SNIC Newsletter

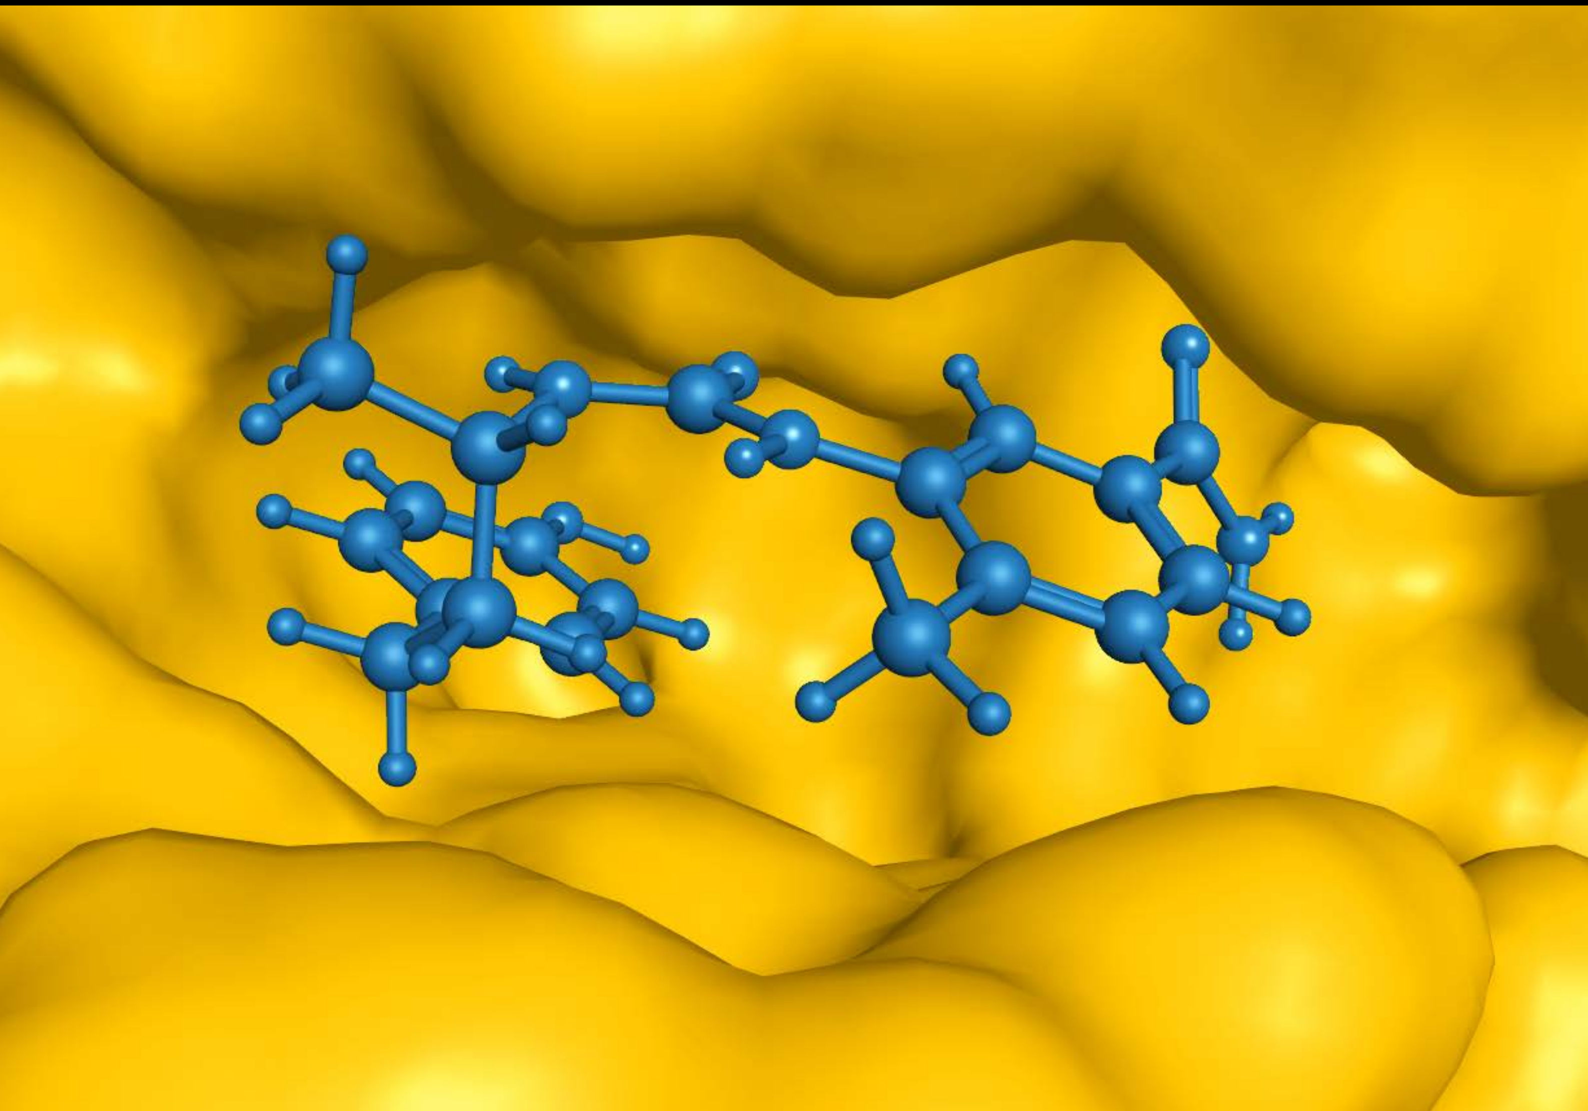

Autumn 2022

Published by the Swedish National Infrastructure for Computing

Contents

View from the SNIC director's desk

The period of transition from the Swedish National Infrastructure for Computing (SNIC) to the National Academic Infrastructure for Supercomputing (NAIS) has begun! NAIS is the new organisation that will be responsible for administering computational and storage resources to support Swedish research from the start of next year onwards. Final approval for NAIS is expected from the Council for Research Infrastructures (RFI), which is part of the Swedish Research Council (VR), in mid-October, and the transition process is already well underway. The general idea is for NAIS to continue the operation of the SNIC resources as seamlessly as possible and for the resources to be gradually adapted to the goals of the new organisation over time.

For researchers who are, or plan to be, using SNIC/NAIS resources in their research, the important thing to be aware of is that SNIC and NAIS are working together to make sure that the transition from SNIC to NAIS will not cause any interruptions in the actual operation of the high-performance computing (HPC) systems, storage, or support systems. NAIS has set up a working group that will manage the transfer, and transition teams are in place. Meetings are being organised over the coming months between SNIC and NAIS personnel to coordinate the process and make sure that the teams are prepared for the transition. SNIC is planning to publish a newsletter in December to provide information about the final aspects of the transition for researchers who use SNIC/NAIS resources – so stay tuned! Meanwhile, the other details we can share at this stage are primarily of interest to the SNIC centres and their staff, rather than researchers.

In parallel with the ongoing coordination work between NAIS and SNIC, NAIS will be holding discussions with all affected organisations. This will involve organisations currently involved in providing or managing SNIC resources in Sweden, or using those resources, or both. The purpose is to exchange information about the needs of the various organisations and how NAIS is planning to work and collaborate with the different entities. To this end, Matts Karlsson, who is one of the Deputy Vice-Chancellors of Linköping University and Chair of the Interim Steering Group for NAIS, will be calling for meetings with all the Swedish universities, as well as inviting the current SNIC centres to participate in discussions with NAIS. In addition, all SNIC and NAIS staff will be invited to a "SNIC & NAIS All-Hands Meeting" in late November to ensure that everyone has a good understanding of the final transfer plans. (Details, including how to register, will be emailed to SNIC/NAIS staff soon.)

The EU has invested heavily in HPC resources through the EuroHPC Joint Undertaking. Swedish researchers are strongly encouraged to apply for this plethora of resources through the open access calls (see https://eurohpc-ju.europa.eu/participate/access-our-supercomputers_en) to ensure that the resources are efficiently used, thereby providing a good return on the investments Sweden has made in EuroHPC. Please contact the SNIC User Support for free help, support, and advice on how to prepare proposals.

Lars Nordström, SNIC Director

Save-the-date

SNIC and NAIS invite all staff who work for SNIC or who are participating in the NAIS working groups, teams and transition processes to the

SNIC/NAIS staff will receive further details, including the registration link, via email soon.

Reminder: SNIC Large Compute Call closes on 14 October at 15:00 CEST!

How SNAC allocates SNIC resources

The Swedish National Infrastructure for Computing (SNIC) provides computational and data storage resources to facilitate Swedish academic research. The Swedish National Allocations Committee (SNAC) is responsible for allocating those resources to researchers.

SNIC provides computer time on a variety of architectures (such as CPU- and GPU-based systems, or systems with extremely large memory nodes), as well as different storage solutions (for example, storage based at individual SNIC centres, and Swestore), and all researchers who are associated with an academic institution in Sweden may apply to use SNIC resources. There are currently four different sizes of allocation that researchers can apply for: small, medium, or large allocations (on the main SNIC systems), or extra-large allocations (for utilising the Swedish part of the LUMI supercomputer, https://www.lumi-supercomputer.eu). The type of allocation that is appropriate for a particular research project will depend on the amount and type of resources that are needed (for example, if time on a GPU-based system is needed for artificial intelligence research). About two thirds of the available time on Swedish computers is reserved for use as large allocations.

Requests for large and extra-large allocations are evaluated biannually by a team of international experts, and then discussed individually by SNAC. Requests for small and medium allocations on some resources are handled directly by the SNIC centres that host the resources, while requests for small and medium allocations on Alvis, Dardel and Kebnekaise are evaluated by the technical experts within SNAC. Making the decisions is a difficult and at times frustrating process as the amount of time that is requested often exceeds the available resources. This means that strict reductions may need to be imposed on the amount of time that is awarded, as compared to what is requested, and that some applications need to be rejected. These decisions are based on the scientific excellence of each proposal, the suitability of the SNIC resources, and the expertise and experience of the research team that submits the application. SNAC always endeavours to find acceptable solutions for as many of the applicants as possible.

SNAC is independent of SNIC and consists of six members who are appointed from amongst researchers active in HPC-relevant scientific fields in Sweden (see https://snic.se/allocations/snac). SNAC considers it to be very important for the committee to remain independent of the SNIC office and centres, and for the evaluations to be based on two independent scientific reviews of the proposals, so that the allocation and provision of resources are handled separately.

In recent years SNAC has increased the amount of feedback to applicants in an effort to help them write better and more competitive applications. In particular, SNAC encourages research teams to apply for extra-large (and eventually PRACE-type European) resources, for which SNIC time for developing, optimising and testing codes and workflows is essential. In fact, SNAC is also responsible for evaluating applications from Sweden to use the Tier-1 European high-performance computing (HPC) systems provided by the PRACE research infrastructure (https://prace-ri.eu); those may be used for academic research or for industrial/business research undertaken by small and medium-sized enterprises (SMEs).

More information about the SNAC policies regarding access to SNIC resources can be found here: https://snic.se/about/governance/policies. As SNIC transitions into NAIS, the intention is to maintain the excellence-driven access model introduced by SNAC.

Latest SNAC allocations

SNIC holds two calls for large and extra-large allocations each year: one call is open in the spring for allocations that will start in July of the same year and the other call is during the autumn for allocations that will start in January the following year. The two SNAC allocation meetings for large and extra-large allocations for 2022 were held in December last year and in May this year. At those meetings, SNAC members discussed applications for compute and storage allocations that had been submitted a few months earlier.

In the SNIC “Large Compute Spring 2022” call, SNAC received 34 applications for large compute allocations and was able to allocate a total of 36 million CPU hours/month on the Tetralith and Dardel systems (at NSC and PDC, respectively) for 29 projects, which addressed research topics ranging from engineering computational fluid dynamics (CFD) to fundamental aspects of planet formation. In addition, 40,000 GPU hrs/month were allocated on the Alvis system (at C3SE) for four projects (out of the seven that were requested resources on that system). Also, a total of four petabytes of large-scale storage – located at PDC, NSC, Swestore/dCache, UPPMAX and C3SE – was allocated to 13 projects.

Thirteen applications were received in the “LUMI Sweden Spring 2022” call, which proved to be very popular. Both the CPU and GPU resources on LUMI (which are referred to as LUMI-C and LUMI-G) were oversubscribed, in the case of LUMI-C by more than a factor of two. SNAC was able to allocate 34 million core hours on LUMI-C and 1.75 million GPU hours on LUMI-G for use over the next six months. Here, the policy that was applied was to give many researchers access to LUMI in order to allow exploration of and adaptation to this relatively new resource for a variety of different research projects and compute needs.

Apply through SNIC or PRACE to use LUMI

The LUMI pre-exascale system, which is hosted by the CSC – Center for Technology in Finland, was inaugurated in early June. Sweden owns a small part of the system, and an equivalent amount of the system’s resources are available to Swedish academic researchers. These resources are currently managed by SNIC on behalf of the Swedish Research Council. Information on how to apply to use LUMI can be found on the SNIC website: https://snic.se/allocations/apply4access.

The EuroHPC Joint Undertaking owns half of LUMI and researchers from Sweden can apply for large allocations to use those resources through PRACE. There are calls for access to LUMI for benchmarking and development purposes. These are open continuously and the cut-off dates can be found here: https://prace-ri.eu/hpc-access/eurohpc-access/eurohpc-ju-benchmark-and-development-access-calls. In addition, there are calls for regular access, with the next call closing in November – see: https://prace-ri.eu/hpc-access/eurohpc-access/eurohpc-ju-call-for-proposals-for-regular-access-mode.

The LUMI User Support Team (LUST) provide support services for researchers wanting to use LUMI – see https://www.lumi-supercomputer.eu/user-support – and the Swedish representative, Peter Larsson, is based at the PDC Center for High-Performance Computing.

New hope for cancer patients through HPC

The World Health Organization's statistics indicate that nearly one in six deaths worldwide are due to cancer, with breast cancer being the most common type of new cancer in 2020 and the cause of the fifth highest number of deaths due to cancer in that year. About 15% of all breast cancer cases are a type of cancer known as triple negative breast cancer (TNBC). The cells of this type of cancer lack receptors for the hormones oestrogen and progesterone and also lack the HER2 protein. Consequently, TNBC does not respond to hormone-based therapies or medicines that target the HER2 protein.

This type of breast cancer, which is more common among women under 40, is usually more aggressive, harder to treat, and more likely to recur than breast cancers that are hormone receptor-positive or HER2-positive. However, recent computational chemistry research – undertaken on SNIC resources by Leif Eriksson and his colleagues at the Department of Chemistry and Molecular Biology at the University of Gothenburg – has led to the identification and development of new compounds which have proven to hold great promise for treating this form of breast cancer. For example, Tip60 is an enzyme that, amongst other things, controls certain types of DNA repair. Using drugs that block Tip60 together with radiation therapy that induces a lot of DNA damage in cancer cells can thus be used as a new and efficient way to treat TNBC. In addition, just suppressing Tip60 has been shown to trigger genetically programmed death (apoptosis) of cancer cells. Are you wondering how the research group found suitable compounds to treat cancer by inhibiting Tip60 levels when it is neither practical nor ethical to test millions of substances on breast cancer patients?

This is an example of where computational chemistry plays a vital role in the development of new pharmaceutical treatments. By using the latest computational chemistry algorithms and software on high-performance and high-throughput computing (HPC and HTC) systems, Leif and his colleagues have been able to screen vast libraries of compounds searching for potential new drugs, as well as creating computational models of proteins that new drugs could target, and exploring protein-protein interactions that are crucial for signalling pathways in cancer cells or other relevant biological processes. Leif's group has been studying several multi-protein complexes as targets for cancer therapy – the purpose is to identify small molecules which could act as inhibitors by blocking modes of action or signalling pathways, and would thus stop or slow the growth of tumours. The group has also been investigating mechanisms whereby cancer cells develop resistance to drugs and trying to identify compounds that would inhibit resistance from evolving. In addition, Leif and his colleagues have been looking into a small protein, known as AGR2, which has recently been shown to contribute to the development of tumours.

The group follows well-established protocols for this type of research. Appropriate models need to be prepared for the proteins being investigated, then calculations are performed looking into how pairs of different proteins could combine or "dock" with each other. These calculations are done according to a "consensus structure" protocol recently developed by the group and are then followed by molecular dynamics (MD) simulations to determine how stable the resulting structures are and how they would interact with other key compounds. This work is highly computationally intensive. For example, the MD simulations that are normally carried out each take about 500 to 1,000 nanoseconds, and are performed in triplicate (to ensure that the simulations provide proper statistics) which places high demands on the HPC resources.

In these drug development projects, the group also performs systematic docking of large databases, which can contain up to a billion compounds. This means performing calculations for the compounds in the database to check which of them could bind (dock) with the target substance. The resulting group of potential "hits" are then subjected to further modifications and more refined docking simulations. In addition, detailed simulations of the resulting complexes are performed – using binding pose metadynamics (BPMD) and MD techniques – to investigate the likely therapeutic potency and other characteristics of the hits. This is known as the "hit to lead" phase in drug discovery where suitable hits (compounds that can bind to the target) are identified and then investigated to see if they can become "leads" (compounds with a therapeutic effect that require some modification before being suitable as potential drugs).

The enormous numbers of compounds in the libraries that are used in this research mean that the necessary analyses can only be performed using massively parallel systems. Leif and his colleagues use around 1,000,000 core hours in total each month, spread over several SNIC systems. Tetralith at NSC and Dardel at PDC are being used for computationally intensive screening and modelling work, while C3SE's systems for Artificial Intelligence/ Machine Learning - the Alvis GPU-based system and the Mimer storage system – are being used to extend the group's work into the area of machine-learning-based "de novo drug discovery". That essentially means using artificial intelligence systems to generate new drug-like molecules that will bind to a given target. Their work on producing new leads in this way has come up with very promising initial data.

All in all, the research group has been highly successful in using these in silico methods for drug development on the SNIC systems: several patented drugs have already been developed, including (as mentioned earlier) some new compounds that hold great promise in the treatment of triple negative breast cancer. The group has also developed new compounds inhibiting the retinoic acid degrading enzyme CYP26B1 that can be used to treat various dermatology conditions with few or no side effects, as well as compounds blocking the extracellular domain of RET kinase involved in thyroid cancer, and molecules targeting the motor kinesin KIF18B which plays an important part in cell division and in repairing damage to DNA. Most recently, the group's use of SNIC's HPC resources has led to the development of a novel compound to combat the fast-growing lethal brain tumour glioblastoma multiforme (GBM), and in vivo tests indicate that the new compound may be able to cure patients completely! This is wonderful news since GBM is the most common type of primary malignant brain tumour in adults, and the average survival time of all patients who are treated with surgery, chemotherapy and radiation is only 15 to 16 months.

If you would like more information about other fascinating discoveries the group has made using SNIC's resources, see their website at https://www.gu.se/en/research/computational-and-theoretical-biochemistry.

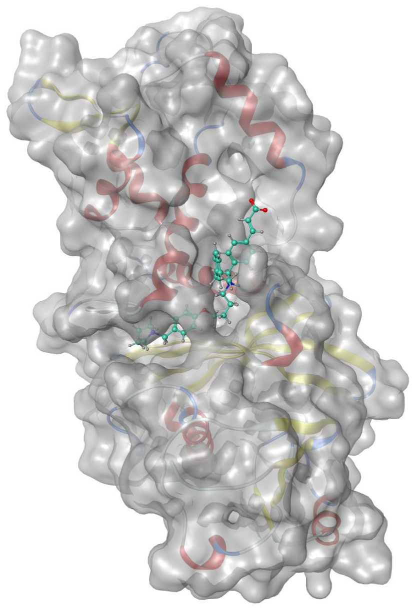

The image on the left represents a Tip60 molecule with one of the inhibitors developed by Leif's group bound to part of the molecule to inhibit its activity. (The inhibitor is the green and red structure in the centre.)

The image on the cover of the newsletter shows part of a protein known as IRE1 with the patented inhibitor compound Z4P bound to it. High IRE1 activity is associated with tumour development in cancers such as GBM, and Z4P (which was also developed by Leif and his team) is showing great promise in the treatment of GBM brain tumours.

Investigating online political communication

Some research areas, like computational fluid dynamics and biomolecular modelling, have been using parallel computing techniques for many decades and have large established bodies of code that researchers can use without having to develop their own routines. Other areas of computational research have arisen more recently and are in the process of developing analytical approaches that can serve as the foundation for widely used future workflows. Computational content analysis is one such relatively youthful area, which has highly practical applications, for example, in the identification of online disinformation and fake news. An interesting current research project in this area is working on identifying latent sentiment in online material, which is a step on the way to working out whether the material is likely to contain disinformation. Since fake news campaigns are often created to provide an advantage to a particular person, group or organisation or to be detrimental to others (for example, during the run-up to elections), we really need reliable ways to identify this kind of information pollution so we can combat it to build a more fair, equitable and just world.

What is computational content analysis?

Computational content analysis is an important research method used in the social sciences. Digital texts (such as organisational documents, interview transcripts, or material from websites) and/or images (for example, digital photos, or videos from YouTube) are collected and then analysed to extract meaningful content features, which are described as being either manifest or latent. A typical example of content analysis would be to determine the sentiment of a collection of social media posts to indicate whether users' attitudes are positive or negative towards a particular topic. In this case, the sentiment could be measured using the number of words with positive or negative emotional values (a manifest feature of the contents). However, the problem becomes more challenging if we also want to assess the latent sentiment in the posts, for example, when the phrase "this is really great" is used ironically to indicate the opposite sentiment. Identifying latent sentiment is vital when it comes to providing accurate measures of overall sentiment, and it is also important to correctly identify other latent features of online material.

Nils Holmberg and his colleagues at Lund University are exploring ways to identify latent content features in digital texts and images using computational methods, such as natural language processing and computer vision techniques. One example of this is using object recognition techniques to locate natural objects (which are another type of latent content feature) in digital images. By identifying people, objects and activities in political advertising, we can begin to understand how these features are related to changes in audience attitudes. The research in the project is also taking things a step further and working on trying to recognise all the objects on individual web pages in order to understand how the objects affect the behaviour of people using those web pages. To illustrate this, the figure shows a user interface (UI) for a particular web page at the top with the various parts of that page identified at the bottom. While identifying parts of

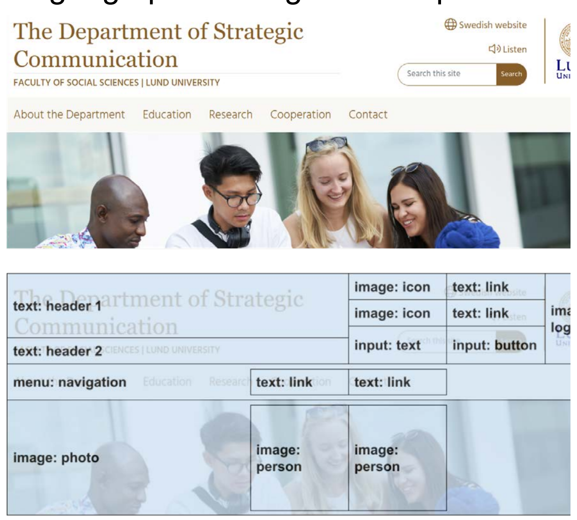

a web page may seem trivial to us as we do that subconsciously, it is non-trivial to develop automatic content analysis processes to reliably perform this type of identification.

Identifying online disinformation campaigns

Computational content analysis can be of great benefit in identifying online disinformation campaigns. The level of informational bias in an online text is an aspect (like natural objects in images) that can be regarded as a latent feature, yet it cannot easily be determined, even by human investigators. However, by leveraging the power of big disinformation datasets such as the Twitter Electoral Integrity Data (https://transparency.twitter.com/en/reports/information-operations.html), it is possible to train machine learning models (containing thousands of textual and pictorial parameters) that will help to infer the likelihood of disinformation content on a large scale. The researchers at Lund are working on doing that using the SNIC artificial intelligence (AI) system, Alvis, which is based at Lund University, along with storage on the Cephyr storage server at the Chalmers Centre for Computational Science and Engineering (C3SE). They will use about 500 graphics processing unit (GPU) hours each month during 2022 on Alvis and around 500 GiB of storage on Cephyr.

Cognitive communication effects

Another aspect of the ongoing content analysis research at Lund involves working towards understanding how people are affected by the way that online information is presented. While content analysis can be used in its own right (for example, to detect and measure sentiment or disinformation), it can also be used to understand, the causal connection between content features on a website and the behaviours (and cognitions) of the people using that website. Since a growing part of people's daily activities occurs through interactions on mobile or computer screens, it is of value for Swedish society to support research on how communication design can affect user decisions.

The traditional method of determining such website communication effects is to simply ask people what they thought, felt or did when visiting a particular website. A significant problem with this approach, however, is its dependence on self-reported subjective answers. In order to remedy this situation, communication researchers are increasingly turning towards objective behavioural observations, such as web analytics that can provide detailed information about individual users' interactions with a website, down to each mouse movement, click and scroll action. This can be expanded into the cognitive domain by the use of eye-tracking that records where users' visual attention focuses on web pages. For example, this image is from the video https://www.youtube.com/watch?v=bC4fVYMjurY which shows results from experiments comparing the eye movements of people looking at an online advertisement for a phone, with and without an image of a girl on the page. The present project envisions that computational methods will help us to first recognise multiple web page object types across multiple user browsing sessions and second predict user behaviour and visual attention based on how graphical objects are presented on web pages. The complexity of the analyses and enormous volumes of online material involved in these steps mean that there is no other way to obtain this information reliably except by using supercomputing techniques and resources. For more information about this research, see the project's website: https://portal.research.lu.se/en/projects/computational-methods-and-behavioral-analytics-in-strategic-commu.

Dardel SNIC system being extended

The first (CPU-based) phase of the SNIC system, Dardel, was installed at the PDC Center for High Performance Computing in August last year (after some unavoidable delays), and the second phase, which is a partition featuring graphics processing units (GPUs), will be installed during the autumn. More CPUs are going to be added to the system too!

As mentioned in the previous SNIC newsletter, SNIC and the Swedish Science Council decided to invest more in Dardel to mitigate the risk of a future lack of adequate capacity for Swedish research. As a result of this, plus compensation awarded for the delays with the installation of the first phase of the system, Dardel will be extended with an extra 412 CPU nodes and the disk capacity will be increased by 50%.

At this stage, it is expected the additional hardware will be delivered in September or October this year. These extra nodes, plus the GPU partition, mean that the system will grow by an extra row, with another two compute cabinets and one cooling distribution cabinet (for controlling the temperature), as shown below. The two phases will operate independently initially as the (interconnect) networks that the CPU and GPU nodes use to exchange data while doing computations work at different speeds.

PDC is currently working on developing a good solution for unifying the two phases – the plan is to do this late this year after the GPU partition has passed its test phase, and the additional CPU nodes have been installed.

SNIC centre round-up

C3SE

The national SNIC services provided by C3SE are now fully consolidated within Chalmers e-Commons, a Chalmers infrastructure that provides a broad span of digital and data support. Within the new National Academic Infrastructure for Supercomputing (NAIS), Chalmers e-Commons will continue to develop and provide the GPU-based resource Alvis - with a focus on artificial intelligence/ machine learning (AI/ML) - for the full lifetime of the system. Here, a specific focus will be on expanding the user base even more, for example, by further developing the existing portal for interactive use. Also, Chalmers e-Commons will continue the national collaborative efforts within the areas of data storage and cloud services. In the near future, Chalmers e-Commons will be expanding with a focus on both e-infrastructure and expertise. Future efforts include building a complete set of data management and storage services for Chalmers, leading the build-up of the Swedish e-infrastructure node for the global Square Kilometre Array (SKA) radio astronomy infrastructure, and building one of the Data Science Nodes in the national Data-Driven Life Science (DDLS) initiative.

HPC2N

HPC2N is assisting SNIC users to migrate from the Kebnekaise system to more

SNIC centre round-up (continued)

modern SNIC resources. Researchers will still benefit from HPC2N's long-term and large-scale storage systems. HPC2N is also contributing to the Swedish Science Cloud and is coordinating the development of an interactive HPC environment for the Dardel system. Introductory, intermediate and advanced-level HPC courses are offered by HPC2N. Upcoming courses include Python for HPC, Introduction to MPI, and Introduction to Git.

LUNARC

LUNARC is preparing for the installation of the new COSMOS system, which will replace the current general-purpose Aurora system later this year. The COSMOS supercomputer will be about four times as powerful as Aurora. There will also be a separate system, called COSMOS-SENS, for researchers who need to perform calculations or simulations using sensitive data (or, in other words, data that must be kept secure). The current plans are for COSMOS and COSMOS-SENS to be available until 2026-27. LUNARC focuses on providing support to make COSMOS highly accessible to researchers in a broader range of disciplines. LUNARC is also actively planning a new modern data centre that will support the needs of both LUNARC and Lund University in the future.

NSC

At present, the plans are for NSC's main systems – the Tetralith cluster and the AI/ML system, Berzelius – to continue to be available for Swedish research until 2023 and 2026, respectively. As the new NAIS organisation will be hosted by Linköping University, NSC will be heavily involved with the upcoming transition from SNIC to NAIS in the coming months.

PDC

PDC is in the process of getting the second phase of Dardel in place; it is a partition with GPU-accelerated nodes. For the first phase of the system, which has already been installed, PDC is focussing on improving its usability for a more diverse user basis and improving the ease-of-use. This includes the introduction of authentication based on secure shell key-pairs, as well as providing support for small jobs, which use only a few cores, and for long-running jobs.

The software offered on the Dardel system continues to be extended in accordance with requests from researchers. Furthermore, tools have been installed to support data transfer from Dardel's parallel file system Klemming to the dCache-based SweStore instances.

UPPMAX

The status of the SNIC systems at UPPMAX is good. Rackham will be available as a general-purpose system and Bianca will be available for sensitive data until the end of 2024. The SNIC Science Cloud East region can be assumed to continue operating until at least the end of 2023.

In the autumn of 2022, Cygnus will be taken into operation as a new storage system for Bianca. Cygnus replaces Castor and will be faster, larger, and more reliable. An extension to Cygnus is already planned for 2023.

Currently, UPPMAX hosts a complete copy of the ZINC15 ligand database. Soon, this database will be superseded by the much larger ZINC22 database, currently being downloaded in its entirety. This resource will be made available primarily to UPPMAX users, but we are investigating the possibility to open up access more widely.

The UPPMAX Application Expert team will welcome a new member in the domain of humanities and social sciences. The role is shared with the Centre for Digital Humanities Uppsala, and aims to improve SNIC usability in this exciting and growing scientific domain.

HPC opportunities & organisation links

HPC in Sweden and Scandinavia

ENCCS: http://enccs.se

eSSENCE: http://essenceofescience.se

NeIC: http://neic.no

SeRC: https://e-science.se

SeSE: http://sese.nu

SNIC: https://snic.se

SNIC centre project participation

BioExcel CoE: https://bioexcel.eu

EBRAINS: https://ebrains.eu

PerMedCoE: https://permedcoe.eu

EXCELLERAT: https://www.excellerat.eu

EOSC-Nordic: https://eosc-nordic.eu

DICE: https://www.dice-eosc.eu

HPC-Europa3: http://www.hpc-europa.eu

European HPC ecosystem

HPC in Europe: https://hpc-portal.eu

EuroHPC: https://eurohpc-ju.europa.eu

PRACE: https://www.prace-ri.eu

LUMI: https://www.lumi-supercomputer.eu

ETP4HPC: https://www.etp4hpc.eu

EOSC: https://eosc-portal.eu

General HPC news sources

HPCwire: http://www.hpcwire.com

insideHPC: https://insidehpc.com

Upcoming SNIC training & HPC events

SNIC zoom-in (Online interactive support and discussion forum)

13 October (and also 17 November) 2022, 14:00-15:00, online

Researchers who are using (or who want to use) SNIC resources are invited to attend the SNIC zoom-ins where SNIC experts are available to discuss computing and storage needs or other matters related to research using SNIC resources: http://snic.se/support/zoom-in.

Save the dates

Here are preliminary dates for some upcoming SNIC training events. For details, visit the SNIC website http://snic.se/support/snic-training, or sign up in the SNIC User and Project Repository (SUPR) where you can register for our newsletter: https://supr.snic.se.

- **Basic Singularity: Running and building Singularity containers**: 28 September 2022, online workshop
- **Cluster architecture and job submission**: 5 October 2022, online seminar
- **Introduction to Alvis**: 12 October 2022, online course
- **UppASD autumn school 2022:** 11-13 October 2022, KTH main campus, Stockholm
- **An introduction to parallel programming using Message Passing with MPI**: 17-20 October 2022, online course
- **Programming Formalism**: 24-28 October 2022, online course in collaboration with NBIS
- **Introduction to PDC Systems**: 27-28 October 2022, online course
- **Julia for HPC**: 8-10 November 2022, online course in collaboration with ENCCS
- **Introduction to Git**: 14-18 November 2022, online course in collaboration with PRACE