Project leaders may also download the grant letter for their project(s).

USAGE STATISTICS FOR NOTUR PROJECTS

In addition to the available services in the project leader graphical interface, UNINETT Sigma2 has also recently developed and made available detailed usage statistics for Notur projects.

Project leaders and their assistants may view detailed breakdowns of the resource usage within their respective Notur-projects. In order to view the usage statistics, follow the *detailed statistics*-link from the project home page at www.metacenter.no. From this entry page, you may filter on computational resource, allocation period or user.

NorStore usage statistics will soon become available through the project leader graphical interface.

In order to log on to the project leader-GUI it is necessary with a FeideID or OpdenIdP. For any questions about using the project leader-GUI, please contact sigma@uninett.no

Gardar is closing down

Gardar is closing down at 30 December 2015

After almost four years of successful operation, the Gardar computer is shutting down at 30 December, 2015, at 12h00 Iceland time. All users must have completed their jobs and removed all data before this time. For continued access to computing resources, the users are directed to the national facilities, since there is no continuation of the Gardar project.

The supercomputer Gardar was installed in November 2011 at Thor Data Center, near Reykjavik, and inaugurated 16 April 2012. The Linux compute cluster was a result of the national high performance computing research infrastructures of Denmark (DeIC), Norway (UNINETT Sigma2) and Sweden (SNIC) placing a joint supercomputer in Iceland. The corresponding project NHPC has been a pilot initiative to understand the organizational and technical challenges of joint procurement, administration and operation of computational infrastructure for science.

Since 2012 computing allocations have been available on the Nordic HPC-cluster, being jointly used by the national HPC infrastructures. For the computing hours on the system available to Norwegian research projects, the resources have mainly been allocated to a few large projects. In particular, chemistry has been the most significant discipline from Norway using Gardar.

Contact: sigma@uninett.no

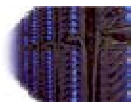

The December 2015 issue of the META magazine will be the last printed issue of the magazine. This marks the end of a ten year era of printed issues reporting the diverse and important use of the national e-infrastructure. Future editions of META will be released in some electronic format and available online.

Earlier editions of the META magazine, published in the period 2006 - 2015, will remain available at the following URL: www.sigma2.no/meta-archive

Photo: Thinkstock

Welcome to www.sigma2.no

Since late summer you have been surfing on a brand new website for UNINETT Sigma2. This new and improved website replaces www.notur.no and www.norstore.no, which will no longer be updated. With this new web site we offer you an interesting journey through the areas of national e-infrastructure. Its presentation reflects for example clearer the ability in which users of Notur and NorStore e-infrastructure facilities may apply for user accounts and reset their passwords at the supercomputers, and how project leaders may administrate their national e-infrastructure projects from UNINETT Sigma2s project leader graphical user interface.

Some pages are still under construction, but we hope you find the information you are looking for. If not, you are welcome to contact us at sigma@uninett.no.

User survey 2015

Photo: Thinkstock

Every year UNINETT Sigma2 carries out user surveys among the research projects that make use of the Notur and NorStore e-infrastructure. The aim of the survey is to determine the overall satisfaction with the resources and services provided by UNINETT Sigma2.

The user survey is an important opportunity for users and project managers to influence current and future services provided to you in Notur and NorStore. It is important for UNINETT Sigma2 to provide an infrastructure that is in accordance with the needs and preferences of the users. The survey gives the user community an opportunity to communicate their satisfaction with the resources and services and point to areas that need increased attention and inform UNINETT Sigma2 about future needs for infrastructure. The results of the survey will be used improve the operations and support mechanisms of UNINETT Sigma2 and guide the further planning of the infrastructure.

An electronic version of the report from the survey will soon be available on www.sigma2.no

Compute and storage projects

Which Norwegian researchers use supercomputing resources?

Supercomputers are critical to the success of researchers from a broad range scientific disciplines across Norway. These scientists use the e-infrastructure facilities to making our lives healthier, safer and better. A list of all research projects being granted computer time and data resources on the Notur and NorStore facilities by the Resource Allocation Committee in the current and recent allocation periods, are now available on www.sigma2.no. By clicking on the actual project, you will be able to read an abstract of the research being executed at the supercomputers.

RETURN ADDRESS

UNINETT Sigma2 AS
P. O. Box 4769 Sluppen
7465 Trondheim
NORWAY

NORGE P.P.
posten

Upcoming events

FEB 3

EUDAT
1st User Forum
3 - 5 Feb, 2016, Rome, Italy
www.eudat.eu

MARCH 1

RDA
7th Plenary Meeting
March 1, 2016, Tokyo, Japan
rd-alliance.org

JUNE 12

TNC16
12 - 16 June 2016, Prague, Czech Republic
tnc16.geant.org

MAY 10

PRACE days 16
10 - 12 May 2016, Prague, Czech Republic
www.prace-ri.eu/pracedays_16

JUNE 19

ISC'16
31st ISC High Performance Conference
June 19 - 23, 2016, Frankfurt, Germany
www.isc-hpc.com

Svanemerket trykksak fra Skipnes Kommunikasjon Lisensnr. 241 731 1215

UNINETT Sigma2