A National High-Performance Scientific Computing Facility

September 2008

This report has been written by:

Henrik Madsen
Mathematical Statistics, DTU Informatics
Technical University of Denmark
hm@imm.dtu.dk

Lise M. Frohn
National Environmental Research Institute
Aarhus University
lmf@dmu.dk

Bernd Dammann
Scientific Computing, DTU Informatics
Technical University of Denmark
bd@imm.dtu.dk

Jostein K. Sundet
Research Computing Service
University of Oslo
j.k.sundet@usit.uio.no

Executive Summary

This report describes the results of a study initiated by the Danish Ministry of Science, Technology and Innovation for the purpose of evaluating the situation for High Performance Computing (HPC) and, in particular, to investigate the need for a large TOP100[1] facility in Denmark.

As an important part of the study we have held meetings with a number of High Performance Computing (HPC) facilities in Canada, US, and Europe. In Denmark we have visited and interviewed a large number of HPC users. Based on these interviews, and on findings during our international visits we have drawn conclusions about an optimal setup for a future HPC infrastructure in Denmark. DCSC accepted our invitations to participate in some of our international meetings, and our findings and conclusions have been discussed during several meetings with DCSC. We appreciate the time spent by DCSC on all issues related to this study.

It is concluded that Denmark needs a combined decentralised and centralised infrastructure for HPC. Supplementing the existing structure with a centralised infrastructure will greatly support the research and education in scientific computing, the internationalisation of Danish HPC as well as the application of HPC in Denmark to new areas and users.

To some user groups the demand for computing requirements is found to be most adequately met by special purpose decentralised facilities located on site, or nearby. Today this applies to some of the 11 DCSC installations. The facility is then administered by and reserved for the local research group. Today the use of HPC in Denmark is narrowed down to rather few areas - mostly physics, chemistry, and (some areas of) biology. However, many research groups would prefer to have access to a large centralised HPC facility with a professional and highly skilled user support. Such a facility should offer help with issues like code tuning, porting, code profiling, selection of algorithms, selection of libraries, parallelisation, and efficiency. It is also concluded that the facility should be linked to a university that offers state-of-the-art research, and education in scientific and high performance computing. Such an infrastructure will enable the synergy potential between scientific applications, and research and education related to scientific and high performance computing.

The central HPC facility must be large enough to be able to solve critical mass computations within a large number of scientific areas. Compared to Denmark, a broader spectrum of HPC users are seen in most other countries, where a wide range of scientific applications within areas like data-mining, statistics, pharma-

[1]The TOP100 list is the 100 most powerful computer systems ranked by the achieved maximal LINPACK performance

ceutics, linguistics, multimedia, etc. are supported and hosted at large central facilities. With an appropriate infrastructure including support on various levels, a new central facility will imply that the demand for HPC, data storage, data analysis, etc. within a much broader spectrum of scientific areas can be facilitated. This includes the governmental research activities, which have recently become the responsibility of the universities, in particular at DTU and the University of Aarhus. A central HPC facility and infrastructure must also be able to meet the demands for mass computations and data storage, which play a vital role in the modeling activities associated with forecasts, scenarios etc, which form a central part of the governmental research activities. Recently the research network (Forskningsnettet) has been upgraded to dedicated 10 GE connections. This means that a researcher will now have the feeling of using a local server in cases where he is in fact using a remote large scale facility.

Today Denmark is not found on the TOP500 list comprising the largest computer facilities in the world. This is contradictory to almost all other countries in Europe (including countries like Slovenia and Cyprus). An entry on the list is a ticket to possibilities for international collaboration and funding. This includes funding from EU. As an example can be mentioned that EU has lately allocated more than 500 Million Euros to HPC to cover the initial phase of the PRACE project. Later on, more funding will be allocated to PRACE activities. One of the primary goals of PRACE is, together with the IT industry, to influence the development of new technologies and components for large scale computer systems.

We see an urgent need for building up a TOP100 facility in Denmark. In order for Denmark to be a candidate for significant EU funding and international collaboration, the facility must be on the TOP100 list of 2009/10. The new facility must be an important part of a new HPC infrastructure with a focus on all aspects of scientific and high performance computing and, hence, enable a healthy and inspiring environment for scientific computing and HPC. The HPC infrastructure in Denmark, and the eScience strategy must be reconsidered, with an increased focus on user support, code migration, efficiency, new technologies, etc. as well as on servicing a broader spectrum of HPC users.

Finally it is concluded that the total TCO of the HPC infrastructure must be visible. During the past seven years only hardware cost has been visible. It is clear that a central facility calls for a significant funding. However, a part of this funding might come from EU and the industry. A new facility linked to education and research within scientific and high performance computing, will also re-invoke and facilitate a technology watch, which is considered a very important aspect. For the time being we see a shift from 'serial-code focus' to 'multi-core CPU focus'. This shift in technology implies a need for new algorithms, HPC software and tools in order to take advantage of those possibilities the new hardware offers.

Contents

1 Introduction

Since the year 2001 the Danish Center for Scientific Computing (DCSC) has been responsible for the allocation of funding for HPC in Denmark, and today 11 facilities are financed by DCSC. Until 2000 UNI-C coordinated and hosted the national supercomputing for the research community. During the 1990's the HPC activities at UNI-C constituted a very small part of the IT services provided, and the courses offered were mainly non-HPC oriented software; examples are courses in Word Perfect and SAS. At the same time UNI-C became responsible for the software and IT network for primary schools.

Looking at its early years, the attitude towards UNI-C was quite positive. This lasted until, say, the decade up to 2001. But it is also clear that because of the shifted focus on HPC, this lead for UNI-C to a situation, which hardly any user of HPC would like to see re-established.

The background for this report are two grant proposals for larger central facilities – one proposal from DTU and one from DMU /Aarhus University. The applications lead to a wish from the Danish Ministry of Science, Technology and Innovation to study the situation for HPC today and, in particular, to investigate the need for a larger (say TOP100) facility in Denmark. In parallel to the above-mentioned applications, research groups with a significant demand for access to HPC, have expressed their frustration related to the lack of focus on efficiency, software, education, and support for the existing HPC structure in Denmark.

It has also been observed that the Danish policy with respect to eScience infrastructure today is in stark contrast to what is the general trend in leading HPC countries as well as in countries with similar research goals as those in Denmark. As a serious consequence Denmark is not really visible in the HPC landscape, and Denmark is very weak with respect to international funding for HPC activities. As an example, this implies that Denmark is one of the only countries in Europe that does not participate in the 500+ Million Euro PRACE project. A primary purpose of the PRACE project is, together with the European industry, to influence the development of new IT technologies.

As already pointed out in the letter from the Danish Ministry of Science, Technology and Innovation, the synergy potential by establishing a large central facility, must be described in detail in our study. This issue will, therefore, be elaborated on in the report.

The main purpose of this report is also to investigate the need for a larger national HPC facility, and the need for building up a new national infrastructure related to scientific and high performance computing. We hope to establish a position from where it will be possible for us to compete and collaborate with researchers in

HPC in countries like Norway, Sweden, Finland, UK, Germany and Holland.

The motivation for this investigation is further described in Section 2. Central in this investigation is a long series of visits to internationals HPC facilities and a series of meetings with HPC users in Denmark. Our approach is presented in Section 3. The synergy potential and the possibilities for collaborative efforts in connection with the establishment of a large central facility are described in detail in Section 4. Section 5 describes the setup for a new infrastructure based on this pre-investigation. A summary of the meetings held with both international HPC facilities and with HPC users all over Denmark is found in the Appendix.

2 Motivation

2.1 The situation today

During the past eight years the focus has been on acquiring as much hardware as possible within the given financial frame, leaving little attention to other HPC issues like, e.g., user-support. However, a few months after the establishment of the consortium behind this investigation, the concept of CAUS (Centre for Advanced User Support) has been established within DCSC, and in June 2008 about DKK 2 Millions was allocated to CAUS (plus DKK 6 Millions were allocated Grid activities) and it is expected that two CAUS centres will be established within a short time frame.

Still, after an implementation of CAUS, the situation in Denmark is that most of the annual DCSC labelled financial support is spent on hardware for the 11 installations in Denmark (DKK 2 Millions for support versus DKK 18 Millions for hardware according to the grants of June, 2008).

Appart from the new initative CAUS there is essentialtially no DCSC funding and focus on software, user support and on optimal operation and code optimisation. These issues are left to the individual researcher. However, it is well known that efficient use of computers calls for an expert knowledge on libraries, software, code tuning, and optimisation. Even the porting from maybe a PC to a large memory system calls for expertise. The consequences of not having this balanced infrastructure are, that experts in physics, chemistry, etc. spend a considerable amount of time on issues which could have been solved rather easily by experts in numerical algorithms, scientific and high performance computing, while experts in physics could have spend time on issues related to their expertise. Today there is essentially no link between the HPC users, and the research and education in scientific and high performance computing.

A further consequence of today's situation is that the use of HPC in Denmark (with very few exceptions) is narrowed down to physics, computational chemistry, and (some areas of) biology. This is problematic as we see a need for HPC in many other areas like finance, data mining, pharmaceutics, visualisation, and even in intelligent storing and accessing of complex and data intensive objects within bio-science, history, culture, linguistics, etc. We have observed that the spectrum of user groups in other countries is much larger than in Denmark, and this indicates that the group of HPC users in Denmark has the potential of significant enlargement.

During 2008 the research network in Denmark is extended with possibilities of 10 GBit/s point-to-point connections. This upgrade of the research net enables a new possibility for centralising the HPC facilities without the risk of a low response time, since only the graphics have to be transmitted to the desktop of the remote user. Such a HPC infrastructure with so-called fat nodes (HPC facilities) and thin clients (the desktop) are becoming increasingly popular. Still, however, for many applications it is important that the data is kept close to the HPC installation, and by enabling a large centralized data storage this critical issue of modern high performance computing can be facilitated. Today simply the time for data transfer may prevent users from moving from one of the 11 installations to another installation for solving a critical task, which, with respect to the hardware and software suite, might have been more adequately solved by another installation.

Today the cooling power of the individual facilities is not reused, e.g., for heating. This is most likely not economically feasible due to the small and diverse installations. One large central facility placed in areas with a demand for heating, e.g., near a large district heating system presents a possibility for using the cooling power for heating. Today the focus on energy prices and the environment implies that new large facilities must be built up with a focus on green computing, which includes focus on efficiency, low power consumption, and reuse of cooling power.

2.2 Synergy

The infrastructure in Denmark today does not facilitate a spread of knowledge related to HPC from one scientific community to other scientific communities. This means that generic knowledge about algorithms or optimisation that might have been established within for instance computational chemistry, most likely is not spread to other sciences, simply because the existing HPC infrastructure in Denmark does not support such transfer of knowledge. For the time being this is very problematic since a worldwide shift from a "serial-code focus" to a "multi-core CPU code focus" will happen in the next few years, and this implies a substantial call for new algorithms and HPC solutions to fully take advantage of the new technical possibilities.

Furthermore, the synergy between applications of HPC within physics, chemistry, etc, and the research in scientific and high performance computing is lacking. The existing infrastructure does not facilitate this synergy. One of the main motivations behind this report is to try to elaborate on the positive effects of such a synergy. This would imply that researchers within HPC user groups, together with researchers within scientific and high performance computing, could be able to exchange not only existing knowledge, but also to discuss problems to be considered in future research projects as well as knowledge on trends and directions in hardware and software. By facilitating this synergy the infrastructure will also facilitate a transfer of knowledge to high level courses on scientific and high performance computing. This will then present a possibility for educating new experts who later on could work on HPC applications, even within areas like physics and computational chemistry.

A main motivation for this project is to study how these synergy possibilities for synergy possibilities are identified and considered in other countries, and to discuss these issues with existing users of HPC in Denmark.

2.3 Industrial and other funding

In most other countries we observe a strong link between industry and the large HPC facilities and their focus. Examples of major industrial companies with a strong collaboration with major HPC facilities are to be found later on in this report.

Entries on the TOP500 or the TOP100 list (detailing the 500 and 100 most powerful computational installations worldwide) also seem to be indicative of the extent of collaboration with the industry. Today countries close to Denmark all have several facilities on the TOP100 list.

An entry on the TOP500 list is also a ticket to funding from, for instance EU. EU recently allocated huge funding in support of large HPC facilities. An example is the FP7 project PRACE. About 500 Million Euro were allocated to HPC to cover the initial phase of the project. At a later stage several hundred million Euro will be allocated. One of the primary purposes of PRACE is, together with the European IT-industry, to influence the development of new technologies. In order to qualify for PRACE membership, the HPC facility must be of a certain size, and none of the 11 DCSC installations is a candidate for even a PRACE entry level installation. Hence, Denmark is one of the only countries in the old EU that does not participate in PRACE – as also seen in Figure 1. One important implication of

Figure 1: Countries listed are participating in the new European PRACE project. The main goal of PRACE is together with the European IT-industry to influence the development of new technologies and components for large scale computer systems.

this is that Denmark will have difficulties getting EU funding for HPC activities.

A motivation behind this study is a wish to establish an infrastructure in Denmark which facilitates more industrial funding and represents a better possibility for international funding and collaboration.

2.4 Governmental research activities

The integration of governmental research activities in universities was one of the major changes to the Danish research infrastructure on January 1, 2007. With this new structure, the universities have a new responsibility with "Research-based Public Sector Consultancy", including the related strategic research. This is also reflected in the new University Law. As per the Law, the universities now have four main tasks: Research, higher education, research-based public sector consultancy and propagation of knowledge. The extent of research-based public sector consultancy includes DKK 1.1 Billion in the Finance Act for 2008.

It would be natural to link a new, High Performance Scientific Computing facility to a new eScience infrastructure for Scientific Computing. Considering the situation today, such a link, however, raises a special problem for the research-based public sector consultancy, because the existing infrastructure does not take into account the special requirements for this type of activities. The current set-up is not satisfactory, taking into account that models and large amounts of data play a vital role in the forecasts, scenarios, etc, which form a significant part of the research-based public sector consultancy.

For the allocation of HPC resources, currently quality criteria are used that are analogous with the criteria traditionally applied in basic research, i.e. the scientific quality alone, while the relevance is not included. Other parts of the Danish research system have responded to this issue: Thus, The Strategic Research Council (DSF) has developed a quality assessment in which at least three equal criteria are applied: The relevance of research, the research impact, and the scientific quality. A given project is thus evaluated on both the research results which are created, as well as the possible economic or social impact it may create. A similar evaluation of quality and relevance is necessary in the allocation of scientific computing resources, in order for the Danish HPC facilities to be able to support the new university tasks of public sector consultancy.

It would also be natural that the Ministry of Science, Technology and Innovation assess the composition of a new eScience board in the light of this new situation. Considerations should especially be taken concerning how the public users (e.g. the sector Ministries) of scientific computing resources can be represented in the strategic decisions in an eScience community, e.g., in line with the current industry representatives.

2.5 Total cost

Today only hardware cost is accounted for, whereas the cost for maintenance, power, consultancy, buildings, manpower, etc. is hidden. Also, the additional cost due to the fact that experts in the research fields need to spend considerable time on porting codes, moving data around, etc., is not taken into account. For instance, the power is paid by the individual universities. A serious consequence of this is that it often happens that research groups from one university are not allocated to a HPC facility at another university, even if the hardware and software suite is more optimal there. In some cases the Head of the university has denied access of research groups due to the hidden cost covered by the university. In most other countries the total cost of HPC facilities is visible.

During our visits to users of HPC in Denmark the issues related to the current cost structure were addressed.

Also, the level of organisation and strategic planning is much higher in comparable countries (e.g., U.K., Norway, Germany and Finland). This became very obvious from many of our visits. The professionalism is seen not only with respect to planning but also in the daily operation (help desk, code porting, code tuning, operations, optimal choice of algorithms, optimal choice of hardware, ...).

The actual infrastructure in Denmark also has the serious implication that a technology watch is lacking. Focus on the entire IT structure in DK is non-existent. During our visits to other countries a priority has been seen to multi-core and parallel computing. However, during our visits to a number of users of HPC in DK it was clear that the knowledge on these aspects and fields is minor or non-existing. One of the motivations behind the visits to HPC users in Denmark was to discuss how such a knowledge can be established and further facilitated in Denmark.

2.6 Procedure for applications

Today DCSC calls for applications once a year; but locally there is some possibility to get access. However, at each of the 11 installations the access procedure is different, which makes it very difficult for new users to get started and it reduces mobility. Another problem today is that many of the users getting financial support are (directly or indirectly) represented in the board that allocates the funding.

In almost all other countries the procedure is different; more frequent possibilities for applications, harmonisation between HPC sites (or a central portal also for smaller applications), etc.

Another important aspect observed in other countries is the seperation between the board allocating the funding and the users applying for funding.

At EPCC in Scotland, they have built up a very large facility, and currently all application for CPU time is granted, independent of research relevance or quality. If there is an acute and very high demand for resources, an application can be made directly to the relevant research council, who will then instruct EPCC on how to proceed. At the University of Oslo in Norway, every research group owns a part of the HPC centre. However, when not in use, other research groups can take advantage of the free capacity due to a very flexible operating system. CSC in Finland is a private company and they perform the allocation of HPC resources themselves. At CSC it is believed that it is impossible that leading scientists allocate computing resources to themselves. If the scientists at a sufficiently high level of research, they are not neutral, and if they are neutral they are not competent.

3 Our approach

In the request from the Danish Agency for Science, Technology and Innovations it is emphasised that our project is to outline the possible synergy effects obtained by establishment of a central HPC facility in Denmark.

The purpose of this project is to prepare the background for a discussion whether it is beneficial to establish a national High Performance Scientific Computing Facility (HPSCF) in Denmark. Besides the synergy effects, also organisational and economical aspects should be addressed.

In order to establish a background for answering these questions we will consider the following main issues in this report:

- An investigation among the Danish universities and research organisations for the purpose of identifying the need and possibilities for establishing a central high performance eScience infrastructure.
- Clarify the scientific, economical, and organisational aspects and synergies in establishing a central facility.
- Identify the potential need and benefits of establishing a TOP100 facility.
- Investigate the possibilities for a collaboration with the existing national facility for allocation of HPC resources, i.e., DCSC.
- Identify how HPC is facilitated in other countries.

In order to address the issues mentioned, we have had a long list of meetings with some of the existing users of HPC in Denmark, and we have organised visits or meetings with a rather long list of international HPC facilities. This list contains HPC facilities within Europe, US and Canada.

In parallel we have had a series of meetings with DCSC; these meetings have been with the Head of the DCSC board, Jens Kehlet Nørskov, several board members, and the CEO, Rene Belsøe. The latter was also invited to join the group at several international visits.

In this section we will briefly summarise these meetings. A more detailed summary is found in the Appendix.

3.1 Meetings and visits at international centres for HPC

In Europe, US as well as in Canada the national HPC setup is rather different from the current setup in Denmark. Foreign HPC centers strongly focus on having a facility that offers support to the users. The main motivation behind this is to ensure that the researcher can spend his time efficiently on activities directly related to his expertise, and it is recognised that experts on scientific and high performance computing are able to solve problems related to HPC much more efficiently. Another motivation is to ensure that the hardware, software and infrastructure are used as efficiently as possible.

It was mentioned several times during our visits that a large HPC facility with an embedded infrastructure related to user support and with close relation to universities which offer research and education in scientific and high performance computing, leads to several possibilities for synergy:

- Synergy between various users of the HPC facility. A central large facility serving different fields of high performance computing facilitates a possibility of absorbing information about new breakthroughs and some of this information is useful in other scientific areas.
- Synergy between the scientific users and the experts on scientific and high performance computing. A close collaboration between users of HPC and experts in scientific and high performance computing implies that the researcher can spend time on his expertise, and experts on HPC can ensure that observed beneficial solutions and research within scientific computing can be used readily within a wide range of scientific areas.
- Synergy between application, research and education on scientific and high performance computing. The link to education is important, since this link ensures that new researchers with an up-to-date knowledge on HPC can be allocated to various scientific areas with a need for HPC.
- Synergy between the host of the HPC facility and the expert knowledge how to build up and run large facilities. In many other countries, large HPC facilities are built up. Technical knowledge how to build up such facilities is important in order to use the power efficiently and in order to have an efficient communication. Also, the location of facilities is important in order to be able to absorb the cooling power in, e.g., district heating systems.

In most other countries we observe a strong link between industry and the large HPC facilities and their focus. In Germany it is well known that, for instance, Porche and Daimler Chrysler have a strong collaboration with some of the three national and state funded large HPC installations (LRZ (Munich), HLRS (Stuttgart), NIC (Jülich)). In Finland, CSC, which during its early years of operation actually copied many of the ideas of UNI-C, prioritises disciplines like telecommunication networks. Today, the Nokia Headquarter is located about one kilometer from the CSC main building. In most other countries, a strong link exists between large HPC facilities and the industrial sector. This seems to be facilitated not only by significant HPC facilities but also by collaboration between the facilities, technical universities, and similar research institutes. Sweden, for instance, has nine entries on the TOP500 list, and six of them are private companies. The remaining three entries are governmentally funded, and found on the TOP100 list.

Today Sweden has 3, Norway 2, Germany 11, and Denmark no entries on the TOP100 list. Even on the TOP500 list Denmark has no entries today.

Apart from all the above-mentioned synergy effects – which today in Denmark are weak or totally lacking - we have also experienced that it is important for a country to have facilities on either the TOP500 or the TOP100 list. Besides being the link to industry, an entry on the TOP500 list is also a ticket to funding from, for instance, EU. EU has actually allocated huge funding in support of large EU facilities and the entire infrastructure. As an example, let us again draw attention to the new EU PRACE project. This project is going to prepare a pan-European HPC service consisting of three to five large HPC centres, similar to the US HPC infrastructure. PRACE is going to establish the tier-0 level of the European HPC ecosystem, and the hosting centres of the planned tier-0 systems are going to provide the expertise, competency, and the required infrastructure to meet the challenging demands of users from both academia and industry. The initial EU funding for PRACE is about 500 Million Euros, to be followed up by annual funds of about 100 Million Euros. PRACE is going to collaborate with the European IT industry to influence the development of new technologies and components for future HPC systems.

In total, 14 European countries have facilities that are members of the PRACE consortium. In order to be a candidate for membership, the facility must be of a certain size. As mentioned previously none of the 11 DCSC facilities qualifies for even a PRACE entry level installation and, hence, Denmark is one of the only countries in the EU, which is not participating in PRACE. The participating countries are indicated on Figure 1 on page 12. This implies not only that Denmark has reduced influence on the future HPC system development, but also that Denmark will have difficulties obtaining EU funding for HPC activities.

3.2 Meetings and visits with Danish users of HPC

Several research groups at Danish universities were contacted. Limited by time, and by the ability to unravel new HPC user groups in the short time at our disposal for preparing this report, the users' conclusions are somewhat biased. They are

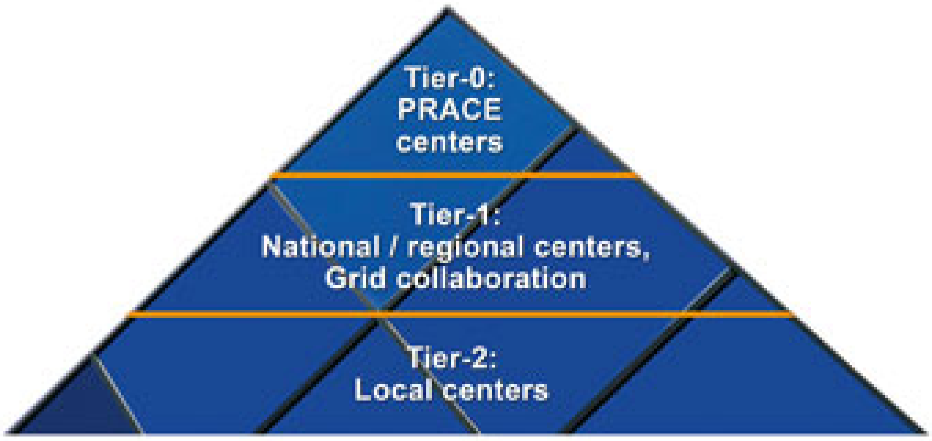

Figure 2: The pyramid structure of HPC facilities in the European funded PRACE project. None of the 11 DCSC facilities in Denmark today are large enough to become a candidate for even a PRACE entry level installation.

biased not only due to the problematic issues in identifying yet unknown HPC research groups, but also due to the fact that DCSC today is in charge of all HPC funding.

We are positive that Denmark has user groups at several universities, university colleges and polytechnic institutes. Groups to whom HPC is new, but who believe that they in the future could make breakthrough research with the aid of HPC. Several groups have contacted us after the meetings (e.g., researchers from CBS in Copenhagen and the Bio Systems Department, Risø) and expressed specific interest in introducing HPC into their field of research. We believe that this is just the tip of the iceberg.

All in all, the visits were positive. Most of the research groups were in favour of the proposed plan presented. However, some groups expressed doubts about the proposals (NBI, Chemistry in Aarhus). Their main objections referred to 'efficiency' of investments, questions regarding the need to involve new user groups, etc. In particular, Chemistry in Aarhus explicitly stated that a local system is an absolute must for their research.

During winter and spring of 2008, meetings were held between the core project Group, consisting of Jostein Sundet, Lise Frohn, Bernd Dammann, and Henrik Madsen, and a number of HPC user groups in Denmark.

Prior to the meetings the following two documents were sent out:

- A pdf-presentation as presented to the Director of DCSC in December 2007 and to the board in March 2008 (see Appendix G.3).

- A preliminary description of the proposed infrastructure for HPC in Denmark. This document was also presented to the board of DCSC in March 2008. The document is included in Appendix G.1.

In order to try to identify synergy effects and possible needs for organisational changes, etc., the above documents contain a rather open invitation to DCSC to discuss the need for a change of the HPC infrastructure in Denmark

At all visits the core Group started by explaining the background and the purpose of this pre-investigation that is to come up with ideas for improvement of the current HPC situation for all research groups in Denmark, including the need for HPC, storage, or visualisation.

A discussion with the research groups was then conducted, where Q&A and comments were recorded and commented. After the visits, meeting minutes were submitted to the visited sites, open for comments and approval up to 6 weeks after the meeting.

For the meeting minutes please see Appendix D.

A list of positive remarks to our proposals:

- A general statement, more professional management of the individual systems.
- More availability of user support and advanced user support.
- More focus on efficiency.
- Availability of courses in scientific and high performance computing, scripting, parallel programming, ...
- Porting of source codes and help for using proper libraries, compiler options between systems in Denmark.
- Introductory course in HPC and scientific computing.
- Advices on new hardware and hardware trends.
- Standardisation of hardware and software at the Danish HPC sites.
- A uniform way to get allocation of CPU, disk space, and support.
- A more flexible procedure for applications.

A list of negative remarks to our proposals:

- Generally, less investment in hardware if the user support part is increased.
- Loss of own influence on use of hardware, on control over funding and on selection of hardware procured.
- Afraid that a new central facility will become another UNI-C and that the focus will change from the researcher to the organisation (aka, UNI-C in the 90's)
- Predictability of turnover time of selected problems
- Loss of influence of DCSC in the Danish HPC landscape

In general, the HPC users visited in Denmark are in favour of changes of the infrastructure.

A common finding is also that HPC moves towards new technology, multi-core moving into many-core with sufficient amount of memory. Many of the HPC applications mentioned are found to be memory bounded. These generally accepted facts make the need for a central, large computing resource more pronounced. Such a resource, coupled with (advanced) user support and professional management of a computing/storage facility, is a natural consequence of this need. But, several users are still afraid that this will become a new UNI-C.

It was, however, mentioned that UNI-C before, say, the mid 80's, was found to be beneficial and useful for Danish HPC users. During that period UNI-C had strong relations to some of the universities and groups considering the theoretical aspects of scientific and high performance computing.

3.3 Meetings with DCSC

Initially we had a meeting with the board of DCSC. At a later stage we had a number of meetings with the Head of the Board and with the CEO. Our ideas were outlined in a presentation. The presentation was also distributed to the entire board of DCSC. The CEO of DCSC, Rene Belsøe (RB) was invited to join us at meetings with EPCC in Scotland and CSC in Finland. We have had fruitful discussions with RB, and we appreciate the time used by RB on all aspect of this investigation. During the entire project, RB has been very positive and open concerning a discussion about changes and a future organisation related to HPC in Denmark.

At an early stage an external expert in HPC, Director Jostein Sundet, was allocated to the project and he has participated in all the meetings since January 2008. The external expert also had a meeting exclusively with DCSC.

In March the project group behind this report formulated a document describing the ideas about needed changes of the HPC infrastructure in Denmark. The document contains a rather open invitation to the board of DCSC to collaborate on defining the future eScience structure. A copy of our document and the letter from DCSC is enclosed in Appendices G.1 and G.2.

4 Synergies and Collaborative Efforts

This section contains a further elaboration on some of the important aspects previously mentioned, namely the possibilities of a centralised and large facility to enable a series of synergies, and some aspects of grid and parallel/multi-core computing.

4.1 Synergy with respect to research in scientific computing

For the past eight years the focus has been on obtaining hardware for a rather limited number of researchers, mostly within physics or chemistry. The focus must, hence be considered as applied with respect to scientific and high performance computing. The feedback and interchange between these applications and the research related to scientific and high performance computing is mostly non-existing. However, some financial support and focus have been devoted to grid computing, but in those cases the massive research related application is very low.

A centrally organised HPC infrastructure will enable a synergy between applications, research and education in high performance and scientific computing. Furthermore it will create synergies between research areas using HPC and new areas will be facilitated.

4.2 Synergy between various scientific applications

During our visits to international HPC facilities we have recognised that today access to high performance computing is needed within a lot of other areas, such as data mining, statistics, informatics, signal processing, biology, finance, and economy. Since many of the scientific problems, e.g., with respect to algorithms, are more or less the same across all scientific areas, a national infrastructure will facilitate a cross scientific transfer of knowledge. During our visits abroad we have, for instance, seen a lot of examples where the classical users of scientific computing are taking advantage of knowledge built up within areas like signal processing and statistics.

Furthermore, the new infrastructure has to be more transparent to new research groups, and in that way facilitate not only the classical areas like physics and chemistry, but also new research areas. Finally, we have noticed, as mentioned before, that in many countries, large funding is allocated to research in scientific computing related to complex and large data within completely new areas like history, multimedia, linquistics, and culture.

4.3 Synergy between HPC and datacentre know-how

It is generally well-known that the energy consumption of datacentres worldwide is considerable and must – due to higher energy prices and an increased focus on emission, etc. – be addressed much more carefully in the future. The need for more experience in this area is also due to the tendency of consolidation of datacentres.

In our neighbouring countries several of the TOP500 datacentres belong to the industry. Hence we see a possible synergy potential in centralising the knowledge how to built large IT centres at a technical university that is already linked to existing know-how on power systems, cooling, heating and civil engineering. By operating a large-scale HPC facility on, e.g., TOP100 level the relevant university will gain valuable information. Such information could be transformed into knowledge regarding the construction of IT centres that focus on green computing.

4.4 Grid computing

During the past years, grid computing has obtained a lot of attention, both with respect to financial support, and politically. The more and more pronounced wish to be able to supply storage and computing in a distributed net reflects this. For a few research areas, this is obvious, but it is important to consider a more balanced support for grid computing in relation to the need for an efficient national high performance computing infrastructure.

The idea of connecting a large number of small computer installations in a grid facility is appealing. However, a number of problems still need to be solved before such a grid computing will be able to serve the mainstream researcher. With the loss of clock cycles between nodes it is also clear that the applications running on grid solutions will only be an optimal solution for a small part of real life scientific problems. Furthermore, some scientific problems, e.g., involving large amounts of readily available input data, are not suitable for grid solutions.

In other countries we have seen a better balance between the focus and support on grid computing and other relevant areas of scientific and high performance computing, such as those mentioned in the next sections. Some of these other areas also seem to be tightly connected to the direction software developers must follow in order to build software for the future mainstream computer.

4.5 Multi-core and shared memory computing

More and more data is being generated and more and more complex problems are being considered, while users at the same time want to process that data and computations in smaller and smaller time frames. The solution to these colliding demands is obvious – parallel computing. This requires a fully parallel infrastructure, ranging from hardware (parallel computers, fast interconnects) to software support (parallel libraries, parallel file systems).

Up to a few years ago, the major manufacturers increased the speed of computing by raising the clock speed of the CPU. This had two side effects: the CPUs needed more power, where most of the power in fact was dissipated into heat, and the gap between the CPU clock speed and the rest of the machine, i.e., memory and bus, was getting wider and wider. The latter means, that the CPU most of the time was wasting cycles waiting for data from memory to be available for the next computation.

Since then it has been realized, that the way to obtain more computing power is to build multi-core CPUs, i.e., to increase the number of processing units on the processor die, keeping the so-called power envelope of the CPUs constant. For the time being, most CPUs on the market run at clock speeds of 2-4 GHz, i.e., no increase of clock speed has taken place during the last years. One exception, though, is the IBM POWER chip series that keeps increasing clock speeds from one generation to the next generation. However, machines with the latest POWER6 chip need to be water-cooled!

Multi-core CPUs share the same benefits as SMP computers with respect to data locality, but introduce new levels of data locality, too, due to the cache structure of the new multi-core CPUs. Already today, four- to eight-core CPUs are available, but in the near future this number is going to explode. An 80-core prototype is under development and tested by Intel. These new multi-core machines offer a single address space, which implies that a software revolution is taking off, and that there is room for both improvement of existing algorithms and for completely new methods.

Besides putting more cores on a single chip, another trend is the so-called CMT (chip multi-threading) or SMT (simultaneous multi-threading) technology. This allows each core to execute multiple threads simultaneously, with (almost) no overhead for a context switch. One such CPU is the Sun UltraSPARC-T2 chip, which is an eight-core chip, where each core can execute 8 threads simultaneously.

In fact, multi-core chips have been known for a long time in specialised hardware, e.g. for graphics cards, the so-called GPUs, or the Cell processor, known from the PlayStation 3. For certain types of applications, e.g., some types of matrix computations, the speed-up that can be achieved by pushing those computations into those add-on cards, is impressive. The fastest supercomputer today, the Roadrunner, is using add-on cards with a modified version of the Cell processor.

We believe that this is just the beginning of a revolution that forever will change not just how high performance computing is done, but also how any kind of computing is done.

At the latest super computing conferences (Reno, November 2007 and Dresden, June 2008) presentations clearly indicated, that within a few years the size of multi-core CPUs is going to be measured in kilo-cores. This type of computers and software calls for a new generation of application software which automatically uses the parallel access to cores and data. We are confident that this will open a new paradigm for software development.

The current Danish HPC infrastructure is not ready for this challenge. However, with the intention to play a difference within the IT sector, we need to change the situation.

4.6 Power consumption and performance

New multi-core CPUs are very efficient with respect to power consumption. Typically, it is the goal of the chip developers to increase the number of cores, without changing the so-called power envelope of the CPU, i.e., the power consumption. As an example it can be mentioned that the new number one on the TOP500 list, the Roadrunner, is approximately twice as fast as the number two on the list, the BlueGene/L, but the power consumption is about the same for both facilities. The tendency that the efficiency to power ratio is much better for multi-core CPUs will increase. Tests conducted at Berkeley labs have shown that the consumption of a 64-ways multi-core UltraSPARC-T CPU is only about the same as the power consumption of a classical, 6+ years old, single core Pentium IV from Intel.

4.7 Software development and user support

The single address space in the new generation of computers implies, that it will be possible to set up flexible clusters with fat nodes, where the operating system and compilers related to each of the fat nodes, essentially takes care of the parallelisation within the nodes. This will significantly reduce the need and work involved in using technologies like MPI to organise coherence across cluster nodes. However, for large-scale computations, it is necessary to combine both technologies, i.e. multi-thread within a node and MPI between the nodes, in order to make the best use of the resources.

Based on the fact that we are facing a new paradigm for parallel computing, we feel an obligation for building up and to consolidate the computing power by building up a flexible super-cluster with fat nodes. This facility is aiming at serving researchers at a national level.

At this facility we also propose to build a support infrastructure, which can be used for exchange of ideas related to supercomputing. Conferences and workshops should be arranged as well. We see the facility as a possibility for creating a fruitful environment, where users of high performance computers and research and education in scientific computing have a mutual benefit of access to the same systems. Such a collaboration has a huge potential not only for enabling an infrastructure for developing the next generation for software for solving large problems, but also existing codes can benefit.

In fact, optimisation of production codes, as well as specialised courses in programming, e.g. OpenMP and MPI, can be of tremendous benefit, as is known from experiences at different sites, e.g. DTU Informatics, CSC in Espoo (Finland), RWTH Aachen (Germany), etc. In an article in 'CSC News 2/2006', about a 3-year code optimisation project in Finland, funded by the Finish funding agency for technology and innovation, Jan Åström from CSC writes: "The general experience of the code optimisation project (at CSC) so far is, that with reasonable effort it is possible to reach a 50%, or more, cut in computing time for most codes. This means that the effective outcome of a code optimisation project is a doubling of the computational resources without investing in any new hardware!"

In conclusion, a new facility will create a national background for obtaining a synergy between applied researchers who demand scale computing, and researchers within scientific computing and software engineering. The facility will absorb and spread information related to scientific computing and code optimisations. This again will ensure that the hardware is used efficiently. Finally, such a facility will ensure that Denmark has the possibility to be among those nations which will be able to build software for the future massive multi-core computers.

5 A national HPC facility - Challenges and Suggestions

In this section some of the various challenging ideas and thoughts with which one may be faced in the process of establishing a national HPC facility, are discussed. In close collaboration with the HPC users in Denmark, the main universities relevant to HPC related research, and with the Ministry of Science, Technology and Innovation, these thoughts need to be further elaborated on.

5.1 Datacentre facilities and expertise

After the shutdown of the central infrastructure at UNI-C in 2000, a large part of the expertise in running and maintaining a large scale installation has been lost. A new facility will, thus, require both the construction of a new datacentre and the hiring of the necessary staff. However, by establishing a tight connection between a new central facility and the universities with expertise in these aspects, it is judged that such an expertise can be reconfigured and tailored to the situation of modern scientific computing.

Since it is suggested that a new facility should be found on the TOP100 list of 2009/10, a major investment into new datacentre infrastructure will be necessary, i.e. new building, power and cooling, networking infra structure, etc. Some of the existing infrastructure at eg. DTU can however be reused. For the HPC facility owned today by DTU, the basic maintenance is carried out by UNI-C at a very moderate cost, whereas the HPC and scientific computing, and support facility is hosted by a department on the university (DTU INFORMATICS).

To justify an investment in new infrastructure, it is also important that the funding of the new facility is ensured for a long-term period (8 years or more, better 12 years).

After a successful start-up phase, sale of services to interested customers (industry, other partners) should be considered. Income from such sales would contribute to maintenance cost (see also the model SARA in the Netherlands uses). The facility could also offer 'integrative hosting', i.e., hosting of hardware for other HPC installations, e.g., Danish universities.

In connection with the facility, it is quite important to establish a setting, that makes it attractive for foreign HPC experts to visit the facility and to contribute to its development.

5.2 Size of the installation

It is suggested that the central facility should be large enough to qualify for inclusion in the TOP100 list of 2009/10. The facility should also be configured in such a way that it re-invokes a possibility for major international (mainly EU) funding, and for international collaboration related to the important areas of parallel computing, which will be the focus in the IT-development in the years to come.

During our visits it became clear that other countries invest seriously in HPC facilities. In Figure 3 some known plans for investment (measured in peak performance) in different European countries are indicated for 2009.

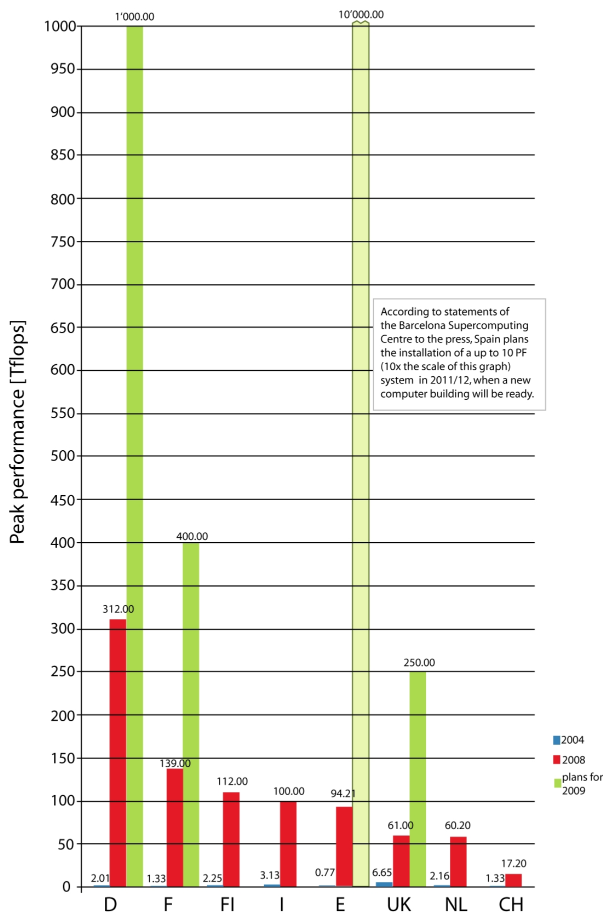

Figure 3: Peak performance of the fastest generally available computers in different European countries for 2004, 2008 and some known plans for 2009 (source: *The Swiss National Strategic Plan for High Performance Computing and Networking*).

5.3 Design process for the HPC facility

In order to get the best solution and to avoid a lot of problems, e.g., regarding the datacentre design, etc, it is suggested to team up with a vendor already from the very beginning of the design process. This has, e.g., been done for other large scale installations, like the TACC Ranger in the US. The vendor and the organisation that wants to host the facility work together on preparing a proposal. The final proposal is then submitted to the funding agency.

5.4 Collaborations - national and European

The new installation can be a step for Denmark to return to the European map of HPC infrastructure, e.g., via participation in the PRACE project. The aim of the PRACE project is to build up a system of a few European large-scale installations (Tier 0), while the national partners provide tier 1 (national) and Tier 2 (local, e.g., university) installations. So far, 15 European nations and Turkey have signed a 'Memorandum of Understanding'.

The HPC facility could, e.g., be a Tier 1 installation in the PRACE project.

5.5 Organisational challenges

Since the facility is supposed to be national, it must be controlled by some national council or ministry. DCSC is one possibility, but this requires a major change of the DCSC organisation and rules.

Another possibility is to set up a parallel eScience structure. In that way the users will be able to decide the preferred infrastructure (and organisation) simply via their applications. This will then lead to an optimal mix between a centralised and a decentralised structure for HPC facilities in Denmark.

No matter which solution is found to solve the problems above, it is necessary to give some thoughts to the very organisations that is to run the facility. This task may be 'outsourced', e.g., to one of the universities, either in a public or closed tender.

5.6 Suggestions

We suggest a national facility that will encourage new research and invite new research groups to use HPC. Such a central facility will leverage synergies between user groups from different disciplines as well as between users and HPC specialists. Furthermore the pooling of money will make it possible to buy one large system which – compared to the sum of several smaller systems – will have the advantage of more optimal operation, maintenance and user support, centralized applied and theoretical IT courses (e.g. on numerics, statistics, informatics, math) as well as advanced user support. A national central facility, in TOP100, will put Denmark back into the international HPC landscape again, with entry to current European projects (PRACE, EGEE-III) as well as future EU projects. Finally, it must be expected, as seen in other countries, that such a national facility will lead to more collaboration with industry. A transfer of knowledge to the industry is important; especially since the IT-industry of today is facing a paradigm shift towards massive multi-core and high memory computers, with an imperative need for changes in software and solutions.

Strategic investments in HPC is a necessity in Denmark, as the funding is limited. As of now the individual researcher is investing in what is his or her preference. However, to get most value for the money for as many researchers as possible, one has to implement a strategy that takes into account the prospects for resarch in Denmark overall in an e-Science perspective.

This again will enable Denmark to participate in international collaborations and move Danish research to international standards. It has been seen in other countries that such large facilities have strong collaborations with industry (eg. EPCC, CSC, and HPCVL).

The facility should host a diverse spectrum of machine architectures in order to ensure optimal hardware for solving dedicated problems. This also facilitates the development of software for the mainstream computers of tomorrow. The facility should be a heterogeneous solution but facilitated in a homogenous way to ensure mobility and a smooth transition towards new trends in supercomputing.

We need a national eScience board which is able to define likely directions of supercomputing. In addition an efficient reporting structure should be set up in order to ensure that the needed flow of information is available at the board level. Reporting from the researchers as well as from the individual site administrators is a must to be able to evaluate the efficiency of use of the different sites and document the value of the investments.

Total cost of ownwership (TCO) is important has to be visible. The increasing cost of energy and CO2 emission demands a focus on the power consumption already at the design stage. Today 11 installations are running under DCSC and the total use of manpower for operation is very difficult to estimate. Also the inhomogeneous setup of the installations makes it rather difficult or even impossible for the users to migrate to another installation. DCSC is reporting the cost of hardware, but the total cost is probably at least double the hardware cost. In general, costs for electricity/cooling only is about the same expense as the hardware costs over a three year period. In addition, one must also include the man months spent operating all the centers.

Other in-kind expense is also probable, e.g. PhD students. It is also very problematic, that professors themselves use time to make code run more effeciently, whereas it is well-known that experts in scientific and high performance computing can obtain better and faster results in code optimization.

Low exploitation of hardware is an issue, since an investment in hardware is more cost-efficient if the exploitation is as high as possible in the entire life-time of the computers. In Denmark the reporting of utilisation of the investment to the funding agencies is weak or non-existing and a more stringent procedure for reporting should be implemented. Utilization of a computer resource show the efficiency of the investment.

As suggested the central facility should host mixture of hardware which, however, are operated homogeneously. This is needed both in order to reflect various optimal hardware solutions, and in order to prepare the HPC community (as well as the local industry) to new IT and HPC solutions.

Examples with motivations

1. Blades – for eg CFD modelling and simulation.
2. FATnodes – for eg. application with large amount of data and data transfer.
3. CELL – this is the architecture used in RoadRunner - the World No. 1 today. The energy consumption is rather low but with some further demands for optimization of libraries etc.
4. Niagara – a rather new massive many-core CPU architecture which seems to be very energy efficient, but with new demands on software, and data handling. Might be a candidate for some green IT-solutions.

The first two examples are conventional CPUs, whereas the latter are new potentially interesting new architectures. A resent study at the Berkeley has shown very promising performance/energy ratios for both the Niagara and the Cell architecture (see the Reference List).

6 Conclusion

This report describes the results of a study initiated by the Danish Ministry of Science, Technology and Innovation for the purpose of evaluating the national situation for High Performance Computing (HPC) and, in particular, to investigate the need for a large TOP100[2] facility in Denmark. The support for this study was offered as a result of two applications for larger centralised HPC facilities – one application from DTU and another from DMU / Aarhus University.

In order to establish a background for the study, our approach was to visit a number of facilities in Canada, US, and Europe. In Denmark we have visited and interviewed a large number of HPC users. Based on these interviews and on findings during our international visits we have drawn conclusions about an optimal setup for a future HPC infrastructure in Denmark. We invited DCSC to join us at some of the international meetings, and we have also had several meetings with DCSC. We appreciate all the time spent by DCSC on issues related to this study.

It is identified, that the best solution is a combined decentralised and centralised infrastructure for HPC in Denmark. This conclusion is mainly based on interviews with a number of existing HPC users in Denmark. For some of the research groups, the demand for computing is found to be most adequately solved by special purpose decentralised facilities hosted at, or nearby, the relevant research group. Today this applies to some of the 11 DCSC installations. The decentralised facilities are today designed, constructed, negotiated, administered, maintained and reserved mainly by the local research group.

However, many research groups would prefer to have access to a large centralised HPC facility with a professional and highly skilled user support. The facility should offer help with issues like code tuning, porting, code profiling, selection of algorithms, selection of libraries, parallelisation, and efficiency. It is also concluded that the facility should be linked to a university that offers state-of-the-art research and education in scientific and high performance computing. Such an infrastructure will enable the synergy potentials between scientific applications and research and education related to scientific and high performance computing.

Today the spectrum of HPC users in Denmark is narrowed down to a few research areas. This is contradictory to most other countries where HPC usages within a wide spectrum of scientific areas are rapidly increasing. These areas include telecommunication, data mining, statistics, pharmaceutics, linguistics, multimedia, virtual reality, culture, etc. Some of these areas are characterised by having a

[2] In the rating of the performance we have picked the entries on the TOP100 – or TOP500 list – since these ratings are based on real measured Tflops. Other references use the peak or theoretical Tflops.

special demand for intelligent storage and treatment of large amounts of data.

Firstly, a centralised large HPC facility in Denmark will make it possible for many research groups – either individually or jointly – to solve much bigger problems than is possible today by the 11 DCSC installations. This will enable all research groups to carry out cutting-edge calculations with a high scientific impact. Secondly, an efficient and centralised facility, with a reasonable amount of storage and data handling capabilities, will open up the possibility for a much broader use of HPC in new scientific areas – with very promising potentials, both regarding research and innovation – as a result.

In other countries it is seen that a large facility motivates collaboration between industry, researchers within scientific computing and the various scientific areas. This leads again to a large industrial impact. In Sweden, for instance, the industry accounts for 6 of the installations on the TOP500 list.

We see a need for building up a TOP100 facility as soon as possible. In order to be a candidate for EU funding and international collaboration, the facility must be on the TOP100 list of 2009/10. Such a facility will open up for collaboration with the industry, significant international funding and international collaboration. Besides international collaboration and funding, a national facility will also make it possible for Denmark once again to be more visible and represented at conferences and international meetings related to high performance computing.

The new facility has to be an important part of a new HPC infrastructure with a focus on all aspects of scientific and high performance computing and, hence, enable a healthy and inspiring environment for dealings within scientific computing and HPC. The HPC infrastructure in Denmark must be reconsidered, so that a real focus on user support, code migration, efficiency, new technologies, etc. can be established.

The facility must be located so that the HPC community in Denmark can take advantage of the long list of synergies, such as documented in further detail in this report. This implies that an obvious location for the facility could be at the Technical University of Denmark (DTU), where the only identifiable education (and most of the research) on scientific and high performance computing takes place today.

The central HPC facility has to be large enough to solve critical mass computations within a large number of scientific areas. With an appropriate infrastructure, including support on various levels, such a facility will also imply that the demand for HPC, data storage, data analysis, etc. within a much broader spectrum of scientific areas can be facilitated. This broader spectrum of HPC users can be seen in most other countries, where a wide range of scientific applications within areas like data-mining, statistics, pharmaceutics, linguistics, multimedia, etc. are supported and hosted at large central facilities.

Today Denmark is not to be found on the TOP500 list of the largest computer facilities in the world. This is contradictory to almost all other countries in Europe (including countries like Slovenia and Cyprus). An entry on the list is a ticket to possibilities for international collaboration and funding, including funding from the EU. As an example, it can be mentioned that EU has recently allocated more than 500 Million Euro to HPC to cover the initial phase of the PRACE project. Later on more funding will be allocated to PRACE activities. One of the primary goals of PRACE is, together with the IT industry, to influence the development of new technologies and components for large scale computer systems.

It is concluded, that the TCO of the HPC infrastructure must be visible. During the past seven years only hardware cost has been visible. It is also clear, that such a central facility calls for a significant funding. A part of this funding might, however, come from EU and the industry. We see several advantages by reporting on all aspects of the costs of HPC facilities. With the most recent TOP500 list, a focus on energy consumption is introduced, although the committee behind the TOP500 list is still discussing the measures. For the new central facility a focus on the energy consumption is one of the most important aspects.

A new facility linked to education and research within scientific and high performance computing, will also re-invoke and facilitate a technology watch, which is considered to be a very important aspect. For the time being we see a shift from 'serial code focus' to 'multi-core CPU focus'. This shift in technology implies a need for new algorithms, HPC software and tools in order to be able to take advantage of the possibilities. During our visits to other countries we experienced that priority was put on multi-core and parallel computing, whereas during our visits in Denmark it became clear that the knowledge of these aspects and fields is minor, as also recognised by a large number of the HPC groups we visited.

Many universities and research organisations are responsible for handling governmental research activities that often require significant HPC resources. It is obvious to link a new large facility to a new national eScience infrastructure with a focus on both HPC related research, education and research-based public sector consultancy. This is also reflected in the new University Law, outlining the responsibilities of the universities. This will also allow a distribution of HPC resources, based not only on scientific quality, but also on relevance of research and possible research impact - economic as well as societal. Such a set of criteria for distribution of resources will correspond better with the criteria applied, e.g., by the Danish strategic research councils.

A Top500

As mentioned by several HPC centres, an entry on the TOP500 of the largest facilities for supercomputing gives a ticket to international (e.g. EU) funding and collaboration, including collaboration with the industry. It is therefore of interest to observe how the nearby countries are placed on the TOP500 list.

Several of the HPC centres visited as a part of this analysis are found on the TOP500 list, and later on, in the section describing the meetings the actual position on the TOP500 is mentioned for many of the centres.

Today the counties nearby Denmark have the following statistics on the TOP100 list. The list is ranked by R_{max} which is the maximal LINPACK performance achieved measured in TFlops.

The No. 1 (RoadRunner) is based on a modified version of the IBM Cell processor known from the PlayStation 3. This is a next generation multi-core CPU, and it is interesting to note that the power consumption of the RoadRunner is similar to No. 2 (BlueGene) although the performance is more than doubled.

By considering the full TOP500 list it appears that several private companies also enter the list. Sweden has, as an example, in total 9 entries on the TOP500 list, and 6 of these are private companies (mostly entitled Financial Institutions and IT Service Providers).

Rank	Country	Site	Computer	Vendor	Cores	R_{max} (TFlops)	Power (kW)
1	US	DOE	Roadrunner	IBM	122400	1026.00	2345.5
2	US	DOE	BlueGene	IBM	212992	478.20	2323.6
3	US	Ar.N.Labs	BlueGene	IBM	163840	450.30	1260.0
4	US	Univ. Texas	Ranger	SUN	62976	326.00	2000.0
5	US	DOE (Oak Ridge)	Jaguar	Cray	30976	205.00	1580.7
6	Germany	Jülich	BlueGene	IBM	65536	180.00	504.0
.	.	.	.	.	.	.	.
9	France	IDRIS	BlueGene	IBM	40960	112.50	315.0
10	France	TEP	Altrix	SGI	10240	106.10	442.0
11	Sweden	Goverm.	Cl.Blade	HP	13728	102.80	?
.	.	.	.	.	.	.	.
18	U.K.	ECMWF	Power 575	IBM	8320	80.32	1329.7
.	.	.	.	.	.	.	.
26	Spain	MareNostrum	PPC 970	IBM	10240	63.83	?
27	Germany	LRZ	Altrix	SGI	9728	56.25	990.2
.	.	.	.	.	.	.	.
29	U.K.	EPCC	Cray XT4	Cray	11328	54.65	?
.	.	.	.	.	.	.	.
39	Sweden	Umeå	Blades	IBM	5376	46.04	173.2
40	Sweden	NSC	Blades	HP	6400	44.46	?
.	.	.	.	.	.	.	.
45	Poland	Gdansk	Galera	ACT.	5336	38.17	?
46	Norway	Univ. Bergen	Cray XT4	Cray	5548	36.87	?
.	.	.	.	.	.	.	.
51	Holland	Univ. Groningen	BlueGene	IBM	12288	35.12	94.50
.	.	.	.	.	.	.	.
53	Slovenia	T.Inst.	Blades	IBM	4096	35.08	191.08
.	.	.	.	.	.	.	.
70	Finland	CSC	Cray XT4	Cray	4032	26.80	?
.	.	.	.	.	.	.	.
90	Italy	CINECA	Blades	IBM	5120	19.91	488.2
.	.	.	.	.	.	.	.
96	Switzerland	CERN	N.Cluster	Self-m.	3760	19.69	?
.	.	.	.	.	.	.	.
100	Germany	RWTH	Pr.Cluster	F-S	2048	18.81	?

Table 1: Part of TOP100 list as of June 2008. On the TOP100 list Germany has 11, UK 7, Sweden 3, and Norway 2 entries. Denmark has no installations on the TOP100 (and also not on the TOP500) list. On the TOP500 list US installations account for 257 entries. The power consumption is a new measure considered by SPEC and hence not available for all installations.

B Meetings held with international facilities for supercomputing

B.1 Canada

B.1.1 HPC Virtual Laboratory (HPCVL)

In Canada we had meetings with HPCVL, only - one of the leading HPC facilities. This facility provides storage, compute facilities and support for more than 130 Canadian research groups. HPVCL's administrative offices and large central compute and storage clusters are based in Kingston.

HPCVL is a CA Dollar160 million public/private collaboration for research and innovation. Public funding is provided by the Ontario Ministry of Research and Innovation, the Ontario Research Fund (ORF), the Canada Foundation for Innovation (CFI), the Ontario Innovation Trust (OIT), and the federal Natural Sciences and Engineering Research Council (NSERC).

HPCVL is composed of Research Chairs in multiple disciplines. More than 1000 Graduate Students and Post-Doctoral Fellows are trained in advanced parallel programming techniques, and expert user support. They have a number of persons who look at codes from scientific users. Presently, 6 persons are dedicated to performance analysis in order to ensure efficient use of the hardware.

HPCVL is accessible from anywhere in Canada through a portal login and, thus, provides all researchers with a common login procedure.

In the latest investment HPCVL is focusing heavily on the power consumption and the massive multi-core CPUs. According to Executive Director, Ken Edgecombe: 'The statement, dubbed "Moore's Law", was revised by Moore and others over time, but expectation that each chip version would result in greater speed persisted in the minds of users. Now, as chip temperatures approach that of rocket engines, it's clear that this expectation cannot be achieved - and that the technology must develop new ways in order to meet the demands of our modern times. In order to solve the problems associated with the chip heat and performance, producers are moving to multiple CPU cores per chip, keeping the cycle speed down by using multiple threads – or processors – per core. Using multiple threads per core also helps keeping the CPUs operating at optimum efficiency. The result for the researcher is the ability to do more work faster at an energy cost of less than two watts of electricity per thread. That is also good for the environment!'.

Presently, HPCVL collaborates with industrial partners in getting new parallel libraries and methods for use in economics, drug discovery, photonic, bioinformatics, and computeraided molecular design.

HPCVL has a Board of Trustees composed of members appointed by member institutions and non-academic sectors. Responsibility for the formulation and implementation of policy and all operations rests with the Executive Director who reports to the Board. The Executive Director obtains advice on scientific and technical matters from the Technical Advisory Committee (TAC). TAC members are appointed by the Board from both the academic and the private sector.

B.2 United States

B.2.1 The TeraGrid

In the United States the money was previously spread, and the various research groups purchased their hardware on an individual basis. According to several researchers in the US, this was very inefficient. Often the wrong hardware was purchased, the computers were NOT critical mass computers and, in general, there were too many small centres.

After publication of a report in 1999 the allocation of money for high performance computing was changed fundamentally. Now the National Science Foundations' support for HPC infrastructure is organised through the TeraGrid. The purpose is to provide a nationally coordinated infrastructure for research. The individual HPC facilities work together in order to establish a smooth infrastructure in the US for the support of both science and engineering.

TeraGrid integrates high-performance computers, data resources, tools, and high-end experimental facilities around the US. Currently, TeraGrid resources include more than 750 teraflops of computing capability and more than 30 petabytes of data storage. TeraGrid is the world's largest, most comprehensive cyberinfrastructure for open scientific research.

TeraGrid is coordinated through the University of Chicago. The University works in partnership with the Resource Provider sites: Indiana University, the Louisiana Optical Network Initiative, the National Center for Atmospheric Research, National Center for Supercomputing Applications, the National Institute for Computational Sciences, Oak Ridge National Laboratory, Pittsburgh Supercomputing Center, Purdue University, San Diego Supercomputer Center, Texas Advanced Computing Center, and University of Chicago/Argonne National Laboratory.

Allocation of resources is based on scientific records, citations, and a total review of the proposal including the potentials. The review process also tries to find the most suitable supercomputing centre for the proposal.

Presently, the biggest centre in the TeraGrid network is the University of Texas' (TACC) Ranger super-computer, which is built in collaboration with SUN Microsystems. TACC is also the largest non-military facility in the world and number three on the TOP500 list.

A common gateway is enabling the entire community of users to use national resources through a common interface. Science gateways are enabled by a community allocation the goal of which is to delegate account management, accounting, certificates management, and user support to the gateway developers. This means that users in Chicago can readily use the compute facilities at TACC or San Diego by using the same web-page or gateway.

The large super-computing installations within the TeraGrid also enable an efficient support infrastructure, which is considered of great importance by almost all researchers. The support infrastructure typically provides porting of codes, code tuning and parallelisation. For the large scientific areas they try to maintain a domain (scientific area) for specific assistance, but keep the host at the central facility.

B.2.2 State based HPC facilities

In some cases the individual states typically establish a midrange centre to match those found at the NSF centres and National Labs.

As an example can be mentioned the Ohio Supercomputing Center (OSC), which is a centralised state based facility. We met with the OSC.

Ohio Supercomputing Center (OSC) OSC recently celebrated 20 years of service. According to OSC they act as a catalytic partner of Ohio universities and industries that provide high performance computing and high performance infrastructure to a diverse state/regional community, including education, academic research, industry, and state government. Furthermore, Ohio Supercomputing Center promotes and stimulates computational research and education in order to act as a key enabler for the state's aspirations in advanced technology, information systems and advanced industries. For additional information, please see http://www.osc.edu.

A so-called Council on Competitiveness (members are from the Government and industry) is assisting OSC in creating a curriculum that addresses business community needs and regional economic potential.

The high performance computing services, coupled with the portals linking to complex modeling packages, are formulated in order to remove many of the barriers that currently prevent companies from utilizing modeling and simulation as a part of their business practices.

At OSC they have established a consulting group in order to ensure optimal use of the computers. In proposals they must address efficiency. They monitor the use of the computers, and sometimes they flag a batch job for further performance analysis. They assist with migration of codes.

Mostly support via email, but meetings may also be called in case of major problems and new users. They try to embed support people with scientific people.

Based on CPU hours, about 50 pct. is running in parallel - and this share is dramatically increasing.

B.2.3 National Labs

NASA Supercomputing Center NASA Supercomputing Center has a Scientific Consulting Group that assists in solving computational problems, porting of codes, optimisation of codes (often a few days optimisation gives a very significant difference), and simple inspection of codes. Benchmarking of codes is considered of great importance.

Also assists in application oriented optimisation of more than a few weeks duration. They now focus on OpenMP, but still runs a lot of MPI jobs.

Lawrence Livermore National Laboratory At Lawrence Livermore National Laboratory they has a development environment group related to MPI, compilers, performance analysis, porting issues, etc.

The larger allocations must be related to missions in a certain number of areas. The final allocation is based on scientific records and potentials.

They try to allocate a scientific support person with some knowledge in the area. The codes are mixed OpenMP and MPI - in the future more OpenMP.

C Visits at some major EU centres for supercomputing

C.1 UK: EPCC in Edinburgh

The total computing capacity at EPCC is about 100 TFLOPS (HECTOR - Cray XT4 - has 63.4 TFLOPS). EPCC is by far the largest system in the UK. Manchester comes in next, but is much smaller. EPCC took part in a collaboration between US and UK in developing the background for the BlueGene computers (the QCDOC project). Today they focus on parallel computers.

EPCC was established in 1990 as a focus for the University of Edinburgh's work in high performance computing. Today the centre serves a large number of scientific groups in the UK. The mission is: "To accelerate the effective exploitation of novel computing throughout industry, academia and commerce". One of the key features of EPCC is the linkage between academic science and industrial projects. During the last five years the collaboration with industry has increased considerably.

EPCC is managed by a Director who is also a Professor, a Commercial Director, a Technical Director, and a Deputy Director. EPCC has an Advisory Board. Very large projects also have a Project Board. All projects have allocated a Project Investigator (PI) and a Project Manager (PM).

Besides its work with industry, EPCC offers advanced services to scientists and the HPC community. They hosted a visitor programme called TRACS for a decade. This programme allowed more than 400 visitors to use their HPC systems thus enabling them to carry out their research.

EPCC also offers training courses on a variety of subjects, and has over the years contributed to standard efforts, such as MPI. In 2001 EPCC started an MSc in HPC which has proven popular and successful. EPCC staff is regularly represented at national and international conferences and events.

Their team of technology experts is trained and skilled in just about every area of technical computing; code profiling, code optimisation, data mining, image processing, data encryption and compression, and Grid and web services, just to name a few. Alison Kennedy (Director) stressed that a science support team is important to ensure that the researchers get the most out of the clock cycles. EPCC has a group who are preparing the scientists for efficient use of the large computers. EPCC spends quite some time on testing new solutions both with respect to hardware and software. An advantage for the central support team is that they in all likelihood have been faced with a given problem before, and then it becomes easier to solve the problem in hand.

EPCC often contacts the top users and ask how their code runs - in order to ensure efficient use of the supercomputers.

Very often a fraction of a science application that includes a request for HPC, is devoted to those dedicated people in the support team who will ensure efficient computing.

The excellent capabilities within scientific computing imply EPCC's participation in a number of the EU's FP programmes on High Performance Computing. Actually, a large share of the financial support is obtained through EU's FP programmes.

Alison Kennedy (Director) said that a lot of money is still spent on grid computing - but it turns out that a HPC consolidation is more successful. Grid is interesting but not for real computing. Kennedy also mentioned that they collaborated with UNI-C until about 8 years ago, but since then they have not had any collaboration with DK.

In order to ensure a minimum use of energy and depiction of the total cost, the users pay for the power consumption as a part of the grant.

An applicant must compile a technical assessment form. Then EPCC investigates whether the machine and software represent an optimal solution for the current problem. This procedure is also used for Nobel Prize groups.

The users may apply at any time, and the Research Council decides the number of CPU hours, data space, etc.

C.2 NL: SARA in Amsterdam

SARA Computing and Networking Services is a non-profit organisation – but also a national computing centre.

They work hard to be recognized as a national centre. SARA's IBM Power 575 facility is No. 73 on the TOP100 list. Holland has 3 facilities on the TOP100 list as of June 2008.

For historical reasons, SARA is very closely connected to other computing infrastructure resources in the Netherlands, e.g., the different networks, both nationally (SURFnet) and internationally (Amsterdam is the major European link node).

SARA's business model has 3 legs:

- HPC & Virtualisation
- High Performance Networking
- Commercial ICT

The last leg is a service for, e.g., industry, and profit from this part helps to finance the other parts of the organisation.

The High Performance Computing & Virtualisation (HPC&V) department at SARA has expertise in various fields. Such expertise is used on a daily basis in the managing of the supercomputer systems and the indispensable support of those systems. The HPC&V department has expertise in the fields of optimisation and parallelisation of computer codes. Using this expertise, the hardware can be efficiently used to shorten the turn-around time.

The expertise can be used for optimisation and parallelisation of programmes that are installed on SARA's systems, but also in connection with programmes that are to be run on one's own systems. The facilities used include: compilers, subroutine libraries, and the large experience available at SARA in optimisation and parallelisation in connection with various kinds of hardware.

For more than 30 years SARA has experience in providing expertise in High Performance Networking (HPN). Designing, installing and managing high performance networks, to which the newest technologies are applied, are among SARA's daily activities. The close proximity to the national and international network backbone is very advantageous to SARA.

For more information on SARA, please see http://www.sara.nl/index_eng.html.

During the visit it was mentioned that the following topics were important when building up and running a HPC facility:

- Hints for building up a new centre:
  - budget: 50% hardware, 50% for everything else!
  - Vendor: good relationship with vendor is important - commitment from both sides! Clear contracts!
  - Datacentre planning: An important step! Get help from the experts.
  - Infrastructure: Close to the fast network backbone. Backup is very important!

- Success criteria:
  - Usability
  - Uptime
  - Capacity (theoretical vs. used)
  - Scheduling - it's hard to move jobs around

C.3 D: CCC at RWTH Aachen

The Center for Computing and Communication (CCC) is both a university computing center and a large installation in the federal state of Nordrhein-Westfalen (NRW). As of June 2008, the largest installation at CCC is No. 100 on the TOP100 list.

As a part of the German university infrastructure, you have standard tasks, like identity management, software licensing and distribution (done by a software portal), backup services, etc. A large part of the installation is dedicated to High-Performance Computing, though there is no real distinction between the activities when it comes to compute time. Access to the resources is granted through the participating institutions, i.e., typically other universities or research institutions in NRW, that hold a kind of share in the system (see also hosting below).

Since 2001 the CCC has had a large installation of big SMP machines, e.g., Sun Fire 6800/E6900 and 15K/E25K, but most of those boxes were recently replaced by smaller 4-core to 8-core boxes, plus a cluster of 20 Sun Fire T5120 machines (8 cores/64 threads). For more information, please see http://www.rz.rwth-aachen.de. In addition to the 'classical' UNIX based HPC installations, CCC also provides a Windows based cluster for HPC activities.

Regarding advanced user support, CCC has a very active group of people of which Dieter an Mey is in charge. This group helps users with porting and performance issues, but is also very active in trying out new technologies. Furthermore, the group participates in the OpenMP language committee work, especially on issues with C++ and memory placement.

As a new part of the business model, the CCC does 'integrative hosting', i.e., other universities/institutions may place their equipment at the CCC datacenter, and CCC takes care of everything (installation, maintenance, power, cooling, licensing, etc). As a reward, the hosted systems are fully integrated into the CCC compute cluster. In that way, the number of CPUs available to CCC users grows with every new hosted system.

Access to those 'hosted resources' is controlled by the batch system thus making sure that the resource owners can get the compute time they are entitled to. The difficulty the CCC faces right now, is to find a way to move running jobs from one machine to another, to make requested resources available, and to minimize waiting time.

C.4 FI: CSC in Helsinki

CSC

CSC was established in 1971 and since then Finnish researchers have had access to the national computing services. From the very beginning experts of various disciplines were hired for user support. When CSC was established it copied many ideas from UNI-C. Later on the activities expanded to also cover telecommunication networks. Nokia's Headquarters is located about one kilometer from CSC's main building.

Today CSC is the Finnish IT centre for science and provides versatile computing, communication and data storage services as well as scientific expertise to the entire Finnish research community.

CSC's mission statement is: 'CSC, as part of the Finnish national research structure, develops and offers high quality information technology services'. According to Director Kimmo Koski, the objective of CSC is to be one of the leading European centres for supercomputing by 2012.

CSC provides solutions for data storage, management, data mining and analysis for a long list of scientific disciplines, such as bioscience, chemistry, linguistics - in general the list covers both science and culture. They maintain huge visual archives.

About 60-65 pct. of the total budget comes from the Ministry of Education (and Research). A big portion of the funding comes from EU activities. Today the share of contractual projects with the industry is growing dramatically. The industrial partners prefer to place the equipment at CSC because the infrastructure and support are in place there.

In Finland the M-grid was established in 2004. Grid is the access to the network, but the grid does not compute anything. It is very important to have a central facility; it is a huge advantage to be able to level up and down the shares on the system. 'Today even physicists call on CSC for support; their expertise is physics - and not computing', Kimmo Koski says.

According to Kimmo Koski, a few physicists want to run their own system. One of the main reasons why a few Finnish groups have their own local installation is to avoid peer review, to avoid experts to keep an eye on the optimal or non-optimal use, and to avoid benchmarking. In Finland only a single university, namely the University of Åbo, is interested in local HPC installations.

The centralised setup offers better deals on software licenses for scientific computing. CSC manages the national licenses for the most popular software. The budget is around 1 mill Euro. They focus more and more on open source, and they contribute intensively to the development of open source software for the academia. A very popular example is the open source software for Finite Element Multi Physics called Elmer (see http://www.csc.fi/elmer), which is now used worldwide. Other examples are: CHIPSTER – micro-array data analysis, SOMA2 – drug development, GPAW – DFT electronic structure code. In general, in-house development can replace very expensive software.

Nano-science occupies 38 pct. of the CPU power at CSC and, basically, there is nothing at the individual universities. According to Director Janne Ignatius, it is very cost effective to have a few large systems, as compared to a lot of small systems. However, both specialised and general purpose machines are needed.

The fact that CSC is on the TOP500 presents a ticket to international collaboration and funding. Today CSC is No. 70 on the TOP500 list.

According to Director Janne Ignatius, those responsible for the resource allocation do NOT come from the university world, since it is important that people are not able to allocate money to themselves. The large allocations are based on international reviews, and the process is transparent; then it is easier to predict the chance for success. The largest projects are published on the web.

CSC offers about 70-80 courses on scientific and high performance computing per year. These courses are complementary to the courses offered at the universities. A large group of people at CSC is collaborating with research groups in code optimisation, porting and benchmarking.

Today about 160 persons are employed at CSC. The turnover is about 15.6 million Euro (exclusive of hardware, but including electricity). The cost of hardware is about 10 million Euro for each 3-year period. In general, the share between hardware cost and manpower is fifty-fifty.

C.5 D: LRZ in Munich

Leibniz Supercomputing Centre (LRZ) under the Bavarian Academy of Sciences and Humanities is the main High Performance Computer in Bavaria. The origins of LRZ go back to 1962, where the 'Electronic Computing Committee' was founded. LRZ provides services to the scientific and academic communities in Munich and to the Bavarian Academy of Sciences and Humanities (BAdW). Together with similar centres in Germany (HLRS (Stuttgart) and NIC (Jülich)) LRZ also acts as a technical and scientific high performance supercomputing centre for all German universities. LRZ also supports the universities in Munich with IT infrastructure.

The HPC centres in Germany also provide resources to industry like Porsche, Daimler Chrysler, and T-Systems. They organize workshops like the recently held at HLRS on 'Supercomputing in Science and Industry'.

The system today is a 62.3 Tflops SGI Altix system with 39 TB memory, and it is located as entry number 27 on the TOP100. This system was opened in the spring of 2007. The system is optimised for high application performance and high memory bandwidth. The storage is about 2500 Tbyte.

The users (about 100.000) come from 5 universities, 30 institutions, and 700 institutes.

As of June 2008, 46 HPC facilities in Germany appear on the TOP500 list.

LRZ recognised the significance of networks early on and has developed a nationwide know-how centre in this field. For instance, LRZ was the first institution in Germany to pilot many new communication technologies, such as the first ATM 34 Mbs-line and DQDB. LRZ also piloted the world's first 2.5 Gbs ATM line using gigabit technology with Wave Division Multiplexing. Many of these new technologies are developed in a collaboration with industrial partners.

LRZ is integrated in research at the universities in the field of High Performance Computing, such as efficient Linux clusters, architectures for high performance computing, tools for parallel computing, numeric programming on supercomputers, etc.

In collaboration with the Bayerische Staatsbibliothek, LRZ is presently working on the design of a library archiving and document supply system. This research will also lead to the development of long-term archiving procedures for many sciences.

According to Dr. Matthias Brehm (Group Leader, High Performance Computing), user support is essential in order to ensure an efficient use of the system. Currently the LRZ's user support group counts 10 persons. They are responsible for solving all kinds of user problems, including porting and tuning, and for installing and maintaining the software stack. They use a long list of tools for tuning and code profiling. In total, about 25 persons are involved in the HPC part of the activities at LRZ. They are all scientists with a HPC related background (University Master or PhD in Scientific or High Performance Computing).

Governmental funding means that the federal government in Berlin typically funds 50 pct. of large investments, while the remainder (including the running cost) is paid by the local state (Bavaria).

For the high end system at LRZ, users can apply at any time (continuous application process). Test accounts (quick access) are provided within two days. 20,000 CPU-hours are allotted to any reasonable application. The application itself to be by handled and reviewed by the steering committee within 6 weeks.

However, the procedure is going to change for the newly formed GAUSS Centre[3]: For (very) large-scale projects (>5 pct. of the available resources per year for the particular project), a situation of competition is established. GAUSS publishes a 'Call for Proposals'. Proposals for large-scale projects are accepted twice a year. After the review process a ranking for all (very) large-scale projects is performed. For normal projects (<5 pct.) the continuous application process will continue to be used.

C.6 NO: USIT in Oslo

UNIVERSITETET I OSLO

The meeting with IT Director Lars Oftedal took place in Reno at the Super Computing Conference (SC07) in November 2007. The meeting with Programme Director, Jostein Sundet (Research Computing Services, University of Oslo), was held at the University of Oslo (UiO) in January 2008. At that time the central facilities for High Performance Computing at UiO were also demonstrated.

In Norway the initiatives and planning of High Performance Computing, Notur, is lead by UNINETT-Sigma (here called Sigma). The four largest universities and Sigma are members of Notur. The board members of Sigma come from these universities and from the Research Council of Norway.

Sigma was established in 2005 as a new subsidiary of UNINETT under a contract with the Research Council of Norway. Sigma is leading the long-term project on high performance computing, Notur, national research storage facility, Norstore, and national grid activity, Norgrid. Notur has a ten year time frame. This is considerably longer than previous HPC projects in Norway, and this enables a much needed long-term planning of infrastructure and services.

Notur replaces the previous NOTUR-I project that has been active since 2000. Ever since 2005, a new national HPC infrastructure has been replacing some of the

[3] The GAUSS Centre for Supercomputing (GCS) was recently established. The main purpose is to trigger a closer collaboration among the three national supercomputing centres. A further objective of forming GAUSS is to enable Germany to take a leading role in European supercomputing. This is reflected in, e.g., significant German participation in the EC-funded project PRACE.

most outdated central installations with the aim of providing a modern, national HPC infrastructure in an international and competitive setting, and of stimulating computational science as the third scientific path.

The Notur project is a collaboration between Norwegian universities and the Meteorological Institute, and with major funding from the Research Council. Regional university colleges may also join through a national grid structure, and industry partners are invited to join. The operation of the HPC installations financed by Sigma is a national core activity in the project. User support at various levels, ranging from help-desk operations to advanced support and technology transfer, training and marketing, will also contribute to reaching the vision of the project.

In several areas, Norway has excellent scientific communities that have developed and combined mathematical models, algorithms and numerical methods into their specific branch of computational science. Construction engineering, geophysics, fluid dynamics and quantum chemistry are good examples. In other areas, computational science is a new leg within which scientific developments are taking place, such as in bioinformatics, medicine, economy, social sciences, legislative politics, linguistics, history, and archeology. The central organised collaboration ensures a synergy between scientific computing and the scientific communities using the HPC facilities.

According to Program Director Jostein Sundet, no individual research group at the university is allowed to buy their own hardware. Money for High Performance Computing facilities is allocated to a central facility (USIT).

It is the view of both Lars Oftedal and Jostein Sundet that Denmark seems to be missing an e-Science program. According to J. Sundet, reports or feedback from DCSC to the funding agencies seem to be missing. The most recent reports found on the DCSC homepage are from 2003 and 2004. Feedback to the decision makers is very important in order to adjust the national e-Science structure and composition.

The purpose of the e-Science program of Norway, EVITA, is primarily to establish a collaboration and synergy between the various scientific communities using HPC, and researchers focusing on algorithms, statistics, data mining, parallel computing, etc., and to spread the use of HPC to areas like social science, culture, etc. An important, but rather new area is the structured handling of a large amount of multimedia data (video, sound, interviews).

Research programmes combining several research areas are the basis for obtaining the synergy within EVITA. Even in applications the three elements, algorithms, scientific application, and IT must be addressed.

EVITA owns the research related part of the infrastructure. The other part of the infrastructure provides services to, e.g., forecast resources at the Norwegian Met-office.

Applications for larger amounts (more than 750.000 CPU hours) for compute cycles are reviewed by an international scientific board. However, according to J. Sundet, it is often rather difficult to compare applications within the various scientific areas. Also, political input is important. Medium applications (more than 100.000 CPU hours) are reviewed twice a year, whereas smaller applications are handled continuously. Research groups with a large amount of resources must write an annual report documenting the obtained scientific results, including the number of peer reviewed publications.

According to J. Sundet, it is often difficult to find the true need for high performance computing of a research group, and the local interest can be in conflict with national interests. A primary goal of EVITA is to discuss, consider and handle this conflict of interests.

The four universities all have their own research profiles:

- Bergen: Methods, algorithms, and informatics.
- NTNU: Mostly related to technical aspects; a major hard core computer group is found at NTNU.
- Tromsø: Applications – Computational chemistry is the most important area related to HPC.
- Oslo: Applications in a rather broad sense. Also some non classical HPC applications are considered, such as legislative politics (election processes in the EU parliament), Archeology (optimal conservation policies), Psychology (reconstruction of 3-d brains).

An important task for the Research Computing Services at UiO is to ensure that the computers are used efficiently. Researchers often need support related to scientific computing, choice of libraries, algorithms, hardware, etc., in order to ensure that the investment (hardware, software, manpower, etc.) is utilised efficiently. Also, the consideration to green computing implies that utilisation and efficiency are of importance. Another important task is continuously to keep track of the choice of hardware for a given code. Very often the service group at UiO assists in porting an application to a more optimal hardware platform.

At the University of Oslo 15 people are associated with the HPC support team. 9 of them have a PhD degree.

Licenses are coordinated among the four sites in Norway. The researcher asking for the specific software is charged with the cost of using it. The national coordination significantly reduces the total cost for software in Norway. It is estimated, however, that more than 80 pct. of the codes are either open source codes or codes written by the individual research group.

A focus on the theoretical peak performance is less interesting than a focus on the TOP500, since the performance measure applied in connection with that list is the actually measured performance.

D Visits at a number of HPC-users in DK

This section contains a condensed summary of a number of meetings held during the spring 2008 between the core project group consisting of Jostein Sundet, Lise Frohn, Bernd Dammann, and Henrik Madsen, and a number of HPC-users in Denmark.

At all visits the core group started by expressing the background and the purpose of this pre-investigation project.

Prior to the meetings the following two documents were sent out:

- A pdf-presentation as presented to the Director of DCSC in December 2007 and to the board in March 2008. This presentation is included in Appendix G.3.
- A preliminary description of the proposed infrastructure for HPC in Denmark, also presented to the board in March 2008. This description is included in Appendix G.1.

D.1 Aalborg University

Professor Erik Lund, Department of Mechanical Engineering, Assistant Professor Ondrej Franek, Professor Gert Frølund Pedersen, Department of Electronic Systems.

(Due to a sudden problem GFP did not participate in the meeting but we have had several fruitful conversations with him, and he supports the comments and suggestions made by EL and OF).

OF: Would like to have access to support and help in connection with porting of codes and code tuning. He is considering a large system of Maxwell Equations. Memory is the main issue. Large multi-CPU systems with huge memory are, thus, the optimal solution for the HPC structure.

EL: The research is mainly performed using the in-house developed code called MUST for finite element modeling of composite structures, such as wind turbine blades. It consists of more than 150.000 lines of Fortran 95. The section collaborates closely with Vestas, Siemens, and LM Glasfiber. Typically, the program requires a lot of memory. They are interested in direct solvers, where the Jacobian can be reused (kept in memory) – and, therefore, access to SMP would be beneficial. Today no courses on parallel computing are offered in Aalborg on a regular basis. They support the idea of having a national series of courses on parallel, scientific and high performance computing. Aalborg University will be glad to participate in setting up such courses, and in defining the optimal setup for a national infrastructure for HPC.

They all agree with the ideas as presented in our presentation.

EL and OF mentioned a list of research groups in Aalborg. These groups do not use HPC today, but would, no doubt, like to have access to a central HPC facility:

- Production Department
- Department of Energy
- Chemistry
- Acoustics

D.2 University of Aarhus: Department of Economics and Management

Martin Møller Andreasen, Jonas Staghøj, Bent Jesper Christensen, Valeri Voev, Henning Bunzel, Jie Zhu.

They have only recently begun to consider the use of HPC techniques. Martin M. Andreasen and Bent J. Christensen have since the startup in September 2007 succesfully established parallel programs for financial modeling. They support the idea of having a national series of courses (introductory as well as more advanced ones) on parallel, scientific and high performance computing. Such courses could also be offered locally. The compute servers could be hosted at any place in Denmark. They would like to have evaluations of applications for computational resources more often (not only once a year). Courses on Fortran and C++ are on a wish list.

Some of the researchers (e.g., Valeri Voev) do have a huge amount of high frequency data. Therefore, access to large memory systems would be attractive. In general, the researchers need help with the writing and optimisation of codes. Tools and some good advice for efficient cross compiling from high level languages like Matlab or R to Fortran or C are also on the wish list. Currently it is difficult to keep up with international standards (similar research groups worldwide) due to the limited access to HPC facilities.

D.3 University of Aarhus: Department of Chemistry

Professor Frank Jensen, Professor Jeppe Olsen, Assoc. Prof. Svend J.K. Jensen, Professor Poul Jørgensen, Assoc. Prof. Ove Christiansen.

After a short introduction JO left the meeting due to his membership of the board of DCSC.

PJ: He did not see much of a difference between our proposal and the new ideas of DCSC in connection with the Center for Advanced User Support (CAUS). He would like to have a large SMP computer placed in Aarhus and is worried that the distance between a central support facility and the individual user groups will be too long, as was experienced previously in connection with UNI-C.

OC: Often he considers rather complex codes. He would like to see changes as to how the resources are allocated. He likes the idea of having a more frequent possibility for getting CPU resources.

SJKJ: A specialist on Gaussian. A huge amount of memory is often needed for Gaussian. The expertise within computations comes primarily from chemists using Gaussian all around the world. This means that it will be too expensive to allocate resources for code optimisation instead of hardware.

FJ: Works with complex simulations of large molecules. He supports the new initiative by CAUS since he finds a need for: 1) code optimisation 2) optimised use of libraries, 3) support related to specific software, 4) user support, and 5) support in selecting the optimal algorithm. He believes that multi-core computers will become the future of high performance computing. The optimal solution is two centers for High Performance Computing in Denmark. People capable of identifying the need for support and for HPC is completely lacking today.

D.4 University of Southern Denmark (SDU)

Professor Ole G. Mouritsen, Prof. Julian Shillcock, Research Assist. Professor Himanshu Khandelia.

MEMPHYS - Center for Biomembrane Physics.

The aim of the center is to conduct parallel experimental and theoretical research into biological and model membranes on the molecular level, the relationships between the physical and physico-chemical properties of membranes and their functional behavior.

JS: Particle based simulations using C++. Previously used Monte Carlo and Molecular Dynamics methods; today he uses Dissipative Particle Dynamics and Brownian dynamics to reach the length and time scales relevant to biophysical processes. He often observes limitations in the memory on the computers used today, and this requires a move to a parallel codebase. Wants to have more experience with parallel computing and the optimisation of such codes, and is interested in participating in courses. A typical job may run for more than 7 days on a single CPU. The possibilities are most often limited by the memory on the single processor. He also thinks that the primary condition for entering a computing cluster is that the group owns a certain share which is always available to the group, and that the group has access to the remainder of the cluster when it is possible.

OGM: Simulations of particle systems. UNI-C provided very useful help in the 80ties, both with libraries and efficient compilers, but in the late 90ties UNI-C was not helpful at all. The new initiative must not end like UNI-C. Because of the interchange of experiments it is very important that the computation power available is predictable. In general, it is very important that MEMPHYS does have a fixed share of the computational resources.

HK: Liquid dynamics, and complex biological systems. A typically CPU time for a single job is 8 weeks today. Code optimisation is not an issue at all, even for new users. Local training can improve the efficiency with which new users implement existing codes to address their scientific problems.

D.5 Technical University of Denmark (DTU): MEK, SPACE

DTU

Professor Ole Sigmund (MEK), Assoc. Prof. Allan Roulund Gersborg (MEK), Head of Astrophysics Allan Hornstrup (SPACE)

OS briefly described the research at DTU MEK, where at least 3 research groups need access to HPC facilities. They have found it rather difficult to get access and reasonable response from DCSC. However, people at IMM have helped occasionally, but according to OS it is not reasonable that this effort is not recognised. The relevant people are more or less helping them after ordinary working hours. They see a need for experts on scientific and high performance computing that can help the research group at MEK with those aspects. Presently, the situation is simply not reasonable. The people at MEK spend too much time on how to solve large systems of equations, and how to pick up the best algorithm and, hence, they spend too much time on issues that are not within their expertise. MEK wants access to a highly qualified support group.

AH is mostly dealing with neural network, ray tracing and parallel applications. He would like to see one large central facility with competences related to various aspects of scientific and high performance computing.

D.6 Technical University of Denmark (DTU) MAT, MEK, Risø

Professor Martin Bendsøe (MAT), Professor Jens Nørkær Sørensen (MEK), Ass. Professor Dalibor Cavar (MEK), Senior Scientist Anders Nielsen (Risø)

MB explained that DTU MAT is a mixed setup of users, who are focusing on new mathematical methods implemented on a computer. They need access to one central facility, and this facility should not be located at MAT. They see an immense need for help to implementation and porting of advanced codes. At MAT they want access to a highly educated group of experts related to scientific computing who can help with various problems. They want a more flexible way of getting resources than provided by DCSC today. MAT wants also more focus on software – not only hardware.

JNS and his group are mostly working with computational fluid dynamics. He expressed that they have had great advantage of DCSC. However, they would like to have one larger computer. He questions whether all numerical problems are solved, but he agrees that expertise on scientific computing and expertise as far as the actual problem is concerned, must be separated. JNS sees a need for central organisation of courses on scientific and high performance computing. MEK wants access to a much larger facility, and they will support a large facility in Denmark with a close connection to the experts on scientific computing, preferably at DTU. JNS sees a need for more flexibility in getting resources from DCSC.

AN and his group are working with fusion energy. AN often uses a computer in Germany. Risø wants to collaborate with specific groups how to solve problems related to algorithms, etc. He sees a huge need for more available expertise on scientific and high performance computing. They also see a need for support, which is not available today.

D.7 Technical University of Denmark (DTU) CBS, MEK (ETH-Zurich), Chemistry

Professor Søren Brunak (CBS), Head of System Admin Kristoffer Rapacki (CBS), Associate Prof. Hans Henrik Stærfeldt (CBS), Assoc. Prof., Dr. Jens H. Walther (MEK and ETH-Zurich), Assoc. Prof. Irene Shim (Chemistry), Assoc. Prof. Flemming Yssing Hansen (Chemistry), Assoc. Prof. Günther H. J. Peters (Chemistry)

SB (CBS) started by explaining the research at CBS, and focussed on the diversity in the large number of software and computer tools used at CBS – several hundreds of programs are used. At CBS they focus on shared memory computers with a large amount of memory. Data is typically loaded directly into memory in order to speed up the calculations. HHS is an expert in data warehouse solutions, which also require large areas for data storage. He is also leading the initiatives to find optimal algorithms. SB explained that CBS has a large amount of competence in-house. However, more training in efficient computing and parallel computing is needed. DCSC has been a big advantage to CBS, but SB agreed that if today support were available for both the decentralized solution in combination with a large central facility, this would be beneficial. He also mentioned that the new courses on High Performance Computing at DTU serves as a very much needed source for new experts.

JHW (MEK and ETH-Zurich) has a longtime time experience in using large HPC facilities; mostly at ETH-Zurich. He works on simulating fluids in complex situations by the use of particles, and he has an expertise in porting codes to large HPC systems. His HPC requirements are small single CPU computers for development and testing, and massively parallel computers with a large distributed memory and thousands of compute cores for production simulations.

In DK JHW finds it very difficult to identify how to get access to large HPC facilities, and that the existing HPC facilities do not offer the necessary size and reliability. He believes these problems are associated with the strongly decentralized structure of HPC in DK. He, therefore, still relies on the Swiss HPC facilities. He finds that DK should host one large national HPC center, with top of the line shared and distribute memory machines. He advocates for an agile, slim HPC centre with a limited core staff for system administration, including user support/guidance in porting and optimisation. Moreover, the hardware architectures are currently undergoing a radical change: from few core (single, dual, quad) to many core systems. It is imperative that research is conducted in this area of scientific and high performance computing in order to secure proper utilisation of these systems in the future. One possibility would be to associate a PhD programme with the national HPC center, or to promote this research through a strategic research programme.

FYH (Chemistry) finds that a centralisation is needed, and that a centre should assist researchers in porting of codes, benchmarking and code optimisation. He agrees with many of JHW's comments. IS is an expert on Gaussian, and it is critical that this software runs efficiently. This application demands large memory computers. Chemistry in general is very interested in courses and assistance related to parallel and high performance computing. The support should be available promptly during working hours.

D.8 University of Copenhagen (KU)

KØBENHAVNS UNIVERSITET

Head of Department, Prof. John Renner Hansen (Niels Bohr Institute), Assoc. Prof., Klaus Mosegaard (Niels Bohr Institute / Theoretical Geophysics), prof. Kurt Mikkelsen (Chemistry / Theoretical Chemistry), Prof Åke Nordlund (Niels Bohr Institute / Astrophysics)

JRH stressed that the resources for HPC must be allocated based on scientific capabilities, and he finds that computational services for the goverment should be totally separated from scientific computing. He finds that some support structure is needed for new researchers, and he finds that PhD schools are the most appropriate instrument for building up expertise on high performance computing. However, he stressed that DCSC should only provide compute cycles and that the financing of the courses should lie within the universities, only. He argued that Denmark should focus on the activities related to Nordic eSciences.

With a reference to www.arcade.eu.org, K. Mikkelsen argued that DCSC has succeeded in getting DK in a leading position within EU in relation to the number of Mflops available for the researchers. He agrees with JRH that support should be based on scientific records. He finds that it is important to have an expert in high performance computing directly related to (and located close to) the departments.

ÅN pointed out that the resources provided by DCSC to participating groups are larger than many of the allocations given to similar groups in the US, that (as also pointed out by K. Mikkelsen) when Arcade*EU (www.arcade-eu.org) measured resources in flops per capita in 2004, Denmark was at the top in Europe, and that it is doubtful if these statistics have been adequately updated recently. In particular, since the Danish resources have followed or exceeded Moore's Law since then, Denmark cannot have fallen significantly behind. ÅN prefers support infrastructure to be mainly (geographically) local, with occasional coordination of lectures and workshops on parallel computing at the national level.

K. Mosegaard stressed that the interaction between the computer and the researcher is important, and he has experienced useful collaboration with external experts on algorithms. He finds that DK is lacking people with an expertise in High Performance Computing and algorithms. He also finds that experts on computing should be used for code optimisation; he finds that within the optimal HPC infrastructure he would be able to focus on computational geophysics while other experts could focus on or assist in code tuning and optimisation of the use of algorithms. K. Mosegaard stressed that a PhD or a Postdoc should not spend the first year learning how to operate the HPC facilities. Such knowledge should be available through, e.g., courses and support functions.

E References

1. PRACE — Partnership for Advanced Computing in Europe, http://www.prace-project.eu/

2. Enabling Grids for E-sciencE (EGEE), http://www.eu-egee.org/

3. U. Suter: *The Swiss National Strategic Plan for High Performance Computing and Networking*, Presentation at CSCS User Day, Sept. 2008.

4. D. Patterson and S. Williams: *The Parallel Computing Landscape: The Berkely View*, SC2007, Reno, Nov. 2007.

5. K. Asanoc, R. Bodik, B.C. Catanzaro, J.J. Gebis, P. Husbands, K. Keutzer, D.A. Patterson, W.L. Plishker, J. Shalf, S.W. Williams, K.A. Yelick: *The Landscape of Parallel Computing Research: A View from Berkely*, Technical Report No. UCB/EECS-2006-183, 2006.

6. *Accelerating Innovation for Competitive Advantage: The Need for Better HPC Application Software Solutions*, The Council on Competitiveness, Washington, July 2005.

7. Høringsuttalelse fra Universitetet i Oslo vedrørende forslag til 'Nasjonal plan for koordinering og utvikling av eVitenskap og anvendelser af eVitenskab i Norge.

8. C. Inglis: *HPC-Europa: Building Bridges in European Computational Science*, Capability Computing, Issue 10, Autumn 2007.

9. A. Turner: *Research Computing Facility: HPC for all*, Newsletter of EPPC – Europe's premier high-performance computing and technology transfer center, Issue 61, Autumn 2007.

10. K. Koski: *Competitive Services and Infrastructure – Keys to Success*, CSC Annual Report 2006.

11. J. Åström: *Code optimization–taming the computational beasts*, CSCnews 2/2006

12. T. Sterling: *Multi-Core for HPC: Breakthrough or Breakdown?*. SC2006, Tampa, Nov. 2006.

13. Standard Performance Evaluation Corporation (SPEC), http://www.spec.org/index.html

14. *Memorandum of Understanding concerting the establishment of a European Tier 0 – High Performance Computing Service*, Confidential, April 2007.

15. *Benchmarking Industrial Use of High Performance Computing for Innovation*, The Council on Competitiveness, Washington, May 2008.

September 11, 2008

SHORT CURRICULUM VITAE for

Henrik Madsen, born 15 October 1955. Married. Two children.

- Employed as researcher at the Section of Mathematics at DIAB, February 1980.
- M.Sc. in Engineering (Civilingeniør) from (DTU) specializing (eksamensprojekt) in Statistics, January 1982.
- Selective military service, August 1984 - May 1985.
- Ph.D. (Lic. Techn.), January 1986.
- Ass. Prof. (Adjunkt) in mathematical statistics at IMSOR, February 1986 - September 1989.
- Received 'Direktør Peter Gorm-Petersens Mindelegat' at the annual celebration of the University as an acknowledgement of my research, May 1987.
- Assoc. Prof. (Lektor) in Mathematical Statistics at IMSOR, September 1989 - 1999.
- External lecture at the University of Copenhagen, Institute of Matematical Statistics (KUIMS), in Time Series Analysis, September 1989 - February 1990.
- Head of the G-bar computer system for students and researchers at DTU. Today the system has 16.000 users. 1991 - .
- Invited external lecture (ECMI) at the Fourier University, Grenoble, 1994.
- Received a prize (*60.000 dkr.*) 'Fabrikant Ulrik Brinch og hustru Marie Brinchs legat' as an acknowledgement of my contribution to the theory for and engineering applications of dynamical stochastic systems. March 1997.
- Professor in *Stochastic Dynamic Systems* in the Section of Mathematical Statistics, Inst. Math. Modelling, DTU. April 1999.
- Head of the Research Committee at Informatics and Math. Modelling, 2000-2005.
- Co-chair of the board for the Danish Society for Theoretical Statistics (DSTS), February 2001 - February 2004.
- Negotiated an agreement with Sun Microsystem, San Francisco, where IMM/DTU is declared a *Center of Excellence in Dynamic Systems and Interval Arithmetic*, June 2001.
- Negotiation of contract and installation of a large HPC facility at DTU. The nominal value was about 100 mill. dkr. During 2001.
- Head of *Center for High Performance Computing* at DTU which was opened by the *Danish Minister of Research* Helge Sander, February 2002 - .
- Major DCSC grant holder, January 2002 - May 2007.
- Member of the Advisory Board for DCSC, May 2003 - 2007.
- Member of Advisory Board for MathWorks (Matlab), May 2006 - .

PUBLICATIONS

More than 305 publications including 118 papers in journals, 11 research monograph, 5 books (latest books: *H. Madsen, Time Series Analysis, Chapman and Hall, 392 pp., 2007* and *P. Thyregod and H. Madsen, General and Generalized Linear Models, Chapman and Hall, 280 pp. - will be published in 2008/9).*

PH.D. SUPERVISION

Supervision of 25 completed PhD projects. Currently supervising 11 PhD students.

FURTHER PROFESSIONAL ASSIGNMENTS

I am or have been the leader of – or participated in – a number of research projects financed by EU, INTERREG-EU, Public Service Obligation (PSO), the Danish Council for Strategic Research (DSF), Danish Agricultural and Veterinary Research Counsil (SJVF), EU, NATO, Ministry of Industry (Erhvervsfremmestyrelsen), Nordic Council (Nordisk Ministerråd), Ministry of Environment and Energy (Energi- og Miljøministeriet), Danish Research Academy (Forskerakademiet), Danish Academy for Technical Sciences (ATV) and a large number of private companies.

The total budget for the research projects I am involved in today is more than 120 mill. dkr. with partners from Denmark, a number of European countries and US. I am or have been referee at a large number of professional journals.

September 28, 2008

SHORT CURRICULUM VITAE for

Lise Marie Frohn, born 19 July 1969. Married. Two children.

- Employed as scientific assistant at the National Environmental Research Institute, June 1995.
- M.Sc. in Physics from the Niels Bohr Institute specializing in experimental high energy physics, August 1999.
- Lecturer at Baltic Universities for the Nordic Research Academy, 2001-2002.
- Ph.D. in physics (air pollution modelling), August 2003.
- Employed as research scientist at the National Environmental Research Institute, September 2003.
- Co-convener of the EGU General Assembly, 2003-present
- Received EUROTRAC-2's Young Scientist Award for research contributed to the sub-programmes GLOREAM, CAPMAN and GENEMIS, October 2004.
- Lectures at the University of Tromsø(Norwegian Institute of Fishery Sciences, 2004.
- Co-convener of the GLOREAM workshop, 2004
- Lecturer for the Danish Agricultural Advisory Service, 2004 - present.
- Employed as senior research scientist at the National Environmental Research Institute, July 2007.

PUBLICATIONS

More than 120 publications including 31 papers in international journals with peer-review and 47 consultancy reports.

FURTHER PROFESSIONAL ASSIGNMENTS

I am or have been the leader of – or participated in – a number of research projects financed by EU, the Public Service Obligation (PSO), the Danish Council for Strategic Research (DSF), Nordic Council of Ministers, (Nordisk Ministerråd), Ministry of Environment and Energy (Energi- og Miljøministeriet), Danish Research Academy (Forskerakademiet), the Danish Wind Mill Association (Vindmøllefonden), the Danish Agency for Science, Technology and Innovation, the Danish Environmental Protection Agency, the Danish Research Centre for Organic Farming, the EU commission and a number of private companies.

I am or have been referee and guest editor at a number of professional journals.

September 25, 2008

SHORT CURRICULUM VITAE for

Bernd Dammann, born 8 July 1966. Not married. One child.

PROFESSIONAL ASSIGNMENTS

- M.Sc. in Physics (Diplom-Physiker) from University of Mainz, Germany (1992). Topic: Statistical and Computational Physics.
- PhD studies at the Dept. of Chemistry, Technical University of Denmark (1992-1996).
  Topic: Model studies of random surfaces, interfaces, and microemulsions.
- Research Assistant and system administrator at Dept. of Chemistry, Technical University of Denmark (1996-2001).
- System Consultant at Informatics and Mathematical Modelling, Technical University of Denmark (2001-2005)
- Assoc. Professor in 'High-Performance Computing', at Informatics and Mathematical Modelling, Technical University of Denmark (since 2005)

PUBLICATIONS

13 publications in international journals and conference proceedings.

FURTHER PROFESSIONAL ASSIGNMENTS

I am or have been in charge of different support units, e.g. the HPC and Unix support groups at DTU. My experience in the field of HPC is often used in consultancy projects by other DTU departments or private companies.

September 30, 2008

SHORT CURRICULUM VITAE for

Jostein Kandal Sundet. Born July 6th, 1965. Married. Two children.

- Military service, personal assistent for head of Navy, Rear Admiral Grimstvedt, 1984 - 1985
- Cand. Scient in dynamical meteorology, November 1993
- Ph.D. (Dr. Scient.), December 1997
- Post. Doc., Norwegian Research council, 1998 - 2001
- Visiting researcher, University of California, Irvine, 1998 - 1999
- Adj. Assoc. Prof. (1. amanuensis II) in meteorology, atmospheric chemistry and climate 2001 - 2004
- Programme Director, Research computing service, University of Oslo, 2004 - dd

PUBLICATIONS
Author and co-author for 43 reviewed publications

Infrastruktur for High Performance and Scientific Computing i Danmark

Jostein Sundet, Program Director, Research Computing Services, USIT, University of Oslo
Lise Frohn, DMU/Aarhus Universitet, Henrik Madsen, DTU, Bernd Dammann, DTU

Danmarks Miljøundersøgelser og Danmarks Tekniske Universitet har fået bevilliget et projekt fra infrastrukturmidlerne med det formål at undersøge behovet for at udvide tilbuddet af High Performance Computing i Danmark.

Projektet udspringer af to parallelle ansøgninger - en fra DMU og en fra DTU - hver med det formål at opnå en større central facilitet, gerne som en del af den eksisterende HPC infrastruktur, samt at tilbyde kurser og (avanceret) brugerstøtte til algoritmer. I forprojektet er DMU og DTU sat sammen for at afdække behovet for en central facilitet i Danmark.

DMU og DTU har gennem de oprindelige ansøgninger givet udtryk for, at der er behov for et nyt center for High Performance and Scientific Computing. Situationen for HPC i Danmark er pt:
- Indtil 2001 eksisterede en enkelt central facilitet (UNI-C).
- I de sidste år frem til 2001 blev der i praksis ikke udbudt support og relevante træningskurser for brugere.
- I 2001 blev DCSC etableret, og i dag er der 5 centre med i alt 11 decentrale installationer
- DCSC finansierer kun hardware.
- Der er ansøgningsrunder en gang om året.
- I praksis er der ikke noget link mellem DCSC's brugere og forskning inden for algoritmeudvikling og scientific computing.

Blandt andet mener DMU og DTU, at der er behov for en opdatering af HPC infrastrukturen (maskiner, brugerstøtte, kurser) i Danmark således, at der også kommer fokus på:
- National adgang til et stort system, gerne som en del af det eksisterende HPC arbejde i Danmark.
- Nem adgang og bred tilgængelighed.
- Et stort og effektivt system med mange kerner, fx 10.000, tæt koblet (e.g via infiniband).
- Noder med stor memory (>100GB) for shared memory applikationer).
- Hurtig adgang til store filservere (fx. Parallelle filsystemer, GPFS, LUSTRE).
- At tilbyde verdensklasse regnekraft til det danske beslutningssystem, også med henblik på at understøtte politiske beslutninger.
- Effektiv strømudnyttelse (grøn strøm super computing) bla. med tanke på at systemet kan indvies i forbindelse med klimatopmødet i 2009.

Mere konkret mener vi, der er behov for:
- Solid brugerunderstøttelse og -oplæring.
- Mulighed for hjælp til portering og tuning af kode.
- Installation af de kommende massive high-memory multicore maskiner, herunder mulighed for hjælp til software udvikling i relation til denne type maskiner.
- Mere flexibel adgang til systemet, fx. gennem løbende ansøgninger, automatisk allokering af jobs på relevante systemer og mulighed for testkørsler uden forudgående peer-review af ansøgninger.

Specielt i forhold til brugersupport:
- HPC brugere er forskere, der er interesseret i de videnskabelige resultater, men som ikke

nødvendigvis har spidskompetencer inden for HPC.
- Computeren er et værktøj, ikke en del af forskningsresultaterne.
- Brugere vil ikke og bør ikke skulle spilde tid på detaljer (som fx compiler options eller portering af koder).

Vi forestiller os, at brugerstøtten bliver håndteret gennem en centralt drevet national facilitet, der kan håndtere simpel såvel som avanceret brugerstøtte hurtigt og effektivt. Den centrale brugerstøtte facilitet kan fx finansieres ved, at en vis procentdel af ansøgte midler automatisk allokeres til brugerstøtteenheden.

Opgaver der bør ligge under en support facilitet:
- Bruger introduktion til HPC systemet.
- Installation og vedligeholdelse af compilere og applicationer.
- Et batch system til fordeling af job.
- Hjælp til portering af kode.
- Hjælp til løsning af compilationsproblemer.
- Uddannelse og workshops.
- Online helpdesk.
- Analyse af job performance og hjælp til kodetuning, både i de situationer hvor brugeren efterspørger kodetuning og i de situationer, hvor en hurtig analyse af koden viser, at den ikke kører optimalt.

Generelt mener vi også, at der skal være fokus på:
- Samarbejde med førende forskergrupper.
- Konferencer og workshops med fokus på fx kommende teknologi og hvordan den bedst kan udnyttes.
- Internationalt samarbejde med det formål at holde den danske infrastruktur ajour, såvel hardware- og softwaremæssigt, som organisatorisk.

I forhold til DCSC har DMU og DTU følgende holdninger:
- Vi anerkender indsatsen fra DCSC på strukturering og organisering af HPC infrastrukturen i Danmark indtil nu.
- Vi vil fortsætte med at understøtte DCSC og finansieringen af DCSC.
- Vi håber, vi kan være med til at videreudvikle DCSC.
- Vi håber, vi kan understøtte DCSC brugere gennem fx diskussion af forskellige emner relateret til HPC inden for DCSC's rammer.
- Vi håber, at den nye centrale installation/infrastruktur, som vi efterspørger, kan blive etableret i tæt samarbejde med DCSC, måske endda inden for DCSC regi.

Planen for forprojektet er som følger:
- Afholdelse af møde med et stort antal internationale HPC centre - dette er primært gjort i forbindelse med SC2007 i Reno.
- Besøg hos store succesfulde Europæiske centre - der er aflagt besøg hos Norsk Regnesenter (Oslo), EPCC (Edinburgh), CSC (Espoo), SARA Computing and Networking Services (Amsterdam) og Center for Computing and Communication (CCC – RWTH Aachen).
- Tilknytning af ekstern HPC ekspert - det er Jostein Sundet, som er leder for videnskabelig databehandling på Norsk Regnesenter, Universitet i Oslo.

- Møder med DCSC brugergrupper og besøg hos danske forskergrupper med behov for HPC faciliteter - dette arbejde er i gang.

- Samarbejde i videst mulig udstrækning med DCSC.
- Udarbejdelse af rapport i tæt samarbejde med den eksterne ekspert.

Vore nuværende visioner for en fremtidig HPC infrastruktur i Danmark baseret på de interview og præsentationer, vi har gennemført indtil nu, kan konkretiseres som følger.

Der bør være fokus på:
- Planlægning af den langsigtede strategi.
- Helpdesk funktionalitet der muliggør hjælp til alle typer systemrelaterede problemer inden for relativt kort tid (fx ~24 timer).
- Allokering af ressourcer (regnetid).
- User support (kurser, kode-portering, profiling, optimering af kode etc).

I en fremtidig struktur kunne vi fx forestille os en central enhed (en slags National e-Science Board), der varetager opgaverne omkring planlægning af langsigtet strategi, samt ressourcefordeling og allokering af regnetid (for ansøgninger på mere end et vist antal CPU timer). Alle brugere rapporterer tilbage til den centrale enhed (opnåede videnskabelige meritter), som inddrager rapporteringsresultater og hardware udnyttelsesgrad i forbindelse med vurderingen af nye ansøgninger. Allokeringen sker efter brugerønske og efter en overvejelse af, hvor de enkelte grupper bedst kan udnytte den tilgængelige kapacitet og derved får de optimale forhold.

En særlig forpligtelse for den centrale enhed er at sikre, at systemerne udnyttes maksimalt og mest omkostningseffektivt med henblik på at sikre "grøn it" energimæssigt set.

Mindre ansøgninger allokeres lokalt på de decentrale installationer, og den kortsigtede planlægning og strategi for disse koordineres med den centrale enhed.

Vi ser det som særdeles vigtigt, at den nye installation er tilgængelig for alle typer brugere - også de der ikke traditionelt anvender HPC i deres forskning, men som kan have glæde af det, såfremt de får muligheden og bliver hjulpet i gang. Med dette for øje skal det være let at komme til systemet, specielt for nye brugergrupper og fagområder.

Derudover mener vi, at den nye HPC installation skal have en kapacitet, som dækker meget større jobs (målt i antallet af CPU'er) end de, der kan gennemføres med de tilgængelige ressourcer i dag.

Den nationale centrale facilitet består af en stor samlet facilitet med fokus på at imødekomme et generelt behov hos mange brugere, hvilket giver sig til udtryk i high-memory multi-core kombineret med behov for stor lagerkapacitet og høj I/O rate. Den nationale facilitet forankres i forhold til undervisning og forskning i scientific og high performance computing samt myndighedsbetjening og kunne fx have egen bestyrelse til varetagelse af ressource allokering for mindre ansøgninger. Denne ressourceallokering kan fx ske i en opdateret version af den allokeringsmekanisme, som i dag anvendes af DCSC.

DCSC
Dansk Center for Scientific Computing

Blegdamsvej 17
DK-2100 København Ø
Danmark

Tel +45 3532 5453
Fax +45 3532 5016
E-post info@dcsc.dk
Web www.dcsc.dk
CVR 29979812
EAN 5798000422490

Danmarks Miljøundersøgelser
Afdeling for Atmosfærisk Miljø (ATMI)
Att.: Lise Marie Frohn (lmf@dmu.dk)
Frederiksborgvej 399
Postboks 358
4000 Roskilde

Danmarks Tekniske Universitet
Institut for Informatik og Matematisk Modellering
Att.: Henrik Madsen (hm@imm.dtu.dk)
Richard Petersens Plads
Bygning 321, rum 019
2800 Lyngby

23. april 2008
/RB

Infrastruktur for *High Performance and Scientific Computing* i Danmark

Bestyrelsen for Dansk Center for Scientific Computing (DCSC) behandlede den 1. april 2008 det fremsendte notat "*Infrastruktur for High Performance and Scientific Computing i Danmark*" af 30. marts 2008.

I forlængelse af DCSC's tidligere udtalelse, i brev af 25. september 2007 om DTU-DMU-projektets sigte, takker DCSC for den opdaterede orientering samt for lejligheden til at kommentere projektgruppens fund og overvejelser. Selvom notatet i et vist omfang udbygger de antagelser og synspunktet som projektets rationale og sigte hviler på, finder DCSC dog ikke, at dette giver anledning til at ændre tidligere fremførte vurderinger.

DCSC mener således, at det fortsat står tilbage for projektet at godtgøre de behov, som blev fremført i henvendelserne af 11. september 2007 og 30. marts 2008. Med relevans for projektet kan det tilføjes at DCSC pt. behandler en række ansøgninger til DCSC's seneste opslag, *Centres for Advanced User Support (CAUS)*[1], som tager afsæt i DCSC's offentligt tilgængelige strategiske målsætninger[2].

Vi ser frem til de fortsatte drøftelser om den mest hensigtsmæssige Scientific Computing infrastruktur til dansk forskning. For at fremme processen inviteres DTU-DMU-projektets interessenter hermed til at søge, at demonstrere på DCSC's nuværende ressourcer, de behov for Scientific Computing arkitektur og kapacitet, som projektet mener findes og bør prioriteres. Vi forstår, at et sådant møde allerede er aftalt med DCSC/AU. Men tilbuddet gælder også DCSC's øvrige ressourcer, hvor de andre fire DCSC driftscentre gerne fortsætter en konkret drøftelse af, hvorledes de fremførte behov bedst demonstreres, så længe det sker indenfor rammerne af DCSC's målsætning og vedtægter.

Med venlig hilsen

Rene Belsø
Centerleder, DCSC

[1] CAUS opslag: http://www.dcsc.dk/grant_applications.html
[2] Ansøgning til Infrastrukturpuljen: www.dcsc.dk/news_current.html#2007-12-14

High Performance and Scientific Computing Infrastructure

Lise Frohn and Henrik Madsen

Danmarks Miljøundersøgelser

HPC in Denmark

- *Until 2001 a single central facility called UNI-C*
- *During the later years of UNI-C essentially no support and training were offered.*
- *In 2001 DCSC (Danish Center for Scientific Computing) were established.*
- *Today 11 installations.*
- *DCSC financing – only for hardware.*
- *Applications once a year*
- *Essentially no link to research on algorithms and scientific computing.*

Danmarks Miljøundersøgelser

Our existing HPC Center at DTU

- *first Sun Fire 6800 servers in summer 2001*
- *officially opened Feb 4 by Helge Sander, 2002*
- *Partly a donation from Sun Microsystems*
- *funding DCSC and DTU*
- *part of the DCSC installation at DTU*
- *Sun CoE for “Interval Arithmetic and Dynamic Systems”*

HPC support - Why?

- *HPC users: scientists who are interested in* ***scientific results***
- *Computer is a tool – not part of the research*
- *Users don't want to – should not – waste time on the details (like compiler options)*
- *Comparison: Experimental Physics*
  - *Large installation, eg. A synchrotron or a collider*
  - *Local scientist/technician who know the technical details.*

HPC support

- *Support tasks:*
  - *user introduction to the HPC system*
    - *compilers, applications*
    - *Batch system*
  - *porting problems*
  - *compiler problems (bug reporting, enhancements)*
  - *trainings & workshops (e.g. in collaboration with DCSC)*
  - *consulting*
  - *performance analysis & code tuning*

Code Tuning Results

- ***Our experience: Many codes (also codes applying for huge funding for hardware) can easily be tuned to run 2 to 10 times faster.***
- ***Tuning means that instead of buying computers for 5.000.000 dkr we could instead consider code tuning and obtain the same computational performance with computers for 1.000.000 dkr!***
- ***Tuning is the often a matter of 1-5 days.***

DTU and DMU's view

Danmarks Miljøundersøgelser

- We acknowledge the support and initiative from DCSC
- We will continue to support DCSC – and support further funding of DCSC.
- We hope to be able to contribute to a further development of DCSC – eg link to research and education on HPC.
- We hope that we can be able to support DCSC users and discuss various issues on scientific and high performance computing within the framework of DCSC.
- We hope that the new central infra-structure can be established in a close collaboration with DCSC – maybe even within the DCSC organization.

DTU-DMU's view

Danmarks Miljøundersøgelser

- We need a new center for High Performance and Scientific Computing with focus on:
  - National access to one large system
  - Easy access and high availability
  - Consolidation – implies a huge and efficient system (say 10.000 CPU nodes – tightly connected)
  - Many nodes with large memory tightly connected to the node.
  - Fast access to large file servers (eg. Parallel file servers)
  - Offer world class computational resources for the government – also for supporting political decisions.
  - Efficient power use (green power super computing) – announced at the Climate Summit 2009 in Copenhagen

DTU-DMU's view

- We need a new center for High Performance and Scientific Computing with focus on:
  - Code tuning, user support and training
  - Synergy between applications, research and lectures on scientific and high performance computing
  - Massive high memory and multi-core computing
  - Software development for the future massive multi-core computers
  - Collaboration with leading research groups
  - Conferences and workshops
  - International collaboration

Danmarks Miljøundersøgelser

Plan

- Meetings with a large number of Centers for High Performance Computing (eg at SC2007 in Reno).
- Meetings and visiting major and most successful European Centers.
- External expert allocated.
- Visiting some Danish research groups which calls for a high performance computing facility.
- Collaboration with DCSC.
- Report

Danmarks Miljøundersøgelser