eUROPEAN nETWORK for aDVANCED cOMPUTING tECHNOLOGY for sCIENCE

DISSEMINATION REPORT

Compiled by

Stavros C. Farantos (FORTH)

Institute of Electronic Structure and Laser
Foundation for Research and Technology, Hellas
And
Department of Chemistry, University of Crete

Table of Contents

1. INTRODUCTION

2. ENACTS- Dissemination

3. WHAT HAS BEEN ACCOMPLISHED

4. CONCLUSIONS FOR THE FUTURE

APPENDIX I

CECAM- ENACTS Workshop 2003

CECAM- ENACTS Workshop 2004

APPENDIX II

OVERVIEW OF ENACTS STUDIES

Report on Grid Service Requirements (pp. 326)

Report on High Performance Computing Development for the Next Decade, and its Implications for Molecular Modelling Applications (pp. 103)

Report on Grid Enabling Technologies (pp. 76)

Report on Data Management in HPC (pp. 146)

Report on Distance Learning and Support (pp. 56)

Report on Software Reusability and Efficiency (pp. 81)

Report on Grid Metacenter Demonstrator: Demonstrating a European Metacentre (pp. 41)

Report on Survey of Users' Needs (pp. 69)

APPENDIX III

Article on Distance Learning and Support by ICCC

APPENDIX IV

LIST OF ACRONYMS

1. INTRODUCTION

In this report we present the actions taken to disseminate the work which has been performed in the eight studies of ENACTS in the time interval of four years, 2001-2004. Furthermore, we collect in an appendix the aims and the conclusions of each separate study from their final reports. We believe that this effort is worth pursuing if we take into account that the reports from the eight investigations amount in a total of 898 pages.

Practically, in the ENACTS studies all important topics of Grid technology have been considered – computational Grids, data Grids and data management, collaborative Grids, Grid enabling technologies, software efficiency and reusability, distance learning and support, plus the construction of a demonstrator for a future European Metacentre. The current middleware and software, which deploy and support a Grid, are reviewed, and in the reports the reader can find the results and analysis of questionnaires addressed to leading hardware specialists and users from all computational sciences. These questionnaires state the present status of Grid computing and most importantly they reveal the future trends.

The purpose of this report is to highlight the main achievements of the work done in these four years and we shall try to foresee the future implications of the rapidly advancing Grid technology in computational sciences, specifically in molecular sciences. The importance of the latter stems in the emphasis that has been given in the twenty first century for studying complex physical systems from first principle, i.e. considering the atomistic composition of matter and the basic laws of physics.

For completeness the following Section remind us from where we started by referring to the main objectives of ENACTS project, and particularly to the dissemination tasks. In Section 3 we summarize the actions taken in the dissemination project and in the finale Section some general conclusions are drawn with the future implications of computational Grids in Computational Sciences. In Appendix I we tabulate the programs of the two CECAM-ENACTS workshops held in 2003 and 2004 respectively, in Lyon. In Appendix II we compile the objectives and the results of each study extracted from the sectoral report of each project. Appendix III contains the article on Distance Learning and Support written by ICCC, which will be presented in EDEN (European Distance and E-learning Network, http://www.eden-online.org/), June 20-23, 2005, Helsinki, Finland. Finally, Appendix IV lists the meaning of the acronyms widely used in the Grid literature.

2. ENACTS-Dissemination

ENACTS is a Co-operation Network in the "Improving Human Potential Access to Research Infrastructures" Programme.

This Infrastructure Co-operation Network brings together High Performance Computing (HPC) Large Scale Facilities (LSF) funded by the DGXII's IHP programme and key user groups. The aim is to evaluate future trends in the way that computational science will be performed and the pan-European implications. As part of the Network's remit, it will run a Round Table to monitor and advise the operation of the four IHP LSFs in this area.

This cooperation network follows on from the successful Framework IV Concerted Action (DIRECT: ERBFMECT970094) and brings together many of the key players from around Europe who offer a rich diversity of High Performance Computing (HPC) systems and services. In ENACTS, our strategy involves close co-operation at a pan-European level - to review service provision and distil best-practice, to monitor users' changing requirements for value-added services, and to track technological advances. In HPC the key developments are in the area of Grid computing and are driven by large US programmes. In Europe we urgently need to evaluate the status and likely impacts of these technologies in order to move us towards our goal of European Grid computing, a "virtual infrastructure" - where each researcher, regardless of nationality or geographical location, has access to the best resources and can conduct collaborative research with top quality scientific and technological support.

ENACTS provides participants a co-operative structure within which to review the impact of Grid computing technologies, enabling them to formulate a strategy for increasing the quantity and quality of access provided.

The scope of our network is computational science: the HPC infrastructures which enable it and the researchers, primarily in the physical sciences, which use it. Three of the participants (EPCC, CINECA and CESCA-CEPBA) are LSFs providing Researchers' Access in HPC under the HCM and TMR programmes. All were successful in bidding to Framework Programme V (FP V) for IHP funding to continue their programmes. In this, they have been joined by the Parallab and the associated Bergen Computational Physics Laboratory (BCPL) and all four LSFs are full partners in this network proposal and plan to co-operative more closely in the Transnational Access programme. Between them, these LSFs have already provided access to over 500 European researchers in a very wide range of disciplines and are thus well placed to understand the needs of academic and industrial researchers. The other 10 ENACTS members are drawn from a range of European organisations with the aim of including representation from interested user groups and also by centres in economically less favoured regions. Their input will ensure that the Network's strategy is guided by users' needs and relevant to smaller start-up centres and to larger more established facilities.

A list of participants together with their role and skills is given in Table 1, whilst their geographical distribution is illustrated in Figure 1.

Table 1: ENACTS participants by role and skills

EPCC IHP-LSF Particle physics, materials science
ICCC Ltd User Optimisation techniques, control engineering
UNI-C LSF Statistical computing, bioinformatics, multimedia
CSC User Meteorology, chemistry
ENS-L Society Computational condensed matter physics, chemistry
FORTH User Computer science, computational physics, chemistry
TCD User Particle physics, pharmaceuticals
CINECA IHP-LSF Meteorology, VR
CSCISM User Molecular sciences
UiB IHP-LSF Computational physics
PSNC User Computer science, networking
UPC IHP-LSF Meteorology, computer science
NSC User Meteorology, CFD, engineering
CSCS-ETH-Zurich LSF Computer science, physics

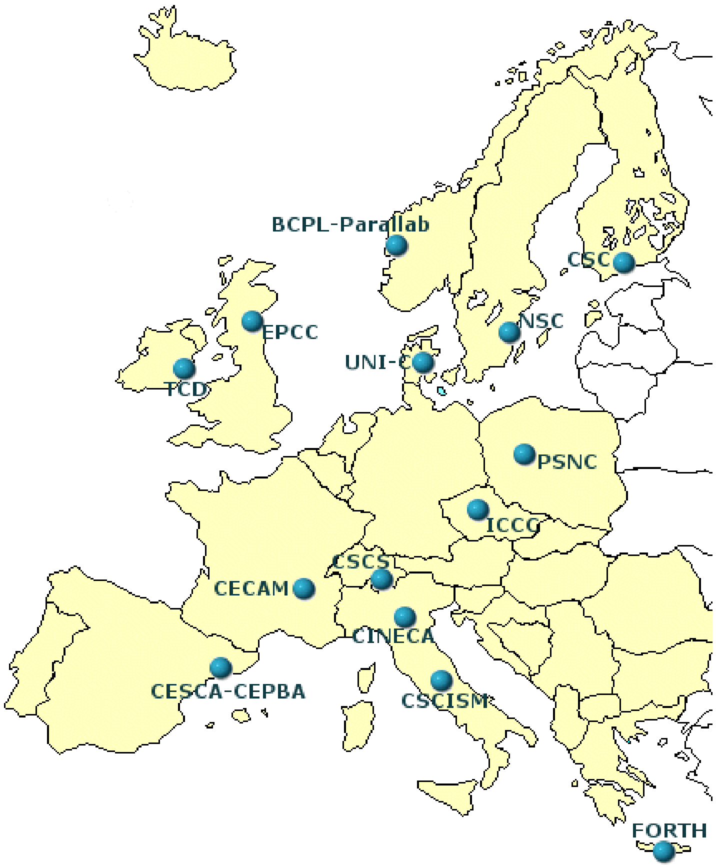

Work Plan

The aim of ENACTS dissemination activity is:

- to raise awareness of the projects and to publicise its activities, particularly its findings and results;

- to provide a mechanism to leverage efforts at a European level;

- to identify, define and undertake exploitation activities which will be beneficial to the operators and users of European Large Scale facilities, at a pan European level.

The main focus in Year 4 will be on the following:

WP1: Workshop – CECAM-ENS-L, one of the project partners, will organise a workshop to disseminate the project results to users of molecular simulation techniques, one of the core user groups of computational science. This workshop will be held in France.

WP2: Workshop/conference - ENACTS will host or participate in a European workshop or conference on Grid Computing. This activity will be organised by ICCC.

WP3: Dissemination Report - FORTH, with input from participants, will produce a report on the dissemination activities with pointers on how the project should be taken forward e.g. RTD

3. WHAT HAS BEEN ACCOMPLISHED

WP1:

According to the Work Package 1 of the dissemination project the following actions have been taken. CECAM (European Centre for Atomic and Molecular Computations, http://www.cecam.fr/) has organized in the year 2003 a four days workshop with emphasis the implications of High Performance Computing to Atomistic and Molecular Simulations. Glenn Martyna, Jacob Schiotz and Gilles Zerah organized this event by inviting key speakers, mainly computer code developers, to give a talk on

``Component Architectures, Open Standards and Parallel Algorithms for Molecular and Atomistic Simulations on Large Grids, Supercomputers, Workstations and Clusters''.

The event took place in Lyon, October 13-16, and the detailed program of this meeting is shown in the Appendix I.

The speakers covered topics such as Ab Initio software for HPC (ABINI, OHMMS, GAMESS), packages for molecular simulations of millions of atoms and the construction of force fields (CP2K, TINKER, CHARMM, AMBER), as well as subjects for managing large amounts of data (XML), molecular informatics and the development of parallel software with PYTHON and BSP. Talks on developing Grid computing algorithms to break fundamental barriers in molecular simulations were particularly interesting. Among a few existing algorithms is the one used in Folding@Home, a project devoted to the protein folding problem. This Grid paradigm is based on internet and volunteers who offer their PCs to run the programs as sceensavers.

A second CECAM workshop was organized in November 16-17, 2004, in Lyon, to collect and evaluate the main results of the ENACTS projects. Representatives from each working group highlighted the main objectives and conclusions of their study. It was interesting to see what happens after the elapse of two years since the completion of the first studies. Here we present in short the events of this meeting. More details can be found in Appendix II, where we have compiled from the published reports the objectives and the conclusions of each project. We hope that this will provide to the interested reader a detailed overview of the ENACTS accomplishments.

I. Grid Service Requirements (2002)

This study was the first ENACTS project carried out by EPCC and PSNC with the aim to establish a firm base on which the following projects could rely. The current Grid technology is reviewed, the terminology is explained and, through a well studied questionnaire, the Grid service requirements are investigated. In their extended report, which amounts to 329 pages, the authors collect important information and draw conclusions for the future trends in the Grid computational technology.
The objectives of the study were to track what the user community expects from the emerging Grid technology. The results will be useful to infrastructure operators, middleware developers, the user community, especially from the physical computational sciences and policy makers and funding bodies.
The well thought and designed questionnaire was addressed to 85 group leaders and examined seven areas of interest

- Key characteristics of user group
- Level of awareness
- User profile
- Application profile
- Infrastructure
- Security and services
- Future needs

Here are some important outcomes:

- 71% expect the "Grid" to help tackle **Grand Challenge problems**
- Grid technology seen as too complicated
  - Often perceived as one further level of complexity from HPC
  - Security concerns – being "in control''

Discussion: Computational Grids for solving Grand Challenge problems are still missing!

II. High Performance Computing Development for the Next Decade, and its Implications for Molecular Modelling Applications (2002)

The second study of ENACTS was conducted by NSC, EPCC and CSCISM. The objectives were

- Determine the likely technology and economic trends, which will prescribe the hardware architectures of HPC systems over the next 5 to 10 years (From 2001).
  - Survey of the technology roadmap for processors, memory, networking (closely coupled and LAN), data storage, custom-built solutions, and software paradigms and standards.
  - Grid-related technologies.

- A case study focusing on the usefulness and implications of the technologies discussed in the technology roadmap.
  - The implications for the molecular science community.

Seven interviewees participated as individuals and not as company representatives.

- Dr. Martin Walker, Compaq
- Dr. Wolfgang Mertz, Sgi
- Mr. Benoit Marchand, Sun
- Dr. Jamshed Mirza, IBM
- Dr. Burton J. Smith, Cray
- Mr. Tadashi Watanabe, NEC
- Dr. David Snelling, Fujitsu

Hewlett-Packard was invited but did not participate.

Here are some of the conclusions from the interviews

- Increased mass market will provide a basis for technological development.
  - Off the shelf components.
  - Linux clusters for HPC.
  - New / non-traditional HPC vendors.
- Future HPC systems are parallel, scalable architectures based on clustered SMPs.
  - COTS as well as proprietary components.
  - Unix with Linux as target.
  - SMP nodes with one to thousands of processors.
  - Price not technology sets the upper limit on size.
- Peak performance as well as price/performance will continue to follow Moore's law.
  - Performance strongly depends on application characteristics.
- Programming models and languages will be on an evolutionary path rather than on a revolutionary.
  - **MPI and OpenMP for parallel programming.**
  - **Fortran, C, C++ and Java.**
  - **Larger efforts than today will be needed to achieve high performance**.
  - Lack of adequate tools for software developments will become a key issue.
- We strongly believe that the Grid will play a key role in the way HPC is used.
  - A layer between the user and the HPC system.
  - Vendors should take the Grid into consideration when developing new systems and make old software Grid enabled.

The case study: Molecular sciences

- Construct four typical model HPC systems five year from now (2001).
  - Two based on projections of a proprietary hardware system (IBM SP Power 4):
    - A 12 M€ system for a large computing centre.
    - A 3 M€ departmental system.
  - Two based on projections of a COTS system (Intel Pentium 3-4):
    - A 12 M€ system for a large computing centre.
    - A 3 M€ departmental system.
- Model potentials based on Force-Field (FF) parameterizations.
  - Scales as *O(M2)* with system size *M*. Accuracy 20 kcal/mol.
  - Excellent scaling and performance on massively parallel systems.
- Density Functional Theory (DFT) derived potentials.
  - Scales as *O(M3)* or *M2logM*. Accuracy 3-7 kcal/mol.
  - Needs very efficient FFT, large memory and extremely efficient interconnects.
- *Ab initio* potentials derived directly from the Schrödinger equation.
  - Scales as *O(M8)*. Accuracy 0.5 kcal/mol.
  - Requires enormous disk storage, large memory, high bandwidth and low latency. Memory-bound, extreme data connectivity.
- Force-Field methods:

    - Will continue to dominate for studies of liquid crystals, ferro-electric neomatic materials and protein-folding. (Up to a few million atoms and time scale up to milliseconds.)
- Density Functional Theory (DFT) methods:
    - Accurate electronic, structural and dynamical reactive properties of systems containing 3000 (10000) atoms 5 (10) year from 2001.
    - Will replace FF-methods in chemiometric applications such as drug-design.
    - Accurate simulations relevant to respiration and photosynthesis.
    - Simulations of nano-scale systems such as molecular engines, quantum computation devices and chemical storage of data.
- *Ab initio* methods:
    - Accuracy 0.2 kcal/mol.
    - Simulations of elementary reactions relevant to the field of atmospheric chemistry and combustion.
    - Computations of cross sections and rate constants of systems of up to 10-50 atoms.
    - Studies of elementary reactions in the interstellar space
    -

Recomendations for HPC centres

- Collaborate and team up with other centres. Share resources.
- Expertise in many different areas of system architecture.
- Involved in open source development.
- Continued emphasis on MPI and OpenMP.
- Increased performance expertise.
- Expertise in Grid infrastructure and middle-ware.
- Application specific centres.

Discussion: Most of the conclusions are still valid (2004)!

III. Grid Enabling Technologies (2002)

This study was carried out by FORTH and CSCS. The main objectives were

1. Report uses the service definitions from the Grid Service Requirements report and evaluates different technologies and middleware which could be used to implement these services.

2. Survey of current status of computational grids in molecular science community – physics, chemistry and biology.

3. Description of authors' experiences in installing Unicore and Globus on a local cluster and thus implementing a small local computational grid.

In the first part of the study, an internet search was carried out (2002) to locate operating testbeds all over the world, to see what kind of middleware is used and what services they provide. Among the middleware for implementing a Grid (Globus, Unicore, Legion) the authors investigated in detail the toolkit approach of Globus, and the abstract job object approach of Unicore, whereas among the established testbed grids they explored the USA based ALLIANCE and NASA Information Power Grid (IPG) and the European EUROGRID.

The two main conclusions of the study were

1. Globus as a Grid deployment technology dominates in both sides of the Atlantic Ocean.

2. Most of the applications concern particle physics, astronomy, environmental and meteorological projects and in less extent biological and chemical applications.

Discussion: Conclusions are still (2004) valid!

IV. Data Management in HPC (2003)

This study was performed by CINECA and TCD. The main objective was to gain understanding of the problems associated with storing, managing and extracting information from the large datasets increasingly being generated by computational scientists, and the technologies which could address these problems.

Principal aim of this study

- Investigate and evaluate currently available technologies
- Explore new standards and support their development
- Develop good practice for users and center
- Investigate platform- independent and distributed storage solutions
- Explore the use of different technologies in a coordinated fashion to a wide range of data- intensive applications domains

This in depth analysis included

1. Basic technology for data management
2. Data models and scientific data libraries
3. Finding data and metadata
4. Higher level projects involving complex data management
5. Enabling technologies for higher level systems
6. Analysis of Data Management Questionnaire
7. Executive summary and recommendation

The conclusions are really interesting

- 60% respondents perceived a real benefit from better data management
- Majority do not have access to sophisticated or high performance storage systems
- **Many computational scientists are unaware of the evolving GRID technology**
- The use of data base management technologies (DBMS) is scarce
- For industry security and reliability are stringent requirements for users of distributed data services
- The survey identifies problems coming from interoperability limits:

  - limited access to resources,
  - geographic separation,
  - site dependent access policies,
  - security assets and policies,
  - data format proliferation,
  - lack of bandwidth,
  - coordination,
  - standardising data formats.
- Resources integration problems arising from different physical and logical schema:
  - relational data bases,
  - structured and semi-structured data bases,
  - owner defined formats, etc.
- Representation of semantic properties of scientific data and their utilization in retrieval systems.

- The questionnaire shows that European researchers are some way behind in their take up of data management solutions.
- However, many good solutions seem to arise from European projects and GRID programmes in general.
- The European research community expect that they will benefit hugely from the results of these projects and more specifically in the demonstration of production based Grid projects.

V. Distance Learning and Support (2004)

The study was done by ICCC and UNI-C. The main objectives were

Part II: Various definitions of distance learning and education. Presentation of basic distance learning tools and models, including a discussion of their features, limitations, and benefits for prospective users. Although this text is relatively general, the authors want to focus on three main user groups which have been identified: HPC centres (service providers), Scientific Grid community (researchers and users) and Industrial Grid community (vendors and end-users). These three target groups are shortly discussed in the end of this section.
Part III: Results from a comprehensive survey distributed among 85 and undertaken among 25 major European research groups. It focuses on different aspects, such as the needs and requirements of various potential target groups and the pedagogical and

organisational approach, which fits best with identified target groups. The survey includes a clear analysis how to ascertain the feasibility, viability and relevance of adapting a proper distance learning strategy to the training requirements and leads into a evaluation and agreement on a framework for collaborative development of suitable distance learning based course material.

Part IV: Conclusions and recommendations. The purpose of the survey presented in Part III is to gain a better understanding of these key user groups' needs and requirements in view of establishing a proper framework for distance learning and support. The analysis of this survey, together with Part II, which presented general concepts and technological issues, will be instrumental in establishing key recommendations for target groups. This section provides a summary of the distance learning features offered by the leading but still rather small groups of Grid specialists and users and it will make recommendations on a possible strategy that supports a successful uptake of Grid technology around larger communities.

Discussion: Grid Service Requirements report highlighted a paradox – little enthusiasm for distance learning from computational scientists, but we need to train a large number of scientists in Grid Computing quickly.

VI. Software Reusability and Efficiency (2004)

The study examines issues surrounding portability of software and suggests how program efficiency can be ensured across a set of heterogeneous and/or geographically distributed computational platforms and data storage devices. Furthermore, it reviews emerging software technologies as these affect the ability of existing and new applications to make best use of HPC facilities.

Here is some of their conclusions:

1. Moving large numbers of application codes between heterogeneous HPC systems or to new software environments is only feasible if well-accepted standards, languages and tools are available.

2. An HPC centre or consortium should therefore promote best practices in portable data management towards its user community and provide appropriate advice and training.

3 A tool that is considered mature enough and provides added value should be promoted to the user community of the HPC consortium by means of providing adequate training material.

4. An HPC user may not rapidly move onto the grid if this requires significant changes to his/her existing application.

5. Adapting to emerging web and grid standards may require regular or frequent changes to the grid-enabled application.

6. The use of open standards should be promoted actively

7. TheWeb has become the user interface of global business, and Web services now offer a strong foundation for software interoperability through the core open standards of XML, SOAP, WSDL, and UDDI. Models and applications that make use of this huge potential are just beginning to emerge.

8. Many of the new technologies require knowledge of programming languages like Java and C++, scripting languages like Perl, middleware like Globus, and XML-based standards. Several of these are fairly new and evolving.

9. Another emerging form of grid computing is the *collaborative* grid. Such a grid enables the creation of virtual organizations in which remote research groups can perform joint research and share data.

10. The use of grid technologies will eventually lead to more remote (distant) collaborations. It is therefore essential that mechanisms are in place for code maintenance by a large research group with multiple programmers modifying a single code.

VII. Grid Metacenter Demonstrator Demonstrating a European Metacentre (2004)

The ENACTS project would seem incomplete without having some experience, at least some of us, working with a real grid. This was materialized by the construction of a Data Grid to run a project on Quantum Chromo- Dynamics. The aims of the project were

- Set- up a data grid across the 3- sites using QCDgrid.
- Use a genuine scientific scenario.
- Ensure all data is described using meta- data.
- Ensure the data is portable between the systems involved.

- QCDgrid was written to manage the QCD data belonging to the UK QCD community (UKQCD).
- The original grid consisted of 6 geographically dispersed sites (UKQCD).
- Around 5 terabytes of data.
- The amount of data is expected to grow dramatically when QCDOC comes online later in 2004.
- QCDgrid is a layer of software written on top of the Globus Toolkit.
  - Uses security infrastructure and basic grid operations such as data transfer
  - also uses more advanced features such as the replica catalogue.

Conclusions and difficulties

- Certificates
  - Some certificate issuers took several weeks to issue certificates.
  - Different policies on issuing certificates, e.g. non- human users (project accounts).
  - Not too many difficulties using multiple certificates.
- Moving to a heterogeneous environment.
  - Installing of Globus 2.x is difficult on Solaris – led to the Solaris node being unable to submit data.
- A few minor problems getting system specific functions to work (e.g. df command).
- Usual minor compilation issues – did require gcc compiler.
- Globus
  - This presented the biggest difficulty!
  - Installation difficulties and firewall issues
    - several months before a “helloworld” job would run from any site to any other.
  - Migrating from GT 2.0 -> GT 2.4
    - Major difficulties!
    - Had to re- write the replica schema.
    - Remove some error- handling functionality.

VII. Survey of Users' Needs (2004)

ENACTS in its final year launched a questionnaire to find out what the European users need for HPC and Grid computing. The final conclusions are

The crucial role of the computing power is the key issue in the involvement in Grid related issues. It is very encouraging that a large majority of the users is willing to share knowledge, tools, data and results. The main challenges Grid computing will face will be political; more agreements will be needed on how to use the data and resources in multinational Grid infrastructure.
The Grid is in development stage and users are willing to contribute to the Grid infrastructure development, if proper help and support is provided. Virtual tools for education and training could enable faster, cheaper and more environmentally friendly communication. To enhance the adaptation of virtual tools, these tools should be more user- friendly.
The growing amount of data that is produced presents challenges. Standardisation and interoperability efforts of the international community are needed for diffusion of the knowledge, cooperation and best exploitation of resources.

Discussion: Good agreement with what has been learnt in the previous years.

WP2:

With respect to the Work Package 2, ICCC will present a paper on Distance Learning and Support in the EDEN annual conference which will be held in Helsinki, June 20- 23, 2005. The article is shown in Appendix III.

WP3:

This report.

4. CONCLUSIONS

The ENACTS project with its well planned and in depth studies has shown that GRIDS of computers are well developed not only at the national level but also as a European multi-national integration. Computer scientists in most European countries have taken initiatives and financial support to create a Grid all over the country to support scientific research, education, economical activities and administration. European Union in its VI Framework for research and technology approved several projects for Grid computing,: EGEE (Enabling Grids for e-Science in Europe), DEISA (Distributed European Infrastructure for Supercomputing Applications), HPC-Europa (High Performance Computing-Europa) and NextGrid (an architecture for the next generation Grid). The middleware to deploy and operate a Grid is well established with GLOBUS taking a leading role.

The key players in hardware development have given their full support in this new technology. IBM, for example, is backing the **World Community Grid** (http://www.worldcommunitygrid.org/) and its human proteome folding project, which uses distributed computing to predict the structure of all proteins found in human genome. IBM has donated the hardware, software, technical advice, and expertise to built the infrastructure for the project which is run by grid-computing specialists **United Devices** (http://www.ud.com/). IBM also contributes with the publications of books and articles available in Redbooks (http://www.redbokks.ibm.com/, ``Introduction to Grid Computing with Globus'', ``Enabling Applications for Grid Computing with Globus'', ``Fundamentals of Grid Computing''), thus supporting education in the new computer technologies (http://www.developer.ibm.com/university/students/contests/Scholars_Challenge.html).

In the last years, tools to deploy and operate a Grid have been developed and it seems that the number of such packages will increase with the time. For example a few of these products are: ``The **Gridbus** Toolkit'' (http://www.gridbus.org/) , **ASSIST** (Application of the Grid and the coordination language ASSIST), **P-Grade** (Parallel Grid Run-time and Application Development Environment, http://www.lpds.sztaki.hu/pgrade/).

What kind of applications do the existing Grids run?

Among the existing Grids (computational, data, collaborative, access-grid, educational) the most successful are the data-Grids with the best representative of them to be the **DataGrid** (http://eu-datagrid.web.cern.ch/eu-datagrid/) project coordinated by CERN and funded by the European Union. The objective was to build the infrastructure for handling data from hundreds of terabytes to petabytes across widely distributed scientific communities. Contrary to this, successful computational Grids similar to the DataGrid have not yet appeared in spite of what the majority of the HPC users want;

``71% of the users asked expect the "Grid" to help them to tackle Grand Challenge problems''

Why such a difference? Let us see what is the situation in Molecular and computational Material and Biological sciences. Admittedly these are among the computational sciences the most demanding in computer resources. Numerical simulations provide a means for studying the properties of materials under equilibrium or non-equilibrium conditions. Starting from physically founded cohesion models or from an ab-initio description of atomic bonding, atomic scale simulation techniques (Molecular Dynamics and Monte Carlo) give insight to the behaviour of the materials studied, biological or not. This can help improving the quality of materials for technological applications and possibly to the development of new materials and/or fabrication processes. Unfortunately, the scales of most properties relevant for applications are fundamentally different from these of atomistic simulations. The challenge therefore for numerical simulations consists in bridging the gaps of space and time scales that separate atomic and macroscopic material properties which may differ more than twelve orders of magnitude.

***Grid (distributed) computing could be the solution in bridging the gaps in time and space*.**

As a matter of fact, computational quantum chemistry was among the first disciplines to try to exploit the idea of networking computers. In early nineties Hartree-Fock calculations were distributed across continents to evaluate the performance in such remotely distributed cluster of computers. It is interesting that the early experiments proved that, i) on a network of computers a computation can be performed faster than on any single machine, and ii) a computation can be accelerated further by introducing heterogeneity in the network. However, in the following years there was no significant progress in this subject and the conclusion of ENACTS was that, whereas computational quantum chemistry was among the pioneers in metacomputing, other areas of HPC have taken the lead. It turns out that this pessimistic conclusion is still valid.

In our days, the most successful computational Grid model is that based on internet by accessing thousands of PCs via running the programs as screensavers. The first such application was SETI@home (http://setiathome.ssl.berkely.edu/) to analyse the data from radio telescopes looking for signs of extraterrestrial life. Taking its inspiration from SETI, a protein folding project called Folding@Home (http://www.stanford.edu/group/pandegroup/folding) has been in operation at Stanford University for four years now.

1In order to use successfully a worldwide distributed computing environment of thousands or even millions of heterogeneous processors such as the Grid, communications among these processors should be a minimum. There are not many molecular simulations using the Grid environment such as to allow us to point out the strategy one should adopt in writing codes for molecular applications. A practical rule is to allow each processor to work independently, even though calculations are repeated, and only if something important happens to one of them, then they communicate. Such a strategy is called **Embarrassing Parallelization**. A. F. Voter (Phys. Rev. B, **57**, R13985, 1998) has shown that a linear speedup can be obtained in calculating the first transition rate constants by integrating replicas of the system independently on M processors. A set of parallel replicas of a single simulation can be statistically coupled to closely approximate long trajectories. This produces nearly linear speedup over a single simulation. The investigators with the Folding@Home have simulated the folding mechanism and accurately predicted the folding rates of several fast folding proteins.

Another active academic project is Predictor@home (http://predictor.scripps.edu/) which is aiming to structures prediction. This is a pilot of the Berkeley Open Infrastructure for Network Computing (BOINC) – a software platform for distributing computing using volunteer computer resources. Predictor@home which is led by Charles Brooks of The Scripps Research Institute, has the final goal of testing new algorithms as part of the sixth bianuual CASP (Critical Assessment of Techniques for Protein Structure Prediction). Another action on the human protein folding problem has been taken by the World Community Grid (WCG) (http://www.worldcommunitygrid.org/) which is supported by United Devices (UD). UD also runs the life-science research hub (http://www.grid.org/) to search for drug candidates to treat the smallpox virus. Some 35 million molecules were virtually screened against smallpox proteins, and 44 strong treatment candidates were identified and handed to the US Department of Defense for further evaluation. The WCG takes the idea a stage further, and aims to create the world's largest public computing Grid and apply it to a number of humanitarian projects

– of which protein folding is the first. In the ``Human Proteome Folding Project'' the program used is called **Rosetta** which computes a scoring function to sort through the possible structures for a given sequence and choose the best. The computers in the Grid try to fold the protein in different ways. This will be done millions of times for each protein, and lowest scoring structures compared by researchers with the Protein Data Bank (PDB) at Rutgers University.

The common ingredient of the above successful applications in Internet Grid Computing is the computational algorithm employed and which guarantees **``Embarrassing Parallelization''**. Thus, the practice up to now demonstrates that the Grid computing requires new approaches in solving scientific problems. Parallelized codes written for large parallel machines may be of no good use when thousands of computers should be exploited connected by relatively slow networks.

``Without these new algorithms and new programming paradigms suitable for distributed computing Computational Grids will not play the role and they will not fulfill computational scientists' ambitions to tackle Grand Challenge problems''.

APPENDIX I

CECAM-ENACTS
Workshop 2003

Component Architectures, Open Standards and Parallel Algorithms for Molecular and Atomistic Simulations on Large Grids, Supercomputers, Workstations and Clusters

Day 1 : Morning Chair : Glenn Martyna

9:00-9:20 **Glenn Martyna**, IBM, YKT, "Welcome and Introduction".
9:20-10:10 **Xavier Gonze**, Universite Catholique de Louvain, "The ABINIT software project".
10:10-10:30 Coffee Break!
10:30-11:20 **Jeongnim Kim**, University of Illinois, "Development of portable electronic structure tools on high-performance computers (OHMMS)'.
11:20-12:10 **Mark Gordon**, Iowa State University, "Enabling high performance electronic structure theory: Models and applications (GAMESS)".

Day 1 : Afternoon Chair : Heather Netzloff

1:40-2:30 **Thijs Vlugt**, Utrecht University, "Parallel Configurational Bias MC".
2:30-2:50 Coffee Break!
2:50-3:40 **Rajiv Kalia**, USC, "Multimillion atom simulation of nanoscale dynamics and fracture".
3:40-4:30 **Jens Jergen Mortensen**, Danmarks Tekniske Universitet, "Real space implementation of the projector augmented wave method in CAMPOS Atomic Simulation Environment".
4:30-5:20 Discussion.

Day 2 : Morning Chair : Ken Esler

9:00-9:40 **Fawzi Mohamed**, CSCS Science Division, "The CP2K molecular simulation package".
9:40-10:30 **Martin Head-Gordon**, UC Berkeley, "Fast methods for electron correlation (QCHEM)".
10:30-10:50 Coffee Break!
10:50-11:40 **Ramkumar Vadali**, UIUC, "A Framework for Multiple Concurrent FFTs".
11:40-12:20 **Glenn Martyna**, IBM, YKT, "Fine grained parallelization of plane wave based DFT".

Day 2 : Afternoon Chair : Gilles Zerah

1:40-2:30 **Xavier Gonze**, Universite Catholique de Louvain, "Introduction to XML".
2:30-2:50 Coffee Break!
2:50-3:20 **Peter Murray-Rust**, Cambridge, "Molecular Informatics for the Grid".
FSatom Forum : **Gilles Zerah** moderates
i. Introduction to FSatom.
ii. Summary of activities of FSatom.
iii. Report from Pseudopotential work group.

iv. Forum and Collective Decisions.

Day 3 : Morning Chair : Michael Shirts

9:30- 10:20 **Jakob Schiotz**, Danmarks Tekniske Universitet, “Python as a glue in atomicscale simulations: Advanced simulation methods and parallel molecular dynamics within the CAMPOS Atomic Simulation Environment”.
10:20- 10:40 Coffee Break!
10:40- 11:30 **Weitao Yang**, Duke University, “(1) O(N) Electronic Structure Calculations with Nonorthogonal Localized Molecular Orbitals (2) Distributed Computation of Chemical Reaction Paths in Enzymes.”.
11:30- 12:20 **Heather Netzloff**, Iowa State University, “Simulating solvent effects and liquid structure with the effective fragment method (GAMESS)”.

Day 3 : Afternoon Chair : Philip Blood

1:40- 2:30 **Konrad Hinsen**, Centre de Biophysique Molecul, “High level parallel software development with Python and BSP”.
2:30- 2:50 Coffee Break!
2:50- 3:40 **Pengyu Ren**,Washington University, “TINKER Polarizable Mutlipole Based Model for Molecular Simulation”.
3:40- 4:30 **Michael Crowley**, Scripps, “CHARMM and AMBER: Palaverous colloquy kills parallel efficiency. How do we cut down on the chatter?”.
4:30- 5:20 **Michael Shirts**, Stanford University, “Folding@Home: Grid computing algorithms to break fundamental barriers in molecular simulation”.

Day 4 : Morning Chair : Ramkumar Vadali

9:00- 9:15 **Philip Blood**, University of Utah, “Grid computing for multiple time and length scales”.
9:15- 10:05 **Ken Esler**, University of Illinois, “Path Integral Monte Carlo for Systems of Heavy Atoms”.
10:05- 10:25 Coffee Break!
10:25- 11:15 **Carlo Cavazzoni**, Cineca, “Computational Material Science Issues on Grid and Large Clusters”.
11:15- 12:05 **Mark Tuckerman**, NYU, “Towards a linear scaling approach in plane wave based DFT: Gauge invariance and localized orbitals (PINY MD)”.

Day 4 : Afternoon Chair : Gilles Zerah

(a) Workshop Summary and Discussion.
(b) **Glenn Martyna**, IBM, Yorktown Heights, “Closing Remarks”.

CECAM-ENACTS
Workshop 2004

NOV 16, 2004
============

10.00 - 10.30 What is ENACTS: our objectives - our methodology

10.30 - 12.30 6 x 15 minutes talks presenting the strong points of each of the 6 studies

12:30 - 14:00 Lunch

14.00 - 14:30 Lessons learn form the ENACTS demonstrator (feasibility)

14.00 - 15.00 Conclusions / most interesting aspects of user survey

15.00 - 15.30 Coffee break

15.30 - 16.00 An overview of trends in HPC and Grid Computing in Europe (policy/funding aspects as well as technological aspects)

16.00 - 18.00 Presentation from the invited speakers

==

NOV 17, 2004
============

9.00- 10.30 users presentations and general discussion (follow)

coffee break

11.00 general discussion and conclusions
12.00
Meeting close

APPENDIX II

Overview of ENACTS Studies

PROJECT 1

Grid Service Requirements
http://www.epcc.ed.ac.uk/enacts/gridservice.pdf

JC Desplat, Judy Hardy, Mario Antonioletti (EPCC)
Edinburgh Parallel Computing Centre

and
Jarek Nabrzyski, Maciej Stroinski, Norbert Meyer (PSNC)
Poznan Supercomputing and Networking Centre

Study objectives

The objective of this study was to specify the level and quality of services users require from a computational Grid. The report was written in consultation with a representative selection of users from different computational science groups across Europe. This top-down approach is complemented by a bottom-up examination of the enabling technologies. There were five work-packages in this project, totalling 6 staff months of effort.

This study is composed of four main parts:
Part I: Presentation of the ENACTS project, its scope and membership.
Part II: Presentation of the different types of computational Grids, with a discussion of their features, limitations, and FAQ for prospective users. Since the material presented in this section is relatively detailed and tackles technical considerations which are mostly irrelevant to end-users, the authors expect this section to be of particular interest to operators of HPC systems and systems administrators who want to find out more about Grid computing. Sections specifically aimed at end-users (such as a FAQ) have also been included in this second part. This section has been written by PSNC.
Part III: Results from a comprehensive survey undertaken among 85 major European research groups. This covered aspects such as awareness and expectations of Grid computing, groups and applications profile, working practices, bottlenecks and requirements, future needs, etc. This section has been written by EPCC.
Part IV: Conclusions and recommendations. The purpose of the survey presented in **Part III** is to gain a better understanding of these key user groups' composition and working practices in view of establishing their key requirements. The analysis of this survey, together with **Part II**, which presented technical and implementation issues, will be instrumental in establishing key recommendations for policy makers and resource providers. This section will provide a summary of the features offered by the leading-edge Grid middleware presented in **Part II**, compare this to the users' requirements established throughout **Part III**, and following these, make recommendations on a

possible strategy for a successful uptake of Grid technology within the user community. This section has been written by EPCC.

Summary and conclusions

This section comprises three main parts:
1. A summary of the features offered by leading-edge Grid middleware technology and higher level components, based on the detailed material presented in Part II;
2. A summary of the key service requirements identified by the user community in the preceding survey; and
3. A proposed strategy to increase a successful uptake of Grid-based computational resources within the user community. This strategy will lead to recommendations aimed at both HPC resource providers and funding agencies.

A. Summary of features available within a Grid environment: comparative study of different Grid models

This study aims to help users of Grid systems that can exploit heterogeneous networked computing resources by discovering under-utilised remote computers and deploying jobs to them. As many of the systems which are currently available tend to have significantly similar functions, a study of the relative suitability of such systems for being extended and used for this purpose is a logical first step. This report presents the findings of such a study that considered the systems presented in the previous section. These are **LSF, Globus, LEGION, Condor, Unicore and Entropia.**
The main features of the systems analysed are divided into four major categories.

System, flexibility and interface: See Table 1. These data lead to the following conclusions.

The majority of the systems are available in public domain. The only system that does not have a public domain version is LSF.

Documentation is excellent for LSF, and good for the majority of the other systems. The majority of the systems are constantly being updated and extended with new capabilities.

The majority of the systems operate under several versions of UNIX, including Linux.

Windows NT versions are under development for Condor and Globus.

Impact on the owners of computational nodes is typically small and configurable by the users themselves.

The majority of the systems have both a graphical and a command-line interface. The only system without a graphical interface is Legion.

Scheduling and Resource Management: See Table 2. We draw the following conclusions.

Batch jobs are supported by all systems, whether centralised or distributed. Interactive jobs are not supported by Condor and Unicore.

Parallel jobs are fully supported only by LSF, although limited support is provided by Globus and Legion.

Resource requests from users are supported by all the job management systems except Legion. In Globus, a special specification language, RSL, is used to specify the job requirements.

The most flexible scheduling is provided by LSF. In Condor and Entropia, one of several predefined scheduling policies can be selected by a global administrator.

All centralised job management systems support job priorities assigned by users.

All centralised job management systems and Globus support job monitoring.

Accounting is available in all centralised job management systems.

Efficiency and utilisation: See Table 3. These data lead to the following conclusions.

Stage-in and stage-out are supported by LSF, Globus and Legion, and to a limited extent by Condor.

Time-sharing of jobs is supported by LSF, Globus, Legion, and Unicore.

Time-sharing of processes is available in Condor.

Process migration is supported by all the centralised job management systems.

Dynamic load balancing is supported only by LSF.

Scalability is high for LSF and all distributed job management systems.

Scalability in Condor is limited by a central scheduler.

Fault tolerance and security: See Table 4. We draw the following conclusions.

System level checkpointing is supported only by LSF and for only a small subset of the operating systems. Run-time library checkpointing is supported by LSF and Condor. User level checkpointing is supported by LSF and Condor.

Fault tolerance is the best in LSF and Condor.

Authentication based on Kerberos is provided in LSF, Globus, Legion and Unicore.

Strong authentication based on SSL and X.509 certificates is provided in Condor, Globus, and Unicore. Additionally, Globus supports hardware authentication tokens.

Strong authorisation is available in LSF, Condor and Legion, DCE.

Encryption is clearly documented in Globus, Legion, and Unicore, but may also be present in several other systems using Kerberos and SSL.

B. Summary of key user requirements within a Grid infrastructure

Advocating the use of Grid computing by the user community invariably yields one reaction:

"Why? Can you demonstrate how the Grid would benefit my work?"

This has to constitute the core preoccupation for people involved in the uptake of Grid computing *before* any of the technical aspects presented in Part III are considered.

The survey presented in Part III has provided a relatively detailed account of user groups characteristics (such as group size and scope, level of co-ordination and composition, etc.), of their expectations of Grid computing and of their current requirements. The most remarkable finding was how much these factors varied across research areas. A summary of these characteristic features is given in Table 25, page 195.

This diversity across the user spectrum suggests that there will be no such thing as an ideal Grid infrastructure. Grid user service requirements must be considered from two different angles to get a better understanding of how best to address them:

1. Consider users' aims and expectations: What do they plan to use computational Grids for? What are the expected benefits for the user community from their own point of view?
2. Consider users' requirements: What do users identify as the key features and services within a Grid environment, and why?

What do users expect from a computational Grid?

One of the most reassuring aspects of the survey was the high level of support from the user community for the concept of Grid computing. Indeed, a stunning 91% of the group leaders participating in this survey believe that their group will benefit from accessing computational Grids. Beyond this though, it is important to highlight exactly what users expect from this new technology.
Summarising the discussion from "Perception and benefits of computational Grids," page 130, in order of importance users expect Grid computing to:

1. provide more cycles;
2. enable the study of larger scale problems;
3. increase opportunities to share application codes and packages;
4. provide access to sophisticated visualisation and data analysis tools; and
5. share distributed data.

Which services will user require within a Grid environment?

"Section 2: Awareness," page 125 and "Section 6: Security and services," page 174 considered the key features users would expect to find within a Grid environment. A summary of these findings is presented in Table 25, page 195 under the categories '*most important factors in a Grid environment*' and '*most important added services in a Grid environment*'.
A detailed discussion of these findings, and in particular an analysis of discipline specific variation and requirements is available in Part III. The core requirements essentially consist of increased levels of conventional hardware resources such as faster CPUs, more processors or greater total memory. Whilst the need for greater total memory and access to more PEs is a clear indication of the users' intention to tackle larger or more complex problems (see IV.2.2.2), the need for faster CPUs also encompasses the issues raised in IV.2.2.1, i.e. more cycles and increased throughput.
It is not surprising that these were the most highly-rated factors, as these are the ones typically put forward during procurements for HPC systems. However, they are not linked to Grid computing per se. In this respect, the following categories are much more relevant i.e. the need for (in decreasing order of importance):

1. longer run-times;
2. ease of use;
3. better throughput;
4. best machine for the job; and
5. access to resources not locally available.

In addition to these requirements, users also regarded the following services to be important within a Grid environment (again, in decreasing order of importance):

1. availability and reliability;

2. ease of use;
3. support (e.g., Technical queries);
4. security of code and data;
5. guaranteed turn-around time; and
6. remote visualisation tools.

Miscellaneous considerations

A number of other important issues also emerged throughout the questionnaire. Although they did not fit into any the categories of sections IV.2.2 and IV.2.3, these issues have nonetheless a definite relevance to this study. They are:

1. The difference in awareness level across research areas is quite noticeable and will have implications in terms of strategy through funding and dissemination activities. Different research groups will adopt the technology earlier than others, based on the following combination of factors:

- the benefits they expect from computational Grids (see “Section 2: Awareness,” page 125);
- the type of services they will require (see “Section 2: Awareness,” page 125, and “Section 6: Security and services,” page 174);
- their structure, level of co-ordination, and existing collaborations at a national and international level (see “Section 1: Start ,” page 121, and “Section 3: User profile,” page 146);
- the type of resources they own and their willingness to integrate them within a Grid (see “Section 5: Infrastructure,” page 170 and “Section 6: Security and services,” page 174); and
- their existing links with HPC centres.

2. There is a much greater diversity of working practices, size, structure collaborative links and applications’ characteristics and requirements across research areas than within any single area (see “Section 3: User profile,” page 146). This suggests the suitability of topical Grids for large, well co-ordinated groups involved in international collaborations.

3. The type of architecture deemed essential within a Grid infrastructure also varies considerably across research areas (see Table 25, page 195 and “Most wanted resources in a Grid environment ,” page 142). In particular, the reliance on large MPP systems has to be better understood as such systems will progressively be superseded by cluster of shared memory systems.

4. The preference for a particular Grid concept varies not only across disciplines but also depends on the level of experience of the users concerned (see Figure 35, page 150 and “Features of mainstream research areas ,” page 147). Thus, portals are clearly favoured by the less Grid literate whilst toolkits (and in particular Globus) are put forward as the best solution by those at the other end of the spectrum. This suggests that the development of portals for mainstream application codes and packages (such as Gaussian) is the most effective strategy for the introducing the concept of Grid computing to groups from the chemistry and bioinformatics areas (see “Section 3: User profile,” page 146). It is likely that the additional features of portals, such as the ability to generate configuration files and perform sanity checks on these latter, will also prove important to these communities.

5. The preferred scope for computational Grids (local, national, topical, etc.) correlated directly with both the type of collaborations groups were involved in and the amount of resources they directly control.

6. The composition or structure of the research groups in term of proportion of black box users and user developers i.e. the distribution of their users' level of computing expertise, suggest that many groups in physics and engineering are often making a low re-use of their codes (see "Section 3: User profile," page 146). These practices result in a plethora of codes 'in development' which are inadequate to be integrated within a portal or problem solving environment. This was illustrated by the high frequency of recompilations required between successive runs (see "Section 4: Application profile," page 156). Steps should be taken to encourage community led initiatives to develop standard skeleton codes following the examples of say, Cactus [27] or DL_POLY [98].

7. Application bottlenecks appear roughly uniform across applications areas: CPU performance, memory to CPU bandwidth, and memory capacity were reported as the top 3 bottlenecks across all areas. However, network performance starts to become an issue for the computing, engineering, and astrophysics communities, whilst data storage capacity is of concern to the chemistry and astrophysics communities (see "Section 4: Application profile," page 156). Note that network performance will be one of the main bottlenecks to any user who makes routine use of computational steering or metacomputing. Mechanisms providing job migration (e.g., for fault tolerance or guaranteed turn around time) and access to distributed datasets will be equally affected by poor network performance.

8. Following recent investments in the national and transnational network infrastructures, the bottleneck for end-to-end bandwidth has now been displaced to within institutions and universities (see "Section 5: Infrastructure," page 170). A successful uptake of services such as remote visualisation will require upgrades to local networks for which funding can prove difficult to secure.

9. Security is a matter for researchers both as users of Grids and as prospective resource providers. Indeed, should a group decide to set up their own local Grid or even join an existing Grid, it is important that the security framework can be fully understood, operated, and trusted by the group's systems administrators and support teams (see "Section 6: Security and services ," page 174).

10. Although it was shown that most application codes could be made fully portable with few modifications, it is important to keep on promoting code portability within the research community (see "Section 4: Application profile," page 156). The same remark also applies to data interoperability.

.

PROJECT 2

High Performance Computing Development for the Next Decade, and its

Implications for Molecular Modelling Applications
http://www.epcc.ed.ac.uk/enacts/hpcroad.pdf

Jan Fagerstroem, Torgny Faxen, Peter Munger, and Anders Ynnerman (NSC)
National Supercomputer Centre

J-C Desplat (EPCC)
Edingurgh Parallel Computing Centre

Filippo De Angelis, Francesco Mercuri, Marzio Rosi, Antonio Sgamellotti, Francesco Tarantelli and Giuseppe Vitillaro (CSCISM)
Center for High Performance Computing in Molecular Sciences

Study objectives

This report presents the results obtained within the HPC Technology Roadmap study within the ENACTS project. The project started in February 2001, and the present report presents the results from the second study within ENACTS: The HPC Technology Roadmap.

The HPC Technology Roadmap study is undertaken by the National Supercomputer Centre (NSC) in Link¨oping, Sweden in collaboration with Center for High Performance Computing in Molecular Sciences (CSCISM) in Perugia, Italy. The objective of the study is to determine the likely technology and economic trends, which will prescribe the hardware architectures of HPC systems **over the next 5 to 10 years**, and to evaluate the effects that this will have on applications software. The work within the study has been shared between NSC and CSCISM in the following way. NSC has provided a survey of the technology roadmap for processors, memory, networking (closely coupled and LAN), data storage, custom-built solutions, and software paradigms and standards. This survey was accomplished by interviews with several major HPC vendors, and is reported in section I. NSC has also coordinated the study. Based on the results of the technology roadmap survey CSCISM has provided a case study focusing on the usefulness and implications of the technologies discussed in the technology roadmap, for the key molecular science community. The case study is presented in section II.

Summary and conclusions

Who can look into the future? We all know this is impossible but nevertheless it is necessary. In order to plan our activities we need to have some model of what we think that the future brings. The interview material we have at hand for this report is made by highly acclaimed, knowledgeable representatives from the most prestigious and well known vendors of HPC systems. It is a very interesting and exciting material to read. All interviewed have stressed that this is their own personal view of future HPC and not necessary that of the vendor, so we have unique personal opinions of HPC future based on experience and knowledge from rich sources, some which is only available to people within the corporations.

Many forward looking statements are extrapolations of today's trend. All material from the interviews are made public so it comes as no surprise that vendor representatives can not share all of their knowledge, especially major alterations in strategies. Nevertheless, having access to such a complete set of in-depth interviews gives a unique snapshot of where the industry believe it is heading.

However, it is hard not to get lost in all the details, especially in an area where there is so much rapid progress as in the computer industry. Everything is moving ahead at a very rapid pace, and some areas faster than others. Below we try to summarise and interpret the interviews keeping the focus on what the implications will be for the end users and the computer centres.

Summary of the interviews

1. HPC market

If we start by looking at how market is viewed, HPC today has a rather small share of the total computer market whereas the commercial (enterprise) and consumer markets are much larger. This trend will continue. This also means that a majority of the investments will be in non-HPC markets. There is a difference of opinion on the implications of this however. Either this market trend will force future HPC systems to be integrated directly from commercial components (processors, networks or even complete systems), or proprietary components can still be developed cost effectively for HPC specific solutions since what the consumer and commercial market develops and invests in can be used also for HPC, especially the facilities and methods for making chips.

There is also general agreement that a majority of the investment will move (if it has not already moved) to the consumer market, but again there are different opinions on how this will effect future HPC system architectures. Some believe this means that future systems are going to based mainly on consumer market components even though this might require a whole new programming paradigm, while others believe it will enable new technologies that can be used in more traditional HPC environments.

Constant price pressure and the presence of low budget solutions continue to put pressure on profit margins. Vendors need to look at where they can cut costs. Software is expensive to develop and maintain and the Open source movement is viewed as a way of sharing and thus cutting cost. People intensive tasks such as extensive benchmarking is another area where vendors are looking to cut cost.

2. HPC systems

Future system architecture actually has both similarities and large differences between the interviewed. From a very general point of view one could actually say that all share a view of future systems with:

Clustered SMP nodes giving a scalable parallel architectures.

The SMP nodes can be anything from a single processor up to thousands processors.

These systems are only limited in size by:

- available money
- cost of building space and utility (power)
- reliability

General purpose systems (but special purpose systems is not ruled out).

Non uniform memory access (but with very large variations in "NUMAness").

This is also reflected by a general agreement on MPI and OpenMP as the parallel programming models for the present as well as the future.

On the other hand though, underneath this very general common view almost everything else can differ: the processors, memory subsystems, node interconnects, compilers, operating systems, power consumption, cooling, square foot usage etc. and it equally fair to say that we really are talking about very different computer systems.
As a complement to general purpose systems, special purpose systems will emerge. These systems will be based on consumer product type of components, giving very high price/performance ratio but requiring very different programming models. Initially they will be targeted towards niche markets with suitable applications such as life sciences.

In addition to vendor specific architectures, low cost cluster HPC systems (Beowulf) has opened up a whole new sector for HPC. From being a rather limited solution it has matured significantly over the last years and continues to evolve, particularly in the high end. Most of the vendors claim that Beowulf is complimentary to their offering. Many vendors now also offer their own Beowulf clusters. It is yet unclear how this market will evolve and who the main players are going to be. It is likely however that Beowulf clusters will blend into the HPC world more and more seamlessly. Just as with traditional HPC we also believe there will be a multitude of various HW and SW configurations for Beowulfs.

3. Building blocks (processors , memory (bw, latency), network, storage, graphics etc.)

Silicon will continue to be used for chips for the next 5–10 years. Moore's law is believed to continue. Important to note however is that this is for the ``original'' Moore's law, the doubling of transistors every 18 months. Several questions can be derived from this fact: How are all new transistors going to be used? How is this translated into sustained performance improvements? Answers to these questions were expressed differently from different persons.

Number of processor architectures continue to decrease, but there is no general agreement as to if this means death to special purpose processors. High volume general purpose processors has to be very general purpose so performance for HPC applications will not be optimal. On the other hand is the question open if the cost for developing special purpose processors will be prohibitive.

Memory subsystems will continue to be a differentiator between vendors. Microprocessor performance will continue to increase much faster than memory performance. Latency is getting to the point where the latency (in terms of clock cycles) will actually start to increase. Instead latency reduction techniques such as pre- fetch, multi- threading, deep memory hierarchies etc. will be applied, and this is an area where many different techniques will be developed. Codes with regular memory patterns are likely to benefit more than others by many of the latency reduction techniques. Vectorisation is only actively pursued by two vendors but could also be viewed as a latency reduction technique. Bandwidth is the main focus for some vendors. It is likely to become one main differentiator between different systems and could replace the whole vector/no- vector debate.

Mass storage systems will increasingly be networked rather than directly attached. Disk density will double every year and tape density somewhat slower. I/O rates are not likely to keep up with capacity increase. No consensus if disks will overtake tapes for long term storage. Alternative mass storage solutions are still far from being commercially ready.

Graphics was touched on in most interviews and there is strong consensus that increased bandwidth and at the same time very limited increase in display resolution will change the way some visualisation is done. Centralised graphics servers will be used with thin clients for display only, even at large distances.

4. Parallel programming models

No one is predicting any major change in the parallel programming models, most vendors seems be of the opinion that whatever the hardware looks like it must be possible to program in either **MPI or OpenMP**. There are also some voices for co- array Fortran and UPC but is very unlikely that they will be as dominant. The hybrid MPI+OpenMP programming model is favoured by some but it remains very much unclear if this is going to be widely used. **HPF seems to to have a future in Japan only**.

5. Programming languages

No surprises here: **Fortran, C, C++ and maybe Java** continue to dominate.

6. Software tools

There seems to be a potential for a conflict when it comes to the necessary software tools for debugging and performance analysis of very large systems. Vendors are more and more shying away from developing the necessary SW due to cost, instead directing customers to commercial 3rd party software or open source. At the same time, even though the parallel programming models are the same across different platforms, the hardware underneath might be very different and thus requiring system specific tools to be able to fully understand and utilise the most out of the system.

7. Pain/gain

The pain versus gain is not going to be improved, probably get worse if not much worse to be able to tame the biggest and most powerful systems. Almdahl’s law is a stern task master, especially with processors counts reaching 10000’s. Maintaining very large systems will also require exceptional efforts. Housing, powering, maintaining and scheduling system with ten’s of thousands of processors is not for the faint of heart. It is no coincidence that public presentations of new very large systems spend a significant time talking about the building size, power required, pictures of cables etc.

8. Operating systems

Linux is the driving force around open source operating systems in the HPC world. No one can afford to stay outside of Linux. How HPC specific features are going to make it into Linux is not yet clear and it seems like the initial euphoria has settled down a bit when vendors are starting to realise that working with the Open source community is very different from doing your own development. One way or the other **we still believe that Linux will be the main HPC operating system 5 years from now.**

9. Price performance (peak vs sustained)

There is general agreement that price/performance will continue to improve at current rates, which is approximately Moore’s law. Important to note however that this is for peak speeds. When it comes to sustained performance, the picture is much more vague and most answers indicate that there will be an even stronger dependence on the type of application being used. Highly parallel applications with low bandwidth requirements will increasingly see a much better price/performance than traditional memory intensive HPC applications that are not trivially parallel.

10. Benchmarks

Almost unanimous agreement that benchmarking for customer procurements will have to change. The cost is too high for this in today's market where margins are getting smaller all the time. Requirements for better general benchmarks is there but the awareness of current efforts such as IDC's was low.

11. Grid

Most believed the Grid is still immature today but will become important in the future, but the opinions differed substantially in how and what way this will happen. Direct experience seemed low and one could argue whether the response might have been different if the people interviewed had more experience of the Grid. Interestingly to note was that no one thought that the Grid and Grid development would have any impact on future system architecture. It is rather something that their systems are used for and an opportunity to sell systems!

12. Future computer centres

HPC computer centres will have a place in the future according to the interviewed. Many different reasons were given but the fact that future very high performance will require even more expertise in various fields such as hardware, applications, performance analysis/tuning, logistics, procurements etc is reason enough to believe in a future for HPC centres. Application specific centres were predicted by many.

Our conclusions of the technology watch report

If anyone believed that future HPC systems all would move towards a unified computer architecture they will be disappointed. New advances in all relevant computer technology made possible by an increasing market as well as the Open source movement instead make the possibilities even larger than before.

We see five major trends for the next 5–10 years:

1. We believe that the increasing mass market for computer and consumer systems will provide the basis for a continued exponential technical development in all areas of computer systems development. This will mean many new and exciting possibilities also for HPC computer systems and we will continue to see a proliferation of systems instead of a convergence towards some common platform. A substantial part of future HPC systems will be provided by non- traditional HPC vendors. A major driving force behind this will be Linux clusters used for HPC. Off the shelf components means much lower profit margins which will create new and different business models than what we have today. New HPC vendors will appear, some will merge and some will disappear.

2. We see future HPC systems are parallel, scalable computing architectures that are based on the notion of clustered SMPs. Scaling is possible by either adding nodes and/or increasing the nodes themselves. A significant portion of the components are based on COTS but proprietary components play an important role. Operating systems is Unix with Linux as target. SMP nodes can be anything from a single processor up to thousands of processors and will have large variations in processors, compilers, memory

subsystems, node interconnect and mass storage. The upper limit in size is set by the prize, not the technology.

3. Peak performance as well as price/performance will continue to improve according to Moore's law. This is not true for sustained performance however where there will be an even more profound dependence than today between performance and application characteristics. Applications that easily can take advantage of advances in consumer product technology might even get a "super" Moore's law performance increase. Life sciences are in that category.

4. Parallel programming models and programming languages will be on an evolutionary path rather than a revolutionary. MPI and OpenMP for parallel programming and Fortran, C, C++ and Java for programming languages. This is good news for the user. The bad news is that to achieve high performance, detailed knowledge about the application as well as the underlying hardware is still required. Getting to extreme performance requires even larger efforts than today due to the sheer size of systems that will be available. Lack of adequate tools for software development will likely become a key issue for HPC.

5. The Grid is evolving fast and we strongly believe it will play a key role in the way HPC is used. In the context of this report, we see one new key role will be to serve as a layer between the user and the HPC computer system. Hiding details of not only geographical location but (maybe more importantly) system architectures and usage will in many cases enable the continued proliferation we see in HPC system architectures. We believe HPC vendors should consider this when developing new systems. This development should be very straightforward as long as we are talking about users of 3rd party codes and straightforward compile and run applications. Many high performance user however still needs to be able to be close to the system and this is an area where the Grid will take longer time to develop.

For HPC centres we believe that they will be needed but the rapid development of ever faster networks, Grid and Grid middleware, increased pain/gain ratio, increased system complexity and the increased infrastructure cost for logistics (power,housing) etc will mean fewer but larger centres. HPC centres will have to compete with other centres and to be successful we have the following recommendations:

- Collaborations and teaming up with other centres. Resource sharing.
- Expertise in many different areas of system architecture.
- Involved in open source development.
- Continued emphasis on MPI and OpenMP.
- Performance expertise, increasing. Vendor support decreasing due to leaner profit models.
- Expertise in Grid infrastructure, middle- ware.
- Application specific centres.

Molecular sciences

Overview

From the interviews to the computer companies representatives, the following scenario appears to characterise the near- future development of computer hardware and architectures:

Almost without exception, the experts interviewed concur that Moore's law will continue to describe accurately the pace of further chip integration and the resulting trend in theoretical peak performance of computers, at least for the next 5 years. Similarly, there is wide agreement on the (short-to-medium term) exponential increase in communication bandwidth. However, among the most important consequences of this for high performance computing is that memory — and more generally communication — latency as measured in CPU cycles will increase rapidly and impede most applications from reaping fully the benefits of Moore's law. The principal architectural features proposed to partly alleviate this problem aim at hiding latency behind an increased depth of data storage hierarchies and/or a moderate-to-high level of hardware multi-threading.

Increased transistor integration will also have the likely consequences that, on one hand, more room for extra logic will be available on-chip, allowing a widespread development of relatively low-cost special purpose processors. On the other hand, high integration will probably mean that not all transistor will be able to sync to one clock and multiple-clock chips will appear.

It is likely that, with various shifts of balance, we will continue to witness the development of both relatively low-cost Commercial-Off- The-Shelf (COTS) computers and proprietary hardware, the latter delivering higher performance at a much higher cost. SMP clustering will presumably be the dominant architecture. The total number of processors comprising a HPC machine will be in the 10000's, with consequent problems of reliability. These will have to be addressed by redundancy and statistical algorithms, eventually merging into true AI management systems.

Disk drives for I/O systems will likely continue to increase in capacity at a much faster pace than I/O bandwidth, so that the latter will become the limiting factor. Networked I/O systems, based on SAN, NAS and Internet SCSI will become widespread.
The most popular programming languages currently used (Fortran, C, C++, Java) will continue to dominate software development. Interestingly, it is generally agreed that the programmer's pain/gain ratio is destined to increase, signalling that compiler and RTE technology development will not keep pace with hardware complexity and speed.

It can be safely assumed that Unix-like operating systems will continue to dominate HPC. Linux will rapidly increase its market share, although proprietary OS will be phased out very slowly, if at all.

Based on the above assumptions we have projected two probably typical HPC computers 5 years from now, within the price range allotted, one exemplifying a proprietary SMP cluster and the other being the natural development of current Beowulf clusters. Based on the interviews and current trends, we have assumed that, within the projected time span, the price per processor will remain roughly unchanged.

Conclusions of the case study report

At the actual state of theory and algorithms, we can try to project the impact on applications in the field of Molecular Sciences that the three approaches to the definition of the inter-atomic potential outlined in the previous section will have in 5 to 10 years. In particular:

Model potentials based on Force-Field (FF) parameterisations will continue to dominate the scene in the field of liquid crystals, ferroelectric nematic materials [21] and protein-folding. Indeed, due to the large dimensions of the systems under study (up to a few million of atoms) and to the very long time-scale of the dynamical phenomena relevant in such fields (up to milliseconds), FF-based methods will still represent the only viable simulation tool up to 10 years.

In 5 (10) years, DFT methods will probably allow the accurate computations of electronic, structural and dynamical reactive properties of systems containing 3000 (10000) atoms. We can therefore predict that DFT methods will substitute FF parameterisations in chemiometric applications, in which a large number of medium-size calculations is needed. This will have a direct impact in pharmacology; indeed, the design of a new drug usually requires a pre-selection operated by computer simulations and data analysis; the advantage of a much higher accuracy in the description of the investigated molecular systems and properties directly translates into a high selectivity of the target system with a significant reduction of the number of laboratory tests, up to a factor of 10. DFT methods will also allow the accurate simulation of small protein systems, or of realistic portions of them, with particular impact on the comprehension of the action mechanism of metalloenzymes, where a reduced model usually neglects the fundamental underlying interactions. To understand the importance of such a field, it is sufficient to mention that both respiration and photosynthesis involve metallo-organic active centres constituted by several thousand atoms; comprehension of the action mechanism of such systems will allow to device efficient synthetic bio-mimetic analogues of the natural systems, with a high impact in the field of energy storage and molecular sensors. Moreover, we can predict that DFT-based methods will allow the accurate simulation of nano-scale systems with a high impact in the design of molecular engines, quantum computation devices and chemical storage of data [18].

Ab initio potentials will reach such a high standard accuracy (ca. 0.2 kcal/mol) so as to allow the realistic simulation of elementary reactions of relevance in the field of atmospheric chemistry (e.g. ozone depletion) [37] or combustion (e.g. nitrogen oxides chemistry) [46], integrating and partially replacing existing experimental techniques. Ab initio simulations will allow the accurate computation of reactive cross sections and rate constants of elementary systems composed up to 10–50 atoms, under extreme conditions (high temperature and pressure) on a quantitative basis; moreover, the study of elementary reactions in the interstellar space [16, 45, 34], e.g. the synthesis of organic compounds from small molecules and atoms (C, N, O, H) which is not usually directly accessible to the experimentalists, will be of great help in the design of new space-aircrafts materials and ultimately in the comprehension of the origin of life. However, according to the previous analysis, the computational resources required for such demanding applications can be achieved exclusively by large-scale computing facilities. In this view, the 12 MEuro platform option discussed above might represent only a starting point for the creation of a large transnational European super- computing resource.

PROJECT 3

Grid Enabling Technologies

http://www.epcc.ed.ac.uk/enacts/gridenabling.pdf

Stavros C. Farantos and **Stamatis Stamatiadis** (FORTH)
Institute of Electronic Structure and Laser
Foundation for Research and Technology, Hellas
And
Nello Nellari and **Djordje Maric** (ETH-CSCS)
Swiss Center for Scientific Computing

Study objectives

This study was accomplished in 2002. The main objective was to evaluate the current technologies for Grid computing as described in the "Grid Service Requirements" study and have been implemented in several projects and Grid testbeds around the world. Academic applications were mainly covered, with emphasis given in **Molecular Sciences**. This sectoral report was aiming:

- To briefly review the most popular software packages for implementing a Grid. Here we examine **Globus, Legion** and **UNICORE** since most applications are related to these technologies.
- To consider the current trend of development for Grid technologies and definition of standards.
- To locate the available test-beds for computational Grids worldwide, mainly focusing on those based on Globus, Legion and UNICORE.
- To review the current applications and projects with the established software packages for the computational Grids mentioned above.
- To investigate the present status of Grid Computing for Molecular Simulations and point out the needs for a new programming design of such applications.
- To report personal experience from building a local Grid based on Globus paradigm and UNICORE.

The reports consists of seven chapters.

After the **Introduction, Chapter 2** briefly reviews the most commonly used middleware for implementing a Grid: Globus, Legion and UNICORE. Their main characteristics and differences are emphasized and it considers next future developments for Grid technologies. In **Chapter 3** the authors collect in tables 5 to 7 test-beds of computational and data Grids as well as applications and projects and give short descriptions of their history and functions. This information, which is principally collected from the World Wide Web, is by no means complete. In the **Appendix A** the URL addresses are given which the interested reader can consult for more information. The types of applications and the categories of scientists to whom current Grids address are presented and it attempts to point out future needs. **Chapter 4** deals with the current status of Grids in Molecular Sciences- Physics, Chemistry, and Biology. **Chapters 5 and 6** present the personal experience of the authors on installing Globus and UNICORE in local computers. Small local computational Grids were constructed, which allowed to run a few applications. Finally, in **Chapter 7** the main conclusions of this study are summarized.

Summary and Conclusions

The study reported in this document is organized around six main points: a general analysis on the principal Grid software packages, the current development of Grid technologies, the location of computational Grid test- beds, the reviewing of the applications/projects running on those test- beds, the status of Grid computing for Molecular Simulations and a summary of our practical experiences on implementing a computational Grid with Globus and UNICORE.

The analysis on Grid software packages in Chapter 2 considers particularly Globus, Legion and UNICORE, since at the moment those products characterize most of the current Grid deployments and the most relevant Grid projects. This part aims to provide complementary information with respect to the first ENACTS report "Grid Service Requirements". Actually, the first ENACTS report offers an extensive presentation of available Grid technologies, solutions and products, and this sectoral report often refers to this document for the technical details and the presentation of less popular Grid products.

Currently the majority of projects and Grid initiatives are based on Globus. Moreover, Globus seams to be dominant also in Grid test- beds, especially in academic test- beds and Grid deployments built on Linux machines and Linux clusters. However, the impression is that the current phase corresponds to an early phase for Grid computing characterized by a difficult interoperability between different Grid technologies, a complex installation and configuration process for Grid products, and finally a very limited number of end users. At the moment there is a growing consensus around the OGSA proposal based on Web Services standards and specifications. Now, OGSA compliant Grid products are available, such as the Globus 3.x. The expectation is that the following phase of Grid computing, characterized by an agreement on a well- defined set of standards and specifications, will enable an easier interaction between different technologies and components, thus enabling the competition and possibly also the specialization of the various Grid products. The final objective is to provide a set of services to the user community for an effective, seamless and easy access to various, distributed and heterogeneous resources. The success of Grid computing will be measured on the benefits that a final user can get in a day by day work by using Grid technologies and products.

The lists presented in this report of the current test- beds and the applications/projects run on them by no means are supposed to be exhaustive. We know that several others exist which may have been just started. However, the following conclusions can be drawn from the internet search we have carried out:

11. Globus as a Grid deployment technology dominates in both sides of the Atlanti ocean.

22. Most of the applications concern particle physics, astronomy, environmental and meteorological projects and in less extent biological and chemical applications.

It is interesting to compare the above lists with a previous review the reference of which is given in the Introduction. The explosive growth of test- beds and applications in Grid computing is remarkable. It is worth mentioning the announcement of Mathematica for its new version suitable for a grid environment, the GridMathematica (http://www.wolfram.com/products/gridmathematica/). Mathematica is a popular software package for symbolic and numerical calculations as well as graphics.

Special attention was given for the Grid applications in the field of Molecular Sciences. Although, quantum chemistry has always kept a pioneer role in applying new computer technologies, it turns out that in the Grid case there is no much progress. This is because distributing computing requires new numerical algorithms and new ways of programming. At present, there are just a few examples which apply the established Grid deployments, but we believe in the near future we shall see them expanding. We may think that parallelized codes are ready for running in a Grid. This is not true since the employment of thousands or even millions of computers requires different programming approaches. At present only problems with trivial ('embarrassing') parallelization can be treated in a Grid.

The last two chapters of the document describe the experiences of the authors in installing Globus and UNICORE. Those chapters intend to give an idea how the models and concepts, expressed in the second chapter, are then implemented in reality. Therefore, the intention is to complement this study with practical experiences rather than provide a user guide for installing Globus or UNICOR.

PROJECT 4

Data Management in HPC
http://www.epcc.ed.ac.uk/enacts/datamanagement.pdf

Giovanni Erbacci, Marco Sbrighi (CINECA)
Inter- University Consortium

And
Audrey Crosbie, Carole McGloughlin (TCD)
Trinity College Dublin

Study objectives

The objectives of this study were to gain an understanding of the problems associated with storing, managing and extracting information from the large datasets increasingly being generated by computational scientists, and to identify the emerging technologies which could address these problems.
The principlal deliverable from this activity is a Sectoral Report which enables the participants to pool their knowledge on the latest data- mining, warehousing and assimilation methodologies. The document reports on user needs and investigate new developments in computational science. The report makes recommendations on how these technologies can meet the changing needs of Europe's Scientific Researchers.
The study focus on:

The increasing data storage, analysis, transfer, integration, mining, etc;
The needs of users;
The emerging technologies to address these increasing needs.

This study benefits the following groups:

Researchers involved in both advanced academic and industrial activities, who will be able to management their data more effectively;
The research centres which will be able to give better advice, produce more research and deliver better services for their capital investments.
European Research Programmes developing more international collaborations and improved support to European Research Area.

The report consists of eight chapters:

Chapter 1 ENACTS Project
Chapter 2 Basic technology for data management
Chapter 3 Data models and scientific data libraries
Chapter 4 Finding data and metadata
Chapter 5 Higher level projects involving complex data management
Chapter 6 Enabling technologies for higher level systems
Chapter 7 Analysis of Data Management Questionnaire
Chapter 8 Summary and Conclusions

Summary and conclusions

The objective of this study was to gain an understanding of the problems and reference the emerging solutions associated with storing, managing and extracting information from the large datasets increasingly being generated by computational scientists. These results are presented in this report which contains a state of the art overview of scientific data managemet tools, technologies, methodologies, on-going projects and European activities.
More precisely, the report provides an investigation and evaluation of current technologies. It explores new standards and supports their development. It suggests good practice for users and Centers, investigates platform-independent and distributed storage solutions and explores the use of different technologies in a coordinated fashion to a wide range of data-intensive applications domains. A survey was conducted to assess the problems associated with the management of large datasets and to assess the impact of the current hardware and software technologies on current Data Management Requirements. This report documents the users' needs and investigates new data management tools and techniques applicable to the service and support of computational scientists. The following main observations come from the analysis of the questionnaire;

- 60% of those who responded to the questionnaire stated that they perceived a real benefit from better data management activity.
- The majority of participants in the survey do not have access to sophisticated or high performance storage systems.
- Many of the computational scientists who answered the questionnaire are unaware of the evolving GRID technologies and the use of data management technologies within these groups is limited.
- For industry, security and reliability are stringent requirements for users of distributed data services.
- The survey identifies problems coming from interoperability limits, such as;
    - limited access to resources,
    - geographic separation,
    - site dependent access policies,
    - security assets and policies,
    - data format proliferation,
    - lack of bandwidth,
    - coordination,
    - standardising data formats.
- Resources integration problems arising from different physical and logical schema, such as relational data bases, structured and semi-structured data bases and owner defined formats.

The questionnaire shows that European researchers are some way behind in their take up of data management solutions. However, many good solutions seem to arise from European projects and GRID programmes in general. The European research community expect that they will benefit hugely from the results of these projects and more specifically in the demonstration of production based Grid projects. From the above observations and from the analysis of the emerging technologies available or under development for scientific data management community, we can suggest the following actions.

Recommendations

This report aims to make some recommendations on how technologies can meet the changing needs of Europe's Scientific Researchers. The first and most obvious point that is a clear gap exists in the complete lack of awareness of the participants for the available software. European researchers appear to be content with continuing with their own style of data storage, manipulation, etc. There is an urgent need to get the information out to them. This could be by use of dissemination, demonstration of available equipment within institutions and possibly if successful on a much greater scale, encourage more participation by computational scientists in GRID computing projects. Real demonstration and implementation of a GRID environment within multiple institutions would show the true benefits. Users' course and the availability of machines would have to be improved.

1. The Role of HPC Centres in the Future of Data Management

The HPC centres all over Europe will play an important role in the future of data management and GRID computing. Information on current state of the art and software should be available through these Centres. If information is not freely available then at a minimum these centres should direct researchers to where they will be able to obtain this information.
Dissemination and demonstrations would certainly be a good start to improving the awareness of researchers. Production of flyers and poster of the current state of the art, along with seminars, tutorials and conferences that would appeal to all involved in all areas of the scientific community. HPC centres can be seen the key to the nation's technological and economical success. Their role spans all of the computational sciences.

2. National Research Councils

National Research Councils play an important role in the research of the future. This report aims to make recommendations for 'national research councils' to address avoiding bottleneck in applications. The ENACTS reports endeavour to find current bottlenecks and eliminate them for the future researchers. This report introduces two national research councils, one from the UK and one from Denmark. These are used as examples to demonstrate promotion activities and how national bodies can encourage the use of new tools and techniques. The UK Research Council states that 'e-Science is about global collaboration in key areas of science and the next generation of infrastructure that will enable it.' Research Councils UK (RCUK) is a strategic partnership set up to champion science, engineering and technology supported by the seven UK Research Councils. Through RCUK, the Research Councils are working together to create a common framework for research, training and knowledge transfer. http://www.shef.ac.uk/cics/facilities/natrcon.htm
The Danish National Research Foundation is committed to funding unique research within the basic sciences, life sciences, technical sciences, social sciences and the humanities. The aim is to identify and support groups of scientists who based on international evaluation are able to create innovative and creative research environments of the highest international quality. http://www.dg.dk/english_objectives.html

3. Technologies and Standards

There is an increasing need by computational scientist to engage in data storage, data analysis, data transfer, integration of data and data mining. This report gives an overview of the emerging technologies that are being developed to address these increasing needs. The partners in ENACTS would support and encourage continued research and technology transfer in data management tools and techniques.

New Technologies

Discovery Net Project. The arrival of new disciplines (such as bioinformatics) and technologies will transform a data dump to knowledge and information. The Discovery Net Project aims to build the first e-Science platform for scientific discovery from the data generated by a wide variety of high throughput devices at

Imperial College of Science, UK. It is a multi-disciplinary project, serving application scientists from various fields including biology, combinatorial chemistry, renewable energy and geology. It is a service orientated computing model for knowledge discovery, allowing users to connect to and use data analysis software as well as data sources that are made available online by third parties. It defines standard architectures and tools, allowing scientists to plan manage share and execute complex knowledge discovery and data analysis procedures such as remote services. It allows service providers to publish and make available data mining and data analysis software components as services to be used in knowledge discovery procedures. It also allows data owners to provide interfaces and access to scientific databases, data store sensors and experimental results as services so that they can be integrated in knowledge discovery processes. http://ex.doc.ic.ac.uk/new/index.php

OGSA-DAI is playing an important role in the construction of middleware to assist with access and the integration of data from separate data sources via the grid. It is engaged with identifying the requirements, designing solutions and delivering software that will meet this purpose. The project was conceived by UK Database Task Force and is working closely with the Global Grid Forum DAIS-WG and the Globus Team. It is funded by DTI e-Science Grid Core Project involving: National e-Science Centre; ESNW; IBM; EPCC and ORACLE. http://www.ogsadai.org.uk/index.php

Data Standards

A push for the standardisation of data will increase the usability of the software that is currently available. There is an ongoing push to provide a standardised framework for metadata including binary data, such as the DFDL initiative. The DFDL (Data Format Description Language) is part of the Global Grid Forum initiative (GGF). http://www.epcc.ed.ac.uk/dfdl/. Currently DFDL is an informal email discussion group, providing a language to describe the way formats for metadata should be written. There is a need for a standardised unambiguous description of data. XML provides an essential mechanism for transferring data between services in an application and platform neutral format. It is not well suited to large datasets with repetitive structures, such as large arrays or tables. Furthermore, many legacy systems and valuable data sets exist that do not use the XML format. The aim of this working group is to define an XML-based language, the Data Format Description Language (DFDL), for describing the structure of binary and character encoded (ASCII/Unicode) files and data streams so that their format, structure, and metadata can be exposed. This effort specifically does not aim to create a generic data representation language. Rather, DFDL endeavours to describe existing formats in an actionable manner that makes the data in its current format accessible through generic mechanisms.
Data interoperability in of great importance especially within a GRID context. The iVDGL project is a global data grid that will serve as the forefront for experiments in both physics and astronomy. http://www.ivdgl.org/. Data interoperability is the sharing of data between unrelated data sources and multiple applications. Creating enterprise data warehouses or commerce websites from heterogeneous data sources are two of the most popular scenarios for Microsoft SQL as an interoperability platform. It preserves their investments in existing systems through easy data interoperability, while providing additional functionality and cost effectiveness that their existing database systems do not provide. It enables easy access of data and the exchange of data among groups.

Global File Systems

Traditional local file systems support a persistent name space by creating a mapping between blocks found on disk devices with a set of files, file names, and directories. These file systems view devices as local: devices are not shared so there is no need in the file system to enforce device sharing semantics. Instead, the focus is on aggressively caching and aggregating file system operations to improve performance by reducing the number of actual disk accesses required for each file system operation.

GFS

The Global File System (GFS) is a shared-device, cluster file system for Linux. GFS supports journaling and rapid recovery from client failures. Nodes within a GFS cluster physically share the same storage by means of Fibre Channel (FC) or shared SCSI devices. The file system appears to be local on each node and GFS synchronizes file access across the cluster. GFS is fully symmetric. In other words, all nodes are equal and there is no server which could be either a bottleneck or a single point of failure. GFS uses read and write caching while maintaining full UNIX file system semantics. To find out more please see http://www.aspsys.com/software/cluster/gfs_clustering.aspx

FedFS

There has been an increasing demand for better performance and availability in storage systems. In addition, as the amount of available storage becomes larger, and the access pattern more dynamic and diverse, the maintenance properties of the storage system have become as important as performance and availability. A loose clustering of the local file systems of the cluster nodes as an ad-hoc global file space to be used by a distributed application is defined. It is called the distributed file system architecture, a federated file system (FedFS). A federated file system is a *per-application global file naming* facility that the application can use to access files in the cluster in a location independent manner. FedFS also supports dynamic reconfiguration, dynamic load balancing through migration and recovery through replication. FedFS provides all these features on top of autonomous local file systems. A federated file system is created ad hoc, by each application, and its lifetime is limited to the lifetime of the distributed application. In fact, a federated file system is a convenience provided to a distributed application to access files of multiple local file systems across a cluster through a location-independent file naming. A location-independent global file naming enables FedFS to implement load balancing, migration and replication for increased availability and performance. http://discolab.rutgers.edu/fedfs/

Knowledge Management

e-DIKT (e-Science Data Information & Knowledge Transformation) is a project which applies solid software engineering techniques to leading edge computer science research to produce robust, scalable data management tools that enable new research areas in e- Science. E-DIKT has been funded through a Research Development Grant by the Scottish Higher Education Funding Council. E-DIKT will initially investigate the use of new database techniques in astronomy, bioinformatics, particle physics and in creating virtual global organisations using the new Open Grid Services Architecture (OGSA). E-DIKT's realm of enquiry will be at the Grid scale, the terabyte regime of data management, its goal to strain-test the computer science theories and techniques at this scale. Presently e-DIKT is looking into the following areas of research:

- Enabling interoperability and interchange of binary and XML data in astronomy
  - tools to provide “implicit XML” representation of pre-existing binary files;
- Enabling relational joins across terabyte-sized database tables;
- Testing new data replication tools for particle physics;
- Engineering industrial-strength middleware to support the data management needs of biologists and biochemists investigating protein

structure and function (as it relates to human disease and the development of drug therapies);

Building a data integration testbed using the Open Grid Services Architecture Data Access and Integration components being developed as part of the UK's core e-Science programme and the Globus-3 Toolkit.

Working over time with a wider range of scientific areas, it is anticipated that e-DIKT will develop generic spin-off technologies that may have commercial applications in Scotland and beyond in areas such as drug discovery, financial analysis and agricultural development. For this reason, a key component of the e-DIKT team will be a dedicated commercialisation manager who will push out the benefits of e-DIKT to industry and business. http://www.edikt.org/

Meeting the Users' Needs

One of the points addressed in the Users' Questionnaire was the ease of use of new data management tools. While a lot of researchers are not content with their own data management, they would not be willing to change unless it was an easy changeover. That is ease of use and quantity and quality of functions would be important issues for researchers when looking at migrating to new and improved systems. The ENACTS partners welcome the European initiatives and projects aiming to develop GRID computing and data management tools. However, this development must be focused at the end user and project results must be tested on real systems to enable applications research to benefit from migrating to new tools.

Future Developments

The principle deliverable from this activity is represented by this Sectoral Report which has enabled the participants to pool their knowledge on the latest data management technologies and methodologies. The document has focused on user needs and has investigated new developments in computational science. The study has identified:

The increasing data storage, analysis, transfer, integration, mining requirements being generated by European researchers;

The needs of researchers across different disciplines are diverse but can converge through some general policy and standard for data management;

There are emerging technologies to address these varied and increasing needs, but there needs to be a greater flow of information to a wider audience to educate the scientific community as to what is coming in GRID computing.

The study will benefit:

Researchers who will gain knowledge and guidance further enabling their research;

Research centres that will be better positioned to give advice and leadership.

European research programs in developing more international collaborations.

PROJECT 5

Distance Learning and Support
http://www.epcc.ed.ac.uk/enacts/DistanceLearning.pdf

Josef Novak, Miroslav Rozloznik, Miroslav Tuma,
Division HPC, Prague, Czech Republic (IHCC)
Institute of High Performance Computing Centre

Study objectives

The main task of this study was to describe and analyze the role of distance learning and support within the field of Grid computing. It often happens that Grid computing tools have a strong feedback on the strategies and efficiency of the distance learning. The original plan of the composition of the study was as follows. It had to be composed of the four main parts:

Part I: Presentation of the ENACTS project, of the Distance Learning and Support Project, the envisaged workplan, technical objectives and benefits. We also give a short description of participating organizations together with the list of research team members.
Part II: Various definitions of distance learning and education. Presentation of basic distance learning tools and models, including a discussion of their features, limitations, and benefits for prospective users. Although this text is relatively general, the authors want to focus on three main user groups which have been identified: HPC centers (service providers), Scientific Grid community (researchers and users) and Industrial Grid community (vendors and end- users). These three target groups are shortly discussed in the end of this section. This section was written by ICCC.
Part III: Results from a comprehensive survey distributed among 85 and undertaken among 25 major European research groups. It focuses on different aspects, such as the needs and requirements of various potential target groups and the pedagogical and organisational approach, which fits best with identified target groups. The survey includes a clear analysis how to ascertain the feasibility, viability and relevance of adapting a proper distance learning strategy to the training requirements and leads into a evaluation and agreement on a framework for collaborative development of suitable distance learning based course material. This section incorporated in the Final Report was written by I.C.C.C. in close collaboration with UNI- C.
Part IV: Conclusions and recommendations. The purpose of the survey presented in Part III is to gain a better understanding of these key user groups' needs and requirements in view of establishing a proper framework for distance learning and support. The analysis of this survey, together with Part II, which presented general concepts and technological issues, will be instrumental in establishing key recommendations for target groups. This section provides a summary of the distance learning features offered by the leading but still rather small groups of Grid specialists and users and it will make recommendations on a possible strategy that supports a successful uptake of Grid technology around larger communities. This section was written by ICCC.

Based on the results of the questionnaire, we have changed the structure as follows. Keeping Part I, Part III and Part IV as they were proposed, replacing Part II by a short introductory text on distance learning definitions and identified user groups. The basic reasons for this change were the following: evaluation of the results implies unexpected conclusions. While it seems that distance learning is relevant to Grid computing, it should be applied in a special form with which we deal below.

Summary and conclusions

This section contains three subsections that cover the following subjects:

1. Related projects and activities
2. Recommended distance learning concepts and tools offered by contemporary technologies
3. Key requirements and needs of users in the Grid community

The first subsection summarizes the relations to related projects and partners' efforts. We decided to put such a subsection here in order to see a few important facts: to see some projects dealing with the issue of distance learning from another perspective; to understand some of the concepts considered here in different frameworks and with different results; to motivate further investigations.
Other subsections are based on answers from Part III, their interpretation and derivation of general conclusions. The way the results were assembled was mentioned above. We believe that the proposed strategies and recommendations should help to establish a proper framework for distance learning and support. Consequently, they might contribute to the successful uptake of Grid technology even in larger communities in Europe and all over the world.

1. Related projects and activities

There are a number of projects and activities closely related to ENACTS. An important contribution will be probably played by the results of the LeGE working group whose main task is to facilitate the establishment of a European Learning Grid Infrastructure. This project is interesting for both its goals and for the way grid technologies attempt to achieve these goals (LeGE-WG: Learning Grid of Excellence Working Group; http://www.lege-wg.org). It clearly distinguishes between primary scientific and technological objectives on one side and operational objectives on the other. The former set of objectives is devoted to basic technological and pedagogical issues, legislative conditions, new European methodologies and standardisation of emerging Grid-aware solutions. The latter set of objectives deals with the practical steps to promote e-learning in fact, they include the whole European Higher Education, European Scientific and Engineering Research communities and the like. In a sense, LeGE might be considered as a very general metaproject with respect to our research.
An important dual view could be provided by a research project which would try to generalize our target groups horizontally. In other words, one may consider a general target groups of students, researchers or computational scientists. It is difficult to predict how important is distance learning for each group. An overview of their interests, needs and requirements can be found at ILIAS Opensource, University of Cologne (http:// www.ilias.uni-koeln.de/ios/index.html).
Another interesting project somewhat related to ENACTS is the EGEE project (Enabling Grids for E-science in Europe). This project was launched on April 1st, 2004 (http://public.eu-egee.org/). Its goal is to describe and interpret current national and regional Grid efforts. The project covers a large portion of the industrial partners.

2. Summary of recommended distance learning concepts and tools

As we have seen above, there is, predictably, a non-negligible interest in **acquiring scientific and technological information**. What we did not expected was that most of the users acquire this information by means of traditional sources. In particular, the users prefer standard research papers and non-research articles, manuals, booklets, hardware, software and software documentation. Yet, there is one modern feature. The ways to access are based on electronic tools (e-mail or internet). The web

environment is the new wrapper that includes mainly the classical sources of information as listed above. The new scientific and technological information can be extracted not only from a paper based agenda. Another important way to obtain it includes workshops, conferences, congresses. As far as the size of such meetings is concerned, there is typically a reasonable limit. It is well-known that meetings with larger number of participants are less effective in passing scientific information and for the learning process. Rather, they play a social role in the scientific society. They are important for celebrating important personalities, awarding prizes and the like.
We have investigated the role of the training type in the overall educational process. Based on the previous study, there are basically two important types of training: **informal and organised training.** The former is individualistic and tries to understand the subject matter from scratch. In such cases, for instance learning a programming language, it is important to start with reading examples and not reading user guides or manuals. These tools are more important later when the informal training transfers smoothly into an organised training. Users often really need to check carefully basic ideas of the new subject. The amount of time spent in these initial exercises which we might call a setup time is rather individual. Then the user might attend intensive courses on some software or hardware products, learn new ways how to cope with new communication tools and how to use grids.
If we try to point out the consequences of the basic user proposals, we can see that they may need a **peer-based** (**community facilitated resource)** rather than organised instructor courses. Such courses or training should come later when the user requires more advanced information. One example from previous times was a community of users of parallel computational tools. The individual access was much more important than an organised learning of various formalisms. Note that an organised training might be in this case not very effective since the techniques of parallel programming and using parallel machines are rapidly changing. In addition, parallel programming tools are typically very individualistic.
Users are interested in "**Open source**" approach to training and information resources and materials. They need to have access to databases in order to be able to **see closely related information.** They need to be able to have enough material in order to extend their knowledge in various directions. The next section will try to analyze the obtained results from the point of view of distance learning. Then a specific distance learning way will be proposed.

3. A proposed strategy and recommendations for establishing a proper framework for distance learning and support in Grid community

As we have seen, five basic features can characterize the mainstream of distance learning (conservative definition):

1. it is a highly structured activity
2. it deals with a highly structured content
3. it is typically a one-to-many process (teacher-centered process) in which the tutor plays really a key role
4. it is characterised by a frequent monitoring of its participants by tests, assignments etc.
5. it is based on combination of various web-based tools.

Consider now questions the users may ask. They may need large data resources that do not precisely correspond to the previous distance learning characterization. Although the activity might be considered as highly structured in both its form and content, it is **not teacher-centered** in the strictest sense of the word. The database-oriented learning serves more for information exchanges between two partners: those who create them and those who use them. Nevertheless, this database-oriented learning has a tendency to develop in the teacher-centered way. While the first encounters with the teaching in grid computing might be very unorganised, they have to change. They have to transform into a fully organized training.

We see another conclusion concerning the role of teachers and students in distance learning. It appears that the delayed form might be further developed by increased activity of the students. Let us try to explain this conclusion more carefully. There are many types of specific grid computations. Just now it would be very costly and inefficient to prepare very specialized experts in the field of Grid Computing that is changing so rapidly. Nowadays, it is more important to increase the general level of knowledge of Grid Computing in particular communities. Deep training of experts might be more useful once Grid Computing becomes a generally accepted form of computation. In such a future scenario, students will play a more active role in the learning process than current distance learning analysis based on the questionnaire suggests. In this context, let us distinguish two basic user groups: experienced users entering the new emerging field of grid computing and students who are getting their first qualification. One important difference between these users consists in their possibilities. While for an experienced user there is no difficulty in attending a couple of workshops per annum, this could be a problem for students because of a lack of funding. Consequently, the students are a very specific class of users with a much more open attitude to **distance learning techniques**.
Now let us take into account a user-friendly environment -- a reasonable compromise of users' demands for a flexible distance learning framework. We will call it a grid training portal. It is a customisable, personalised web interface for accessing services using distance learning and education tools. It would provide a common gateway to resources with special attention to the tools mentioned above. That is, **personalised and delayed distance learning forms** must be preferred.
We do not aim to present a fixed form of what we call the distance learning portal. Instead, we would like to summarise our conclusions and subsequent recommendations in terms of a **flexible** tool "under construction". Of course, it should reflect available technologies as well. We will describe its basic structure putting an emphasis on its most important features.

1. Control (management) unit
2. User resources
3. Information resources
4. Communication subsystem

Let us now describe these parts of the proposed portal system. The **control unit**, or more accurately, the **management** unit contains basic institutions and individuals jointly with tools that run the portal. Although there should be various specific rules how to handle the portal organisation, it is important to solve the problems of its technical updates, financial support, technological development, software upgrades etc. Some of them might need a rather sophisticated strategic decision. The control unit should contain two specific layers: **service (maintenance) layer** for implementing the control mechanisms and **evaluation layer.** One of the most important tasks of the control unit is to balance two basic functions. First, there will be a strong pressure of technological developments on hardware tools which will include both the **node demands** and the **network demands**. Technologically, these demands will present themselves in the need to make the nodes more powerful and to make the network with ever-larger bandwidth. Nevertheless, there is a strong gap between very fast processors and relatively slow connections. In other words, the network technology is lagging behind while increasing the network bandwidth is a real technological challenge. Second, the management unit must take care of financial resources including:

1. The start?up development costs,
2. The cost of the hardware and software tools connected to the portal including regular
updates and upgrades,
3. The operational cost of the technology,
4. The management cost and
5. The technology remediation costs.

The cost function must be carefully evaluated and balanced with the technological requirements that increase the overall portal cost. Once these two items are balanced, we should take care of the overall efficiency. Mr. Soren Nipper in [14] proposed and presented a picture showing the user groups on top of the pyramid with its large base corresponding to the overall costs. This might be a figure which is temporarily valid now.

However, it does not need necessarily correspond to the future development. Therefore, when taking into account such models, a realistic **forecast** has an important role. As of now, this figure might represent large startup costs since we are still in the **start**? **up period** of the Grid technology.

By **user resources** we mean the groups of portal users. Their leaders will be engaged in strategic decisions. The learning mechanism will not be strongly teacher- centered but more or less teacher- student balanced. The users will come with new initiatives in order to offer overall improvements. As far as the target groups are concerned, they should not be very large assuming the results of the questionnaire. On the other hand, we do not have an exact idea how they will develop in the future. Some hints, however, suggest that the target groups may increase, such as parallel computational tools two decades ago. After a long period of relatively small user groups, we can see large teams collaborating over a net on the development of large- scale HPC applications using distributed tools like SourceForge and dealing with powerful version of synchronising software that enables collaboration of tens of developers.

Information resources are the third part of the overall portal system. By that we mean the technological content of the portal covering both its hardware and software parts, particularly information databases with papers, lectures in written or recorded forms, simulation software and technological tools to present all these various materials. As of now we do not have very large multimedia resources for Grid Computing in our field. This will likely change in the future, however. In any case, the development of mechanisms how to store, protect, develop, update and clearly organise these data items is a more challenging problem. There will be a specific layer in this item. A specific feature of the information resources will be its hierarchical nature. In our case, the grid content will be first discussed on the level of HPC centers then on the national level (if there is some) and finally within the European grid community. In fact, this subsystem is exactly the one which must be organized hierarchically.

As far as the portal content of the information resources is concerned, it can be stratified into independent layers. The first layer may contain databases of written and electronically distributed information. The users accessing the portal have different needs. Therefore, the documents contained in its databases should be sorted out according to their requirements. Other data, such as test codes, video and audio material, should be stratified in a similar manner, thereby creating an user-friendly environment for gird users.

The final item is the **communication subsystem --** a mechanisms that of exchanges information among the three previous items. The communication patterns do not need to be uniform for the whole portal. In fact, various ways of communication can be supported. For some types of information exchanges, the synchronous connections, such as videoconferences, audio conferences or Access Grid, are preferable. Sometimes asynchronous mechanisms are preferred. In general, we distinguish two types of information signals: control and service ones. The first type serves for keeping the portal

in good shape and supporting its development. The second type (which should prevail by a large margin) serves the users.

While creating a portal is our main recommendation, the question remains open as to whether this is not something we had in mind even before evaluating the answers. In other words, we need to ask whether the answers do not hide some other possible and even completely different solutions. We have aimed at minimizing this hidden risk by our methodology in which we give to some standard technologies a new, well-defined content.

PROJECT 6

Software Reusability and Efficiency
http://www.epcc.ed.ac.uk/enacts/softwareefficiency.pdf

Jacko Koster, (Parralab, BCCS)
Parallab, Bergen Center for Computational Science,
University of Bergen (Norway)

Study objectives

The objective of this study was to determine the implications for new, and more importantly, existing programs of using a pan-European computing metacentre. Specifically, the study looks at how portability between facilities can be ensured and how this will affect program efficiency. The report consists of six chapters.

Chapter 1 Introduction to the ENACTS project.
Chapter 2 Software reusability
Chapter 3 Standards programming languages
Chapter 4 Software efficiency,
Chapter 5 Grid and distributed environments, middleware
Chapter 6 Summary and conclusions

Challenges and Potential Benefits. The uptake of computer simulation by new groups and the constant quest for greater efficiency of computer utilisation means that new techniques and approaches are always required. The study addresses emerging software technologies that address these challenges and how they affect the efficiency and reuse of existing applications. This study aims to review the current position. The study has two beneficiaries: the users who will be able to work more effectively, and the computer centres which will be able to give better advice and produce more research for their capital investments.

Scope of the study
During the study, we became aware of the difference in the way various software developers and development groups approach software reuse. For an academic researcher or research group, software reuse is typically the use of software that is in public domain, open source, or otherwise freely available to the researcher in some form. Reuse of software is often a way to acquire competence and improve and build upon methods, techniques and algorithms (and the software itself) that are developed elsewhere by colleagues or other research groups. Whether or not to reuse software is often a psychological and/or individual issue. In an industrial environment on the other hand, software reuse includes different concepts. Obviously, commercial considerations (like licensing, development cost estimates, market potential) play an important role, but also other issues like quality assurance, reliability, maintainability, and the support organization of the software are more critical. In all cases, software reuse aims at exploiting previously acquired competence and to reduce the development cost of new applications. In the report, we primarily address the HPC user community that is connected to the ENACTS consortium, and this is mostly academic. However, we touch upon commercial aspects of code reuse as well.
The concept of reuse and efficiency of software is closely related to many other concepts in software engineering, including (but not limited to) the following:

• software architecture
• end-user and application programming interfaces
• software life cycle (design, prototyping, implementation, verification, testing, and maintenance)
• software quality, reliability, fault-tolerance
• software development environments

It is beyond the scope of this study to review all these in detail. There is a vast amount of methodologies and literature and on-going research in these areas. However, during the study we found it impossible to avoid addressing some of the aspects related to these concepts. For the writing of this sectoral report, we have tried to stay as close as possible to the original objectives of the study (efficiency and reusability) and minimized the inclusion of other software engineering concepts.
Software reusability in community-led initiatives deals a lot with the design of standards for protocols and languages. We have not attempted to describe recent and emerging language standards in detail. Many of these are still under standardization and therefore subject to change in the near future. Describing the temporary status of evolving standards in detail would make this report obsolete in a relatively short time.

Summary and conclusion

In this report, we have reviewed some of the aspects related to efficiency and reusability in modern and emerging software technologies for computational science. We hope that the document has been useful to the reader in a variety of ways, for example

• to understand some of the major factors that impede software efficiency and successful software reuse in traditional HPC environments and in distributed and grid environments
• to recognize the relationship between software reuse and various other software engineering concepts and techniques
• to recognize the complexity related to software efficiency in modern (multi-level) hardware and software technologies
• to understand the importance of software standards and software interoperability and their impact on many facets of software design and application development and on computational science and engineering in general.
One objective of this study was to put forward recommendations for the establishment of a pan-European HPC centre with respect to software efficiency and reusability. We believe the following conclusions and recommendations are valid.

1. Moving large numbers of application codes between heterogeneous HPC systems or to new software environments is only feasible if well-accepted standards, languages and tools are available. In a pan-European HPC consortium, this will be facilitated if the user work environment is standardized in some form and best practices for software development and job execution are established. Potentially troublesome hardware-specific or site-specific dependencies should obviously be eliminated or be dealt with in a transparent manner.
2. Traditionally, data management connected to a numerical simulation was limited to the use of the programming bindings provided. The application typically ran on only one architecture and the local file systems would ensure that I/O to the application could be achieved by using the primitive bindings provided by programming languages like C and Fortran. The archival, retrieval, and exploration of remote (off-site) data is of growing importance as computer systems enable the rapid generation of extremely large, complex, and remote data sets. Increasing the ease and

efficiency of data transfer and access will greatly enhance the amount of science that can be performed on HPC facilities, and allows for inter- disciplinary sharing of data. An HPC centre or consortium should therefore promote best practices in portable data management towards its user community and provide appropriate advice and training.

3. Tools are important in enhancing the efficiency of the software engineering process and of the software itself, and permit rapid code debugging, performance profiling and analysis, and subsequent software re- engineering. HPC centres should actively keep users
up- to- date with the latest information on emerging tools, monitor their development and report on those that are deemed most effective and useful for the user community. Promising tools should be evaluated for robustness, functionality, and performance, and their applicability in real- life applications assessed. A tool that is considered mature enough and provides added value should be promoted to the user community of the HPC consortium by means of providing adequate training material.

4. An HPC user may not rapidly move onto the grid if this requires significant changes to his/her existing application. Ideally such move must be seamless and require little software revisions. Getting users onto the grid will require that end- user interfaces towards the grid middleware are simple or are similar to what the user uses in a traditional HPC environment. The computational scientist need not have to go through a lengthy learning process before he/she feels comfortable using grid technology. Ideally, the sheer technological complexity of the grid should be hidden to the user as much as possible. The HPC centres will have to make a serious effort in achieving this by choosing the right middleware technologies, provide friendly end- user interfaces (portals) to this middleware, and provide adequate training and support for the new methodologies.

An emerging problem in grid computing is the sheer amount of portals that are being developed world- wide. A grid portal may be constructed as a Web page interface to provide easy access to grid applications and provides user authentication, job submission, job monitoring, and results of the job. Many grid projects are developing their own interfaces to software packages that have a well- defined interface (e.g., the chemistry application Gaussian). There appears to be quite some overlapping (and also incompatible) activity in this context that could be avoided by merging portal development projects and reusing previously developed (generic) portal framework software.

5. Standards for grid services, web services are rapidly evolving. This is good in the sense that it shows a significant progress in (and need for) the methodology. On the other hand, it makes it harder for a user of the grid (not necessarily a code developer) to keep up- todate with the latest developments and keep the application ’alive’. Adapting to emerging web and grid standards may require regular or frequent changes to the grid- enabled application. This is undesirable for production software. An additional complication is that HPC centres typically do not support the same (recent) version of a rapidly evolving standard or tool. Moreover, different (e.g., vendor) implementations of evolving standards may not all support the latest changes to the standard definition. Regular synchronization between the HPC centres is needed in this context to reduce this kind of portability problems.

6. The use of open standards should be promoted actively. Open standards have several advantages, see Sections 2.5 and 3. Also the contribution by users or user groups to the standardization process of open standards should be promoted. Open standards are not only used in the area of high performance computing or grid computing, but also in many other areas of science, like for example visualization (OpenGL),

7. TheWeb has become the user interface of global business, and Web services now offer a strong foundation for software interoperability through the core open standards of XML, SOAP, WSDL, and UDDI. Models and applications that make use of this huge potential are just beginning to emerge.

It can be expected that Web services will play an increasingly more important role in integrating new initiatives in HPC consortia. Web services standards permit application- to- application interoperability, but the coordination of a set ofWeb services working towards a common end is still an open issue. Several XML-based protocols are under standardization to target specific needs of businesses and application domains. These standards clearly illustrate the momentum behind the Web services computing community.

The large majority of HPC user community is not currently familiar with Web services. We therefore strongly recommend that HPC centres provide good training programs for their users to get acquainted at an early stage with the basic concepts and tools in Web services and show the users the impact these services can have on their research activity.

8. Many of the new technologies require knowledge of programming languages like Java and C++, scripting languages like Perl, middlewares like Globus, and XML-based standards. Several of these are fairly new and evolving. Computational science groups and (senior) scientists that traditionally have used HPC in their activities are not always familiar with these recent languages and technologies. This lack of familiarity hinders the uptake of these new technologies by traditional HPC user groups. It is therefore important that HPC centres not only provide adequate training programs on the latest technologies but also on the basic components (and concepts) that underpin these technologies.

9. In its early days, grid computing was often thought of as a *computational* grid, an infrastructure that combines a set of distributed computing resources into one big computing resource on which one can run a large scale application that solves a computational problem that is too large to fit on any single machine. However, application efficiency on such a grid remains an issue. The interconnect between the individual resources can be slow, faulty, and insecure, and hence, the efficiency and reliability of the overall distributed application may not be what one would like. Moreover, these limitations in the interconnect may lead to under- utilization of the individual (but expensive) resources. These days, one can see a shift towards grid computing as being a *data* grid, an infrastructure for distributed data management that is transparent to the user and application. Such infrastructure greatly facilitates the reusability of existing applications. For example, a data grid allows the physical location of the data to be decoupled from the physical location of the application. The application need not be aware of this since the run time environment will make sure that the data will be transferred to a place where the application expects the data to be. The data grid provides mechanisms that form the glue between remote applications, devices that generate data, and data bases, and thus enables the creation of smart coupled applications.

Another emerging form of grid computing is the *collaborative* grid. Such a grid enables the creation of virtual organizations in which remote research groups can perform joint research and share data.

10. The use of grid technologies will eventually lead to more remote (distant) collaborations. It is therefore essential that mechanisms are in place for code maintenance by a large research group with multiple programmers modifying a single code. It is our experience that many scientists have little knowledge of the wealth of tools available to assist in this. An HPC centre or consortium should be committed to provide advice and training in code development and maintenance; application of these techniques leads directly to increased productivity and enhanced code portability. This will become even more apparent when grid- enabling technology will have become mature enough to become an acceptd tool for remote (global) collaboration.

Finally, we note that several newly established European projects in Framework VI address some of the issues addressed in this report. These include HPC- EUROPA [33] and DEISA [24].

PROJECT 7

Grid Metacenter Demostrator

Demonstrating a European Metacentre: Feasibility Report
http://www.epcc.ed.ac.uk/enacts/demonstrator.pdf

Chris Johnson, Jean-Christophe Desplat (EPCC)
Edinburgh Parallel Computing Centre
and
Jacko Koster, Jan-Frode Myklebust (BCCS)
Parallab, Bergen Center for Computational Science,
University of Bergen (Norway)
and
Geoff Bradley (TCD)
Trinity College Dublin

Study objectives

Here we describe the ENACTS Demonstrator activity itself, starting with the objective stated in the Technical Annex [ENA]. We then go on to explain how the objectives evolved as new trends within the user community were followed. We then give a description of the deliverables expected. We also describe the way in which this "Demonstrato " activity fits into the ENACTS project as a whole.

Objective [Demonstrating a European Metacentre]: *To draw together the results from all of the Phase I technology studies and evaluate their practical consequences for operating a pan-European metacentre and constructing a best-practice model for collaborative working amongst facilities.*
Since the time at which the objective was written, trends and technology have moved on and a fast and unpredictable pace. In particular, large scale scientific communities have necessarily become more concerned with handling the large amounts of data they produce, rather than simply concerning themselves simply with getting the most out of compute cycles on large machines. For this reason, we concentrated on the data aspects of the "pan-European metacentre". The intention being to demonstrate that "The Grid" is able to solve many of the problems of data-sharing across what are becoming known as "virtual organisations", something which will become increasingly important over the coming years (see [EU]).
The Demonstrator activity began on 30th June 2003 following on from a kick-off meeting in Dublin earlier in 2003 and involved three European Partners, previously described:
Centre Role Skills & Interests
EPCC leading activity Particle Physics & Globus
Parallab participating Physics & Globus
TCD participating & providing users Physics & Globus

In contrast to the earlier ENACTS activities, the main deliverables of the Demonstrator activity do not consist of reports, but of the actual demonstrator itself. This document describes the work done during the ENACTS Demonstrator activity. It is intended primarily as a "Feasibility" report for those interested in setting up a pan-European metacentre based on our findings in setting up such a centre

Summary and conclusions

Successes

This demonstration project was successful in a number of areas:

- A data Grid has been successfully deployed on three clusters/supercomputers with QCDgrid running across the Grid.
- An XML Catalogue is operating with an XML Schema to describe the MILC metadata, which is stored in an eXist database.
- The MILC code has been altered to produce machine independent XDR output.
- The MILC code has been altered to readin/writeout XML datafiles.
- The test users have certificates and accessed the data Grid.
- All components of the Grid have been demonstrated to the users.

Problem areas

The users had conceptual problems with the Grid and understanding the purpose of the Metacentre. Their limited knowledge of Grid technologies automatically led them to believe that they were gaining access to a computational Grid as opposed to a data Grid.

User feedback

The users were two QCD scientists in the School of Maths, Trinity College Dublin.
Q. How familiar were you with Grid technologies prior to this project ?
A. Only aware of a user working on this type of project at Edinburgh University.
Q. Would you be more/less likely to get involved in a Grid project now ?
A. It would be nice to use such resources, but I think it would need a huge number of nodes before it would be useful as other groups would also need access to the machine. We would rather run on local machines with guaranteed resources.
Q. What functionalities were missing from the Grid test- bed ?
A. Ideally, the Grid would have one central node that allocates jobs depending on the load of each cluster/supercomputer. Effectively the Grid should appear as one large cluster. Additional web based documentation of resources etc. required (Note: this was provided).
Q. Would access to this Grid improve your productivity/efficiency and how would it alter your work practices ?
A. Having loads of CPUs is helpful, but it would only really be usable if I didn't have to go searching for idle machines. A well organised file server would be useful

Q. Would you engage in collaborative activities more readily via the Grid ? Would you share datafiles, results, etc. ?
A. Yes !

Additional user feedback

The following additional user feedback was provided

- Having machine independent data is useful.
- It is useful to have the binary data accessible on all sites.
- The metadata for each jobs is very useful because if you're sharing data with other users you can easily find out the parameters they used to generate their results.
- Automatic upload of XML & binary output after run has completed would be useful. The MILC code was modified to do this.

It would be desirable to have the entire jobs submission process available through a web interface.

A global queueing system, global ‘qsub’ would be very useful.

The large XML Schema used for MILC has many blank fields for this application which is confusing and messy.

The whole Grid system seems to be a bit delicate !

Where’s the computational Grid ?

In the longer term it would be better to take software written by the TCD QCD group and Grid enable it.

The whole Certificate issue was very frustrating !

Feasibility

Our ENACTS “Demonstrator” activity has shown that it is feasible to set up a datagrid across geographically dispersed sites using available technology provided that considerable effort is first put in to set up Globus and it’s associated packages. Given more time we could have increased the number of users to our system which would have provided us with more feedback.

One could also ask the question *“How would such a Grid scale up to use by more sites and users?* There does not appear to be an obvious limitation of the number of sites or number of users within QCDgrid itself although in the future it would be useful have more control over the ownership of files stored in the Grid. For example, it is not presently possible for users to delete files once they are on the Grid - this has to be done by the administrators.

The “demonstrator” itself was intended to take present technology and evaluate it, and we believe this has been successful. The activity has also motivated those involved in the QCDgrid project to generalise the software for use beyond their initial UK Grid. Readers are encouraged to try the software for themselves.

PROJECT 8

Survey of Users' Needs
http://www.epcc.ed.ac.uk/enacts/userrequirements.pdf

Giovanni Erbacci and Claudio Gheller (CINECA)
Inter- University Consortium
and
Satu Torikka (CSC)
Centre for Scientific Computing

Study objectives

The present report aims at presenting the results of the "Survey of Users' Needs" scientific/technological activity, whose objective are to determine users' requirements for access to HPC facilities and Datastores and assess the implications for changes in their working patterns if these were provided within a metacentre model. Furthermore, the report should outlines how users perceive emerging technology affecting their research and looking at the technological barriers to mobility of researchers. The results are based on the opinions of both large user groups and of individual users of high performance and distributed computing facilities in Europe.

Work Plan
One of the key tasks for ENACTS is to collect information from existing and potential users of High Performance and distributed computing about their future requirements. While much of the qualitative, in- depth information on requirements can be collected via ENACTS participants (both HPC centres and user representatives), ENACTS also aims to collect quantitative information from a wider range of groups via a web based questionnaire. This will enable the Network to check the requirements of a far wider cross- section of the computational science community than would otherwise be
possible.
The questionnaire (see the details from section 2) has been designed to gain information in a range of areas including the value placed on current services, the limitations and applications bottlenecks, the level of user experience and expertise and future requirements.
In addition, detailed information have been collected from key users groups, represented or identified by ENACTS participants, by means of an interview. Each user group was asked the same series of open ended questions about their requirements.

Main Tasks
The "Survey of Users' Needs" activity consists in five workpackages, totalling 6.6 staff- months of effort. The elapsed time for the activity is 6 months. The workpackages are summarised below.

WP1 User survey 2.0 CINECA, CSC Network Co- ordinator
WP2 Results analysis 1.0 CINECA, CSC Network Co- ordinator
WP3 In- depth interviews 2.0 CINECA, CSC Network Co- ordinator
WP4 User requirements report 0.5 CINECA, CSC Network Co- ordinator
WP5 Dissemination 0.6 CINECA, CSC Network Co- ordinator

WP6 Project Management 0.5 CINECA

WP1: This is a survey of the current experience and requirements of users and potential users of HPC facilities. It involves the design and promotion of a web-based questionnaire for users of European LSFs and HPC centres. This workpackage comprises 2 staff-months of effort and has been performed by CINECA and CSC. The other HPC centres in the ENACTS network has promoted the questionnaire to their users and contact organisations.

WP2: The questionnaire returns have been analysed and data-mined to look for significant trends. This workpackage has taken 1 month of effort and will be undertaken by CINECA and CSC.

WP3: Eleven in-depth phone or face-to-face interviews have been conducted with representatives of significant computational science research groups or organisations in Europe, to solicit their views and opinions on future requirements for HPC. This activity has been led by CINECA and CSC, with assistance also from other ENACTS participants, who have identified target interview groups and conducted interviews. This workpackage took 2 months of effort.

WP4: The output from the analysis of the user survey and the completed in-depth interviews has been used to produce the User Requirements report (this report). This report details the differing requirements of significant user groups in Europe and summarises user requirements for Grid Computing. CINECA and CSC have written this report. The workpackage had a duration of 0.5 staff-months.

The report consists of 7 sections

- Chapter 1 introduces the ENACTS project and the objectives of the specific activity.
  Chapter 2 describes the structure and the content of the questionnaire. Furthermore it presents the data collection and analysis procedure adopted.
  Chapter 3, 4 and 5 present the details of the various sections of the questionnaire and the results obtained for each section.
  Chapter 6 presents the result of in-depth interviews with several selected users.
  Chapter 7 draw the conclusions and propose some recommendations and suggestions for the future development of users needs driven research tools.

Summary and conclusions

In this report we have presented the results of a research which aims at analysing the requirements and desiderata of European HPC users with respect to high-end computing resources, applications and data management tools. Moreover the survey outlines how users perceive emerging technologies and how these can affect their research/development work.
The research is based on a questionnaire dedicated to users of HPC facilities and datastores and Grid users, in particular researchers and scientific software developers who require medium/large computing resources, so that they can provide a meaningful feedback on high-end CPU devices and infrastructures.
Eleven in-depth interviews to selected representatives of computational science research groups or organisations in Europe have been collected and analysed in order to provide a wider perspective vision of how technological resources should evolve in order to fulfil the expectation of the research community.

Conclusions on the user questionnaire

The questionnaire was answered by 125 users from eighteen European countries, mostly university researchers, representatives of a much larger scientific community. Most of the participants are part of medium (2-10 members) or large (11-50 members) research groups. It must be noticed that the majority of these collaborations are local; only a third of the research groups account for international collaborations. This is mainly due to the difficulty of having everyday remote working sessions. The collaborative environments technologies can be crucial to improve the diffusion of transnational working groups and the spread of skills and knowledge across Europe.

The average present-time HPC user which emerges from the Questionnaire Part 2 shows that there is still a "traditional" approach to computer sciences. Most of the participants use local resources (workstations, departmental servers) with small (1 to 4 processors) or medium (8 to 32 processors) configurations. The main concerns are related to the speed of a single CPU and the memory size, rather than having plenty of distributed resources. The access to the computing platforms is mostly via ssh connection, rather than more sophisticated methods like web portals. This can be due mostly to security concerns. The most common operating systems are Linux and various proprietary Unix flavours – Aix, Irix etc. Windows is getting a growing success, even though it is still quite little diffused in the research community. Other products like MacOS are quite uncommon. A large fraction of the researchers computing related work is dedicated to code development. Commercial or freeware codes are not common in the scientific community. This is due both to the high specificity of many problems, which require dedicated algorithms and codes, and to certain scepticism towards commercial software. Self-made or self-modified codes, starting from previously home-made programs, are the common choice for most of the researchers. Very little space is given to commercial or freeware applications. Specialised scientific libraries are frequently used, since they are highly optimised, precise and accurate tools to perform standard tasks (like array and linear algebra operations, FFT etc.). A traditional approach to numerics is once more confirmed by the choice of Fortran and Fortran 90 as programming languages. However, also other high-performance languages, in particular C and C++, are starting to diffuse, emphasising the openness toward the experimentation of different and new technological opportunities. This is confirmed also by the interest toward open source products, which usually do not represent completely stable and easy to use tools, but can provide the basic components to develop new applications.

The data seem not to be a major concern for researchers. Most of the applications are computing intensive, but the amount of data that are produced is rapidly growing. However their storage, management and even analysis are considered as a secondary problem. Results are usually stored in files, with no particular organisation and often with no standard format (like CGNS, FITS, HDF). Usually either raw binary files or ASCII tables are used to save data. This can represent a strong limitation for data exchange even inside the same research group. But, first of all, it can be a critical challenge for the standardisation and interoperability effort of the international community and, more generally, for diffusion of the knowledge, cooperation and best exploitation of resources.

The crucial role of the computing power emerges also from the analysis of the users' interest and involvement in Grid related issues. In fact, the main attractive is to have a large aggregate computing power available, with much less concern of its architecture. This is confirmed also from the other points the users have indicated as the most critical in a distributed system, that are high network bandwidth, principally to download data, and efficient schedulers, to maximise the throughput of the workflow. Other opportunities, like portals, shared file systems on distributed platforms, large storage capacity, data management via databases, collaborative working session, are perceived as much less interesting and attractive. The Grid, since it is still in a development phase, is seen as an unfriendly environment, mainly due to lack of stable API, incomplete documentation, difficulties in the management and in the development of suitable applications. Nevertheless, a large fraction of the users is willing to contribute to the Grid infrastructure development, if proper help and support is provided. It is also interesting to notice that more than 27% of the users have already been involved in Grid related research projects. Finally, it is very encouraging that a large majority of the users is willing to share their codes and, under proper conditions, their data and results. However, those who do not want to share resources justify their choice with the fact that these are too specific of their work and therefore useless for the community.

Conclusions on the in-depth interviews

The in-depth interviews have outlined the following results.

The need for more CPU power was emphasised by most representatives. Some of them considered the need for more storage capacity as important as CPU power, and some saw that data storage and data management have become more important than CPU power. The importance of data management tools is likely to increase and the development of data analysis, data mining and visualisation tools will continue. The future "winning" HPC architecture for the scientific community will be clusters for the throughput volume computing. Clusters, possibly combined with Grid access, will comprise common, relatively cheap processors and efficient networks between processors. Also the evolution in the software area will be towards clusters distancing from platform-specific applications. Shared memory systems, efficient parallel machines especially, will remain as solutions for selected sub-communities who need
peak performance. In some answers the price of the shared memory systems and vector machines was considered too high for the future systems. The amount of data will increase in the future. Many research groups will use huge data sets stored in distributed storages. Easy and reliable access to the data sets is important and it requires international collaboration. Technological solutions have to be developed to handle very big data sets.

The Grid is here to stay and developing according to all the interviewed, but opinions of its development vary. Grids can be useful for the scientific community, or Grids may remain as limited solutions used by traditional high performance computing and people needing manageable huge data sets distributed over several places. For a successful future of Grid infrastructures, a lot of development is needed to make the Grids more reliable and easy enough to use. There are interesting pilot projects going on all around EU. The interviewed agree that the scientific community is willing to accept the sharing of knowledge, resources, tools and data, because it is the prerequisite for having a successful Grid architecture. However, there will be both open and commercial software tools available. Databases will be open to a large extent, but commercial interests may restrict the access to data. The real challenge is collaboration on both the political and technical level. More agreements will be needed how to use the data and resources.

Virtual organisations, distributed workshops (e.g. Access Grid), or distance learning are already useful tools in educational and training activities. Today, the technology is still difficult and requires usually a set of non-standard equipment and software and special skills to operate. In order for these tools to be more useful for the scientific community, both technological and human resources are needed, e.g. to have easy access from your desktop. These tools can enable faster, cheaper and more environmentally friendly communication.

The developments of cluster computing and data storage solutions are seen as the most important expectations of the scientific community with respect to the near future technological development. The HPC centres will have an important role in helping the scientific communities build cluster solutions and giving users advanced application support, e.g. in code optimisation and Grid usage. HPC centres can become the backbone of a true Grid infrastructure offering the scientific community an easy and sound access to HPC and datastores.

Closing remarks

The user requirements, representing a common view of the HPC users' community and the representatives of significant computational science research groups or organisations in Europe, can be summarized as follows:

The crucial role of the computing power is the key issue in the involvement in Grid related issues. It is very encouraging that a large majority of the users is willing to share knowledge, tools, data and results. The main challenges Grid computing will face will be political; more agreements will be needed how to use the data and resources in multinational Grid infrastructure.

The Grid is in development stage and users are willing to contribute to the Grid infrastructure development, if proper help and support is provided. Virtual tools for education and training could enable faster, cheaper and more environmentally friendly communication. To enhance the adaptation of virtual tools, these tools should be more user- friendly.

The growing amount of data that are produced puts challenges. Standardisation and interoperability efforts of the international community are needed for diffusion of the knowledge, cooperation and best exploitation of resources.

APPENDIX III

LIFELONG Distance Learning and Support Study in the ENACTS Project

Josef Novák, I.C.C.C. Group, a.s, Miroslav Rozložník, *I.C.C.C. Group, a.s.,*
Miroslav T?ma, *I.C.C.C. Group, a.s*

Abstract: - The paper is devoted to the joint scientific and technological study on distance learning and support within the ENACTS project, which is a network of leading institutions from around Europe offering High Performance Computing systems and services together with key developments in the area of Grid computing. The first part describes the ENACTS project with an emphasis to the methodology for evaluation of distance learning techniques and support used in applications related to Grid Computing. The second part of the paper presents the main results achieved in the study on current experience with Distance learning in Grid Computing community. We will emphasize especially the lifelong aspects of distance learning.

1 The ENACTS Project

The ENACTS Project is a Co-operation Network in the 'Improving Human Potential Access to Research Infrastructures' Programme. It is running since 2000. This Infrastructure Co-operation Network brings together High Performance Computing (HPC) Large Scale Facilities (LSF) funded by the DGXII's IHP programme and key user groups. Its aim has been to evaluate future trends in the way that computational science would be performed and to cover the pan-European implications as well. As a part of the Network's remit, it runs a Round Table to monitor and advise the operation of the four IHP LSFs in this area, EPCC (UK), CESCA-CEPBA (Spain), CINECA (Italy), and BCPL-Parallab (Norway).

This co-operation network follows on from the successful Framework IV Concerted Action (DIRECT: ERBFMECT970094) [1] and brings together many of the key players from around Europe who offer a rich diversity of High Performance Computing (HPC) systems and services. In ENACTS, our strategy involves close co-operation at a pan-European level – to review service provision and distil best-practice, to monitor users' changing requirements for value-added services, and to track technological advances. In HPC the key developments are in the area of Grid computing and are driven by large US programmes.

In Europe we urgently need to evaluate the status and likely impacts of these technologies in order to move us towards our goal of European Grid computing, a 'virtual infrastructure' - where each researcher, regardless of nationality or geographical location, has access to the best resources and can conduct collaborative research with top quality scientific and technological support. ENACTS provides participants with a co-operative structure within which to review the impact of Grid computing technologies, enabling them to formulate a strategy for increasing the quantity and quality of access provided.

The principal objective of the project is to enable the formation of a pan-European HPC metacentre. Achieving this goal requires both capital investment and a careful study of the software and support implications for users and HPC centers. A part of the latter goal is the core objective of this study. The project is organised in two phases. A set of six studies of key enabling technologies has been undertaken during the first phase:

1. **Grid service requirements (EPCC, PSNC)**
2. **the roadmap for HPC (NSC, CSCISM);**
3. **Grid enabling technologies (ETH-Zurich, Forth);**
4. **data management and assimilation (CINECA, TCD);**
5. **distance learning and support (ICCC, UNI-C);**
6. **software efficiency and reusability (UPC, UiB).**

2 Joint Scientific and Technological Study on Distance Learning and Support

This section describes basic issues of a study devoted particularly to distance learning and support. Objective starting points (sources) of this study were: users accessing facilities through a pan- European metacentre, typically, require training and support from remote centers in order to make best use of the available facilities. The WWW-based technologies are emerging to accomplish this, but are largely untried outside intranets. Here, ENACTS aims to determine the most appropriate support and training methods and the enabling technologies. The advent of Grid computing will make it ever more likely that users will be using facilities remotely. This means that the same networked technologies must be used to provide training and support. Currently, there is little standardization in the technologies used to develop training courses and none in the area of distributed support.

Let us describe now the technical objectives and conditions for this study. They are implied by the necessity to agree on a framework for collaborative development of distance learning based course material. The objective is to make it easy for participants (and other institutions in Europe) to develop or customise re- usable training material. The widespread availability of distance learning material will increase the accessibility of HPC systems. One of the aims of the metacentre is to make training in appropriate tools and techniques available to researchers who are remote from the facilities they are accessing. The provision of an appropriate framework for course development makes this more feasible.

Distance learning is thus of interest to all infrastructure operators and research groups in Europe, but the cost of developing and maintaining training material has discouraged most centers from committing time and effort to it. Collaborative projects offer the potential for sharing the costs and the effort, but it is vital to select an appropriate development environment. The result of this research is a set of papers and focused studies reviewing the state of the tools, the standards and the methodology will be of practical benefit.

The definition of distance education may seem straightforward enough, but there is an ongoing debate as to what is involved in the process and concept of distance education. Glenn Hoyle's Distance Learning on the Net (http://www.hoyle.com/distance.htm) provides a list of definitions of Distance Learning from various sources. His own summary is: *"Distance Learning is a general term used to cover the broad range of teaching and learning events in which the student is separated (at a distance) from the instructor, or other fellow learners."*

There are many different distant learning definitions. They will be discussed in the talk. With few notable exceptions, the actual beneficiaries and users of tomorrow's Grid technology have not yet established a dialogue regarding standardization of technologies and tools used in education. Remarkably, formal remote training and distance learning are not rated highly in our survey. Yet, it is appropriate to raise a question about what would be the most relevant and effective distance learning method in the context of European and international Grid communities.

3 Grids and Distance Learning: Results and Recommendations

3.1 Peer- based versus organized courses

This subsection is devoted to basic overview of our main results. First, in the Grid community there is a non- negligible interest in **acquiring scientific and technological information**. What we did not expected was that most of the users acquire this information by means of traditional sources. In particular, the users prefer standard research papers and non- research articles, manuals, booklets, hardware, software and software documentation. Yet, there is one modern feature. The ways to access are based on electronic tools (e- mail or internet). The web environment is the new wrapper that includes mainly the classical sources of information as listed above.

The new scientific and technological information can be extracted not only from a paper-based agenda. Another important way to obtain it includes workshops, conferences, congresses. As far as the size of such meetings is concerned, there is typically a reasonable limit. It is well-known that meetings with larger number of participants are less effective in passing scientific information and for the learning process. Rather, they play a social role in the scientific society. They are important for celebrating important personalities, awarding prizes and the like.

We have investigated the role of the training type in the overall educational process. Based on the previous study, there are basically two important types of training: **informal and organised training.** The former is individualistic and tries to understand the subject matter from scratch. In such cases, for instance learning a programming language, it is important to start with reading examples and not reading user guides or manuals. These tools are more important later when the informal training transfers smoothly into an organised training. Users often really need to check carefully basic ideas of the new subject. The amount of time spent in these initial exercises which we might call a setup time is rather individual. Then the user might attend intensive courses on some software or hardware products, learn new ways how to cope with new communication tools and how to use grids.

If we try to point out the consequences of the basic user proposals, we can see that they may need a **peer-based (community facilitated resource)** rather than organised instructor courses. The latter courses or training should come later when the user requires more advanced information. One example from previous times was a community of users of parallel computational tools. The individual access was much more important than an organised learning of various formalisms. Note that an organised training might be in this case not very effective since the techniques of parallel programming and using parallel machines are rapidly changing. In addition, parallel programming tools are typically very individualistic.

Users are interested in "**Open source**" approach to training and information resources and materials. They need to have access to databases in order to be able to **see closely related information.** They need to be able to have enough material in order to extend their knowledge in various directions.

The next paragraphs will try to analyze the obtained results from the point of view of standard distance learning. Then a specific distance learning way will be proposed. As we have seen, five basic features can characterize the mainstream of distance learning (conservative definition):

1. **it is a highly structured activity**
2. **it deals with a highly structured content**
3. **it is typically a one-to-many process (teacher-centered process) in which the tutor plays really a key role**
4. **it is characterised by a frequent monitoring of its participants by tests, assignments etc.**
5. **it is based on combination of various web-based tools.**

3.2 Role of teacher in distance learning for grid computing

Consider now questions the users may ask. They may need large data resources that do not precisely correspond to the previous distance learning characterization. Although the activity might be considered as highly structured in both its form and content, it is **not teacher-centered** in the strictest sense of the word. The database-oriented learning serves more for information exchanges between two partners: those who create them and those who use them. Nevertheless, this database-oriented learning has a tendency to develop in the teacher-centered way. While the first encounters with the teaching in grid computing might be very unorganised, they have to change. They have to transform into a fully organized training.

We see another conclusion concerning the role of teachers and students in distance learning. It appears that the delayed form might be further developed by increased activity of the students. Let us try to explain this conclusion more carefully. There are many types of specific grid computations. Just now it would be very costly and inefficient to prepare very specialized experts in the field of Grid Computing that is changing so rapidly. Nowadays, it is more important to increase the general level of knowledge of Grid Computing in particular communities. Deep training of experts might be more useful once Grid Computing becomes a generally accepted form of computation. In such a future scenario, students will play a more active role in the learning process than current distance learning analysis based on the questionnaire suggests.

In this context, let us distinguish two basic user groups: experienced users entering the new emerging field of grid computing and students who are getting their first qualification. One important difference between these users consists in their possibilities. While for an experienced user there is no difficulty in attending a couple of workshops per annum, this could be a problem for students because of a lack of funding. Consequently, the students are a very specific class of users with a much more open attitude to **distance learning techniques**.

Now let us take into account a user-friendly environment -- a reasonable compromise of users' demands for a flexible distance learning framework. We will call it a grid training portal. It is a customisable, personalised web interface for accessing services using distance learning and education tools. It would provide a common gateway to resources with special attention to the tools mentioned above. That is, **personalised and delayed distance learning forms** must be preferred.

3.3 Longlife Distance Learning Portal for Grid Users

In this subsection we do not aim to present a fixed form of what we call the distance learning portal which we mentioned also above. Instead, we would like to summarise our conclusions and subsequent recommendations in terms of a **flexible** tool "under construction". Of course, this tool should reflect available technologies as well. We will describe its basic structure putting an emphasis on its most important features.

Control (management) unit
User resources
Information resources
Communication subsystem

Let us now describe these parts of the proposed portal system. The **control unit**, or more accurately, the **management** unit contains basic institutions and individuals jointly with tools that run the portal. Although there should be various specific rules how to handle the portal organisation, it is important to solve the problems of its technical updates, financial support, technological development, software upgrades etc. Some of them might need a rather sophisticated strategic decision. The control unit should contain two specific layers: **service (maintenance) layer** for implementing the control mechanisms and **evaluation layer.** One of the most important tasks of the control unit is to balance two basic functions. First, there will be a strong pressure of technological developments on hardware tools which will include both the **node demands** and the **network demands**. Technologically, these demands will present themselves in the need to make the nodes more powerful and to make the network with ever-larger bandwidth. Nevertheless, there is a strong gap between very fast processors and relatively slow connections. In other words, the network technology is lagging behind while increasing the network bandwidth is a real technological challenge. Second, the management unit must take care of financial resources including:

1. **Start-up development costs**
2. **Cost of the hardware and software tools connected to the portal including regular updates and upgrades**
3. **Operational cost of the technology**
4. **Management cost**
5. **Technology remediation**

The cost function must be carefully evaluated and balanced with the technological requirements that increase the overall portal cost. Once these two items are balanced, we should take care of the overall efficiency. Mr. Soren Nipper in [14] proposed and presented a picture showing the user groups on top of the pyramid with its large base corresponding to the overall costs. This might be a figure which is temporarily valid now. However, it does not need necessarily correspond to the future development. Therefore, when taking into account such models, a realistic **forecast** has an important role. As of now, this figure might represent large startup costs since we are still in the **start-up period** of the Grid technology.

By **user resources** we mean the groups of portal users. Their leaders will be engaged in strategic decisions. The learning mechanism will not be strongly teacher-centered but more or less teacher-student balanced. The users will come with new initiatives in order to offer overall improvements. As far as the target groups are concerned, they should not be very large assuming the results of the questionnaire. On the other hand, we do not have an exact idea how they will develop in the future. Some hints, however, suggest that the target groups may increase, such as parallel computational tools two decades ago. After a long period of relatively small user groups, we can see large teams collaborating over a net on the development of large- scale HPC applications using distributed tools like SourceForge and dealing with powerful version of synchronising software that enables collaboration of tens of developers.

Information resources are the third part of the overall portal system. By that we mean the technological content of the portal covering both its hardware and software parts, particularly information databases with papers, lectures in written or recorded forms, simulation software and technological tools to present all these various materials. As of now we do not have very large multimedia resources for Grid Computing in our field. This will likely change in the future, however. In any case, the development of mechanisms how to store, protect, develop, update and clearly organise these data items is a more challenging problem. There will be a specific layer in this item. A specific feature of the information resources will be its hierarchical nature. In our case, the grid content will be first discussed on the level of HPC centers then on the national level (if there is some) and finally within the European grid community. In fact, this subsystem is exactly the one which must be organized hierarchically.

As far as the portal content of the information resources is concerned, it can be stratified into independent layers. The first layer may contain databases of written and electronically distributed information. The users accessing the portal have different needs. Therefore, the documents contained in its databases should be sorted out according to their requirements. Other data, such as test codes, video and audio material, should be stratified in a similar manner, thereby creating an user- friendly environment for Grid users.

The final item is the **communication subsystem --** a mechanisms that of exchanges information among the three previous items. The communication patterns do not need to be uniform for the whole portal. In fact, various ways of communication can be supported. For some types of information exchanges, the synchronous connections, such as videoconferences, audio conferences or Access Grid, are preferable. Sometimes asynchronous mechanisms are preferred. In general, we distinguish two types of information signals: control and service ones. The first type serves for keeping the portal in good shape and supporting its development. The second type (which should prevail by a large margin) serves the users.

4 Conclusions

In the four sections of our paper we summarized the basic points of the distance learning- related part of the ENACTS project of the EU. After describing its basic goal we devoted to the methodology and results. Embedding of our results into the distance learning and support framework is also interesting from the theoretical point of view. On the other hand, creating of the learning portal seems to be very practical implication of our research. In this paper we covered the basic issues related to its form and contents. While creating a portal is our main recommendation, the question remains open as to whether this is not something we had in mind even before we started the research. In other words, we need to ask whether this answer do not hide some other possible and even completely different solution. We have aimed at minimizing this hidden risk by our methodology in which we give to some standard technologies a new, well- defined content.

References:

[1] Glenn Hoyle: Distance Learning On the Net. (http://www.hoyle.com/distance.htm).
[2] Distance Education Clearinghouse
(http://www.uwex.edu/disted), The University of Wisconsin-Extension, 2003.
[3] Distance Learning Resource Network (http://www.dlrn.org), Star Schools Program, U.S. Dept. Of Education, 2000.
[4] Virginia Steiner: "What is Distance Education?", DLRN Research Associate, 1995. (http://www.dlrn.org/text/library/dl/whatis.html).
[5] Michael Moore, Greg Kearsley: *Distance Education: A Systems View*, Wadsworth Publishing Company, 1996.
[6] Instructional Technology Council, Washington DC, 2003. (http://144.162.197.250).
[7] The Institute for Distance Education, University of Maryland University College,

(http://www.umuc.edu/ide/modlmenu.html).
[8] Distance Education: A Consumer's Guide, Western Cooperative for educational telecommunications, Pub. 2A300, 12 pp., 1997.
(http://www.wcet.info/resources/publications/conguide/index.htm).
[9] Derek Stockley, Chirnside Park, Melbourne, Victoria, Australia, 1996-2003.
(http://derekstockley.com.au/elearning-definition.html)
[10] I. Stanchev, E. Niemi, N. Mileva. Teacher's Guide *"How to develop Open and Distance Learning Materials"*, University of Twente, The Netherlands.
[11] Kit Logan and Pete Thomas. Learning Styles in Distance Education tudents Learning to Program. Proceedings of 14th Workshop of the Psychology of Programming Interest Group, Brunel University, June 2002, pp. 29-44. (www.ppig.org)
[12] Karen Bradford: Deep and Surface Approaches to Learning and the Strategic Approach to Study in Higher Education; Based on Phenomenographic Research. (http://www.arasite.org/guestkb.htm).
[13] The Centre for Teaching and Learning, University College, Dublin
(http://www.ucd.ie/teaching/good/deep.htm)
[14] Soren Nipper: talk in Dublin, annual meeting of the Enacts project, 2003.
[15] (www.online-colleges-courses-degrees-classes.org/q-and-a-online-accredited-college.html
[16] VOD – The Online ITV Dictionary (http://www.itvdictionary.com/vod.html)
[17] Institute of Educational Technology at The Open University (http://iet.open.ac.uk/courseonline/)
[18] Stanford Online and the Stanford Center for Professional Development (www.scpd.stanford.edu)
[19] TheAccess Grid Project(http://www.accessgrid.org)
[20] The European Access Grid
(http://euroag.accessgrid.org)
[21] EPCC's Access Grid Node
(http://www.epcc.ed.ac.uk/computing/grid/accessgrid/)
[22] (http://tecfa.unige.ch/edu-comp/edu-ws94/contrib/peraya.fm.html)
[23] Interwise - enterprise communications platform for Web conferencing (http://www.interwise.com)
[24] Guide to Web Conferencing, Online Conferencing, e-Conferencing, Data Conferencing...
(http://www.thinkofit.com/webconf/).
[25] http://publish.uwo.ca/maandrus/Table.htm
Blackboard.com (http://www.blackboard.com).
[26] WebCT (http://www.webct.com).
[27] Fle3. Learning Environments for Progressive Inquiry Research Groups, UIAH Media Lab, University of Art and Design Helsinki (http://fle3.uiah.fi)
[28] ILIAS Opensource, University of Cologne (http:// www.ilias.uni-koeln.de/ios/index.html)
[29] COSE VLE, Staffordshire University Enterprises Ltd. (http://www.staffs.ac.uk/COSE).
[30] The Centre for Educational Technology Interoperability Standards (CETIS)
(http://www.cetis.ac.uk/).
[31] LeGE-WG: Learning Grid of Excellence Working Group. (http://www.lege-wg.org).
[32] Distance learning projects in EU.
(http://www.know-2.org/index.cfm?PID=1).

APPENDIX IV

LIST OF ACRONYMS

ACM Association for Computing Machinery
AI Artificial Intelligence
API Application Program Interface
ASP Application Service Provider
ATM Asynchronous Transfer Mode
B2B Business- to- business
BEOWULF Cluster of PCs connected by a fast network
BLAS Basic Linear Algebra Subprograms
BSP Bulk Synchronous Parallel computing
CAVE CAVE Automatic Virtual Environment
CI Configuration Interaction
CMOS Complementary Metal Oxide Semiconductor
COTS Commercial Off- The- Shelf
CP Car- Parrinello
CPU Central Processing Unit
CRM Customer Relationship Management
CSCISM Center for High Performance Computing in Molecular Sciences
DAS Direct Attached Storage
DFT Density Functional Theory
DRAM Dynamic Random Access Memory
DRMAA Distributed Resource Management Application API
DSP Digital Signal Processing
DTF Distributed Terascale Facility (IBM HPC installation in the US)
DWD Deutche Wetterdienst
EJB Enterprise JavaBeans
ENACTS European Network for Advanced Computing Technology for Science
EPCC Edinburgh Parallel Computing Centre
ERP Enterprise Resource Planning
FAQ Frequently Asked Questions
FeRAM Ferrorelectric RAM
FF Force Field
FFT Fast Fourier Transform
FeDFS Federated file system
FM Fast Multipole
FP- CMOS Flexible Parameter CMOS
FPGA Field Programmable Gate Array
GFS The Global File System
GGF Global Grid Forum
GSN Gigabyte System Network
GT2 Globus Toolkit 2
GUPS Giga Updates Per Second
HPC High Performance Computing
HPF High Performance Fortran
HSM Hierarchical Storage Management
IA- 64 Intel Architecture 64 bit
IDC International Data Corporation
IEEE Institute of Electrical and Electronics Engineers
IP Internet Protocol
ISV Independent Software Vendor
ITRS International Technology Road- map for Semiconductors
J2EE Java2 Enterprise Edition
JSP Java Server Pages
JXTA JuXTApose

LAN Local Area Network
LAPACK Linear Algebra PACKage
LTC Linux Technology Center (at IBM)
MCSCF Multi- Configuration Self- Consistent Field
MD Molecular Dynamics
MPI Message Passing Interface
MPP Multiple Parallel Processing
MRAM Magnetoresistive RAM
MTA Multi Threaded Architecture (Cray HPC system)
NAS Network Attached Storage
NAVi Network Animated View
NSA National Security Agency
NSC National Supercomputer Centre
NUMA Non- Uniform Memory Access
OGSA Open Grid Service Architecture
OpenMP Open Multi Processing
P2P Peer- to- peer
PBLAS Parallel Basic Linear Algebra Subprograms
PC Personal Computer
PP Pseudo Potential
PW Plane Wave
QOS Quality Of Service
RAM Random Access Memory
RFP Request For Proposal
RSL Resource Specification Language
RTE Run Time Environment
SAN Storage Area Network
SCSI Small Computer System Interface
SCSL Source Code Software Licensing
SHMEM Shared Memory (access library)
SIA Semiconductor Industry Association
SMP Symmetric Multi Processing
SPEC Standard Performance Evaluation Corporation
SPP Special Purpose Processor
Sun ONE Sun Open Net Environment
TCO Total Cost of Ownership
TCP Transmission Control Protocol
TSMC Taiwan Semiconductor Manufacturing Company
UMA Uniform Memory Access
UPC Unified Parallel C
W3C World Wide Web Consortium
WAN Wide Area Network
WDM Wavelength Division Multiplexing
XC Exchange Correlation
XML Extensible Markup Language
ZPL Z (level) Programming Language