The Research Council
of Norway

Norwegian eInfrastructure Roadmap

Prepared by the eInfrastructure Scientific Opportunities Panel

Programme
eScience – Infrastructure, theory and applications (eVITA)

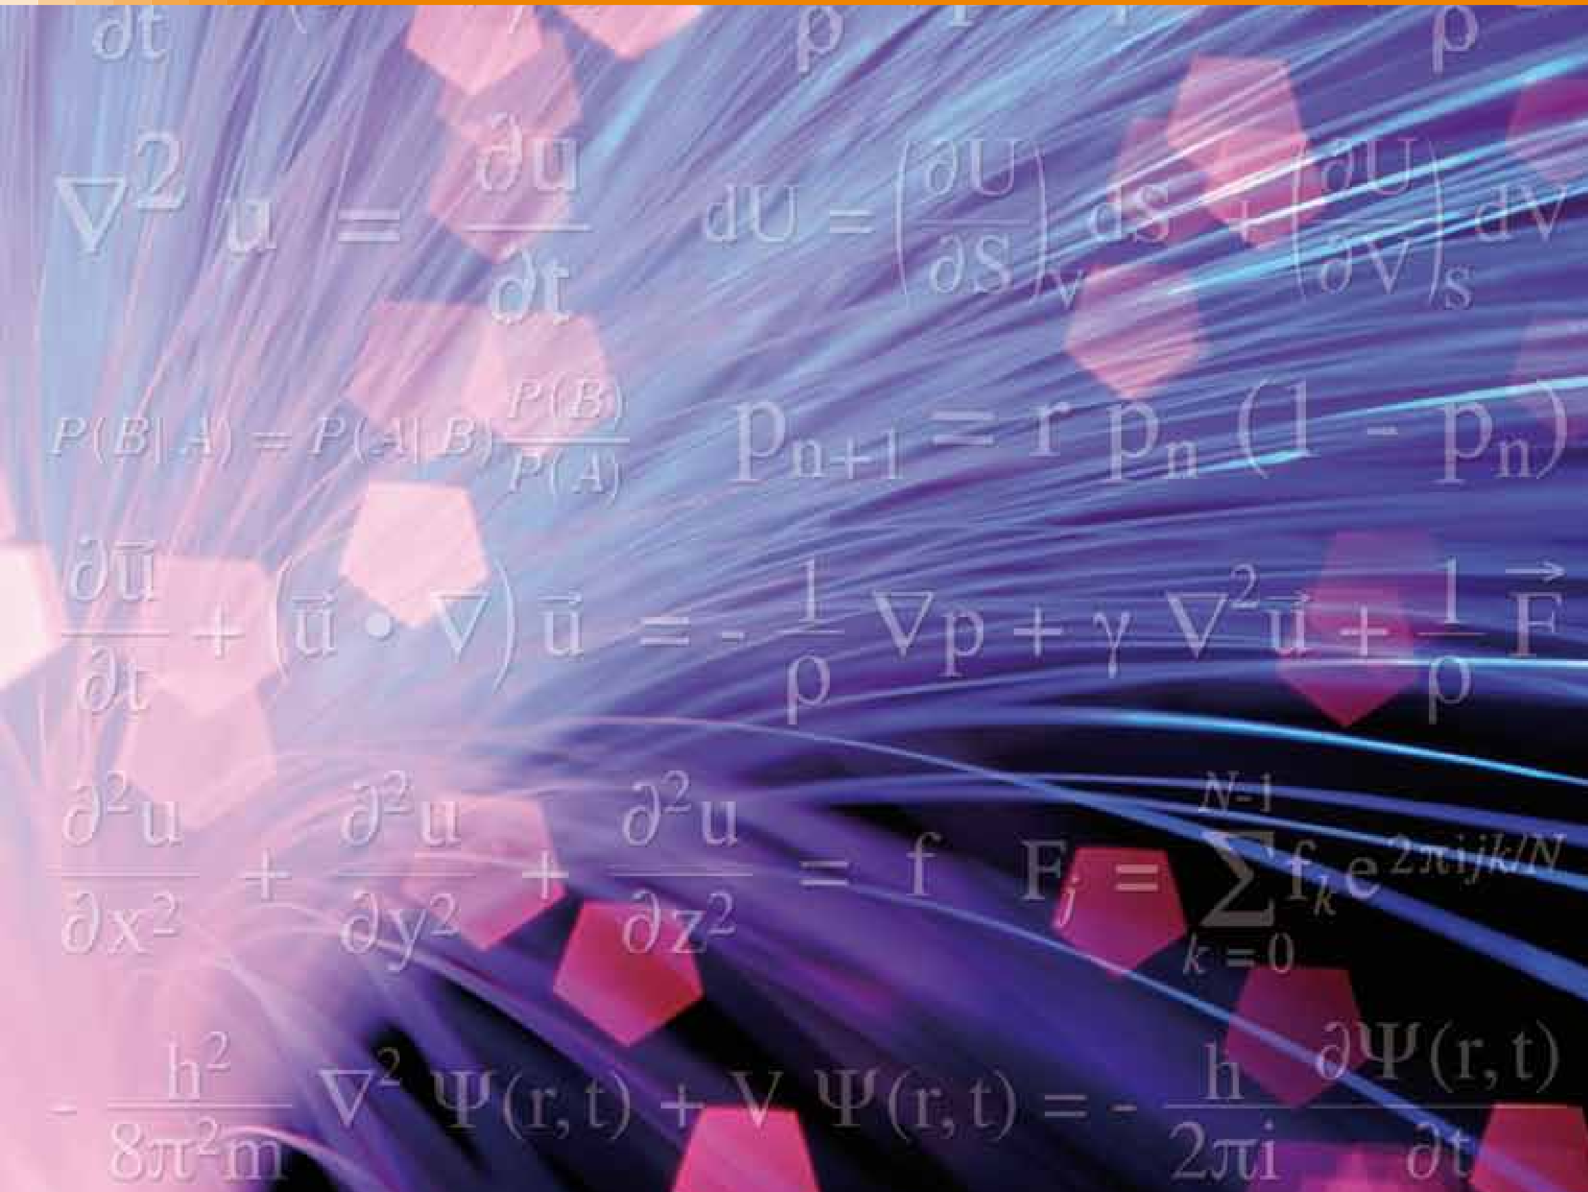

About the programme
eScience – Infrastructure, theory and applications (eVITA)

eVITA is a research and infrastructure programme designed to address computing- and data-intensive challenges in science, technology and medicine. By promoting research on methodologies, competence development and investment in new *eInfrastructure*, eVITA will work to ensure that Norwegian research in the *eSciences* achieves a high international standing, and seek to address important national challenges in the national priority areas of energy and the environment, oceans, food, and health.

The objectives and measures described in this programme require a budget framework of NOK 110 million per year starting in 2008, of which NOK 50 million is reserved for investment in *eInfrastructure*. However, the planning group recommends increasing the budget framework to NOK 170 million per year, of which NOK 70 million is reserved for investment in *eInfrastructure*. The planned programme period is ten years (2006–2015).

The report is prepared by the eInfrastructure Scientific Opportunities Panel

Galen Gisler (chair), University of Oslo
Dag Bjørge, met.no
Elena Celledoni, NTNU
Trygve Helgaker, University of Oslo
Colin Jones, SMHI
Anna Lipniacka, University of Bergen
Arvid Lundervold, University of Bergen
Stephan Oepen, University of Oslo
Nils Reidar B. Olsen, NTNU
with the assistance of
Jacko Koster (secretary), UNINETT Sigma
Gudmund Høst (observer),
the Research Council of Norway

Table of contents

Foreword

We are in the midst of an unprecedented development in which the technologies of information, communication, and computation have merged into what is called an electronic infrastructure, or *eInfrastructure*.

This development poses great challenges as well possibilities for our society.

To describe this the eVITA Programme Committee, under the auspices of the Research Council of Norway, appointed the eInfrastructre Opportunities Panel (eSOP) with two charges:

(1) Develop the scientific case for the eInfrastructure that can best serve Norwegian research groups and operational forecasting from 2015. eInfrastructure in the present context covers electronic resources such as large data collections, large-scale computing resources and high-speed networks, as well as the tools and services enabling efficient use of these resources.

This first charge was fulfilled in the document published in 2010, *The Scientific Case for eInfrastructure in Norway*[1].

(2) Produce a first version of a Norwegian eInfrastructure Roadmap. This Roadmap should cover current and new scientific areas for eInfrastructure use, taking into account the opportunities offered by existing and emerging large-scale international collaborations. The Roadmap should also make international comparisons and, in collaboration with the NOTUR project leader, match application areas against services.

The arguments presented here are intended for all stakeholders in Norwegian eInfrastructure including users and support staff, while the recommendations made are intended for decision makers at universities, the Research Council, and the Ministry of Education and Research. This document frequently refers to its predecessor.

The eVITA Programme Committee is very grateful to the chair of the eSOP, Prof. Galen Gisler of the University of Oslo, and his team, for their dedication in carrying out this task. We are convinced that this document will be very useful for decision makers and scientists alike.

Helge Holden
Chair, the eVITA Programme Committee

Sammendrag

Teknologiutviklingen innen informasjon, kommunikasjon, beregninger og datalagring har stor og økende betydning i hverdagen, og gjør oss i stand til å møte de store samfunnsutfordringene, herunder riktig bruk av energi, transport og mobilitet, helse og aldrende befolkning, miljø, produktivitet og sikkerhet. I dag har disse teknologiene utviklet seg til en egen infrastruktur, ofte omtalt som *eInfrastruktur*, som samfunnet ikke lenger kan være foruten.

Norges avanserte teknologi og eInfrastruktur støtter opp under forskning og utvikling slik at vi skal kunne opprettholde vår posisjon i en konkurranseutsatt verden. Imidlertid henger vår finansiering av eInfrastruktur etter de øvrige industrialiserte land. Finansieringen må økes, og den må være bærekraftig, stabil og forutsigbar for å kunne utnytte den pågående teknologiske utvikling.

I dette dokumentet oppsummeres de samfunnsmessige og økonomiske utfordringer som vil bli adressert av en effektiv eInfrastruktur. Vi beskriver noen forventede teknologiske fremskritt og hvordan de kan utnyttes på en best mulig måte. Vi gir anbefalinger om hvordan en oppnåelig og samfunnsnyttig visjon for Norge etter 2015 kan realiseres. Våre anbefalinger fokuserer på forskjellen mellom situasjonen i dag og forventede fremskritt innen beregnings- og kommunikasjonsteknologi. Vi foreslår også tiltak for hvordan vi kan ligge i forkant av denne utviklingen gjennom god planlegging og et riktig offentlig investeringsnivå.

Optimal utnyttelse av ny teknologi vil alltid være vanskelig, men gevinsten omfatter mulighet for å studere tradisjonelle problemer i større detalj og med mer nøyaktighet. En kan også forske mer effektivt enn før, men viktigste er at en får muligheter til å løse problemer som ikke kunne løses tidligere. Utviklingen innen tungregning går meget hurtig i Europa, Japan, Kina og USA. Norge er teknologisk avansert og konkurrerer på et høyt internasjonalt nivå. Imidlertid står vi i fare for å havne i bakleksa innen flere fagfelt med mindre finansieringsmekanismene for eInfrastruktur blir radikalt forbedret og organisasjonen strømlinjeformes og effektiviseres.

I det følgende gis anbefalinger samlet under generelle temaer. Innenfor hvert tema gis referanser til de deler av dette dokumentet der anbefalingene er nærmere begrunnet.

Finansiering og innkjøp av maskinvare: *Norge må ha egne tungregneanlegg på toppnivå, og minst ett av disse må hele tiden være blant verdens 500 kraftigste.* De aller beste ressursene vil finnes utenlands, men avstanden mellom regnekapasitet tilgjengelig i Norge og i utlandet må ikke bli for stor. Innkjøp av et tungregneanlegg hvert eller annethvert år vil bidra til å holde tritt med den teknologiske utviklingen. Stabil og forutsigbar finansiering er mer nyttig enn store uforutsigbare tilførsler av midler. Hvert større innkjøp bør bringe oss inn på omtrent 30te plass på «Topp500-listen», noe som vil sikre at Norge alltid har to anlegg på Topp500-listen. Energieffektivitet bør være blant kriteriene som benyttes ved innkjøp, og forskning og utvikling på "grønn IT" må stimuleres. *Norge må ha et stabilt, forutsigbart og bærekraftig finansieringsnivå for eInfrastruktur* for å bidra til innovasjon, strategisk planlegging og å møte våre samfunnsutfordringer. Denne finansieringen kan komme fra Forskningsrådet eller direkte fra Kunnskapsdepartementet. *Vi anbefaler et årlig finansieringsnivå på 100 MNOK.* Se kapitlene 1, 3.1, 3.4, 4.1, 4.2, 5.4, og EU[2], NSF[3].

Data: *Norge må ha en permanent nasjonal infrastruktur for lagring av forskningsdata, med enhetlig struktur, nomenklatur og protokoller for tilgang.* Data bør være godt sortert og verifisert, tilgjengelige, sikre og kunne gjenbrukes. Offentlig finansierte data må gjøres tilgjengelig for allmenheten etter en passende tidsperiode og metadata må gjøres robuste mot endringer i lagringsformater. Internasjonalt samarbeid innen håndtering og lagring av data bør stimuleres, og Norge bør

vurdere å tilby sikre og miljøvennlige sentre for datalagring til internasjonale forskningsorganisasjoner. Se kapitlene 3.4, 5.3.

Organisering: *Det trengs en enkeltstående nasjonal eInfrastruktur for tungregning, inklusive drift, brukerstøtte, forskning, utdanning og samarbeid*. Denne eInfrastrukturen bør fungere som et metasenter med et tilstrekkelig mangfold av ressurser til å håndtere ulike mengder og typer oppgaver, og med sentre som er knyttet sammen i en lagdelt struktur med et felles miljø og harmoniserte standarder. *Beslutningsstrukturen må inkludere styrer med uavhengige brukere på alle nivåer og aktivitetsområder.* Sentre på de lavere nivåer med brukerstøtte bør samlokaliseres med de større universitetene, men øverste nivå bør samle spissressursene i ett nasjonalt senter på et nøytralt og miljøvennlig sted. Norges deltakelse i tungregning bør fremmes på nordisk og europeisk nivå. Norge bør investere i større ressurser sammen med andre land. På grunn av vår grønne energiprofil bør Norge vurdere å tilby vertskap for større maskiner, gjerne et felles europeisk system i verdensklasse. Se kapitlene 1, 3.1, 3.4, 3.5, 4.1, 6.

Personale: *Tilstrekkelig bemanning må tilbys for å møte behovene til både uerfarne og avanserte brukere*. Brukerne bør ha tilgang til kompetent støttepersonell for å legge til rette for effektiv bruk av de største nasjonale og internasjonale systemene. De ansattes kompetanse må oppdateres kontinuerlig, og de bør eksperimentere med nye arkitekturer for å forberede seg til anskaffelse og bruk av ny teknologi. Støttepersonell bør plasseres nær forskere. Universiteter og forskningsinstitusjoner må vektlegge utvikling og vedlikehold av kompetanse for å kunne dra nytte av nye programmeringsmodeller, algoritmer og programmeringsspråk, og for å kunne flytte programvare til nye arkitekturer. Behovet for energieffektivitet og trender i retning av heterogene beregninger og hybrid programmering gir nye utfordringer. Utvikling av åpen kildekode bør stimuleres. Se kapitlene 1, 3.1, 3.4, 3.5, 5.2, og NSF[3].

Utdanning og formidling: *Den nasjonale eInfrastruktur må skape og videreutvikle et robust system for utdanning, opplæring og dokumentasjon for både erfarne og nye brukere*. Opplæringen bør knyttes tett opp mot pensum ved universitetene for å engasjere studentene. Tverrfaglige kurs i beregningsmetoder for medisin, lingvistikk og samfunnsvitenskap må opprettes. Utdanning av lærere i den videregående skolen bør omfatte elementer av beregningsorientert matematikk og informatikk for å eksponere ungdom for fordelene ved å bruke og å utvikle eInfrastruktur. Formidling på tvers av fag og dialog med allmenheten er meget viktig for å sikre at verdien av eInfrastruktur blir godt forstått og verdsatt av storsamfunnet. Se kapitlene 1, 3.4, 4.1, 5.1, og NSF[3].

Tilgang: *Et brukervennlig og harmonisert brukergrensesnitt til ressursene må opprettholdes*. Den nasjonale eInfrastruktur og dets metasenter bør være tilgjengelig for all forskning og utdanning. Brukerkontoer, brukernavn og identiteter bør harmoniseres på tvers av alle systemer, med standard prosedyrer for tilgangskontroll. Personell som yter brukerstøtte bør være direkte tilgjengelig for alle brukere uavhengig av hvor brukerne befinner seg. Der det er mulig og praktisk bør Norge samarbeide med tilsvarende organisasjoner i andre land for å fjerne hindringer for levering av tjenester over landegrensene. Se kapitlene 3.5, 4.1, 4.3, og NSF[3].

Executive summary

The technologies of information, communication, computation, and data storage have a tremendous and growing impact on everyday life, enabling us to counter the grand challenges facing our society, including appropriate use of energy, transport and mobility, health and the ageing population, the environment, productivity, and safety. Today, these technologies have developed into an infrastructure, *eInfrastructure* as it is called, that society can no longer do without.

Norway's advanced technology and eInfrastructure support the research and development that maintains our position in a competitive world. However, our funding for eInfrastructure lags behind that in other industrialised countries. This funding needs to be increased, and it should be sustainable, stable, and predictable in order to take advantage of new technological advancements.

In this document are summarised the societal and economic challenges that will be addressed by an agile eInfrastructure. We describe some of the advances anticipated in technology, and how these might best be used. We make recommendations that will help realise a laudable and achievable vision for Norway for the years beyond 2015. Our recommendations focus on the gaps between our present situation and the advances anticipated in the technologies of computation and communication, suggesting ways in which to stay abreast of these advances through appropriate government investment and oversight.

Making the best use of new technology will always be difficult, but the rewards include the ability to study old problems with greater fidelity and accuracy, to carry out research more efficiently than before, and most importantly, to solve problems that could not be solved before. Europe is moving aggressively forward in high-performance computing, as are Japan, China, and the United States. Norway is technologically advanced and competes well on the international level now, but Norway will fall behind unless the funding mechanisms for eInfrastructure are radically improved and the organisation streamlined and made more efficient.

We begin with our recommendations collected under general rubrics. Each rubric is followed by references to the sections of this document in which the background for those recommendations is developed.

Hardware funding and procurement: *Norway must have its own systems within reach of the top level, and must maintain a continuous presence among the 500 fastest computers in the world*. The very largest parallel resources exist internationally, but the gap between performance available in Norway and performance available in Europe and the rest of the world must be kept from getting too large. Procurements staggered with a frequency of every year or two will keep pace with the development of the industry. Steady and predictable funding is more helpful than sudden injections of capital. Each major procurement should be at a level above or near number 30 on the Top 500 list, ensuring that Norway always has two systems among the Top 500. Energy efficiency should be among the criteria used in procurement, and green computing research and development should be encouraged. *Norway must have a stable, predictable, and sustainable level of funding for eInfrastructure*, to allow innovation, to permit better strategic planning, and to meet enduring social needs. This funding could come either from the Research Council or directly from the Ministry of Education and Research. *The level we recommend is 100 MNOK annually*. See Sections 1, 3.1, 3.4, 4.1, 4.2, 5.4, and EU[2], NSF[3].

Data: *Norway must have a permanent national infrastructure for storing scientific data, with unified structure, nomenclature, and protocols for access.* Data should be well curated, accessible, secure, and reusable. Publicly funded data should ▶

▶ be accessible to the public after an appropriate time, and metadata should be made resilient to changes in storage formats. International cooperation in data archiving and storage should be encouraged, and Norway should consider offering secure and environmentally friendly sites for permanent data repositories to international scientific consortia. See Sections 3.4, 5.3.

Organisation: *There must be a single national eInfrastructure for high-performance computing including operations, support, research, education, and collaboration.* As a Metacentre, this should have a diversity of resources to handle different kinds of workloads, with virtually linked centres in a tiered structure with a common environment and harmonised standards throughout. The governance must include independent users' boards at all levels and in all aspects. Lower-tier centres with user support should be colocated with the major universities, but the highest tier with the highest performance machines should be a single national centre at a neutral and environmentally favourable location. Norway's participation in supercomputing should be promoted on the Nordic and European levels. Norway should invest in larger resources jointly with other countries, and should consider offering to host very large machines, even a joint European world-class system, using our green energy infrastructure. See Sections 1, 3.1, 3.4, 3.5, 4.1, 6.

Staff: *Adequate staff must be provided to address the needs of both inexperienced and advanced users.* Users should have competent support to facilitate efficient use of the largest national and international systems available. The skills of the staff must be continually updated, and they should experiment with emerging architectures, to prepare for the procurement and use of new technologies. Support staff should be located close to researchers. Universities and research institutions should develop and sustain competence to take advantage of new programming models, algorithms, and languages, and to port software to new architectures. New challenges arise from the need for energy efficiency and trends towards heterogeneous computing and hybrid programming models. Open source software development should be encouraged. See Sections 1, 3.1, 3.4, 3.5, 5.2, and NSF[3].

Education and Public Relations: *The national eInfrastructure must create and maintain robust education, training, and documentation for both experienced and beginning users.* The training provided should be tied closely to the educational curricula at the universities to encourage student participation. Interdisciplinary courses in computational methods for medicine, linguistics, and social sciences should be created. Training for secondary school teachers should include components in the disciplines of computational and computer science to expose young people to the advantages of using and developing eInfrastructure. Dissemination of information across disciplines and engagement with the public are essential to ensure that the value of eInfrastructure is well understood and appreciated by society at large. See Sections 1, 3.4, 4.1, 5.1, and NSF[3].

Access: *Friendly and harmonised user interfaces to the resources must be maintained.* The eInfrastructure and its Metacentre should be available for all research and education. Accounts, user names, and identities should be harmonised across all systems, with standard procedures for identity assurance. Staff providing user support should be directly accessible to all users regardless of their location. Where possible and practical, we should cooperate with similar organisations in other countries to eliminate barriers to delivery of services across borders. See Sections 3.5, 4.1, 4.3, and NSF[3].

1

Benefits to society, science, and industry

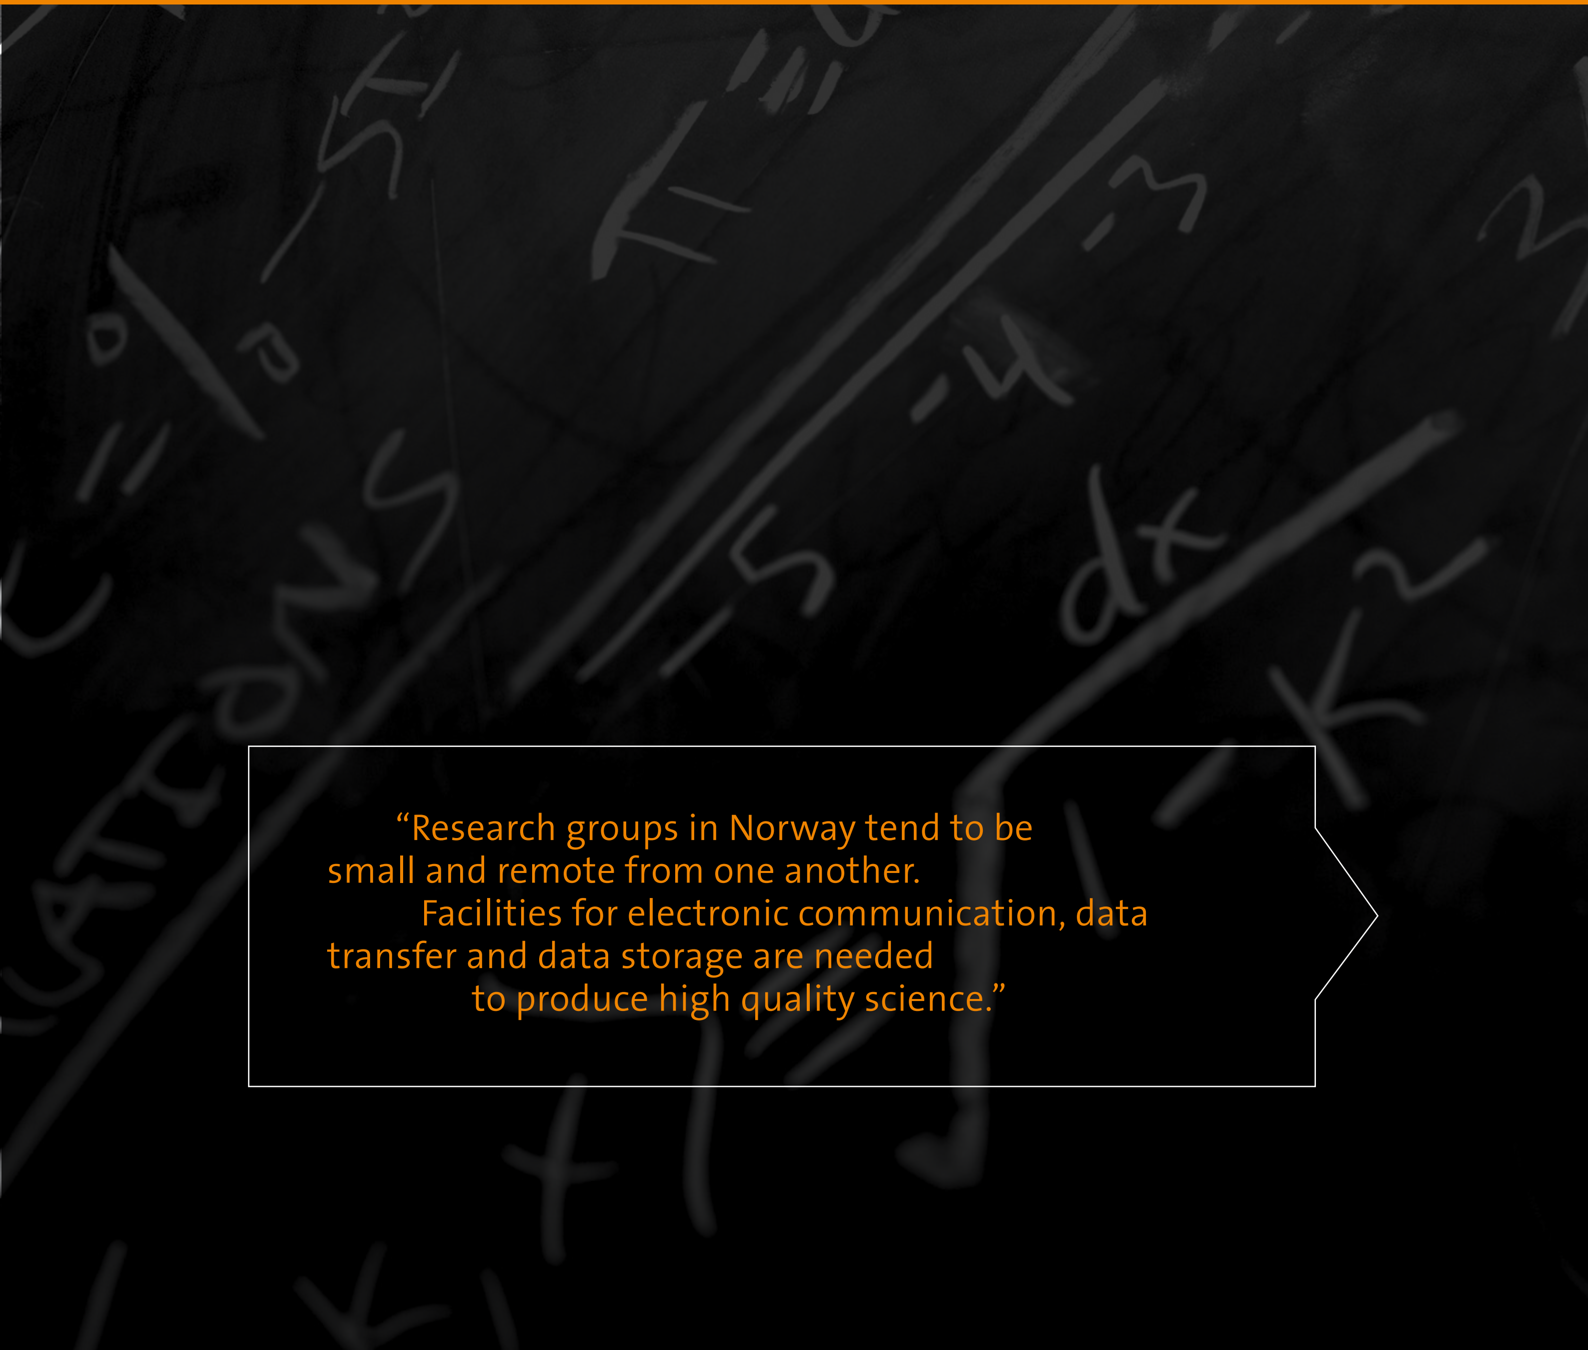

> “Research groups in Norway tend to be small and remote from one another. Facilities for electronic communication, data transfer and data storage are needed to produce high quality science.”

1 Benefits to society, science, and industry

A robust, reliable and efficient eInfrastructure provides benefits for society on many levels. Solving the energy challenge, providing the framework for health research and drug-testing, increased industry productivity, safe and fast transport, and the use of eInfrastructure to protect life and property are on the top of the list.

Improving the infrastructures associated with the technologies of information and communication will enable us to face important societal challenges. Advances in these technologies have led to mobile communication, computer-aided design and manufacturing, medical imaging, drug development and synthesis, climate simulation and weather prediction, the internet, and so on. In the future, new advances in these technologies will change our lives in an even more dramatic manner, as illustrated by the following list of scenarios, the first one recently realised.

1.1 Scenarios

(A) A computer system participates in the popular television quiz show “Jeopardy!” The machine understands questions posed in English, probing a broad range of common knowledge. With considerable skill, the machine evaluates its probability of having the correct answer, and demurs if that probability is low. Two all-time “Jeopardy!” champions battle the machine for two days but are defeated by a good margin.

As is well known, and contrary to widespread predictions by experts, this actually happened[4] in early 2011. A team of IBM researchers used natural-language information processing to cover vast quantities of textual material and designed the hardware and algorithms to take on this grand challenge; they named the system Watson, after IBM’s first president, Thomas J. Watson. Although not a true intelligent agent, Watson demonstrates the power of combining natural language processing, machine learning, reasoning, parallel computing, and hardware. Watson would likely rank among the top 100 of the most powerful supercomputers in the world. Watson was not connected to the internet during the game. Work at IBM is underway to adapt this technology to more practically relevant tasks.

(B) Setting out on a trip from Oslo to Bergen, a commuter consults the holographic display in her vehicle. The grid recognises her and instantly retrieves her preferences and pertinent information. It reports that there is still snow in the mountains, with an 85% chance of more snow combined with strong winds across the Hardangervidda. The descent past Vøringsfossen could be especially hazardous. The final choice of route need not be made until closer to Gol; forecasts and recommendations will be updated along the way. Because she prefers the magnificent view of the Vidda over the tunnel’s gloom, she will wait to make her decision. However, the display also reminds her that her risk-weighted insurance contract will probably force aesthetics to yield to economics, unfortunately.

This may be possible in the future with a detailed regional weather forecasting model, initiated with data from satellites, radar, communication systems, conventional weather stations, and “crowd-sourced” input from other travellers. Boundary conditions come from a global model, and a regional ensemble system produces uncertainty estimates as well as alerts and thresholds, updated continuously. Parts of this system are already in place, but low-resolution and inadequate processing power limits its accuracy and operational usefulness.

(C) A well-travelled patient visits a doctor’s office with unusual and serious symptoms. A diagnostic station takes histories, measures vital signs, and performs a blood test. The machine searches a database of similar tests made in locations the patient has visited recently, and collates the results of the search with the medical literature. A pathogen is identified and its genome sequence is analysed. Moments later, an effective and specifically targeted countermeasure is recommended and initiated.

Technologies for this diagnostic and treatment system are now in development in the biomedical community. Data archived in the system require an appropriate mix of security and transparency, maintaining both patient anonymity and source integrity. Health is a public good: society benefits greatly from the rapid diagnosis and effective suppression of epidemics.

(D) A circle of conferees around a meeting table consists of some present in person and others participating remotely through telepresence. At each site of the conference holographic displays give in situ participants the impression that those in remote locations are physically present. Gestures, facial expressions, and body language convey much of the information transferred, making the meeting much more efficient than a traditional video or telephone conference, and much less expensive and time-consuming than a face-to-face meeting involving travel.

High-resolution three-dimensional displays, multi-view camera systems, and low-latency high-bandwidth connections, managed by a communications infrastructure and software for three-dimensional rendering, will make this scenario possible.

(E) A starship hovers near a planet. An analyst at a console on the bridge announces an impending existential doom facing the planet's inhabitants. Resources from around the galaxy are quickly marshalled to prevent the apocalypse; the inhabitants are moved to safety or the threat is otherwise neutralised.

This classic science-fiction formula assumes computational capabilities far beyond what we have today. The interesting details of the diagnosis and analysis are hidden; the story's writer has only a vague idea of how they might work. The analyst's console must be fed by a computer system that integrates data from hundreds of multi-band sensors, compares these data with data from other star systems at other times, runs dozens of forward and backward full-solar-system models and uses these to choose the most effective plan of action.

Each of these scenarios, from the recently accomplished **(A)** to the very futuristic **(E)**, requires what we call eInfrastructure. The computers to analyse sensor inputs, process the data, and build models to predict consequences; the networks for fast retrieval and access of data obtained at other times and places, and the technologies of archival, storage, and retrieval. These are all part of eInfrastructure, as are the people, institutes, and services that keep these systems in operation.

1.2 Socio-economic benefits

The benefits of eInfrastructure affect society at all levels, from a family's reliance on weather forecasts, road conditions, mobile communications and the internet, to an industry's planning for future sources of energy, to a government's use of massive sociological databases, to public health and responses to natural disasters[3,5,6,7,8,9]. The increasing complexity of our technology-dependent society demands increasingly sophisticated tools to manage the generation, flow, and storage of vital data. A tool first used by researchers in academia, eInfrastructure has become a fundamental societal infrastructure, financed by public money[2].

We have become dependent upon eInfrastructure to manage society's complexities, and we have accomplished a great deal in keeping up with the rapid growth of the world economy. In the future, eInfrastructure will be even more vital for addressing important challenges, including energy, health and ageing, the environment, transport and mobility, productiv-

▶

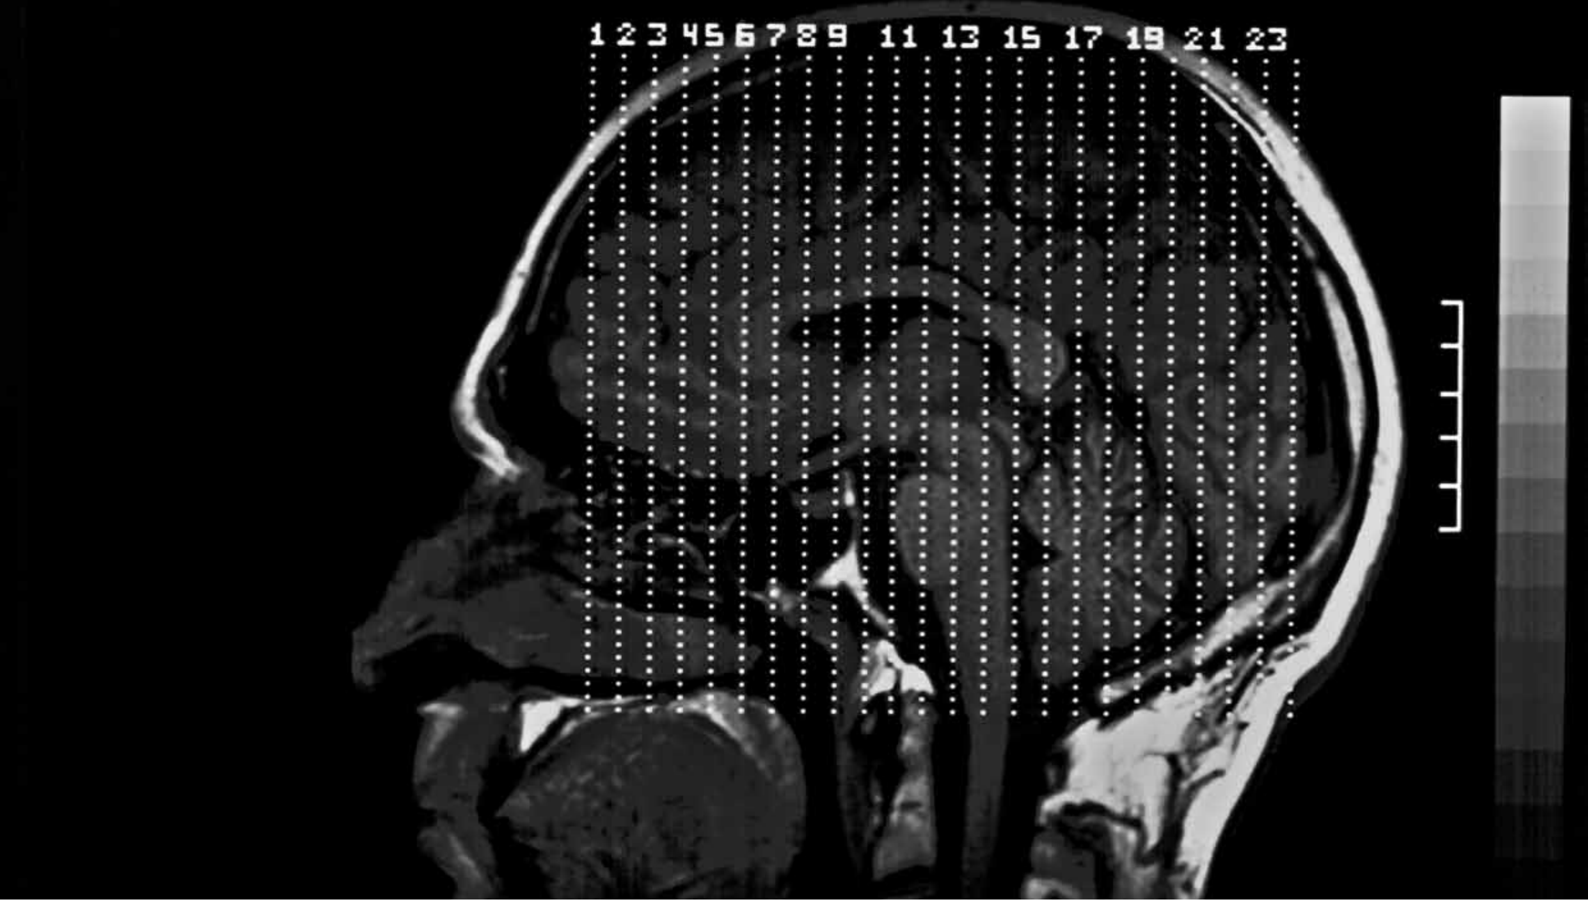

▶ ity, and safety. We must continue to support the growth and expansion of eInfrastructure to keep up with the rest of the world and maintain Norway's prosperity and competitiveness[10]. We summarise some of the benefits in this section; more details are found in *The Scientific Case for eInfrastructure in Norway*[1].

Energy
Our society uses more energy than ever before, and the majority of the world's current energy sources are non-renewable, with a significant detrimental impact on the environment. Solving the energy challenge requires two distinct approaches. We need research into safe and sustainable alternatives to our current energy sources, but we also have to reduce our overall energy consumption. While computers use power, they also contribute to energy savings globally when they are used to improve business processes, to substitute for power-consuming design tests, and to help design more efficient materials and processes. Since they enable on-line media, e-commerce, video conferencing and teleworking, they also reduce the need for physical transport and business trips.

Health and ageing
Sophisticated computer-aided devices that monitor health and assist healing processes, and that can effectively identify diseases at early stages, are in rapid development. New techniques arising from eScience accelerate drug design, control epidemics, enable personal genome mapping, and economically viable health monitoring. As populations worldwide get older, maintaining good health is increasingly dependent on advanced biomedical techniques. The sensitivity of health-related data places stringent requirements on the security and integrity of data storage. An individual's rights to privacy and society's needs to monitor the origin and spread of communicable diseases must be balanced through anonymisation and access protocols.

Environment
The world's growing population and economy create demands on the environment that our finite Earth cannot sustain. The ecological footprint of humans must be reduced. By continuously monitoring environmental parameters, by optimising the efficiency of engines, by reducing or optimising traffic flows, and by otherwise controlling and optimising our impact, eInfrastructure can assist in protecting the environment. Research into renewable energy sources and environmentally-friendly materials is also improved through eInfrastructure.

Transport and mobility
Modern society depends critically on inexpensive, safe, and fast modes of transportation. In many industrialised areas of the world mobility is a nightmare: it is an environmental hazard, it kills thousands of people every year, and is very inefficient. Automation and optimisation of traffic would save energy, reduce air pollution, increase productivity, and improve safety. Both public transportation systems and personal vehicles can be designed with computerised traffic control to provide efficient and rapid routing. Accident-avoidance sensors and autonomous negotiation of rights-of-way can improve safety.

Productivity
Significant improvements in the productivity of multi-institutional collaborations, in all sectors of the economy, result from better eInfrastructure in research and industry. In order to remain at the forefront of global competition, and produce more and better goods at a lower price or more quickly, businesses and nations must continuously improve the efficiency of all processes.

> "While computers use power,
> they also contribute
> to energy savings globally."

Safety
Advancements in eInfrastructure enable individuals to gather information necessary for the protection of life and property. Critical systems can be monitored by controllers that recognise failing components, and deal with them under constraints of timing and functionality. Law enforcement is improved with more sophisticated analysis and forensic means, and national defence is enhanced through more intelligent data gathering and improved logistical systems.

1.3 Implications for Norway

The considerations above all apply to Norway as they do to the rest of the world, of course, but the particular circumstances of our location, our resources, and our social structures make certain aspects more important. Norway's energy use per capita is among the highest in the world. Even though much of our electrical energy comes from renewable sources, greater efficiencies achieved through eInfrastructure could contribute to reducing global dependence upon non-renewable sources through freeing up more Norwegian renewable power for export.

Research groups in Norway, whether industrial or academic, tend to be small and remote from one another and they require that facilities for electronic communication, data transfer, and data storage be more robust and flexible than research groups in large urban centres.

Human lifetimes in Norway tend to be long, and the problems associated with ageing and health care correspondingly severe. The remoteness of much of the Norwegian population provides both a need and an opportunity to develop innovative new forms of remote diagnosis and care, aided by a robust, reliable, and efficient eInfrastructure.

The Arctic region will very likely be affected more rapidly and more significantly than other regions under global warming. It is vitally important for Norway to understand how and when these changes will occur, and what investments to make to adapt to the world we will inherit.

The challenge of climate change is also a golden opportunity for Norway to take a leading role. Its geographical location and topography, its access to plentiful renewable energy and the opportunity to make use of waste heat will enable it to maintain and expand its eInfrastructure to address these challenges. A high level of knowledge of, and use of, high-performance computing and associated technologies must be supported in Norway, particularly as our transition from a resource-extracting economy to a knowledge-generating one continues and accelerates. Only by expanding Norwegian use of — and state support for — eInfrastructure will Norway continue to be economically competitive in the world of the future.

1.4 Computational science

As discussed in our previous document, *The Scientific Case for eInfrastructure in Norway*[1], computational science is the bridge that connects observation, theory, and experiment. Detailed mathematical models simulate physical phenomena from chemical reactions, to the behaviour of biological systems, seismic waves, stars, and even people and financial markets. The value of these models is limited by the available computing power: with greater power, more detailed models lead to more accurate and reliable results.

In global climate modelling, for example, results become more accurate as more subsystems are modelled: the entire atmosphere from troposphere to exosphere; the hydrosphere of oceans, lakes, and rivers; the cryosphere of ice sheets,

▶

▶ glaciers and icebergs; the biosphere of animals, plants, and cities; and the solid earth with its mountains, volcanos, and desserts. The application of these coupled models requires computing power that is not yet available in the world's largest supercomputers. In the next decade, as we discuss in Section 3, supercomputers may increase in power by a factor of a thousand, which will help in addressing the most complex problems facing society, including the development of accurate long term weather predictions, better understanding of climate change, personalised drugs, and predictive health care based on detailed DNA screening.

Scientists are concerned with the correctness, validity, and usefulness of their models, and spend their time using computers as tools to solve their problems. Programmers and computer scientists are concerned with the programmability, portability, efficiency, performance, and getting the most out of the resources available. The bridge between a scientist's equations and the final application are the algorithms and abstract programming models of computational science.

Extensive collaboration among scientists, programmers, and computational science is needed to ensure the best use of the supercomputers of the future. This collaboration should develop domain-specific frameworks and toolboxes for expressing the algorithms and making them more readily portable between different systems. These frameworks will speed up program development, and hide the intricacies of parallelising computational kernels. Current methods are inadequate to deal with future Exascale[11] systems with millions of cores, especially considering the likelihood of component failure during the execution of a program[9].

With the development of domain-specific frameworks, scientists in many other disciplines will be able to take advantage of high-performance computing, and the field of computational science will become ever more important.
The recent dominance of multicore processors and the widespread use of heterogeneous architectures have forced a shift towards hybrid programming models. This and the rapidly growing importance of energy efficiency make parallel programming more challenging than in the past. Developing domain specific languages makes it easier for scientists to accomplish programming tasks, but harder to take best advantage of parallelism and diverse architectures. Parallel programming is today more diverse, but also less stable than it was a decade or two ago. Research projects aiming at maximum performance must therefore include computer scientists in joint efforts to address these challenges.

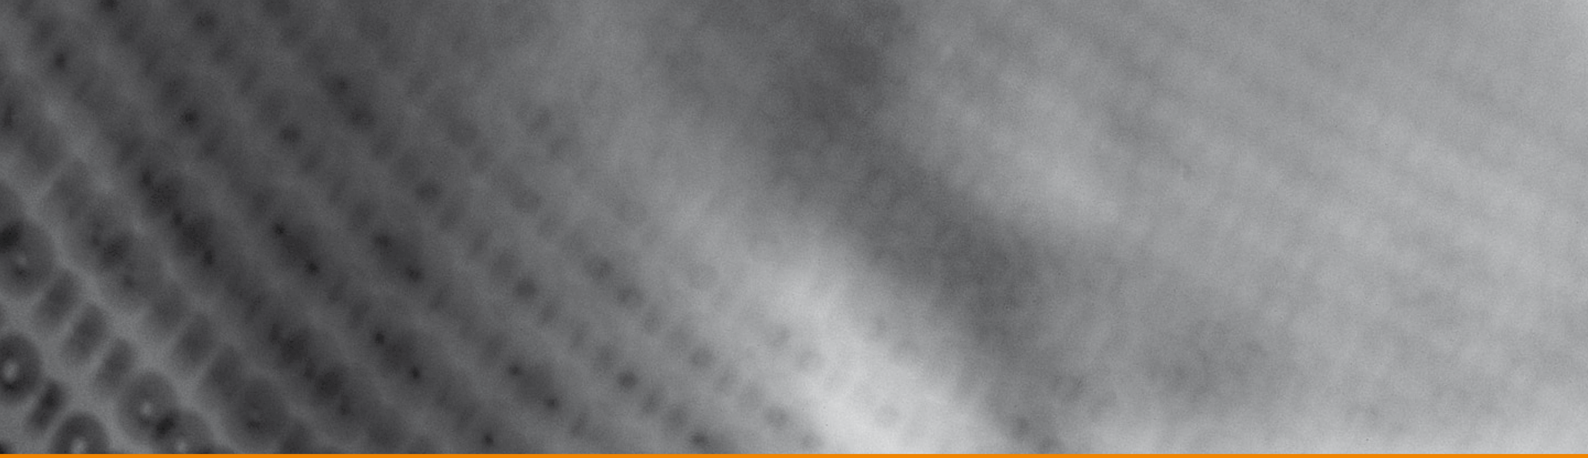

2

Definition of eInfrastructure

> “eInfrastructure is a broad-based public service.”

2 Definition of eInfrastructure

Facilities which grant access to networks, grids, data resources, software and support are defined as eInfrastructure, and the scientific research enabled by it is known as eScience.

The European Commission defines eInfrastructure as the "new research environment in which all researchers — whether working in the context of their home institutions or in national or multinational scientific initiatives — have shared access to unique or distributed scientific facilities."[12] The conduct of scientific research that is enabled by eInfrastructure is known as eScience. Another phrase used to describe these developments is Europe's "Digital Agenda", the idea that Europe should "build its innovative advantage in key areas through reinforced e-Infrastructures and through the targeted development of innovation clusters in key fields."[13]

The technologies of eInfrastructure include computer facilities and peripherals; high-performance and high-capacity networks; grids and collaborative environments; support for software development and life-cycle management; tools to manage and share resources, data, and on-line content; and the applications that produce research. The services to install, manage, and maintain these technologies are also part of eInfrastructure.

Using eInfrastructure, researchers share access to data collections, advanced tools for data analysis, computing resources, and high-performance visualisation. New opportunities arise from remote access and new scientific communities emerge; researchers working in different fields but on similar challenges attain new levels of collaboration and new ways of sharing data, with sophisticated new simulation tools and virtual environments.

We regard eInfrastructure as a broad-based public service, not affiliated with particular communities. The direct users of eInfrastructure are researchers and students who run their applications on the supercomputers, collect and distribute data over the networks, and use storage systems for archiving, maintaining, and retrieving their data. The indirect users are the entire public at large: those who use weather forecasts for planning purposes, who benefit from medical advances made possible through research, whose lives or property are saved through a better public understanding of geological hazards, whose prosperity is affected by the nation's economic performance, and so on.

Because the ultimate beneficiary of eInfrastructure is the public, these resources are funded generally by society, through taxes paid to the state, just as power infrastructures, airports and transport are.

For a more complete discussion of eInfrastructure, please see our document *The Scientific Case for eInfrastructure in Norway*[1].

3

Growth

> “From 1995 and onwards, Norway has been significantly outperformed by her neighbouring countries when it comes to supercomputers.”

3 Growth

Norway's supercomputers should always be counted amongst the 500 fastest supercomputers of the world. As the capability of supercomputers keeps growing at tremendous speed – outdating the most powerful computer in a mere seven years – this demands significant investments in the years to come.

The growth in capability of supercomputers shows no sign of slowing, and the position of two Japanese and two Chinese computers within the top 5 of the Top 500 list of supercomputers[14] shows the ambition of their computer scientists and the determination of their governments to become dominant in this area. Until the present decade, the United States had long dominated the top of the list. Europe has been developing strong regional centres and placed supercomputers in the top ten on most of the recent editions of the biannual listing.

The chart below, from the Top 500 website (with points from Norway added in by hand) shows the evolution of computational power over the past two decades. The vertical (performance) axis is logarithmic; each step represents a ten-fold increase in power. The extraordinarily rapid development of technology is apparent: supercomputer power is exponential with time, increasing by a thousandfold per decade. The fastest supercomputer in the world at any given time falls to number 500 in only seven years, so today's number 500 would have been top on the list in 2004. The trend lines on the diagram are extrapolations of the past growth into the future. If these trends hold, Exaflop (1000 Petaflops) capability will be reached within the present decade, possibly as early as 2018. Today's number 1 has a performance of 10.5 Petaflops, and it is therefore reasonable to expect that number 500 in 2018 will likewise be operating at a handful of Petaflops, tens of times faster than any computer presently running in Norway.

Top 500 performance projection of the most powerful computer systems in the world.

Where should Norway place among the world's top supercomputers? The red dots on the figure above show the historical performance of the top Norwegian computer, for the years in which we scored within the Top 500. We struggle to be on that list. Norway's fastest supercomputer, Hexagon in Bergen, registered at number 488 in June 2011, but it dropped entirely off the list in November 2011. By way of comparison, the top computers in Sweden, Denmark, and Finland sit, respectively, at numbers 44, 70, and 202 on the November 2011 list. Periodic injections of funding give us capability, but then we are left behind.

The standing of the top computer in each of the Nordic countries as a function of time is shown in the graph below. During most of the time from 1995 onwards, Norway has been significantly out-performed by our neighbours. Sweden and Finland, in particular, have managed to stay within striking distance of the top of the chart for all but a few years in this interval, while Norway and Denmark consistently fall behind.

▶

Performance of the top Nordic computer systems on the Top 500 list of most powerful computer systems in the world.

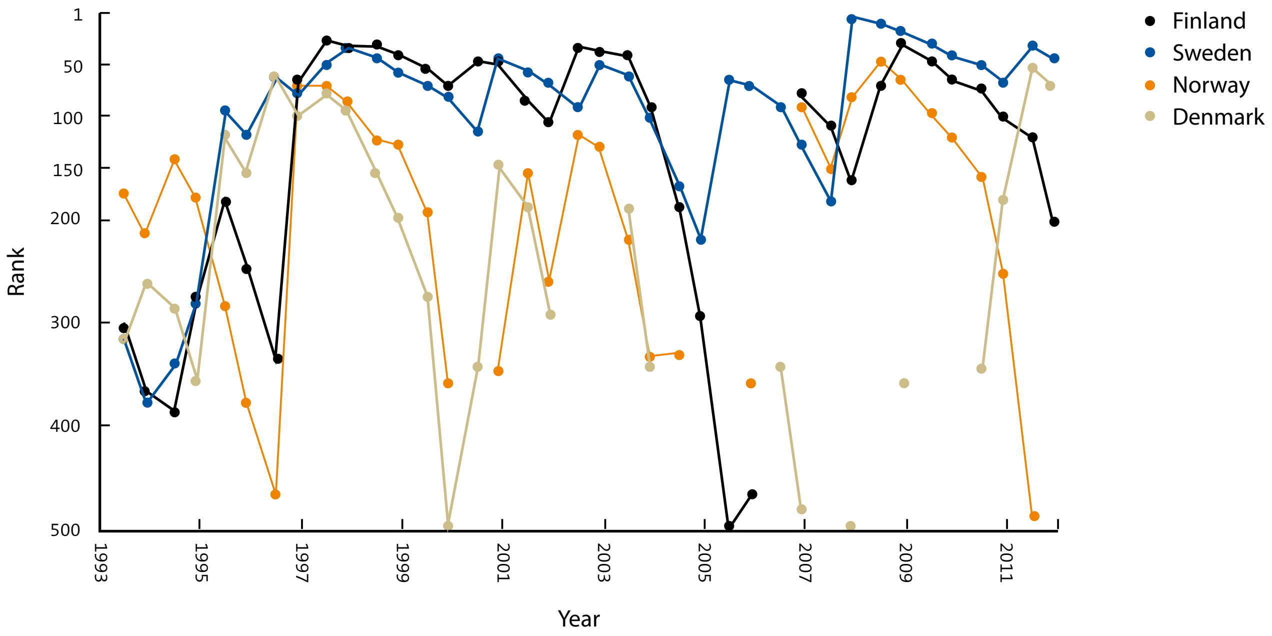

▶ Norway should do better than this. But how much better? Norway is a small nation with one of the richest economies in the world. Our scientists, industries, and the public sector are among the world's most productive; but to retain our competitive edge, we cannot risk falling behind in our use of eInfrastructure. There is now a substantial gap in performance between the supercomputing that is available here in Norway and abroad. *If this gap continues or increases, scientists trained on domestic facilities will fail to be competitive in the broader world.* We must narrow the gap and move Norwegian computing up to a world-class level.

The EU has a gross domestic product that is 10% greater than that of the United States, yet North America has more than twice as much computing power as Europe. Presently the United States has 52% of the computing power in the Top 500, while Europe has 20% and Eastern Asia 21%. Europe will probably boost its share: the European Commission proposes to increase spending in eInfrastructure by a factor of 4. Eastern Asia will no doubt increase its capacity as well. Norway's gross domestic product is 3% that of the EU, while its population is only 1% that of the EU. Norway is highly productive and has both unique advantages and needs. We have green energy, a technologically sophisticated workforce, and good uses for waste heat. The needs we have – understanding the sensitivity of our region to climate change, Arctic ecology, energy exploitation, and natural hazards – represent unique challenges worthy of high-performance computing. Siting a very large supercomputer here, perhaps shared with our international partners, could attract some of the world's top talent.

At present, Norway is significantly under-represented in the supercomputing arena, with no computers among the world's Top 500. To keep pace with Europe alone, based on consideration of gross domestic product, we need to quintuple our efforts. Our present eInfrastructure investment is approximately 20 MNOK annually. Keeping pace with present spending in Europe thus requires investment of order 100 MNOK annually. To keep pace with North America, we would need to double that number. If Europe indeed increases its investment by a factor 4, we would need something approaching 400 MNOK annually to keep pace.

The top few systems in the world are 50 to 100 times more powerful than most of those on the Top500 list, as shown in the chart at right for three separate lists, published in November 2004, 2007, and 2010. The advancement is so rapid that to stay on the list for more than 3 years requires that a system be at number 30 or better when procured, almost an order of magnitude faster than number 500. We must therefore aim to procure systems that are within the top 30 of the Top500 list at the time of purchase. ▶

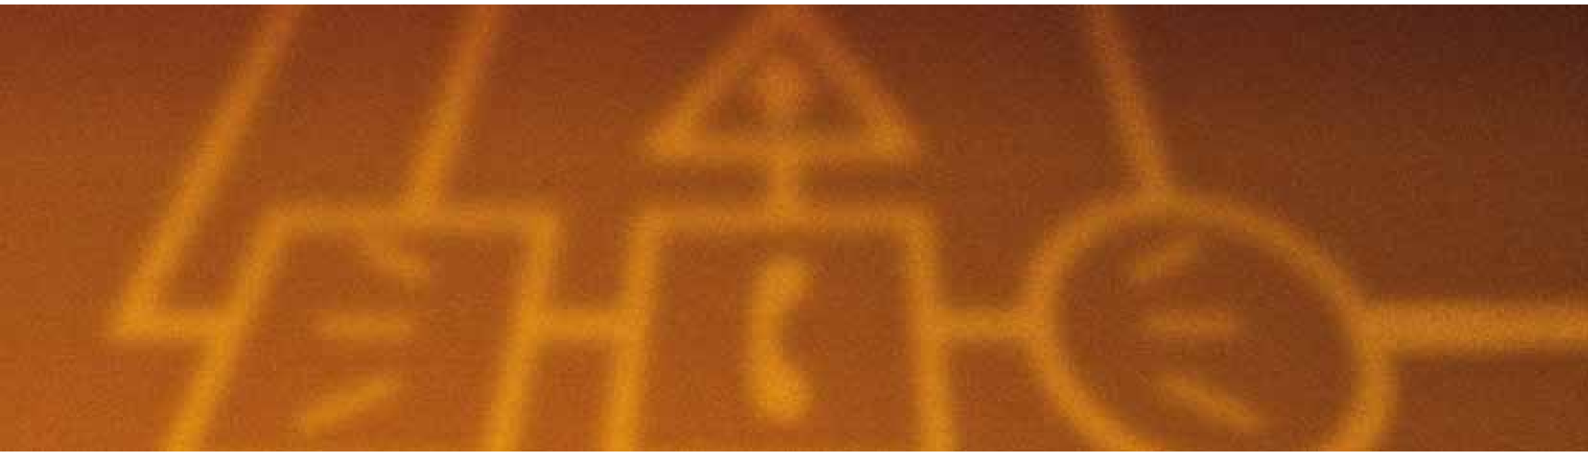

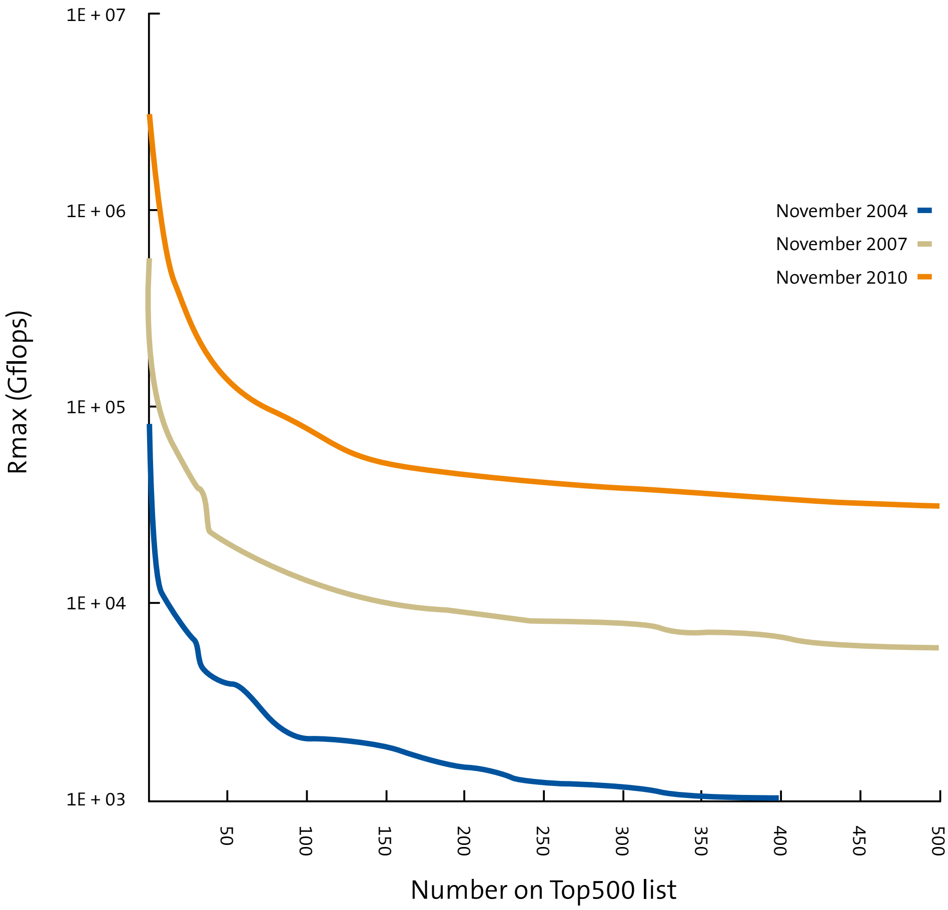

The computational power of the largest computer systems in the world for three separate Top 500 lists, published in November 2004, 2007 and 2010. The top few systems are 50 to 100 times more powerful than most of the other systems in the list.

3.1 Trends

Recent advances in addressing societal challenges through computing have been made possible because of the exponential growth in performance over the last decades. The pursuit of increased performance at lower power consumption and lower cost has led to multi-core parallelism. But parallelism has made computers harder to program. Tools must be developed so that applications written in domain-specific languages are automatically compiled into code that best suits parallel and heterogeneous architectures. As systems grow to much larger sizes, it will be necessary for systems to repair themselves when components fail, and adapt themselves to varying workloads and environmental conditions[3,6,9].

Trends in computing technology
Processor clock frequency has stopped increasing because of power constraints. Performance increase within a processor is accomplished by introducing more cores and more threads per core. *Accelerators* integrated into the processor packaging present an opportunity for improved performance. *Memory subsystem performance* is increasing through non-uniform memory access (NUMA). Efficient utilisation of a compute node will require increased focus on data locality. *Performance of communication networks* in multi-node compute systems will increase significantly, while the fractional cost of the network should decrease. *Energy efficiency* is improving.

Trends in applications
Researchers will have *universal access* to all their data wherever they are, using desktop computers, laptops, tablets, and smart phones. Information will be interchanged and synchronised among all these devices. Data will be stored and computations done at data centres, within access networks, or distributed among other locations. *Personalised services* will become more specialised with preferences automatically taken into account. Traffic advice, location- and context-specific searches, media fitting our personal taste and format, and accessibility adaptations for the handicapped will be available.

Satellites, probes in hostile environments, security systems, and robots are all limited by their ability to understand their surroundings. With more *intelligent sensors* capable of reacting to surrounding events in real time, such devices will provide better science and better services, and lead to enhanced public safety. Consumer electronics employ high-performance multi-core *embedded systems* as microcontrollers. Future applications in safety and security will perform complex analyses on data gathered with intelligent sensors, and initiate appropriate responses to avoid danger or to extract further information. Among the users of such embedded systems will be the automotive, aerospace, and avionics industries, and scientific research.
In all these areas, real-time high performance will be achieved at low cost, low power, low temperature, and high dependability.

Business trends
Low cost microprocessors and integrated circuits surround us everywhere, and semiconductors have become commodities, creating *horizontal markets*. Interfaces are standardised, and tools, foundries, and software are shared, resulting in lower costs. The industry is converging towards modular systems built from a relatively small number of commodity components: standard nodes with standard interconnects. Unfortunately, optimisation at the chip level may not improve overall system optimisation, and requirements of interoperability and communication may induce performance penalties. The convergence of *software standards* among different devices renders the hardware platform irrelevant. Common software

development helps define applications that can efficiently use the hardware resources available on a device.

New communities and opportunities for *collaboration* have arisen. People contribute time and share expertise as never before. The development of Linux and the gnu compilers and utilities, built and offered under free licenses, re examples, the result of work by hundreds of specialists. Wikipedia, social networking sites, YouTube and other portals make knowledge, expertise, and entertainment available to all at little or no cost.

Infrastructure as a service. Processing power and storage services, rather than hardware, are increasingly offered to end users. This began with the cloud computing offered by Amazon and Google, and others have followed suit: data and resources for the end user are stored on a company's servers. Providers that require large infrastructures for peak needs can serve large numbers of customers at non-peak times. Public eInfrastructures could operate in this mode to the benefit of society.

3.2 Technological constraints and challenges

The increase of the number of transistors on a chip has until recently been accompanied by a reduction in supply voltage, keeping the power envelope fairly stable. But further reducing the supply voltage leads to increased power leakage, offsetting the savings. So more cores, instead of denser and faster cores, are used. But running all of them at full pitch also creates power problems, so it is an advantage to make some of them different, to be used in different stages of a calculation. Thus the future of computing is multi-core parallelism and heterogeneity[15].

The economic lifetime of software is longer than that of hardware. Porting software to completely new hardware platforms is expensive, and leads to errors and instability. In some cases, hardware is kept in operation for a much longer period than planned, simply to run legacy software. Identifying concurrency in legacy code is extremely tedious and difficult, and it is almost always preferable to construct a new code from scratch. Newer hardware includes features that only the latest software can use. Language abstractions already exist for this functionality, but are inefficient. The design of better such abstractions, with associated compilation, debugging, and run-time support, is required.

Parallelism adds to the complexity of programming, but cannot be avoided for all users. Higher abstraction levels in programming have become available through tools tailored to specific problem domains, enhancing developer productivity. Concealing the parallelism has been a goal of such domain-specific solutions, but there is a trade-off between higher abstraction levels and higher performance. The increasing use of multicore and heterogeneous architectures and the consequent requirement for data locality has made parallel programming much more challenging. It has therefore become both more important and more difficult to develop abstractions that hide parallelism and other architectural details while achieving high performance. It will become increasingly necessary to foster close collaborations between computer scientists and discipline researchers.

System complexity increases alongside performance, not only because systems are composed of more hardware and software components of various origins, but also because of interconnections to other systems. Impacts of local modifications are felt at the system level, and understanding all the implications of a modification is difficult. *Communication* within a supercomputer system is a fundamental limitation. As proces- ▶

▶ sors become smaller relative to the distances between them, the cost of data transfer increases, and data locality becomes essential. Managing the locality is expensive, however. *System components are unreliable.* Extremely small feature sizes imply that transistors and wires are less robust. Failures will occur, and there will be high variability of global system parameters over time. Sporadic errors will become more frequent and new techniques must be developed to handle them.

Cost-per-performance and power-per-performance have become extremely important for computing centres. Petaflop clusters typically cost of order 50 million euro and use Megawatts of power. Bigger computers use more components, but costs of components have gone down over the years, so the fastest supercomputers have not increased greatly in cost. Similarly, the energy efficiency of components has improved, so power usage has not increased as rapidly as performance has. Nevertheless, the minimisation of both these ratios is critical in the design and placement of new large computing systems. For Exascale supercomputers that are expected to be available around 2020, energy efficiency will be the main design challenge.

There are quantum-mechanical limits to the reduction of transistor size and speed-of-light limits to communications. If growth is to be sustained, the industry will move beyond silicon to radically new technologies such as nanotubes, molecular computing, quantum computing, optical computing, or biological cells. Possibly several of these will be exploited for different tasks.

3.3 Norway and Notur

Today, some 680 researchers spread among more than 130 projects are actively using the computational resources managed by Notur. These users expect the national infrastructure to provide them with competitive resources for computing and storage, including high-quality user support and access to international collaborations. The projects vary in size, ranging from groups of more than 50 researchers to a single senior researcher, and ranging from allocations of 12 million core hours in a year (corresponding to sustained use of more than 1400 cores) to allocations for testing or benchmarking purposes. Current and future research activities are often parts of collaborations that are organised nationally or internationally. Because collaborations implicitly trust that eInfrastructure services will be dependable, it is crucial that funding for them be sustainable and predictable.

The data network for research is used by roughly half a million people, inside and outside of academia, who need to access data stored on the system. Network traffic has shown significant growth through the lifetime of Notur. Between 1995 and 2001 traffic was roughly doubling every year, but since then the exponential growth has slowed to a doubling every four years. Even at this slower growth rate, periodic boosts in network capacity will be required.

3.4 The computational ecosystem and pyramids

High performance computing, essential to scientific research, is an important component of eInfrastructure. It must provide adequate computational performance for all scientific disciplines, fundamental and applied. In *The Scientific Case for eInfrastructure in Norway*[1], we detailed how use is spread among traditional and emerging disciplines. High performance computing must cope with a wide range of usage patterns, ranging from ensembles of smaller jobs with weak coupling to tightly coupled, highly parallel jobs. Special features such as high-performance interconnects or

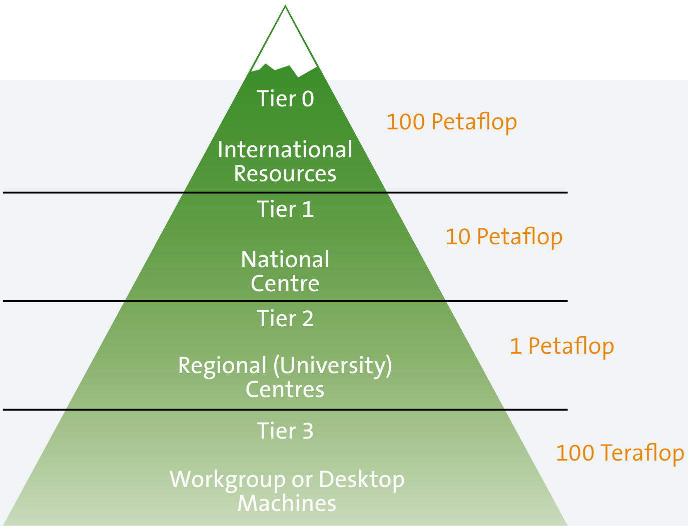

High performance supercomputing pyramid.

large node memories are needed by some jobs, while many others can be run efficiently on clusters built on commodity components. Some groups must perform similar computations periodically, often under time constraints, while others need continuous long-term access.

A sustainable infrastructure for high-performance computing consists of systems of different sizes and architectures to support different application loads in a cost-efficient manner. This leads to a vision for the overall infrastructure as a pyramid. The top represents massively parallel computing for tightly coupled, latency-bound problems. Scalability and overall application performance are of vital importance. Facilities are large in size and expensive. The next lower layer in the pyramid represents throughput computing and includes a larger number of smaller and less expensive high-performance systems for calculations with more modest requirements.

The performance levels indicated at the right of the diagram above are indicative of what we might expect Norwegian scientists to have access to in the next decade. Users gain experience at lower levels of the pyramid before stepping up to systems that are ten times more powerful at the next level, where the scarcity and high value of resources demand stringent control of performance and competence. The diagram is shaded greenest towards the top, because energy efficiency is expected to be paramount in the selection of systems that run at the top tier. Nevertheless, energy efficiency is important at all levels of the pyramid because of the very large number of systems at the bottom. The snow-capped peak on this pyramid indicates the desirability of having the very largest systems in locations where cooling is not problematic.

The depth and breadth of the pyramid grow with time, adding entry levels to the tiers of massively parallel and throughput computing. As disciplines without a computational tradition turn towards eScience, entry levels must cater to users with varied backgrounds and levels of algorithmic and technical expertise. These users require a shift of emphasis towards a service-oriented eInfrastructure, with support in acquiring, using, and adapting software and data. Inclusion of entry-level services in the national eInfrastructure is important for reasons of scalability, upward mobility, and training. At the base of the pyramid are desktop computers, multi-core and sometimes multi-processor, and equivalent to the supercomputers of a previous generation. Even embedded systems such as smartphones and tablets may be considered at the bottom of the pyramid. Those who use such devices may not be aware that things called threads are generated by the applications they use, and that these threads are farmed out to various parts of the system for processing.

Parallelism is everywhere now, and cannot be avoided. Even experienced scientists are turning more and more frequently to packaged applications that hide much of the complexity of concurrent programming.

3.5 Hybrid architectures & hybrid applications

Among the trends in the computer industry is an increase in the heterogeneity of architectures within even a single processor, with distribution of memory from multi-level caches to system-wide memory[3,6,9]. A scientist writing a program in a high-level language is insulated from that complexity, and needs the program to run efficiently on whatever architecture it lands on. It is therefore desirable that compilers and operating systems be made smart enough to recognise and avoid potential mismatches between code and machine.

▶

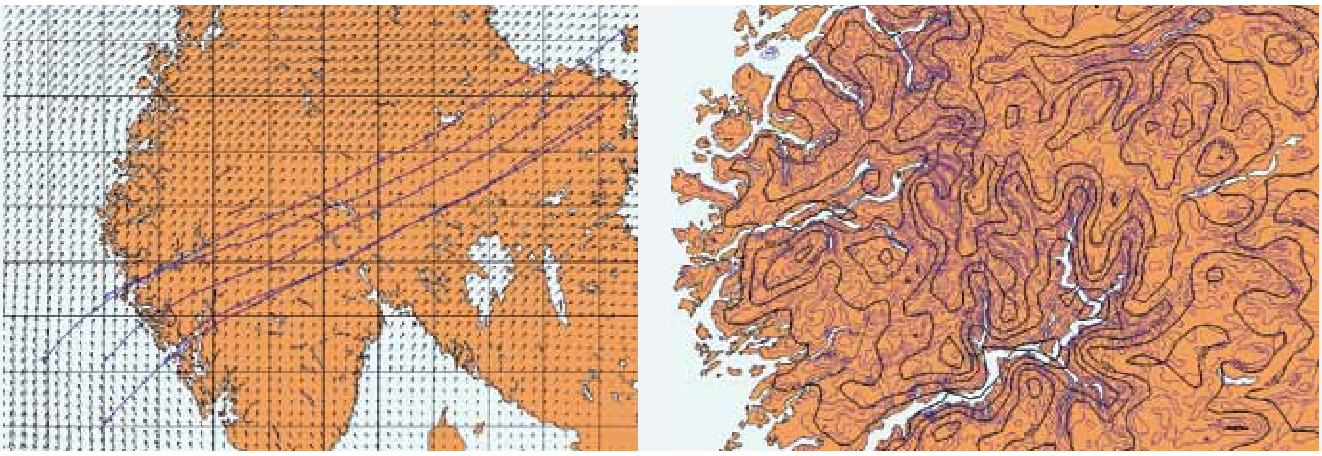

Weather and topography maps. Left: a domain decomposition of a model at 4 km horizontal resolution where each subdomain contains 20×20 grid points. The blue trajectories indicate the air flow through the subdomains. Right: a model topography at 1 km resolution.

Complex computer programs invoke different types of algorithms for different purposes. In an incompressible fluid, like water in a hydropower tunnel, the pressure variations at one point instantly affect all other points. Information is exchanged among all cells in the grid, requiring considerable capacity for data transfer in the computer. On the other hand, pressure changes in compressible air flow (as in the intake of the jet engine illustrated), cause only local variations, and the need for data exchange is correspondingly less. The best computer for modelling incompressible flow will therefore have very fast data transfers among processors, while for computing compressible flow it will simply have as many processors as possible.

Multi-physics codes require different degrees of coupling as the calculation proceeds. A jet engine designer uses equations with local coupling for the airflow in the intake and exhaust, equations with global communications for the reactive flow and heat diffusion in the combustor, and for the stress on turbine blades. Then the calculation needs to take global sums to monitor calculation quality. These global sums require synchronisation signals, so some processers must wait while others finish. Further, the code must produce results, and these often come from large dumps that are later analysed by a post-processor with the eventual production of visualisation data. An ideal system would have compilers, operating systems, and run-time dispatchers that adjust to the real-time balance between local and global communications, and farm out tasks accordingly. Hybrid architectures may fit well, although it will take substantial reprogramming to use them efficiently.

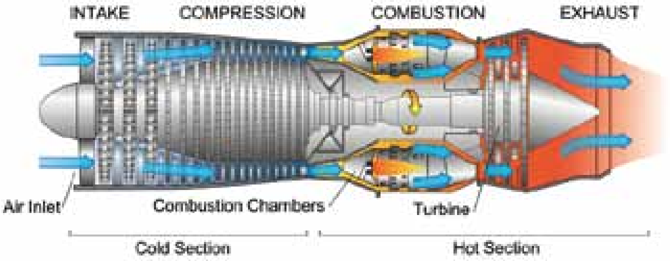

Example of compressible air flow: the intake of a jet engine.

Weather and climate models, and the more complex “Earth System Models” are rich examples of the complexity that future programming must deal with. Different types of equations are used in the different sub-models of the calculations, which produce prodigious and frequent dumps of data that then become the basis for detailed analysis and post-processing.

A numerical weather prediction system consists of a number of tightly coupled sub-processes that depend on each other. The main parts are: observation handling, assimilation of observations into a first guess of the atmospheric state, models for the atmosphere, surface water, ice, and so on. The mere presence of some of the physics routines causes a load-balancing issue for parallel computing. For example, when convection occurs in one part of the grid, the processors for which no convection is diagnosed are idle until the convection calculations finish.

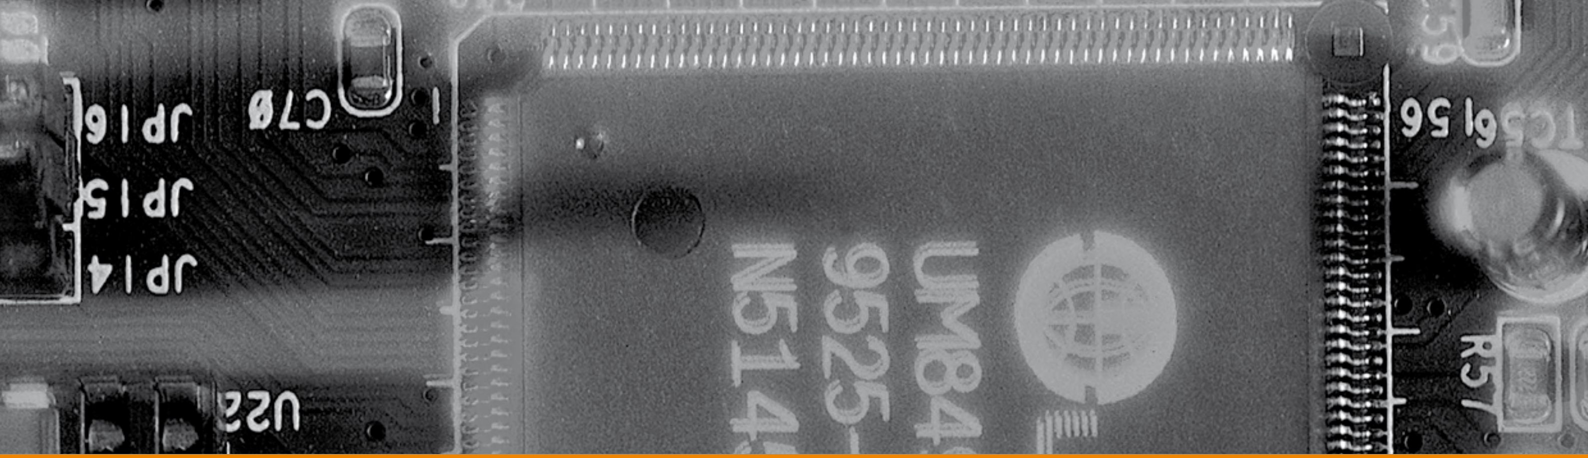

4

International comparisons

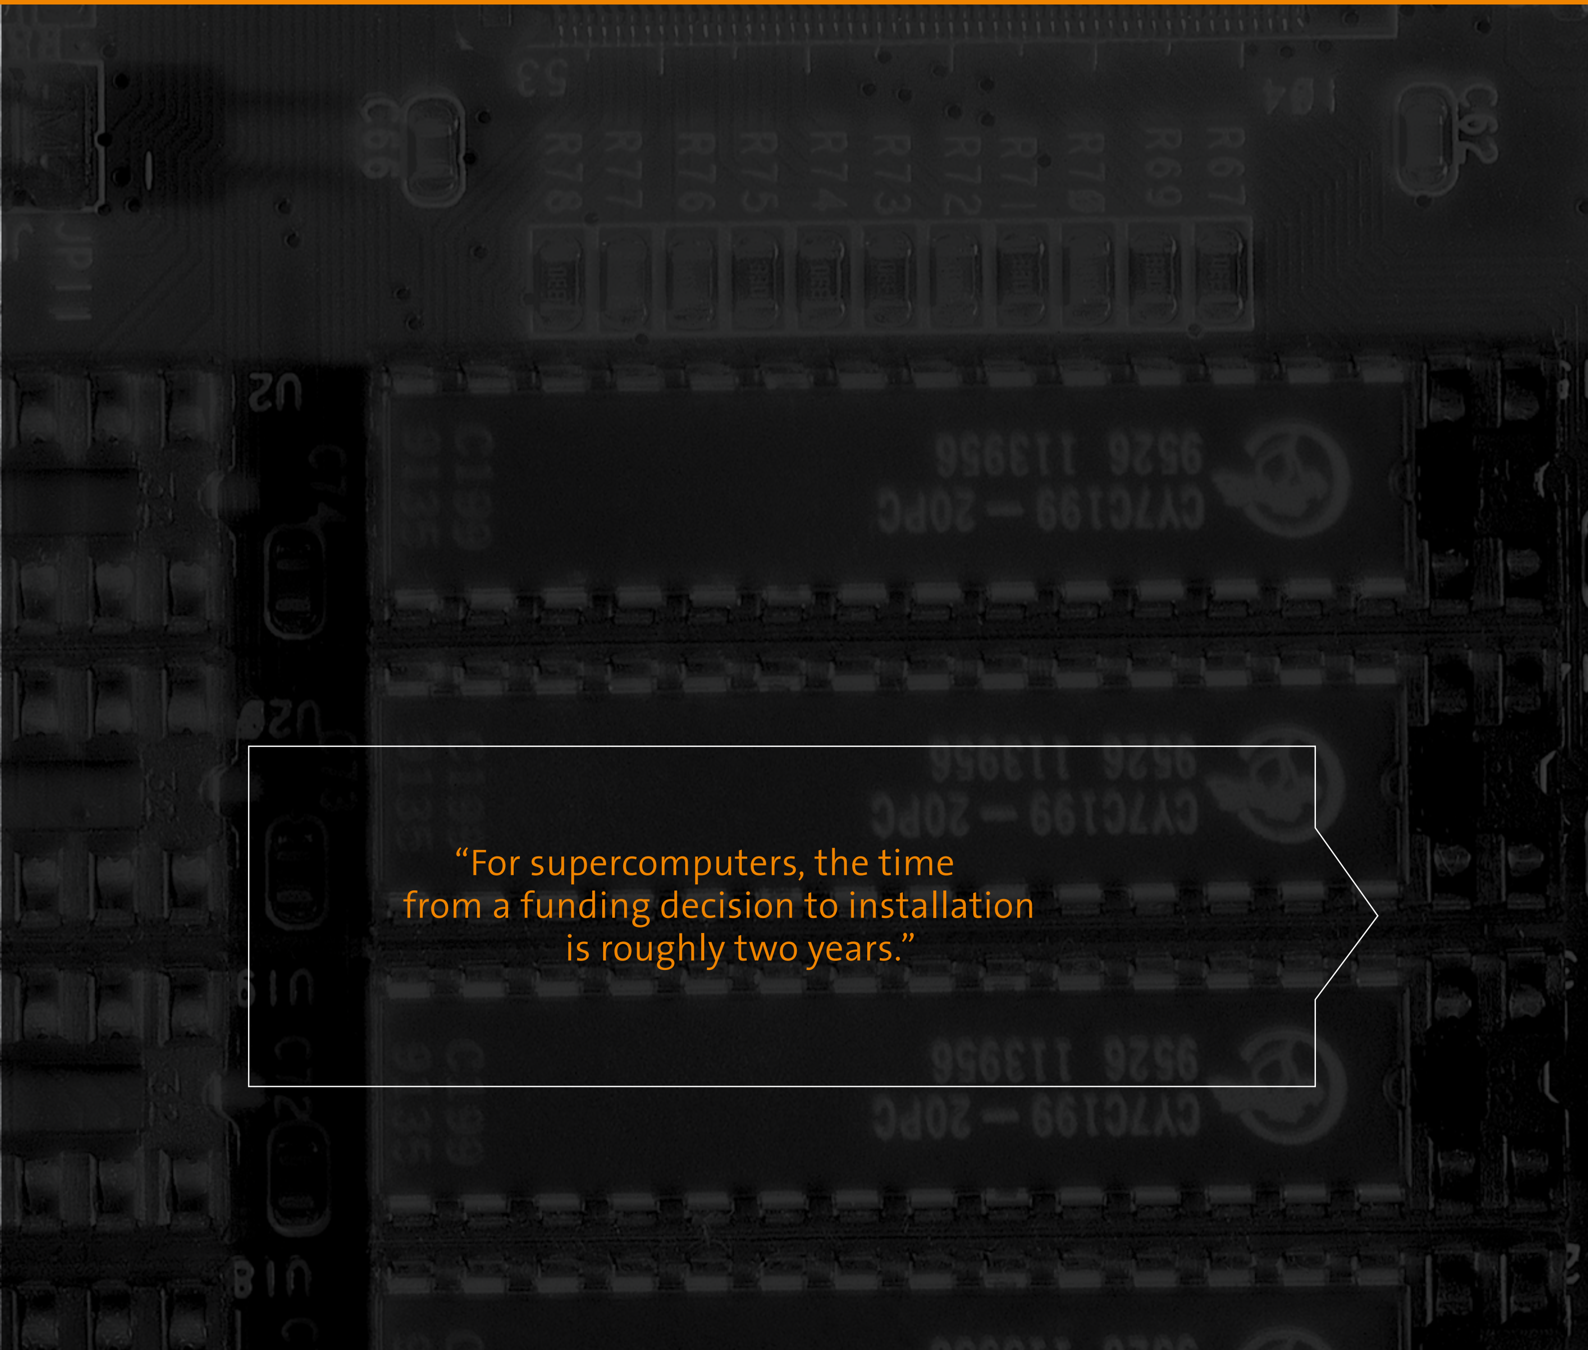

> “For supercomputers, the time from a funding decision to installation is roughly two years.”

4 International comparisons

A greater degree of consolidation of Norwegian eInfrastructure is recommended. This would ensure economy of scale, green power and efficiency and data intensity.

4.1 Organisation

There are several conceptual models for organising eInfrastructure, and different countries have taken different approaches.

In Finland, a single entity, the CSC IT Centre for Science[16] manages the bulk of computational and data storage resources in a single location. It is funded directly by the Ministry of Education, Science and Culture. Economy of scale, lack of redundancy, a critical mass for doing significant research, and a stable and predictable funding stream are significant advantages. Users from all over the country have a single access point and know precisely where to go for resources or assistance. On the other hand, the lack of flexibility and diversity of architectures, and a lack of a direct link to academic programs at the universities may be disadvantages.

In Germany[17], and in the Department of Energy[18] complex in the US, state-funded interdisciplinary research centres include supercomputing facilities as a part of their portfolio. These may be, but are not always, collocated with or near universities. Having a built-in set of researchers as users and developers, and funding that is fairly stable and predictable are advantages. Because there are several such centres, diversity in architectures and approaches can flourish. When these centres are not sited with universities, the lack of an academic link for training is a disadvantage.

Also in the US, the National Science Foundation[19] establishes and funds a number of competitive supercomputing institutes scattered around the country, usually sited at universities but independent from them fiscally and administratively. The available flexibility and diversity of architectures and reasonably good connections to academic programs are advantages. Because these centres must re-compete for new funding every few years, the quality is high, but the uncertainty and instability makes it difficult for them to retain good staff over the longer term. There is also some redundancy of offerings, and great inconsistency from one centre to another regarding availability of support.

In Sweden and Norway, supercomputers are sited at universities and are managed jointly by the university departments in computer science or informatics and by a national entity, Notur[20] in Norway and SNIC[21] in Sweden. The available diversity of architectures and good links to academic programs are strengths of this system, but the disparity of documentation practices, software availability, and user authentication leads to confusion on the part of users.

A tiered structure is used formally by CERN[22] and PRACE[23], and is realised de facto, but without formal agreements, around the world. In an ideal tiered system, the most powerful centre, denoted Tier 0, would be located where access to clean energy and efficient cooling exist, and where competent IT services are available (though not necessarily collocated). Smaller centres, denoted Tier 1, Tier 2, and so on, would be geographically distributed among universities and public or private research centres. See the pyramid diagram in Section 3.4, for example. State funding would fully support the Tier 0 system and contribute to the support of the other Tiers, which would get majority funding from their host organisations or local authorities. Fast networks and grid services would connect all tiers, and credentials and documentation would be uniform across the system.

Comparing the present Norwegian structure with its counterparts in other countries, this panel believes that a greater degree of consolidation is desirable. Economy of scale, green power and efficiency, and data intensity seem to argue in this direction; even ensemble-like calculations (as for climate and weather) would benefit from closer coupling.

> “It is clearly necessary that funding be stable and predictable.”

▶ A single centre (as in Finland) would make high-performance computing more visible to the nation, and make public funding easier to come by. On the other hand, the existing centres at the universities get significant in-kind contributions from the universities in the form of staff, space, utilities, and the reuse of waste heat. Each of the four existing supercomputer centres in Norway have systems for the reuse of waste heat, which would not be possible in a remote location.

An option would be to keep four (or more) centres of competence with smaller systems, user support, and training at the universities while housing the biggest systems where they would be most cost-efficient. Various sites have been proposed[24]. Implementing a tiered structure for Norway is an interesting possibility. Local systems sited at universities can be used as entry-level resources to promote eInfrastructure use for beginning and new users and to train and retain support staff. One or two large state-funded systems, in good locations for power and cooling, would be designed to be accessible for users who have mastered the use of the local entry-level resources, whose requirements exceed the available capacity on the entry-level resource, and whose software satisfies certain criteria. Support staff for the top tiers should remain at the universities.

Harmonising the user experience of support, documentation, batch system protocols, use of scratch space, archiving, and credentials is essential if any consolidation is to take place, and is desirable even if no consolidation occurs. Users must be involved in deciding how resources in Norway should be organised. This is also strongly advocated by the European eInfrastructure Reflection Group[25].

4.2 Financing

State financing of a single entity (as in Finland) or state-supported multiple entities additionally funded from other sources (as in Norway and most other countries) are both possible. But it is clearly necessary that funding be stable and predictable. The time from decision to purchase a supercomputer to its installation, roughly two years, is longer than the annual funding cycle that universities operate under; and obsolescence time is typically only four to five years. Thus a buyer must be ready with the available funds as soon as a new technology is available in order to make the best use of it. This cannot be done if the funding stream is unpredictable. ▶

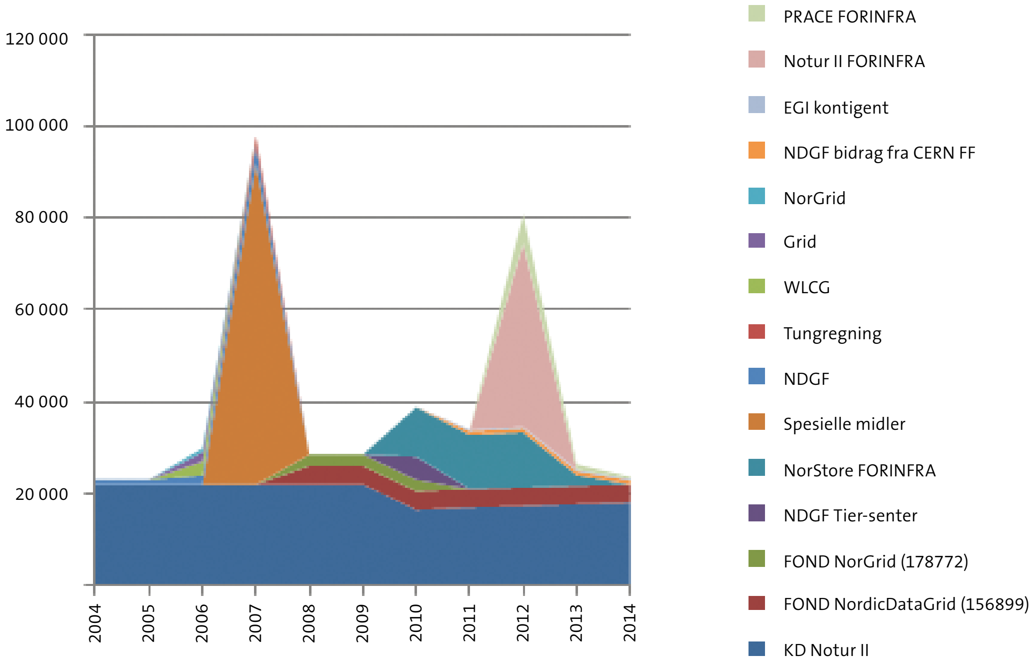

Financing of eInfrastructure projects by the Research Council of Norway.

- The sudden introduction in 2007 of a lump sum funding to Notur of 70 MNOK made possible a leap in the capability of Norwegian supercomputing (see graph above), and an increment of 40 MNOK in the present cycle helps to update the systems purchased then. But such sudden and unpredictable increments do not make for sustainable funding, and wreak havoc with planning. A sensible strategic system for planned rational replacement and upgrading of the Norwegian supercomputing infrastructure requires a steady supply of predictable funding.

The level of financing for eInfrastructure must also be discussed. If Norwegian researchers are to be competitive at the international level, they must be competent enough to run codes on the largest machines in the world. In Europe, the fastest machine at the time of writing, at Jülich in Germany, runs at 2 Petaflops. The fastest in Norway runs at 40 Teraflops, a factor of fifty less. A few researchers in Norway are presently trained to run code at Jülich, but most find the entry threshold too high. This gap must be bridged. While Norway is not likely to have a machine in the top 10 anytime soon, it should aim to purchase machines that are in the top 30 at procurement, requiring a funding ambition of roughly 60 MNOK annually for hardware. To maintain staff at the centres, both for support and maintenance, another 30 MNOK is required, and an additional 10 MNOK for the data storage infrastructure, for a total of roughly 100 MNOK, just to stay abreast of international developments.

Sustainability of supercomputing centres requires that those who staff the centres have permanent jobs, that there is a visible recruitment path for the staffing, and that there are stable funding streams for staff development, for providing documentation and user support, and for purchasing new machines and storage facilities.

4.3 Access

At present, users with Notur projects can have accounts on four different supercomputers and on the Norstore data storage machines as well. These accounts might all have different passwords, different policies on the use of disk space, different batch commands, different file system structures, different support systems, different standards of documentation. Users spend unreasonable amounts of time searching for the right way of doing something on one machine after being accustomed to another, and password policies are inconsistent.

These and other aspects of the user experience need to be harmonised. In Europe, the eInfrastructure Reflection Group has come out strongly in favour of adopting uniform protocols for Authentication and Authorisation Infrastructure. We should do the same in Norway.

5

Necessary steps

> “By school-leaving age, everyone should know what an algorithm is and have some idea of how to construct one.”

5 Necessary steps

We are moving into the era of eInfrastructure, and children must be prepared for a life dependant on it. Mathematics and computational science must find its way into the schools. Specialisation courses should be offered university students. Scientists must receive necessary training, and there must be sufficient specialists making up an adequate staff.

5.1 Education

Children growing up in the twenty-first century are exposed to computers and the internet from the very beginning, and use them for education and entertainment. With growing sophistication, they use them to explore the world, and gradually a few of them become aware that computers can be used not merely to retrieve knowledge, but also to generate new knowledge. As kitchens and garages and backyards have traditionally been the incubators of young chemists and inventors and entomologists, so home computers become the first instruments of young mathematicians and computational scientists.

But the equivalent educational sophistication is not yet in place. You can buy children's chemistry sets, building sets, microscopes, and telescopes, and have been able to do so for a century or more. Where are the equivalent "kits" that would enable a child to master the computer on the desk – and not merely "play with" it? Where are the programming environments that speak to a child's curiosity and willingness to explore? Where are the child-friendly tutorials that encourage innovation?

Before those kits, programming environments, and tutorials can be produced, there needs to be a revolution in education from the ground up. Computers are becoming nearly universally available in the schools, where they are used as instruments for reference, exercises, drills, and examinations. Starting quite early – say as early as third or fourth year, when some elementary mathematics has been mastered – they should be opened up for tinkering with. Starting with simple task automation, progressing to control structures, then conditional execution. By school-leaving age, everyone should know what an algorithm is and have some idea how to construct one. The concept is as fundamental to our technological world as elementary logic has been to previous centuries[3,6].

An analogy can usefully be made with transport. At secondary school age, most students begin to learn to drive a car. At a much younger age, they learned to ride bicycles or scooters and were instructed in the basic rules of the road. In parallel, they learned to use public transport and the manners and conventions thereof. Before the industrial revolution, none of this education would have been relevant because the infrastructure did not exist. Our society evolved to create the necessary education to take advantage of the infrastructure of transport as it was needed.

We are now moving into the era of eInfrastructure, and need to invent the education to prepare children for a life dependent upon it[10]. If Norway were to invest substantially in educating the young for a future dependent on eInfrastructure, Norwegian science and industry would have a substantial competitive advantage in the international science race and in the world economy. All areas of modern life are becoming more and more dependent on the tools of eInfrastructure. Education in the use and development of these tools needs to be started as early as possible.

Young people need to be prepared to enter a world increasingly dominated by computers and eInfrastructure. We are confronted with an unfortunate paradox: just as eScience influences wider and wider aspects of our lives, interest in and preparation for the disciplines enabling eScience (like mathematics, computer science, and the sciences in general) is flagging at the secondary school level. Most high-school teachers have little knowledge or experience in computational science. In teacher-training programs, it would be useful to add units on eScience to training in pedagogy and the liberal arts.

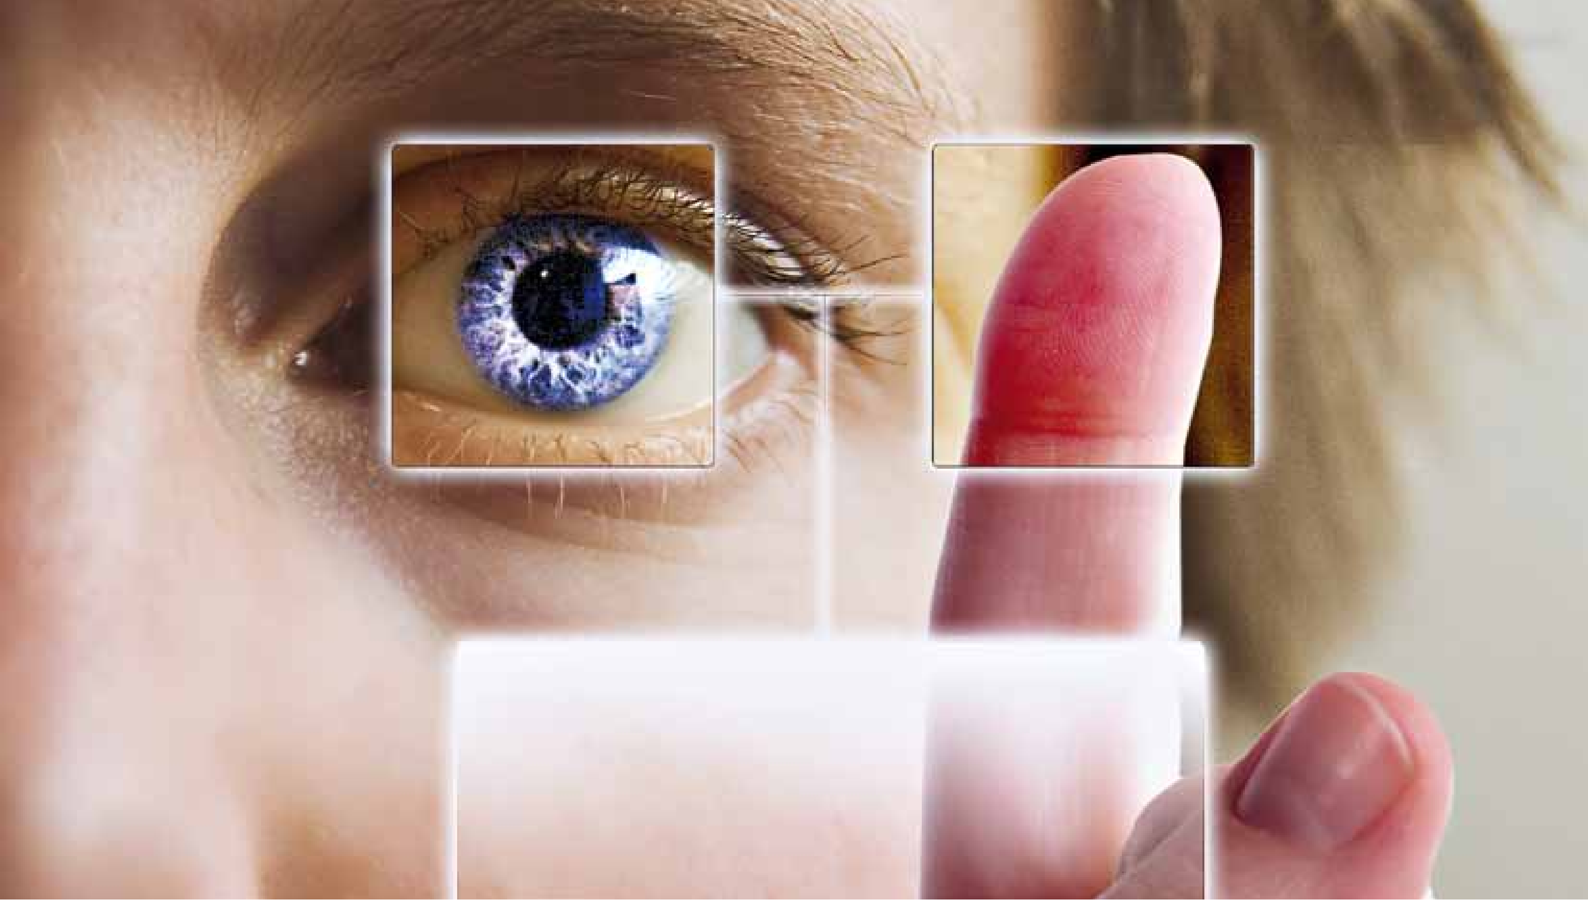

In secondary and tertiary education, specialisation courses in computer science, computational science, and various areas of application should be offered to students who are interested. Students with backgrounds in the natural sciences and engineering are better prepared to develop skills in computational science than their counterparts in the humanities. For the latter, a culture gap as well as a skills gap separates the researchers for whom the computer is a tool and those who program. Experience with computers as appliances for communication, writing, and accounting does not help with numerical analysis or data handling. Attempting to use high-performance computing as an appliance leads to disappointment and frustration.

Effective use of eInfrastructure requires consideration of, and investment in, appropriate education and outreach activities. Intensive courses covering computational tools adapted for different disciplines, like computational medicine or computational linguistics, should be offered by those university departments, possibly through faculty members holding joint appointments in the computer science faculty. There could also be special interdisciplinary Masters degree programs between computational and disciplinary faculties.

The high-performance computing centres themselves should teach specific techniques necessary to make effective use of the computers of the day, through workshops, on-site consultant services, and up-to-date online documentation, while universities should teach the skills in mathematics, numerical analysis, and algorithms that outlive given architectures. Because time on the fastest computers is a valuable resource, the data centres have an interest in assuring that it is used effectively. They therefore need a means of evaluating the skills of their users while deciding resource allocations. This could be done by requiring users to take specific courses on parallel processing concepts and techniques and on specific tools for constructing scripts for job submission and post-processing. In a tiered system, users could use the experience gained at one tier to qualify for entry to the next.

In the early days of computing, there was no distinction between scientists who used computers to conduct research in physics or other fields and scientists who devised algorithms and wrote programs to run on those computers; they were one and the same. Educating young scientists involved simultaneously developing the skills of using computers and conducting scientific research. As computers have become more sophisticated, the gap between those scientists who use computers as tools and those who develop techniques for using computers has grown enormously. Developing the skills necessary to do research in the rapidly changing environment becomes increasingly difficult. New skills that are needed for programming massively parallel machines (for example) are not easily learned, and compete for a student's attention with the scientific background that must be learned for the conduct of research.

5.2 Software

Early computers had single processors and computer programs were therefore written to run in serial mode, executing one instruction after another. Modern high-performance computers have many more processors; to make use of them, programs have to be rewritten to execute instructions in parallel. Legacy codes, inherited from the early days of computing, containing perhaps several million lines of code, are difficult to adapt and must often be entirely rewritten from the ground up. As architectures change over time, rewriting is nearly a continuous process[3,6,26].

If a program runs 100 times faster on a computer with 100 processors, it is said to scale linearly. Such perfect scaling is

▶

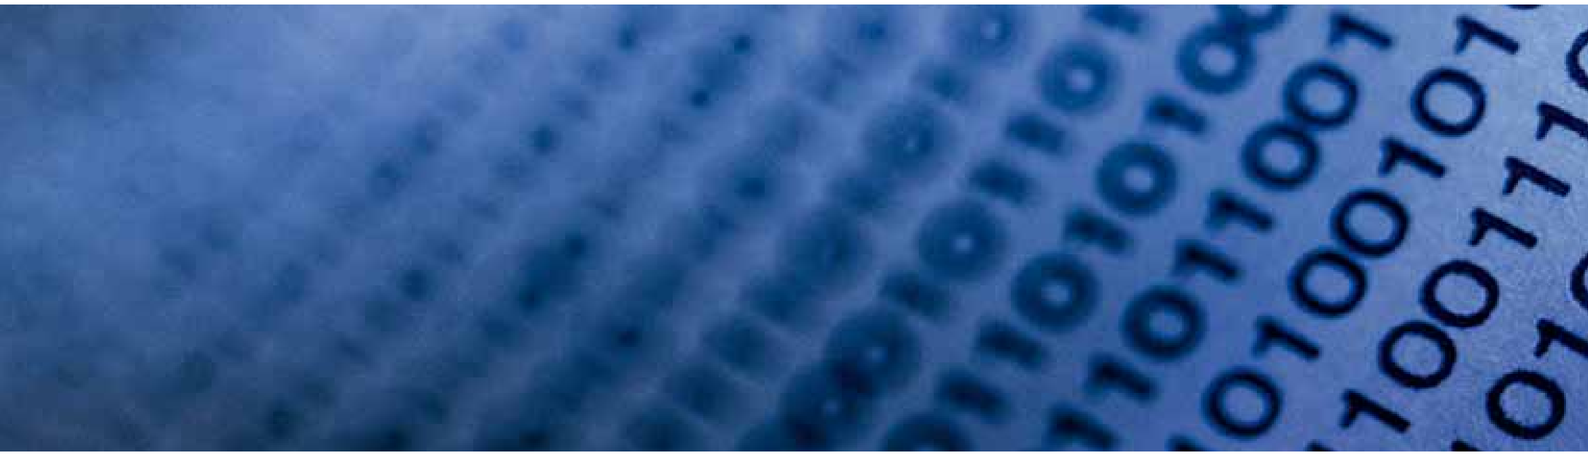

▶ rarely achieved, however. Bottlenecks within the code, latency in interprocessor communication, and unequal distribution of tasks among processors, all limit parallel scalability. The table below illustrates how the open source CFD program OpenFOAM scales for a given simple problem on different numbers of processors on the Hexagon cluster in Bergen.

The effectiveness of parallelisation depends on the algorithms used, the skills of the programmer, and the computer itself. The same code may scale well on some computers and poorly on others, depending on the speed of communication among the processors.

Number of processors	Efficiency (%)	Wall-clock time (seconds)
4	100	4000
8	100	2000
16	100	1000
32	94	530
64	88	280
128	78	160
256	61	100

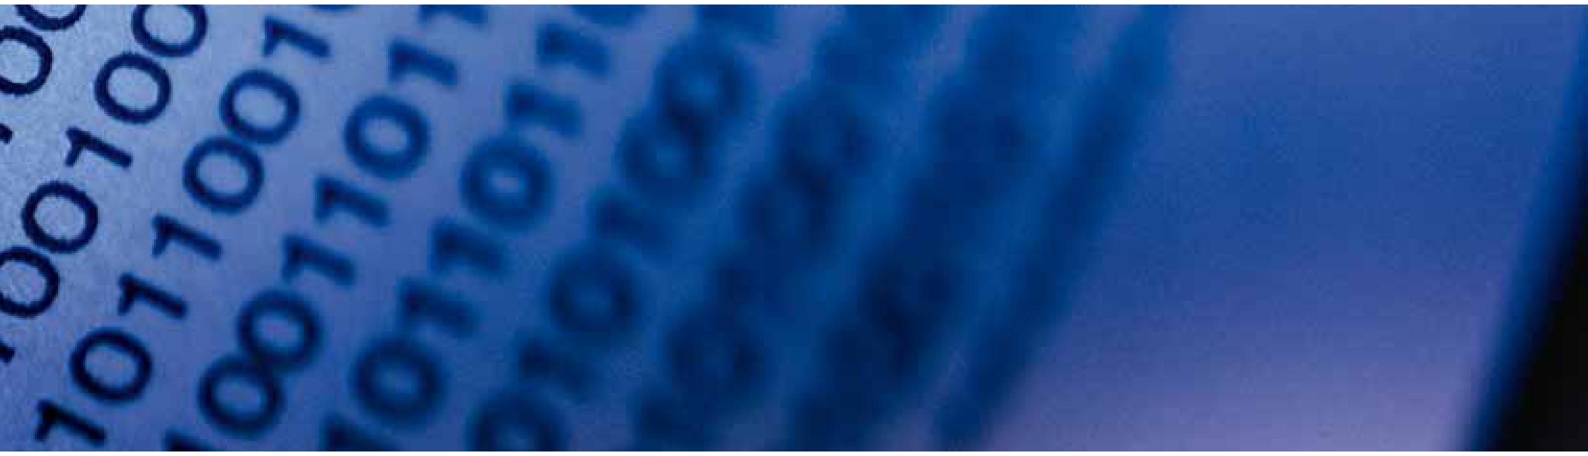

Migrating a code from a serial version on a single processor to a parallel version running on thousands or millions of processors is beyond the ability of most research scientists, and they therefore require assistance from computer scientists. Adding new physics to a parallel code is similarly not a task to be undertaken lightly.

There has been a gradual change in supercomputing over a long period, as greater levels of parallelism have been introduced and exploited. Supercomputers have naturally reflected the trend for increased parallelism as shown by the number of processors in the Top 500 systems. In 2006, systems with less than 1024 processing cores dominated the list with over 300 entries. Only three years later there were no systems with so few cores in the Top 500 at all.

One approach to getting more concurrency is to increase the size of the problem being computed. Such an approach is called weak scaling. The problem size per processor is held constant, and the number of processors is increased. Weak scaling applies for problems involving limited and local communication among processors. In strong scaling, an increased number of processors results in shorter time to solution for a problem with a given size. Putting more processors to work on a code with only weak scaling will result in more communication among the processors and thus less efficiency. In a problem with a given physical size, weak scaling amounts to increasing the resolution of the numerical model. Such an approach would in principal give the advantage of greater fidelity and greater accuracy for the result. However there are situations in which weak scalability is difficult to achieve. In weather forecasting, for example, increasing resolution does not lead to weak scaling because increasing resolution leads to stiffer matrices and less efficient solvers. Novel algorithms for vital parts of numerical weather forecasting codes are being developed through close collaboration between experts in fluid dynamics and software engineering. Ensemble prediction systems, where tens of similar forecasts are initiated and forced with slightly different inputs, use high numbers of cores with today's codes to enable probabilistic forecasts.

Because memory bandwidth and capacity are not keeping pace with the exponential growth in the speed of computers, algorithms of the future will have to minimise data movement, rather than number of operations. Smart and compact representations for numerical solutions of partial differential equations must be developed.

5.3 Data storage and management

Knowledge is preserved and transmitted through digital content. Modern science demands increasingly advanced levels of data curation, that is, maintaining and adding value to digital information for contemporary and future use. Scientific data collections are not merely stored or archived, but are subject to frequent revision and enhancement[2,3,10,27,28].

The needs of data generation, data security, data archiving, and data access differ greatly across the fields we surveyed in *The Scientific Case for eInfrastructure in Norway*[1]. Some of the more traditional scientific projects use high-powered computers to generate a very few numbers, while others, more increasingly now, generate high-volume time-resolved datasets in two and three dimensions for visualisation and insight. Weather and climate calculations produce detailed outputs that must themselves be interrogated for validation and prediction.

The greatest growth in data storage needs, however, is likely to come from outside the realm of traditional high-performance computing. Observational scientists – geologists, astronomers, and biologists, among others – who formerly used film cameras to record data, have now moved largely

▶

to digital media. Much of these data are presently archived in private collections, in stacks of DVDs in offices or on hard drives under desks, where they are inaccessible to other researchers and vulnerable to theft or destruction. Science and society are best served by warehousing these data in a secure public eInfrastructure.The researchers who generated them would have exclusive rights to the data for some period of time, after which access should be broadened under terms defined by the agencies that funded the research and with appropriate regard for data sensitivities. For the data to be useful to others than those who generated them, they need to be tagged with information (metadata) appropriate to the subject in a keyword-sortable catalogue or database.

Migration of data from old formats to continually changing new formats demands constant maintenance and upkeep, making the job of data curation a very important one. Metadata including information regarding data provenance, access rights, anonymised summaries, and so on, must be maintained and migrated as well.

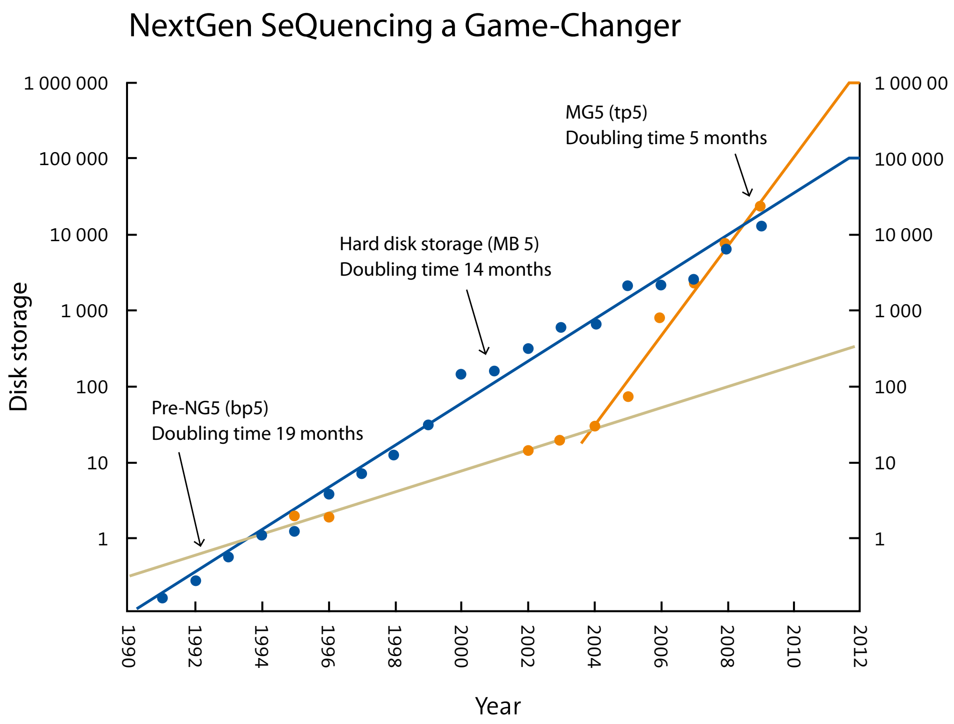

Growth in sequencing data per dollar compared to storage per dollar. The current growth in data has a doubling time of only 5 months.

> “A cost-efficient infrastructure for scientific data will be critical to the progress of research and industry.”

Genomicists have developed machines that rapidly spit out the base-pair sequences containing the genetic code for an organism. These data are used for studies in evolutionary biology, in epidemiology, in the treatment of disease, and in forensics, and it is vital that their storage be secure, robust, and readily searchable by those in need of the information. The current growth in their data needs is prodigious (see figure at left), with a doubling time of only 5 months, far outstripping the doubling time of hard disk storage. These data need to be stored close to the computational resources, to spare network traffic during processing.

The humanities and social sciences also have rapidly growing data storage needs. Projects are underway to digitise the contents of major libraries and museums. Manuscripts are scanned, recorded as images, and transcribed into symbols of languages ancient and modern. Speakers of endangered languages are recorded and the data stored for analysis and preservation. Anthropologists and ethnologists make digital video recordings of cultural practices and need to store these, suitably tagged with geographical and temporal information and privacy limitations.

Policies on access, privacy, anonymisation, data integrity, and data lifetime must be established and rigorously adhered to. These policies differ across and within discipline areas. Some data lose usefulness with time because they are more easily regenerated by future computers than stored; while other data become more valuable with time because they contain information from conditions that cannot be repeated or provide baselines against which to measure changes.

The development of new technology, in particular, intelligent sensors, global positioning systems and the increasing storage capacity of digital databases as well as the availability of high-speed computers has led to an unprecedented ability to collect information and process it. However, there is a considerable gap between the amount of data and our ability to extract the necessary information. An inherent challenge is achieving significant data compression. Dealing with large amounts of data requires good search algorithms and methods for pattern recognition, data analysis, reconstruction of information from sparse data, and data assimilation.

Spatio-temporal datasets with sparse coverage are used in trying to understand the history of the global and regional temperatures of the Earth’s surface. Currently there are more than 7000 stations where temperature, precipitation, and solar radiation are recorded. Despite this large number, computing apparently simple quantities like the uncertainty of the annual mean temperature is a challenge. Techniques for dealing with the sparseness and non-stationarity of the data are under development.

A cost-efficient infrastructure for scientific data will be critical to the progress of research and industry. The data repositories must be actively curated to ensure that data fits its purpose and is available for reuse. They must have a unified structure and nomenclature, and maintain services that can be used where transparent national and international access to scientific data is needed. Data in various forms from raw data to scientific publications will need to be stored, maintained, published and made openly accessible. Trust in data can be enhanced by the establishment of qualified domain specialists who curate the data and deal with issues of security, confidentiality and privacy, ownership, provenance, authenticity, permanency, integrity, interoperability, as well as the quality of the primary data and associated metadata. ▶

5.4 Power requirements

Energy costs are a substantial part of the budget of a high-performance computing centre. Processing, bandwidth, and memory access all cost power, and keeping the components from overheating requires a massive cooling system. As system size increases, the proportion of the budget required for power also increases. Present-day Petaflop systems use on the order of MegaWatts. To build a computer of Exaflop class would require, with linear scaling, a devoted electrical power plant to service it alone. Energy efficiency must be improved to reach higher performance. Dealing with the waste heat is also a problem. This leads to consideration of what is becoming known as green computing[3].

In parallel with the Top 500 list that ranks supercomputers by their processing power, there is a Green 500 list[29] that ranks them according by their power efficiency. The November 2011 list is topped by four prototype IBM BlueGene/Q systems, that run at better than 1600 Megaflops/Watt and rank within the top 65 on the Top 500 list. Number 10 on the Green 500, the Japanese Tsubame, ranks number 5 on the Top 500 list, indicating that energy efficiency is taken very seriously by makers of the world's fastest supercomputers. The Japanese Riken K computer, first on the November 2011 Top 500 list, makes a respectable showing at number 32 on the Green 500 list. Power efficiency of the greenest computers has more than quintupled since the first Green 500 list came out in 2007.

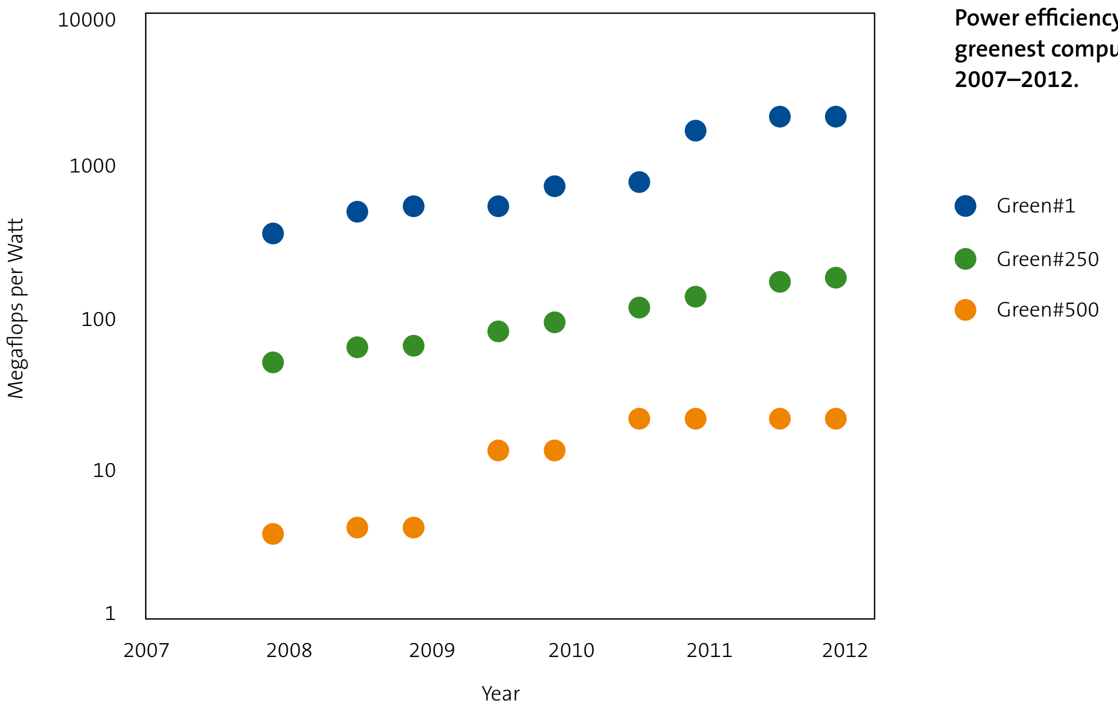

Power efficiency of the Top 500 greenest computers in the years 2007–2012.


75
5434 29
£
3
$

> “Efficient use of hardware and software can make a significant difference in the energy consumption required for the completion of a given computational task.”

▶ The figure on previous page is constructed from the Green500 lists and illustrates a slow exponential trend in the improvement of power efficiency. At current efficiency, an Exascale machine will consume 600 MegaWatts of electrical power. Extrapolating the trend exhibited by the Green number 1 to 2018 brings Exascale down to 16 MegaWatts, which is far more reasonable. While efficiencies are expected to improve, there may be an intrinsic limit to how much computing power can be done for a Watt of electrical power, so efficiency may not continue exponential growth indefinitely.

Efficient use of hardware and software can make a significant difference in the energy consumption required for the completion of a given computational task. Hardware options for greater efficiency include reducing processor clock speeds, new materials for processor design, and multi-core processors. In software, more attention can be paid to improvements in algorithms and parallelisation techniques, and in power management at the operating system level.

The overall energy consumption in a computer centre is determined by the siting of the centre and its available infrastructure, the nature of the computer and other hardware, and how the system is managed. A metric used to measure this is the Power Usage Effectiveness (PUE), defined as the ratio of the total power input to the data centre to the power delivered to the computers. The PUE of a typical data centre is about 2 although there are data centres with PUEs exceeding 4. New data centres are being constructed with the aim of consuming only 10–20% additional energy for the infrastructure corresponding to a PUE in the range 1.1 to 1.2. A small number of centres are attempting to reduce their PUE further by re-using their waste heat. This is not possible for all data centres; it depends on the temperature of the waste heat, its location in relation to potential users, and the viability of transporting the heat. Also, heat is produced all year but may not be wanted in the summer. Siting data centres in cold climates, as in northern Norway or Svalbard, would provide distinct opportunities for the utilisation of waste heat. The prototype supercomputer being installed through Nordforsk funding at the Thor Data Centre in Iceland will very likely have the lowest PUE of any centre now operating anywhere in the world.

New breakthroughs in research on improving CPU speed may also lead to faster and less power-consuming hardware. An example is the use of optical photodetectors to transfer data between CPUs, see figure on previous page. Research done by IBM suggests this technology can be very important for achieving Exaflop performance on a cluster[30]. The data transfer among large numbers of CPUs can easily be the bottleneck for many high-performance systems. An optical device can increase the efficiency by orders of magnitude.

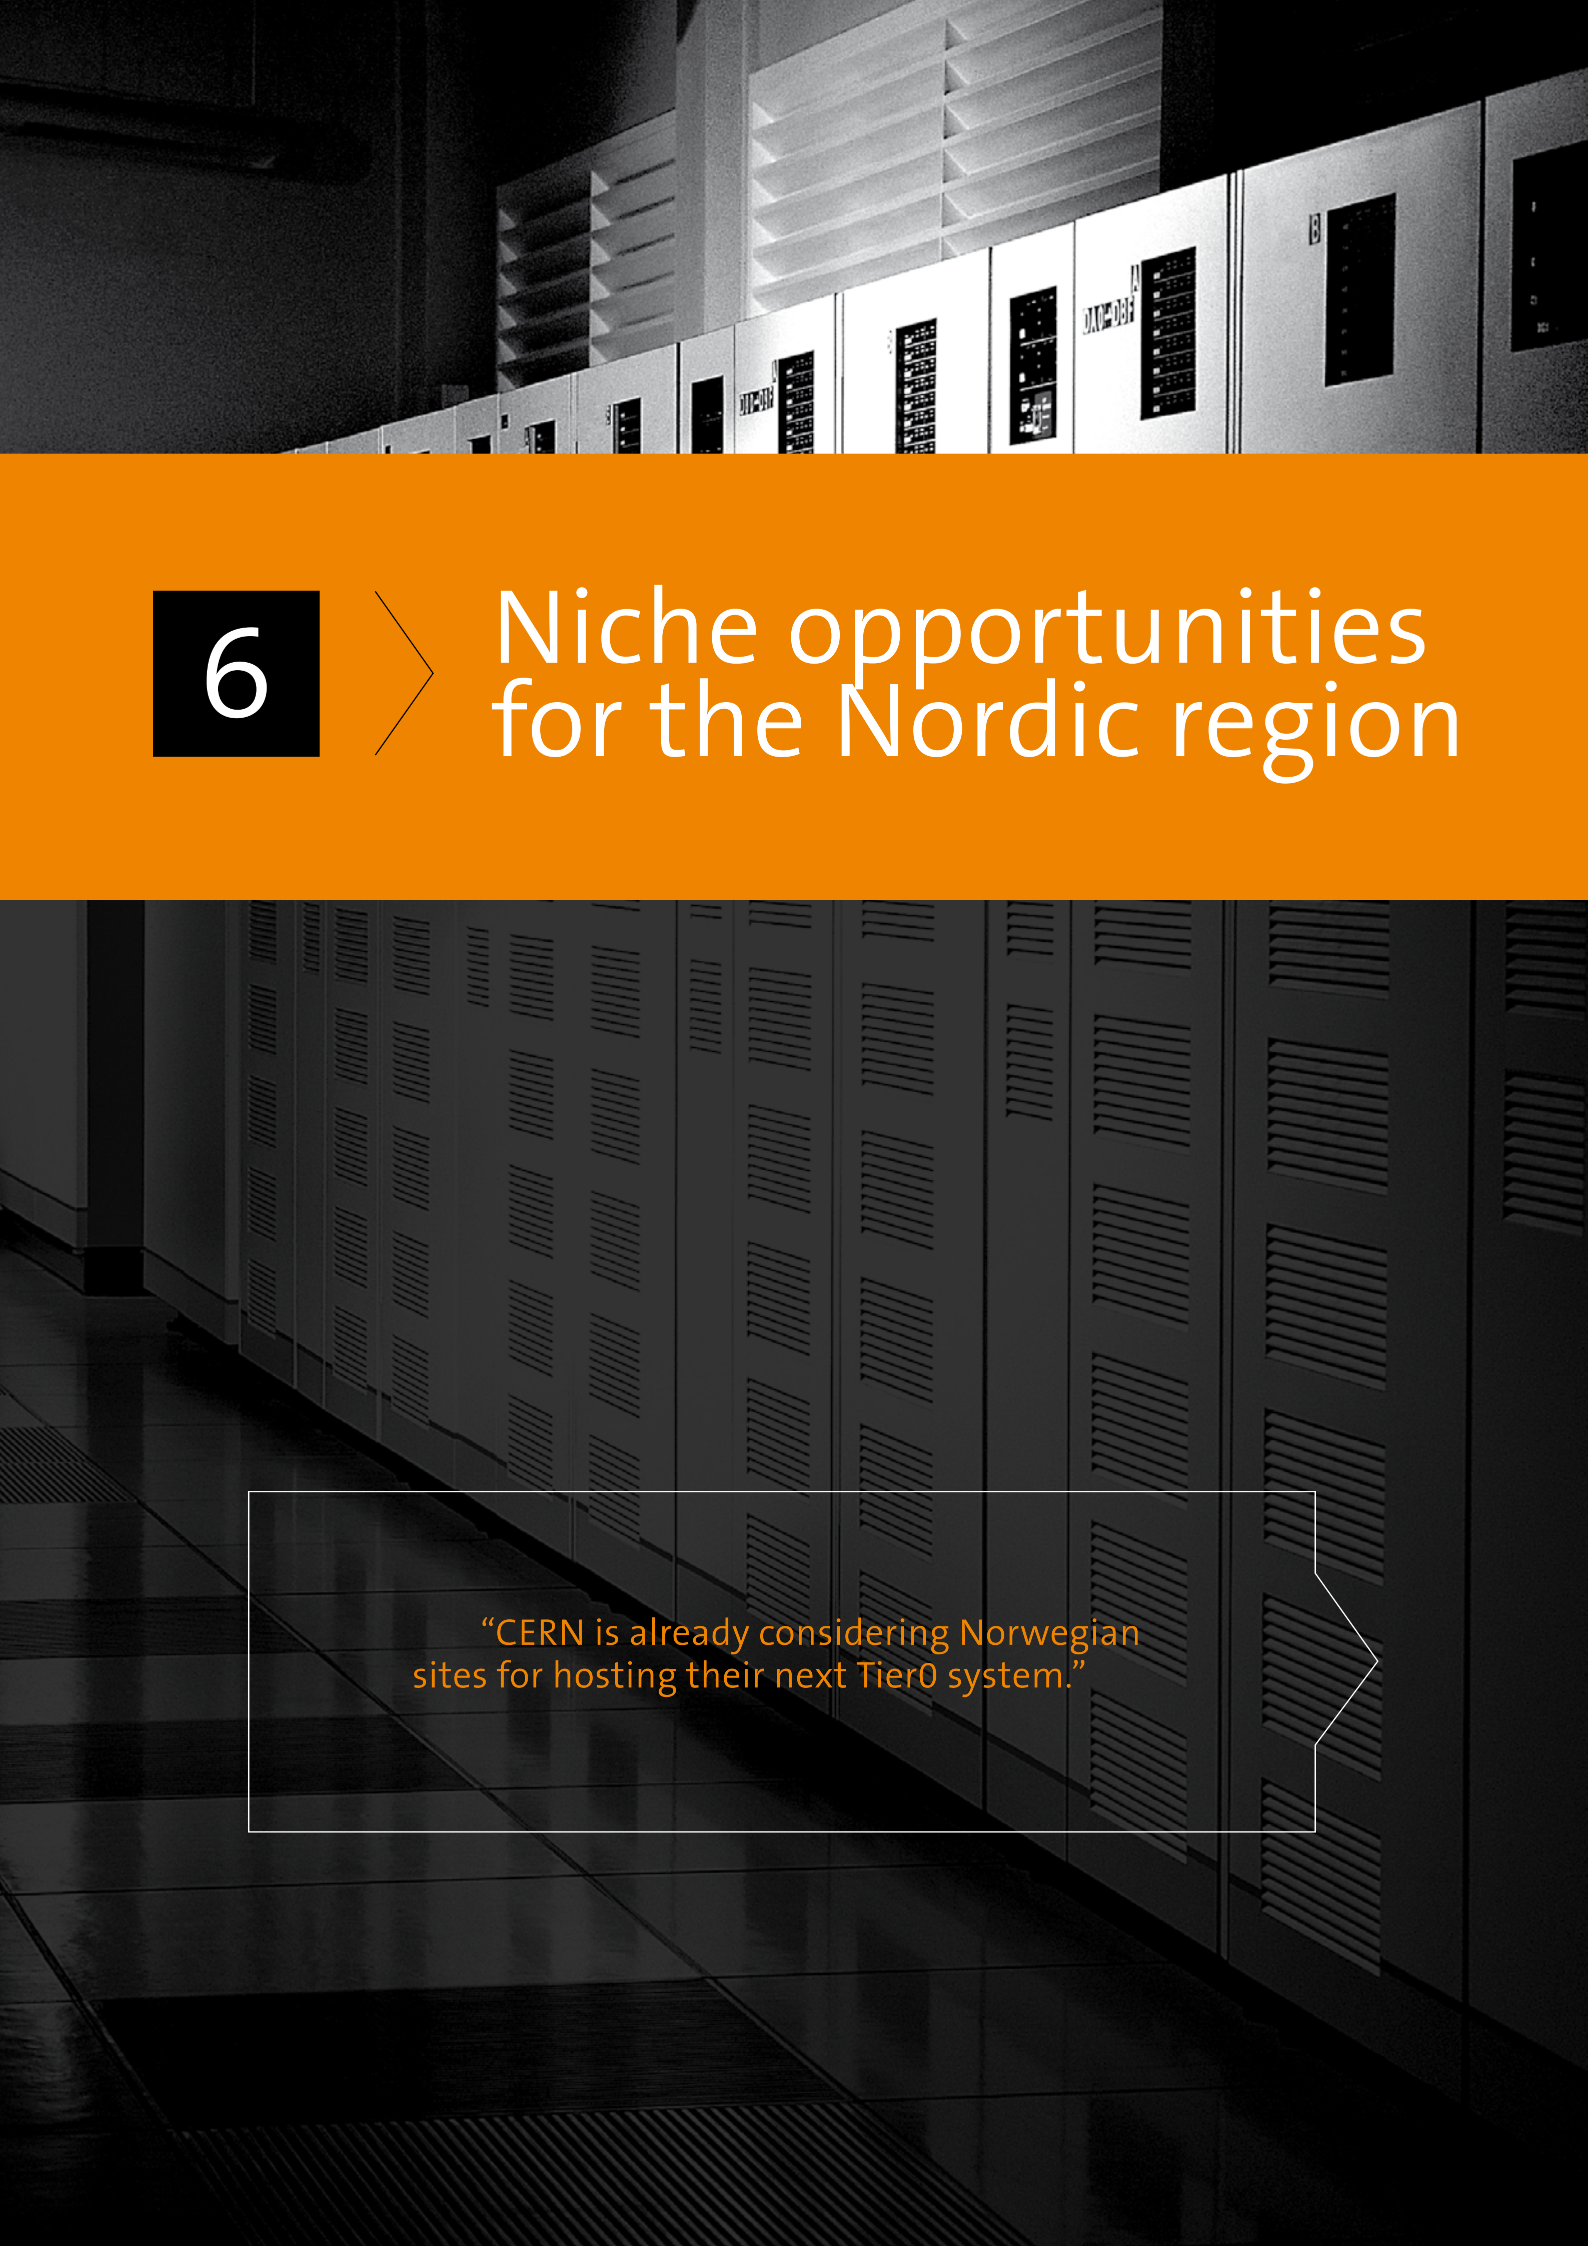

6

Niche opportunities for the Nordic region

> “CERN is already considering Norwegian sites for hosting their next Tier0 system.”

6 Niche opportunities for the Nordic region

All the Nordic countries could exploit advantages to position themselves as attractive partners in international eInfrastructure-projects. Norway has a vast network of mines, tunnels and man-made caverns in geologically stable bedrock which could host large datacenters powered from renewable energy sources.

The Nordic countries, situated at the borders of the Arctic, provide excellent opportunities for the siting of large data or supercomputing centres. Cooling is not as great a problem as it is in lower latitudes, and there are good opportunities for utilising waste heat for public buildings and maintenance of ice-free roads and walkways. Inexpensive and carbon-free electrical power is also available in several Nordic countries through hydropower, wind power, wave power, and geothermal power.

Iceland, with its ample supplies of both hydropower and active geothermal power, would be an excellent site for a new pan-European computer centre, which would be a boon to the Icelandic economy and a means toward their tighter integration with the European mainland. A pilot project to install a modest supercomputer in Iceland at the Thor Data Centre in Hafnarfjörður[31], funded by four of the Nordic countries, is presently underway, and scheduled to be running by 2012. Norway's close ties to Iceland should be used to encourage and facilitate Icelandic ambitions in this regard. Local demand in Iceland for supercomputing resources is relatively low, but siting a very large supercomputer there in the future should not be ruled out. While it is difficult to export Iceland's abundant supply of cheap and clean energy, doing data processing there and exporting the results makes a great deal of economic and ecological sense.

In Finland, the recent closure of a paper mill in Kajaani offers a unique opportunity to install a data centre close to sources of clean hydropower and biofuels. With electrical power exceeding 30 MegaWatts capacity and free cooling from a nearby river, this data centre could supply the infrastructure for the biggest supercomputer on the planet with a near-zero carbon footprint. The data centre is expected to be in operation in 2012.

Norway has not yet made use of its potential for hosting a world-class international data centre. This country has a vast network of mines, tunnels and man-made caverns in geologically stable bedrock that will be secure against natural and man-made disasters, and where cooling and waste heat are unproblematic. A number of localities in Norway have come forward with ideas for using unused mines or tunnels for just such a purpose[24]. Since Norway's electricity production is almost entirely from renewable sources, a data centre sited almost anywhere in Norway will have a lower carbon footprint than in most other places in the world.

A very interesting idea is to offer a site in Norway to our partners in the PRACE collaboration with the aim of hosting the largest supercomputers in Europe. It is noted that CERN is already considering Norwegian sites (among others) for hosting their next Tier0 system. We believe that Norway should promote itself as *the* place in Europe where high-performance computing can be done with inexpensive, carbon-free power, where waste heat can either be redistributed favourably or dissipated harmlessly, and where a well-trained data force exists to maintain, improve, and make good use of the system. With such a data centre, and the work force it would attract, Norway could truly be "at the top of the world" in supercomputing. Spin-off high-technology industries would likely also be spawned.

Unique problems associated with the Arctic climate lend themselves to collaboration among the Nordic countries. As an example, the meteorological offices of Norway and Sweden have recently agreed to begin common operational numerical weather forecasting in 2014. These organisations plan to buy new computer systems every other year, utilising the newest hardware for operational runs while the older system serves as backup. There is room for greater exploitation of the common characteristics of our region.

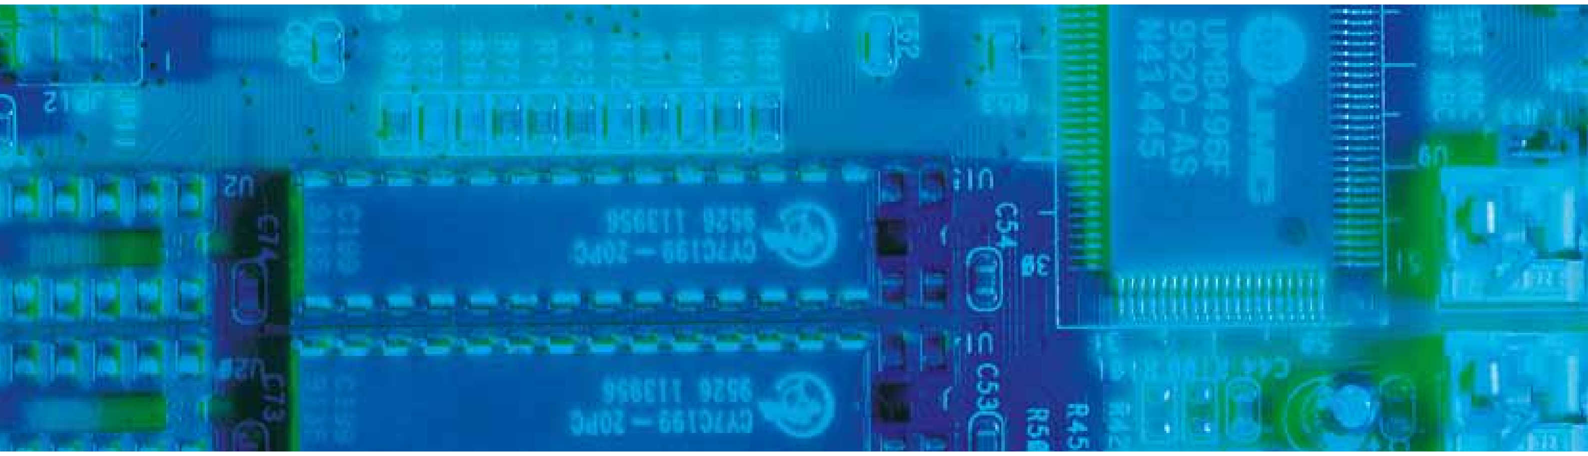

The Arctic Region Supercomputing Center (ARSC) in Fairbanks, Alaska has long supported Arctic-area research in areas such as geomagnetic storms and aurora, Arctic ecology, climate science, volcanology, seismics, and tsunamis. The United States Department of Defense has recently reduced its support for ARSC, providing an opportunity for Norway and the Nordic region to step in with international supercomputing centres devoted to typically Arctic concerns. Nordic supercomputing centres could in the future form a trans-Arctic collaboration with cryogenic, perhaps superconducting, cables linking them under the Arctic Ocean (see figure below). Eventually the governments of Canada, Greenland, and Russia may wish to join as well.

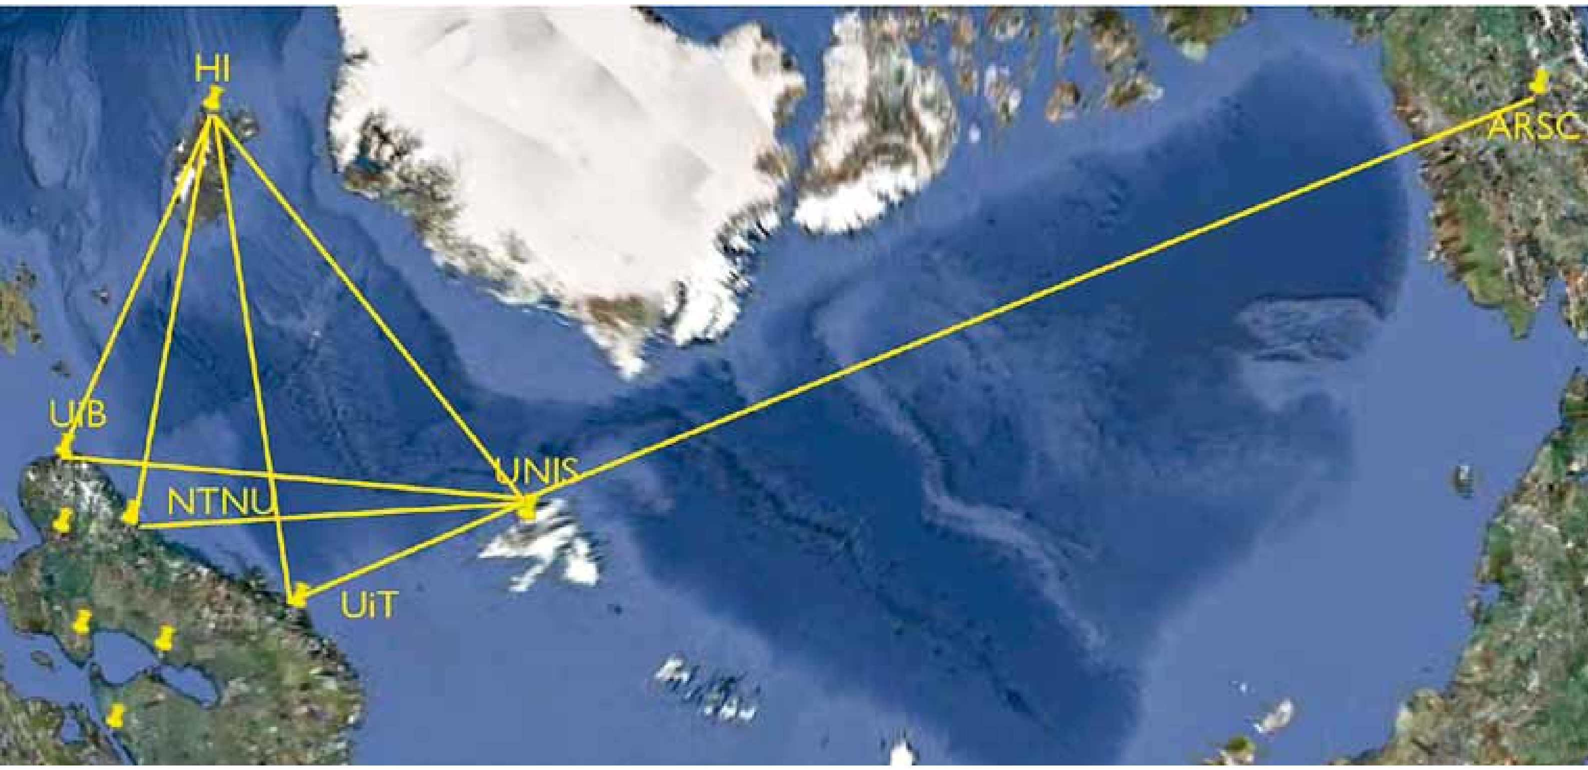

Map of a hypothetical Arctic computing network.

7

Conclusions

> “Investment in eInfrastructure must be increased to at least 100 MNOK annually to keep Norway competitive.”

Illustrasjon: Jon Solberg, Forskningsrådet

7 Conclusions

This Roadmap describes steps towards an upgraded national eInfrastructure to generate a wide range of benefits to society, including increased competitiveness for science and industry. The public investements in eInfrastructure must be increased to at least 100 MNOK annually to keep Norway from falling behind the rapidly accelerating growth in eInfrastructures world-wide.

In this Norwegian eInfrastructure Roadmap, we have attempted to highlight some of the benefits to society that come from investment in eInfrastructure. From our consideration of the worldwide growth in the technologies of computation and information, and comparing the situation in Norway with what we know of the rest of Europe and the world, we have arrived at a set of recommendation for how eInfrastructure should be funded, organised, and managed in the future. First priority among these recommendations is that present annual investment in eInfrastructure must be quintupled, to 100 MNOK or greater, in order to keep Norway from falling behind the rapidly accelerating growth in European and world eInfrastructure.

We believe that our recommendations, if followed, will usher in a bright new era in Norway's use of the technologies of information and communication, and will keep Norwegian science and industry competitive at the highest level with the rest of the world.

We hope this document and the recommendations herein will be useful for the eVITA Programme Committee, for the Research Council of Norway, and ultimately for the Ministry of Education and Research.

References and notes

1. Available from the Research Council of Norway, ISBN 978-82-12-02831-9.

2. *Riding the Wave: How Europe can gain from the rising tide of scientific data*, final report of the High Level Expert Group on Scientific Data, submitted to the European Commission, October 2010.

3. *National Science Foundation, Advisory Committee for Cyberinfrastructure, Task Force Reports on Campus Bridging, Cyberlearning and Workforce Development, Data and Visualization, Grand Challenges, High Performance Computing, and Software for Science and Engineering*, March 2011. http://www.nsf.gov/od/oci/taskforces/.

4. http://www-03.ibm.com/innovation/us/watson/what-is-watson/the-future-of-watson.html

5. *From Challenges to Opportunities: Towards a Common Strategic Framework for EU Research and Innovation funding*, European Commission, COM(2011) 48, Green Paper, Brussels, 9.2.2011

6. *International Assessment of Research and Development in Simulation-Based Engineering and Science*; World Technology Evaluation Center, Inc. – WTEC Panel Report; April 2009

7. *NSF-Supported Research Infrastructure: Enabling Discovery, Innovation and Learning*. National Science Foundation, 2009.

8. *Shaping Europe's Future Through ICT*. Report from the Information Society Technologies Advisory Group (ISTAG). March 2006, www.cordis.lu/ist/istag.htm.

9. *THE HIPEAC Vision*. Marc Duranton, Sami Yehia, Bjorn De Sutter, Koen De Bosschere, Albert Cohen, Babak Falsafi, Georgi Gaydadjiev, Manolis Katevenis, Jonas Maebe, Harm Munk, Nacho Navarro, Alex Ramirez, Olivier Temam, Mateo Valero. Network of Excellence on High Performance and Embedded Architecture and Compilation. 2010. www.hipeac.net.

10. *Empowering the Nation through Discovery and Innovation*, NSF, April 2011, available at http://www.nsf.gov/news/strategicplan/nsfstrategicplan_2011_2016.pdf.

11. Exascale refers to computing power in excess of 10^{18} floating point operations per second.

12. http://cordis.europa.eu/fp7/ict/e-infrastructure/overview_en.html

13. *A Digital Agenda for Europe*, Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee, and the Committee of the Regions, Brussels, 26.8.2010, Com(2010) 245 final/2.

14. http://www.top500.org/lists/2011/11

15. *Preparing for Exascale*, Oak Ridge National Laboratory, ORNL/TM-2009/308.

16. http://www.csc.fi/english

17. http://www.helmholtz.de/en/helmholtz_centres/

18. http://www.energy.gov/sciencetech/index.htm

19. http://www.nsf.gov/dir/index.jsp?org=OCI

20. http://www.notur.no/

21. http://www.snic.vr.se/

22. http://public.web.cern.ch/public/

23. http://www.prace-ri.eu/

24. “Green Datacenters” in http://www.notur.no/publications/magazine/pdf/meta_2011_2.pdf

25. *e-IRG White Paper 2011*, available from http://www.e-irg.eu/publications/white-papers.html.

26. *The Opportunities and Challenges of Exascale Computing*, Summary Report of the Advanced Scientific Computing Advisory Committee Subcommittee. U.S. Department of Energy. Office of Science; Fall 2010.

27. *Sustainable Economics for a Digital Planet. Ensuring Long-Term Access to Digital Information*, Final Report of Blue Ribbon Task Force on Sustainable Digital Preservation and Access. National Science Foundation; February 2010.

28. *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Edited by Tony Hey, Stewart Tansley, and Kristin Tolle. Microsoft Research, October 2009.

29. http://www.green500.org

30. http://domino.research.ibm.com/comm/research_projects.nsf/pages/photonics.index.html

31. http://www.norden.org/en/news-and-events/news/joint-nordic-supercomputer-in-iceland

Figure attributions

Page 15. Top 500 performance projection taken from http://www.top500.org/lists/2010/11/performance_development, with points for Norway added in by hand.
Page 16. Performance of the top Nordic supercomputers on the Top 500 list, constructed by Jørn Amundsen for publication in META.
Page 17. Top 500 statistics constructed from data taken from http://www.top500.org/.
Page 21. High performance supercomputing pyramid constructed by Galen Gisler.
Page 22. Jet engine schematic cutaway from Wikipedia.
Page 23. Weather and topography maps courtesy of met.no
Page 26. Financing of supercomputing by the Research Council of Norway, courtesy Research Council.
Page 31. Growth in sequencing per dollar compared to storage per dollar, courtesy Inge Jonasson.
Page 33. Power efficiency of the greenest computers, constructed from data taken from http://www.green500.org
Page 34. Cutaway of new optical routing 3-dimensional chip design, from http://domino.research.ibm.com/comm/research_projects.nsf/pages/photonics.index.html
Page 36. Map of a hypothetical Arcitc computing network, constructed from Google Earth.

Acknowledgments

The eInfrastructure Scientific Opportunities Panel wishes to acknowledge the contributions from many individuals, including Jørn Amundsen, NTNU; Petter Bjørstad, UiB; Xing Cai, Simula Research Laboratory; Mats Carlsson, UiO; Finn Drabløs, NTNU; Hans Eide, USIT, UiO; Gaute Einevoll, UMB; Bernhard Fabianek, European Commission; Bjørn Hafskjold, NTNU; Erik Hagersten, Uppsala University, Sweden; Svenn A. Hanssen, UiT; Morten Hjorth-Jensen, UiO; Eivind Hovig, UiO; Kjetill Sigurd Jakobsen, UiO; Inge Jonassen, UiB; Martin King, UiB; Petter Kongshaug, UNINETT; Kimmo Koski, CSC, Finland; Marcin Krotkiewski, UiO; Arne Morten Kvarving, SINTEF ICT; Hans Petter Langtangen, Simula Research Laboratory; Sigbjørn Lien, UMB; Torleif M. Lunde, UiB; Arne Løkketangen, Molde University College, Lasse Natvig, NTNU; Michel d.S. Mesquita, Uni Research; Kjell Petersen, Uni Research; Kenneth Ruud, UiT; John Towns, NCSA, USA; Espen Uleberg, UiO; Nils Peder Willassen, UiT;

About the eInfrastructure Scientific Opportunities Panel

The eInfrastructure Scientific Opportunities Panel was charged by the Research Council of Norway to develop the scientific case for the eInfrastructure[1] that can best serve Norwegian research from 2015 and produce a first version of a Norwegian eInfrastructure Roadmap (this document).

Galen Gisler, chairman of the eInfrastructure Scientific Opportunities Panel, has been a senior researcher at Physics of Geological Processes, University of Oslo, since 2006. Trained in astrophysics at the University of Cambridge, he spent twenty-five years at the Los Alamos National Laboratory applying techniques of computational physics to a variety of problems of scientific and societal interest. He is presently studying the physics of explosive volcanic eruptions.

Dag Bjørge, member of the eInfrastructure Scientific Opportunities Panel for the last two years, has been a senior scientist at the Research Division of Norwegian Meteorological Institute (met.no) since 1993. He has been engaged in model development for climate and weather prediction applications, and is currently responsible for parts of the operational numerical weather prediction systems at met.no.

Colin Jones is head of The Rossby Centre at the Swedish Meteorological and Hydrological Institute. He was previously professor and director of the Canadian Regional Climate Modeling Network at the University of Quebec. He is co-chair of the World Climate Research Project (WCRP) Task-Force on Regional Climate Downscaling, a member of the WCRP Working Group on Coupled Modelling and co-leads the WCRP Project CORDEX. He is a member of the EC-Earth Steering Group and coordinates the FP7 project EMBRACE (Earth system Model Bias Reduction and Abrupt climate change). His research focuses on the development of global and regional climate models and their application to climate prediction and global climate change.

Nils Reidar Bøe Olsen has been professor in Hydraulic Engineering at the Department of Hydraulic and Environmental Engineering at NTNU since 2002. He has his PhD from the same department in 1991, and worked at SINTEF, The Norwegian Hydrotechnical Laboratory for seven years before returning to NTNU. His research is in the field of computational fluid mechanics applied to river and hydropower engineering, in particular 3D modelling of sediment transport.

Trygve Helgaker is a professor of theoretical chemistry at the Centre of Theoretical and Computational Chemistry, Department of Chemistry, University of Oslo. He works in the area of quantum chemistry, mainly on the development of new methods for quantum-mechanical simulations of molecular electronic structure. He has spent sabbatical years at the Universities of Cambridge and Durham, UK.

Stephan Oepen is professor in Language Technology at the Department of Informatics at the University of Oslo. His research addresses automated computational analysis of the grammatical and semantic structures of very large collections of natural language texts, for example Wikipedia, thus enabling computers to ‘understand’ (within certain limits) human language. Prior to his appointment at Oslo, Oepen worked at DFKI in Germany and Stanford University in the US.

Anna Lipniacka has been professor at the Department of Physics and Technology, University of Bergen, since 2003. She is a particle physicist involved in experiments at CERN since 1983, when she started analysing millions of neutrino interactions in iron. With career spells at the University of Warsaw, University of California and Stockholm University, she participated in several high energy, high statistics experiments at CERN, Fermilab and DESY. She is presently a member of the ATLAS Collaboration at the Large Hadron Collider and analyses billions of events spawned by LHC collisions using GRID powers around the world.

Elena Celledoni is professor of mathematics at the Norwegian University of Science and Technology, Trondheim (NTNU). She got her PhD in Computational Mathematics at the University of Padua, Italy and she held post doc positions at the University of Cambridge, UK and at the Mathematical Sciences Research Institute, Berkeley, California. Her research is on numerical methods for differential equations and numerical linear algebra.

Arvid Lundervold has a BSc in pure mathematics and philosophy and is MD from the University of Oslo. He earned his PhD (Dr.med) on quantitative MR imaging at the University of Bergen. Since 2005 he has been full professor in medical information technology at the University of Bergen, Department of Biomedicine, and head of the Neuroinformatics and Image Analysis Laboratory. Lundervold has published more than 100 papers and conference reports related to medical image analysis, pattern recognition, and neuroinformatics.

MILJØMERKET
07 OSLO - 241 344

This publication can be ordered at:
www.forskningsradet.no/publikasjoner

The Research Council of Norway
P.O.Box 2700 St. Hanshaugen
NO–0131 OSLO

Telephone: +47 22 03 70 00
Telefax: +47 22 03 70 01
post@forskningsradet.no
www.forskningsradet.no/evita

February 2012
Design: Fete typer
Cover photo: Shutterstock
Print: 07 Gruppen AS
Number of copies: 450

ISBN 978-82-12-03055-8 (printed version)
ISBN 978-82-12-03056-5 (pdf)