ERCIM NEWS

European Research Consortium for Informatics and Mathematics
www.ercim.org

Number 45

April 2001

Special Theme:

GRIDS: e-Science to e-Business

Next Issue:
July 2001

Next Special Theme:
Human Computer Interfaces

CONTENTS

Cover image courtesy ESA. Grid techniques can improve extraction and processing of data provided from satellites (see page 11).

**Lord Sainsbury of Turville,
UK Minister for Science and Innovation.**

From April 2001, the UK will be investing an additional £ 118 Million over three years into e-science, in addition to the funds for upgrading the research network. This is a considerable sum of money – which begs the question "Why?"

A short history lesson brings the first clue. In 1989, a computer system was developed which has fundamentally changed the way business is conducted around the world. Supply chains are now linked electronically; products are designed by international teams on different continents; goods and services are ordered online. This one system, the World Wide Web, has saved companies billions of pounds and has generated billions more in new business. It has also fundamentally altered people's personal lives, with numbers having internet access at home growing by the day.

So was the World Wide Web created by a businessman spotting this huge potential market opportunity? As I am sure all readers of ERCIM News are well aware, the answer to this is an emphatic "no". It was created at CERN, the European Particle Physics Laboratory near Geneva, as a way of sharing information between scientists in different parts of the world. But the story does not end there. Seeing the potential of the Web, the creators went around Europe trying to sell the idea to European industry. They were unsuccessful. Then they went to the United States, and were welcomed with open arms.

In the UK we are determined to learn the lessons from the Web – firstly, that scientists with large IT problems develop technologies which are later of great potential industrial benefit, and secondly that European industry has to be involved from the outset in any new major developments. This view is held not only by UK; it is widespread throughout Europe.

It is our belief that a new major development is right now on the horizon. Scientists in a number of disciplines are coming across limits in their ability to process, communicate, store and access exponentially increasing amounts of data. In tackling these problems, there are the beginnings of a new form of information technology – known as the Grid – which has the potential to have as profound an effect on the world, if not more, as the Web.

The Grid concept derives its name from the analogy to a power grid. When you switch on a kettle to make a cup of tea (which we do a lot here in the UK!), you don't worry about which power station the electricity is coming from, you just plug it in and the electricity arrives. Contrast that with information technology, where you have to specifically search through different databases (which are not compatible with each other), travel to high performance computing facilities, and know exactly where on the system everything is and how best to use it.

With the ever increasing amounts of data being generated by current experiments, scientists are looking towards Grid computing as a solution to several problems. Imagine a world where all databases are interoperable, where huge amounts of data can be transported around the world, where computing power is available "on tap" when it is needed, without the user knowing where any of these resources are but simply paying a service fee. The implications for science are immense. It will give us a potential mechanism to use huge emerging data sources, such as the decoded human genome, in multi-variable analyses with other data elsewhere to develop cures for diseases such as cancer. It will allow international research on very large data sets, such as that from CERN's Large Hadron Collider, which will open in 2005. And it will greatly enhance the ability to mine data again and again by comparing existing data sets collected for one purpose with new and previously unrelated information, so generating new knowledge.

But now imagine the potential business implications of the same technology. Ubiquitous computing, with information sources available uniformly and processing power provided as a service. Search engines which will not only search but look in several locations for data, then look for available computing power, arrange data processing and communication of the results back to the enquirer. Real linkage of the supply chain and all the IT facilities it contains. If the Grid takes off, it has all the potential to change the world even more than the Web has done in the past decade.

The UK wants to take a leading role in this new revolution, both delivering solutions for our scientists and generating opportunities for our companies. We want to work together with partners across Europe and the world to meet this vision. There is a huge amount of activity in Europe, but it needs some coherence, across countries and disciplines, if we are really to make the vision a reality.

The will is there. The people are there. The opportunity is there. In the UK, the money and national political commitment are there. Other European countries, and the European Commission, are thinking along the same lines. Now is the time for the revolution.

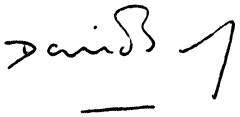

SCHOLNET and CYCLADES: Two New Digital Library Projects

by Donatella Castelli

The intensive activity of the ERCIM Digital Library Working Group (1996 – 2000), coordinated by IEI-CNR, has contributed greatly to the growing presence of DL projects and applications within the Fifth Framework Programme of the European Commission. In addition to the DELOS Network of Excellence for Digital Libraries, the ERCIM Working Group has also provided the framework for the setting up of two innovative IST 5FP funded DL projects: SCHOLNET and CYCLADES.

SCHOLNET and CYCLADES, both coordinated scientifically by IEI-CNR, Pisa, were presented at the EU-DL All Projects Concertation meeting held in Luxembourg in February (see following article). Both projects are directed towards satisfying the needs of scholars and scholarly communities. However, they view this issue from different perspectives: the goal of SCHOLNET is to develop a digital library testbed providing a set of specialised services to meet the requirements of networked scholarly communities, while CYCLADES is focussed on the need to develop a service environment on top of large heterogenous and multidisciplinary interoperable archives.

SCHOLNET

The aim of SCHOLNET is to build a test-bed able to provide not only the traditional digital library services but also support for non-textual data types, hypermedia annotation, cross-language search and retrieval, and personalised information dissemination. The testbed will enable the immediate dissemination and accessibility of technical documentation (and the underlying ideas) within a globally distributed multilingual community. It will also contribute to the creation and diffusion of a new model for scholarly production and use by providing functionality to permit annotation on digital objects in any format by authorised users, to support personalised information dissemination, and to access federated repositories of related material. The enhanced digital library infrastructure should produce significant benefits for a scholarly community, providing it with additional credibility and visibility and encouraging its expansion.

From the architectural point of view, the project will build on and extend the functionality provided by an existing digital library for grey literature: the ERCIM Technical Reference Digital Library (ETRDL) (http://www.ercim.org/delos). This digital library has been developed by ERCIM Institutions to provide their scientists and librarians with on-line facilities to manage their technical documentation and make it electronically accessible worldwide. SCHOLNET will augment the ETRDL test-bed functionality by adding the following services:

- Multimedia data support: the repositories will contain not only textual documents but more generic multimedia documents, in particular videos of tutorials, seminars and lectures, training sessions, demos.
- Hypermedia annotation support: annotation and reference linking features will be fully integrated into the digital library infrastructure.
- Multilingual search and retrieval: functionality will be provided for monolingual querying in English (the common language) or in any of the languages used by the institutions participating in the project (currently Dutch, French, German, Greek, Hungarian, Italian and Swedish). Cross-language search facilities will also be provided.
- Automatic personalised information dissemination support: the test-bed will support a pro-active facility to alert users, on the basis of their system-maintained profiles, of potentially interesting new documents.

An extensive experimentation activity will be conducted in order to assess the degree to which SCHOLNET meets the requirements of a large scholarly community. In particular, this activity will involve the ERCIM Working Groups communities. However, the SCHOLNET infrastructure has been designed for use by scholarly communities working in any domain.

CYCLADES

The main objective of CYCLADES is to support scholars both individually and as members of networked communities when interacting with large interdisciplinary electronic (e-print) archives. Such archives are important vehicles for the dissemination of preliminary results and non-peer reviewed 'grey literature'. Most focus on information dissemination within disciplinary or institutional communities. However, scientific research is now oriented towards an interdisciplinary approach. Scientists thus need to easily retrieve information from diverse sources, and to communicate and collaborate across traditional community boundaries. CYCLADES aims at supporting the transition of e-print systems into genuine building blocks of a transformed scholarly communication model by developing a set of leading edge technologies providing innovative methods for information access, dissemination, sharing and collaborative work.

CYCLADES will thus build an open archives environment consisting of two main architectural components: the archives and the services. The former will be developed by the Open Archives Initiative (http://www.openarchives.org). This US initiative aims at guaranteeing interoperability among e-print archives. It has established a set of relatively simple but potentially powerful interoperability specifications that facilitate the development of services implemented by third parties.

The CYCLADES consortium has established a co-operation agreement with the Open Archives Initiative. This agreement guarantees technical co-operation between the two initiatives and permits the adoption of the Open

Archives Initiative environment as the CYCLADES archive environment.

CYCLADES will develop a leading service environment compliant with the Open Archives Initiative specifications. In particular, a core set of cross-archive value-added services will be developed to constitute a federation of independent but interoperable services. The main CYCLADES services will be:

- Mechanisms for dynamically building meaningful virtual collections (from the prospective of a given community) on top of the archive environment
- Browsing mechanisms based on the concept of multilevel hypertext to deal with heterogeneous archive schemas
- Models for representing and acquiring user information needs (user profiling)
- Algorithms for automatic categorisation
- Models for collaborative recommendation/filtering
- Tools to support knowledge sharing communities.

Both projects had their kick-off meetings at the beginning of this year and both are scheduled to last for 30 months. It will be interesting to compare their activities and results as they progress.

Links:
http://www.ercim.org/scholnet
http://www.iei.pi.cnr.it/cyclades

Please contact:
Donatella Castelli
E-mail: castelli@iei.pi.cnr.it
Tel: +39 050 3152902

EU-DL All Projects Concertation Meeting

by Vittore Casarosa

A concertation meeting among the digital library projects funded by the European Union, held in Luxembourg on 7-8 February 2001, was attended by more than 60 people. The meeting was organized by the DELOS Network of Excellence, in cooperation with the Cultural Heritage Applications Unit of the EC Information Society Technologies programme.

Planned as part of a series of meetings to bring together all the projects funded by the EC IST Programme in the field of Digital Libraries, as well as representatives of other relevant initiatives in Europe and the US, the objectives were to exchange information about the projects, to identify areas of synergy and possible cooperation, to jointly promote standardization and dissemination activities, and to provide to the EU IST Programme an appreciation and a global view on the evolution of Digital Libraries technologies.

The meeting was opened by Bernard Smith, head of the Cultural Heritage Applications Unit, who briefly presented the major programs and initiatives of the European Commission in the field of Digital Libraries. Then each participating project gave a short presentation on the main objectives of the project, the technologies developed/used, results achieved, and areas of interest for potential cooperation. The following projects were presented: PRESTO, ECHO, BRAVA, AMICITIA, SCHEMAS, RENARDUS, COVAX, MIND, SCHOLNET, CYCLADES, ARTISTE, COLLATE, ARION, LEAF, ONE-II, ETB. In addition, three presentations were given by representatives of DL initiatives in the US and Europe. Howard Wactlar (representing the US NSF Digital Library Initiative), Rudi Schmiede (representing the German DL Initiative DL-Forum, formerly Global Info), and Rachel Heery (representing the UK DL Initiative UKOLN) gave presentations on the history, projects, scope and goals of their Digital Libraries Programmes. In particular, Howard Wactlar stressed that international cooperation has recently become a major, rather than peripheral, motivating goal for the NSF Digital Library Initiative.

EU-DL All projects concertation meeting.

Following a request from the Commission, Dieter W. Fellner, (Technische Universität Braunschweig, Institut für Computergrafik) summarized the main research themes and challenges which emerged during the first day, for presentation and discussion the following day. In the morning of the second day, two parallel workshop sessions were held, one on 'Building Digital Library Portals with Harvested Metadata' chaired by Thomas Baker (GMD) and another on 'Evaluation of Digital Libraries' chaired by Norbert Fuhr (University of Dortmund). During each of the sessions there was also a discussion on the main emerging research themes, led by Dieter Fellner.

The session on Metadata examined methods for building digital library portals with harvested metadata. Simple, standards-based metadata descriptions in

XML or RDF could in principle be harvested by a diversity of gateways and portals for multiple purposes (general and domain-specific) and at many levels (international, national and intra-organizational). After a brief introduction by Thomas Baker (GMD), Marianne Peereboom (Koninklijke Bibliotheek) described the Renardus Project, which seeks to integrate access to subject gateways by normalizing local metadata to a shared schema. Rachel Heery and Manjula Patel (UKOLN) described how such schemas can be published in a machine-understandable RDF format in the SCHEMAS Forum Registry. Mariano Sanz (ETB project) explained how the 'push' model of the European Schoolnet – mapping local metadata records to a common schema in RDF for inclusion in a central repository – gives participating repositories the full control to disclose their metadata selectively. Carl Lagoze (Cornell University) described how the Open Archives Initiative has adopted a technologically simple template and protocol allowing repositories to publish their metadata for selective harvesting in a broad range of services, from searching and summarisation to current awareness and reference linking. The point was made in discussion that the harvesting model is relevant not just to digital libraries of scholarly content, but to the global scalability of commercial applications.

The session on Evaluation of Digital Libraries focused on test beds and evaluation methods for DLs. Pasquale Savino (IEI-CNR) presented the evaluation approach taken in the ECHO project, where the test collection consists of more than 200 hours of documentary film, and where the final evaluation will focus on user satisfaction. Daniel Faensen (Free University Berlin) talked about requirements for a DL test bed in general, emphasizing the need for the consideration of domain-specific issues, heterogeneous collections and structured documents. Lars Edinger (DBC) described a testbed for interoperability testing of Z39.50 services. Istar Buscher (SWR) presented the Amicitia project, where a distributed video collection is used for end-user testing. Leona Carpenter (UKOLN) presented some aspects of the Renardus project, which focus on usability in relation to task model. Howard Wactlar (CMU) gave a brief statement concerning his experiences with evaluation during different US-DLI projects: evaluation criteria should include system performance, usability, interoperability, and user satisfaction; parameters to be considered are collection size and granularity of metadata.

Norbert Fuhr (University of Dortmund) presented the activities of the DELOS working group on DL test suites (www.sztaki.hu/delos_wg21). The group has developed a description scheme for DLs and is currently building a metalibrary of available test collections. Subsequent steps will be the identification of needs for additional test suites and the building of such collections. During the final discussion, several project representatives declared their interest in participating in the activities of the working group. Also, a possible joint working group with the members of the US DLI working group on metrics was envisaged. Concerning the acquisition of additional test collections, it was pointed out that multimedia collections pose severe problems due to intellectual property rights issues.

In the afternoon of the second day, the session chairs gave a summary of the results of their session. Costantino Thanos, Director of the DELOS Network of Excellence, then opened a brief discussion session to gather comments and suggestions from the participants on how this type of meeting can be made more effective for the participants. He then closed the meeting, announcing that DELOS plans to continue this initiative of holding DL All Projects meetings; both for European projects and (jointly with NSF) with US projects. The next EU-DL meeting will take place next year, in connection with the ECDL 2002 conference; the first EU-NSF DL All Projects meeting will take place in Europe sometime towards the end of 2001.

Complete details about the meeting (programme, presentations, list of participants) can be found at the DELOS Network of Excellence home page at http://www.ercim.org/delos/.

Please contact:
Costantino Thanos – IEI-CNR
Tel: +39 050 3152910
E-mail: Thanos@iei.pi.cnr.it

Working Group on Soft Computing proposed

by Petr Hájek

Soft Computing is an establishing domain and researchers from ERCIM insitutes propose to form an working group in this field.

Soft Computing has emerged as an attempt to formulate a new paradigm of computing. It can be discribed as association of computing methodologies centering on fuzzy logic, neurocomputing, genetic computing, and probabilistic computing. The guiding principle of soft computing is: exploit the tolerance for imprecision, uncertainty, partial truth, and approximation to achieve tractability, robustness, low solution cost and better rapport with reality. One of the principal aims of soft computing is to provide a foundation for the conception, design and application of intelligent systems employing its member methodologies symbiotically rather than in isolation. Scientists from ERCIM insitutes propose to form an working group in this field, comprising the following research topics:

Mathematical and logical foundations of Soft Computing

- systematic development of the theoretical basis of non-standard (ie approximate, nondeterministic or uncertain, etc.) reasoning, with a special focus to approaches to fuzzy logic, probabilistic and possibilistic logic, Dempster-Shafer theory and related approaches
- research in fuzzy logic programming in predicate calculus and with wider class of connectives (conjunctors and aggregation operators) which appear in practical application where one works with approximation of connectives
- design of the respective calculi for quantification and processing of uncertainty, imprecision and vagueness, analysis of their theoretical properties both from logical and algorithmic points of view; mutual confrontation of these calculi from the viewpoint of their expressive and processing power

- development of the theory of complexity of feed forward and recurrent neural networks and their learning algorithms based on complexity measure corresponding to various implementation possibilities.

Algorithmical Foundations of Soft Computing

- design and investigation of models suitable for realisation of formal calculi studied/proposed in the previous item
- design, development and analysis of formal, abstract machine models embodying the ideas of soft computing, inspired by biological or genetic models, with special regard to massively parallel models, distributed models, and neuromorphic models
- extending the theoretical basis of neural network-based computation with a special regard to approximation theory and related development of new neural computation paradigms
- the design and analysis of efficient algorithms for the fundamental problems in soft computing, both for internal need of the theory, as well as for various application areas, especially in the field of data mining.

Experimental Applications of Soft Computing

- identifying new soft computing information processing application areas particularly in the field of fuzzy and neuro-fuzzy systems, hybrid (ie, analog and neural) systems, data/knowledge bases, data warehouses, data mining, etc.
- identification and formalisation (modelling) of paradigmatical problems of Soft Computing
- The analysis, design, and development of formal methods and algorithms for the inconsistency, conflict resolution in the process of data/dnowledge bases and warehouses, integration, including the problems of different kinds of the fuzziness and of the uncertainty
- experimental implementations of new systems of Soft Computing.

In all those domains the aim of creative synthesis of various approaches to Soft Computing is stressed.

Please contact:
Petr Hájek – CRCIM (UI)
Tel: +420 2 6605 3760
E-mail: hajek@cs.cas.cz

ERCIM offers Postdoctoral Fellowships

The second round of the 2001/2002 ERCIM postdoctoral Fellowship Programme is now open with a deadline of 30 April. The fellowships are usually of 18 months duration, to be spent in two ERCIM member institutes.

The ERCIM Fellowship Programme was established in 1990 to enable young scientists from around the world to perform research at ERCIM institutes. For the second round of the 2001/2002 Programme, applications are solicited with a deadline of 30 April 2001.

Topics

This year, the ERCIM Fellowship programme focuses on the following topics:

- Database Research
- Constraints Technology and Application
- Control and Systems Theory
- Digital Libraries
- Formal Methods
- Electronic Commerce
- User Interfaces for All
- Environmental Modelling
- Health and Information Technology
- E-Learning.

Applications for other topics, eg Networking, Programming Language Technologies, Robotics, or Bioinformatics are also welcome.

Objectives

The objective of the Programme is to enable bright young scientists to work collectively on a challenging problem as fellows of an ERCIM institute. In addition, an ERCIM fellowship helps widen and intensify the network of personal relations and understanding among scientists. The Programme offers the opportunity:

- to improve the knowledge about European research structures and networks
- to become familiar with working conditions in leading European research centres
- to promote co-operation between research groups working in similar areas in different laboratories, through the fellowships.

Selection Procedure

Each application is reviewed by one or more senior scientists in each ERCIM institute. ERCIM representatives will select the candidates taking into account the quality of the applicant, the overlap of interest between applicant and the hosting institution and the available funding.

Conditions

Candidates must:

- have a PhD degree (or equivalent), or be in the last year of the thesis work with an outstanding academic record; the thesis must be held by the time of the beginning of the fellowship
- be fluent in English
- be discharged or get deferment from military service
- start the grant before October 2001.

Fellowships are usually of 18 months duration, spent in two of the ERCIM institutes. ERCIM offers a competitive salary which may vary depending on the country. Costs for travelling to and from the institutes will be paid. In order to encourage the mobility, a member institution will not be eligible to host a candidate of the same nationality.

Detailed description and online application form:
http://www.ercim.org/activity/fellows/

Please contact:
Aurélie Richard – ERCIM Office
Tel: +33 4 92 38 50 10
E-mail: aurelie.richard@ercim.org

GRIDs in ERCIM

by Keith G Jeffery

GRIDs is both a new and an old concept. Many of the components have been the subject of R&D and some exist as commercial products. The GRIDs concept represents many different things to different people: metacomputing, distributed computing, advanced networking, distributed database, information retrieval, digital libraries, hypermedia, cooperative working, knowledge management, advanced user interfaces, mobile and pervasive computing and many others. More importantly, end-users see the GRIDs technology as a means to an end – to improve quality, speed of working and cooperation in their field. GRIDs will deliver the required information in an appropriate form to the right place in a timely fashion. The novelty of GRIDs lies in the information systems engineering required in generating missing components and putting the components together.

In the Summer of 1999, the UK Director General of the Research Councils (John Taylor, whose article follows on page 10) was considering the Research Councils' Strategic Review – the plan for the next few years. I was asked by the CLRC Chief Executive to produce a brief IT Strategy paper. This paper proposed a 3-layer architecture of computation / data grid with networking as the base, an information grid for integrating heterogeneous information into a homogeneous presentation above and on the top a knowledge grid for extraction of knowledge from information and assists for user control of processes (see Figure 1); and it suggested a 'business proposition' that scientific application requirements could require the development of middleware and metadata which would be 'technology transferred' to commercial e-business.

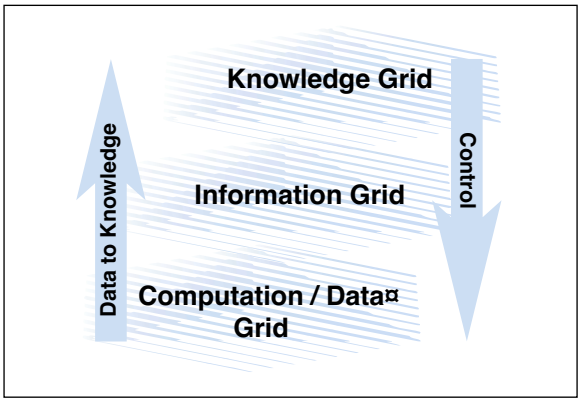

Figure 1: GRIDs 3-layer Architecture.

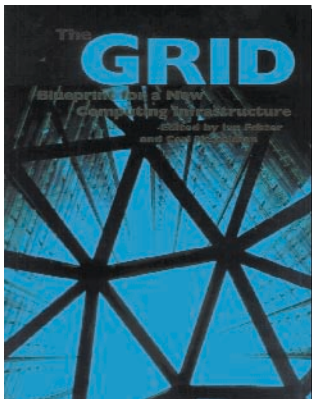

Figure 2: The GRID Bible.

Of course, the basic ideas in that paper were not new and came from previous work my department (and the sister department in CLRC, Computational Science and Engineering) had done in a European context with industry, academia and especially ERCIM colleagues over many years, and especially our work with W3C. We also noted the US work on a computation / data GRID (equivalent to the lowest of the 3 layers in the UK proposed architecture) and especially the book 'The GRID' by Ian Foster and Carl Kesselman known colloquially as 'The GRID Bible' (see Figure 2).

Knowing of work on GRIDs component technologies among colleagues in ERCIM Institutes over several years, I proposed to the ERCIM Directors in November 2000 that we set up a GRIDs Task Force. It is not a traditional Working Group because it relies on contributing middleware and metadata technologies from Working Groups on Digital Libraries, (including Cultural Heritage), Database, User Interfaces for All, Environmental Modelling, Health and Information Technology, e-Commerce, e-Learning and even from other teams who have not yet formed an ERCIM Working Group (eg metacomputing). Each community has its own view of GRIDs technology, emphasising different aspects or layers.

This Task Force initiative was stimulated partly by the emerging ideas on ERA (European Research Area) from Commissioner Philippe Busquin (http://europa.eu.int/comm/commissioners/busquin/index_en.htm) from January 2000 to the resolution of June 2000 and beyond (see ERCIM News no. 42, page 5). Part of the push came from the

developments in the USA and also from the EC-funded Datagrid Project – mainly for particle physics but involving several ERCIM institutes and associated universities and industries.

The Task Force has been set up as follows: the ERCIM member institutes have proposed a contact person or persons and we are busy now collecting data for an up-to-date picture of GRIDs in the ERCIM-represented countries. We are also collecting plans for the future. The general aim is to coordinate activity in ERCIM and associated academic and industrial partners such that we maximise European potential in GRIDs development. Discussions with EC (European Commission) officials indicate that perhaps they are considering measures to leverage national efforts in GRIDs which in aggregate are much larger than the resources the EC has at its disposal for this activity.

The following articles present a flavour of the dynamic GRIDs scene in Europe: John Taylor, Director General of Research Councils sets out the UK national initiative that he has championed for the community. There are articles on GRIDs in some representative countries, which indicate the breadth of approach yet the commonalities that exist. There are articles on metacomputing and related technologies in the computation/data GRID layer. Applications of the GRIDs technology in the Information layer show a rich diversity touching almost all areas of scientific research and development using advanced information technology. The article on a postgraduate course in GRIDs technology – surely a demonstration of abstracted knowledge for the knowledge layer - demonstrates that the technology is coming of age. The other major technologies for the knowledge layer – data mining (to generate knowledge) and e-publication (to disseminate knowledge) – have been documented elsewhere.

Lord Sainsbury – the UK Minister for Science and Innovation – in his keynote article (page 3) provides the rallying call. We in ERCIM have the framework to coordinate and co-develop with industry and academia the GRIDs technology on a 3-layer architecture (as distinct from the current US GRID technology concentrating primarily on the Computation / Data layer). The Information and Knowledge layers require traditional European strengths and competencies including multicultural multilinguality. They have commonality with Tim Berners-Lee's vision of a semantic web (http://www.w3.org/2001/sw/) and the web of trust (http://www.w3.org/Talks/9704WWW6-tbl/slide22.htm), but go beyond and emphasise further the importance of middleware and metadata technologies. These technologies are prerequisites not only for successful e-business but also for successful e-science in a joined-up academic and commercial world such as in bioscience, information technology, materials science, concurrent engineering, healthcare, environment and others. They are no less prerequisites for e-Learning, e-Culture and other disciplines to complete the e-Society.

Europe, working with our colleagues worldwide, should make its contribution such that European science, business and society all benefit. ERCIM has a key role.

Links
http://www.e-science.clrc.ac.uk/
http://www.w3.org/

Please contact
Keith G Jeffery – Director IT, CLRC
Rutherford Appleton Laboratory
Tel: +44 1235 44 6103
E-mail: K.G.Jeffery@rl.ac.uk

Articles in this section

UK Grid: Management and Architecture

by John Taylor, Director General of Research Councils

The days of client/server computing are numbered. Science applications are pushing at the very limits of current technology, and in solving their problems are developing a new generation of information technologies. Increases in several orders of magnitude in processing power, bandwidth and data storage are facilitating this process.

Scientists are facing various challenges, including:

- the acquisition and curation of extremely large, extremely valuable collections of primary data.
- the need for large in-silico experiments.
- visualisation at the desktop.
- the globalisation of scientific collaborations.

Developments are moving towards an information utility, a Grid, where processing power, data storage and retrieval, data analysis software, plus all the networking to link these tools together, are 'behind the wall'. New middleware will allow the scientist to take data and use resources from across a linked, global system without going to each resource individually or even knowing that they exist.

The emerging Grid concept sounds very attractive. But how do we go about building a Grid? How do we ensure that it meets the needs of many different user communities, without building several, incompatible Grids? How can we work within Europe and across the world to bring the new technologies into being?

In the UK, we have set up a specific e-science programme, with £ 118 Million over three years, plus extra funds to upgrade the academic network. We have deliberately divided the programme into two main sections. Firstly, there is £74 Million in specific allocations to the different Research Councils in the UK to tackle individual application science problems for which Grid solutions could provide the answer. Areas of application science include:

- Particle physics, with petabyte data flows and extremely large data storage and communication issues
- Genomics, with huge emerging databases of the human genome, and where metadata, visualisation and interoperability are crucial
- Climate change, where large in-silico modelling with many variables is particularly important
- Engineering systems design.

The second part of the UK programme is the 'Core e-science' programme. This will be managed by one of the Research Councils on behalf of the whole scientific community, and will have £35Million of Government support, covering both scientific and industrial pump priming. The programme will be expected to pull in at least £20 Million of additional industrial support.

This core programme will invest in generic technologies – scaleable networks, shared environments, access/security, metadata and so on. It will need to work iteratively with the application programmes, learning from them, providing them with new generic technology, and disseminating best practice across the disciplines. Progress in the applications will feed through to the generic, and vice-versa. Within the Core e-science programme, academic and industrial researchers will be brought together in open source joint research.

The core programme will also be responsible for expanding linkage with Grid developments internationally. It is essential that the UK programme forms part of the European and global efforts towards developing the technologies. The core programme will be actively involved in establishing international collaboration, although individual application programmes will also have specific collaborations (eg the UK particle physicists will be working with CERN).

The third strand of the UK e-science programme is a £9 Million investment in high performance computing. Added to other baseline funding, this will be used to bring sustained teraflop computing capability to the UK, and make it accessible via the Grid infrastructure.

Parallel to these developments, the UK's academic research network, JANET, has been upgraded. The backbone (Super JANET) has been raised to 2.5 Gbit/s, rising to 20 Gbit/s in 2002 and with additional increases planned. Metropolitan and Local area networks are also being upgraded, and a major boost to university infrastructure funding will allow them to increase their own internal networks.

The e-science programme is one involving many people over a number of disciplines, tackling many different and difficult problems. The overall administration of the programme is complex. The director of this Core e-science programme is an absolutely key individual. He or she will need to be an architect, a manager, a negotiator and an advocate all rolled into one. A Steering Group will bring together the interests of the application programmes and the Core e-science programme to ensure coherence across the whole UK e-science activity.

In this way, we hope to ensure that the core programme and the applications programmes have sufficient flexibility to tackle the problems, but interact in a dynamic way to ensure that we develop compatible technological solutions. It is a difficult balance, and it is ultimately up to the Chief Executives of the Research Councils and myself to ensure that it works.

Link:
http://www.dti.gov.uk/ost/

Please contact:
Gavin Costigan – UK Department of Trade and Industry,
The Office of Science and Technology
Tel: +44 20 7271 2075
E-mail: gavin.costigan@dti.gsi.gov.uk

UK e-Science Grid Programme

by David Boyd, Paul Jeffreys, Julia Goodfellow, Lois Steenman-Clark and Peter Allan

In October 2000, the UK Government's Office of Science and Technology (OST) announced a £98 Million programme over 3 years to develop and deploy Grid technology to meet the data and computing challenges arising in several areas of the physical and life sciences.

The rationale behind this decision is that, today, science is posing some of the most extreme computing challenges. These challenges represent the cutting edge in developing the next generation of distributed IT systems. Just as happened with the Web, however, the solutions to these challenges will quickly find application in the commercial marketplace. By supporting the Grid developments required by science now, the belief is that we will get to commercially useful systems quicker.

This new funding is being channelled through the current structure of the UK Research Councils which support basic and applied research in engineering and physical sciences (EPSRC), particle physics and astronomy (PPARC), the natural environment (NERC), biotechnology and biological sciences (BBSRC), medical science (MRC) and economic and social science (ESRC). CLRC, the seventh Research Council, provides major scientific research facilities in support of the work of the other Research Councils and also supports UK scientists in their use of large scale facilities elsewhere in the world. Together, all these facilities will generate increasingly large quantities of increasingly complex data over the next decade. Grid technology offers a potential solution to the challenge of extracting useful science from this complexity.

Some of the developments which are proposed as part of the UK e-Science Grid programme over the next 3 years are summarised below.

Bioscience

Molecular simulations are an essential tool in understanding macro-molecular function at an atomic level. Researchers at Birkbeck College in London, supported by BBSRC, have been using molecular simulation techniques, such as molecular dynamics, to study the stability and unfolding of specific proteins where the onset of partial unfolding is correlated with disease. One example of this is the protein human lysozyme (see figure) for which two mutations have been found clinically associated with amyloidosis

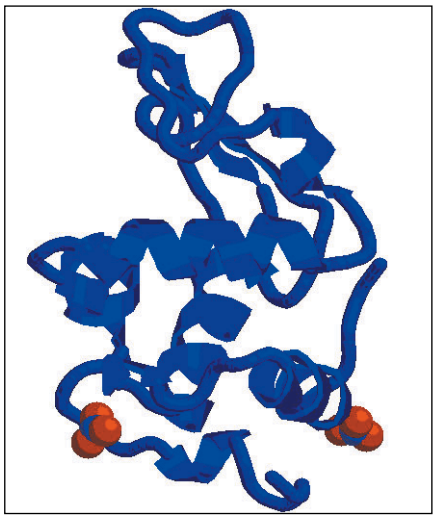

Structure of a human lysozyme mutation.

(formation of the protein into fibres resistant to temperature and enzymes). These calculations are computer intensive and have been traditionally carried out on supercomputers and more recently on in-house departmental multiprocessor workstations. Recent work using a Beowulf cluster at the Rutherford Appleton Laboratory (RAL) of CLRC has demonstrated this architecture to be a very cost effective solution to the considerable computing requirements of these techniques. A project funded by BBSRC and the Wellcome Trust is establishing a prototype Biosciences Grid linking five molecular simulation groups at Oxford, Southampton, Birmingham and York Universities and Birkbeck College with RAL who will provide technical support and computational and data resources.

Earth System Science

One of the possible proposals for a major new NERC initiative has the goal of predicting environmental changes on timescales of up to 100 years on the basis of:

- a sound scientific knowledge of the processes within, and the interactions among, the components of the coupled earth system
- an understanding of the uncertainties of our knowledge of key processes
- the inherent limitations of the predictability of the system.

The intention is to develop a modular modelling environment in which a hierarchy of models can be supported in a common framework. The project would involve collaboration between a broad spectrum of UK researchers from many areas of climate research and earth observation. The primary aim of this prototype climate research Grid is to provide the Earth System Science community with:

- a Data Grid providing high performance, integrated access to distributed massive data resources
- a Simulation Grid supporting large scale and distributed modelling capability spanning process studies to fully coupled climate change experiments
- a Model Grid providing a modular model code framework supporting multi-model experiments.

Astronomy

The UK astronomical community is planning to set up a Grid based project called AstroGrid. This will involve the astronomy centres with major data archives at RAL, the Mullard Space Science Laboratory and Edinburgh, Cambridge and Leicester Universities, and will bring in computer science expertise from Queen's University Belfast. Naturally there are important astronomical resources elsewhere and all of these will eventually be accessible to the whole UK astronomical community via the Grid. The concept of AstroGrid is

to create a Virtual Observatory. This will enable researchers to obtain new knowledge about a particular class of astronomical objects by extracting and analysing information from many existing data resources distributed throughout the world, rather than by applying for scarce telescope time at some remote mountain-top observatory to collect new data. These federated resources will contain data, with appropriate metadata to support efficient search and retrieval, from a wide range of high resolution, all-sky, multi-waveband surveys. Among the data sources feeding into the AstroGrid will be the UK Infra-Red Telescope Wide Field Camera (UKIRT-WFCAM) on Mauna Kea in Hawaii, the Sloan Digital Sky Survey on Apache Point in New Mexico, the XMM-Newton X-ray satellite and the Cluster mission of four satellites measuring the effect of emissions from the Sun on the earth's magnetosphere (see front cover picture). As an example of the productivity of these facilities, UKIRT-WFCAM can generate a substantial fraction of a terabyte of data every clear night.

Please contact:
David Boyd – CLRC
Tel: +44 1235 446167
E-mail: d.r.s.boyd@rl.ac.uk

Italian Initiatives in Grid Computing

by Bruno Codenotti and Domenico Laforenza

There is increasing interest in Grid Computing in Italy. This is reflected by the growing number of projects in this area. We describe some projects now underway in two major Italian research institutions: the Italian National Research Council (CNR) and the Italian Institute for Nuclear Physics (INFN).

Ongoing work at CNR

The actions related to Grid Computing involving CNR are a natural evolution of the PQE2000 project. This project, started at the end of 1995, has involved well-known groups working in three Italian research agencies (CNR, ENEA, INFN) and one industrial company (Quadrics Supercomputers World Ltd). A list of mature proposals includes:

- Computational Grids and Applications: this project, coordinated by CNR, involves 13 research groups belonging to: CNR (4), Universities (8), other Italian research institutions (1), respectively. It represents the first step toward the creation of a National awareness aroud Grid technologies; duration: 1Year; budget: 410KEuro; status: At the end of the evaluation process by CNR.
- GRID Computing – Enhanced Technologies and Applications for E-Science: this project, coordinated by CNR, involves 20 research groups belonging to: CNR (10), Universities (6), Public Administration (2), and Industry (2), respectively. The main goal of this project is to study, evaluate, and exploit grid computing technologies for the development of a new generation of scientific and business applications. A further goal is the creation of a kernel of skilled individuals on this research topic, opened to researchers, service providers and end-users; duration: 2.5 Years; budget: 2MEuro; status: submitted on October, 2000-under evaluation by the Italian Ministry for Research.
- ASI-PQE2000: this project by the Italian Space Agency (ASI) represents a joint effort among academic and industrial Italian partners led by the Department of Computer Science of the University of Pisa. The project mainly aims at building ASSIST (A Software development System based on Integrated Skeleton Technology). Only a part of this project is related to grid computing; in particular, a work package is targeted to specify how to extend ASSIST in order to address heterogeneous platforms (eg, Grids); duration: 2 Years; budget: 2MEuro; status: at the end of the evaluation process by the CNR.
- The DataGrid Project: CNR is associated partner in the DataGrid Project, a project recently approved by the European Commission and led by CERN. CNR is responsible for dissemination and exploitation activities; duration: 3 Years; status: Start-up phase.

The INFN GRID Project

INFN has recently approved a 3 year project, involving a total of 70 FTE's between physicists and computing professionals and a budget of ~ 9MEuro, to develop a Grid infrastructure that will allow INFN users a transparent and effective usage of the computing and storage resources distributed in the 26 INFN nodes of the Italian research network GARR-B.

The INFN national GRID infrastructure will be integrated with similar testbeds being established by ongoing parallel efforts in all major European countries, in US and Japan.

The overall scale of the computational, storage and networking capacity of the prototype INFN GRID testbed is determined by the needs of the LHC experiments. These include experimental activities for physics, trigger and detector studies, and the run of applications at a large enough scale to test the scalability of the distributed computing approaches for solving problems requiring a very large amount of distributed data (PetaBytes), a very large amount of CPU's (thousands) and many of users.

The development of the new software components underlying the GRID technology will be done by the INFN GRID project, whenever possible, in collaboration with international partners. Recently the DATAGRID international project, where INFN is one of 6 main contractors, was approved and funded

with 10 MEuro by the European Union IST program EU-RN2.

The INFN GRID project will investigate the current ideas for the computing models of the LHC experiments, based on a multi-tier architecture of Regional Centers. The investigation will be performed using real applications on real center prototypes, fulfilling at the same time the effective computing needs of the experiments.

The project will be instrumental to the diffusion and the common development of the GRID technologies with other Italian scientific research sectors (ESA/ESRIN, CNR, Universities), addressing the following points:

- develop collaborations with those Italian scientific partners to address problems which are common to different research areas
- promote the integration of different computing resources into one national research GRID.

In parallel to these activities, an additional budget line of ~20 MEuro is being established to implement a multi-experiment INFN Computing Center that will act as the main node of the INFN Grid.

The INFN-GRID Project provides the framework for efficiently connecting:

- the development of the middleware performed by participants in the DATAGRID European project
- the extra development activities needed for INFN specific purposes
- the set up of the prototypes of the Italian part of the distributed computing system for the 4 LHC experiments Alice, Atlas , Cms, Lhc-b; the most relevant components of this system are the prototypes of the Regional Centers for the 4 LHC experiments
- the step by step testing of the middleware with the 'real-life' application of the experiments, which is needed in order to make sure that the middleware will truly respond to the practical requirements of the experiments. The only possible testbeds for such activities are provided by the prototypes mentioned above. An extension of the middleware tests will provided by APE, the well known INFN parallel computer development team.

The currently approved DATAGRID project only required EU funds for the middleware development. The "testbed" and "HEP application" aspects are taken care of in two separate Workpackages. These workpackages require no EU funds for hardware and are just meant to provide for the human resources needed for the integration within the collaboration-wide activities. The INFN GRID will make available both the h/w and the extra manpower required for setting up and running the local prototypes and physics applications.

Please contact:
Bruno Codenotti – IMC-CNR
Tel: +39 050 315 2401
E-Mail: codenotti@imc.pi.cnr.it

Domenico Laforenza– CNUCE-CNR
Tel: +39 050 315 2992
E-mail: domenico.laforenza@cnuce.cnr.it

grid-ireland

by Brian Coghlan and Michael Manzke

In the future, as many people have said, computing will become a ubiquitous strategic resource, which users may avail of any time, anywhere. The subject is in its very earliest stage at present, with a very great number of basic research and practical issues to be investigated. Initial seed funding has been provided by Enterprise Ireland to guarantee the establishment of a working grid between compute clusters at three partner sites in Ireland, that is, the Departments of Computer Science in Trinity College Dublin, University College Cork and NUI Galway.

This project has one primary objective: to establish grid-ireland. This requires:

- Hardware: Compaq will donate a 4-way symmetric multiprocessing gateway machine per site
- Software: initially the Globus services
- Management: the Dept. of Computer Science at Trinity College will do this
- Interconnect: the Irish academic network services will be used.

A secondary objective is to begin co-operative research work between the three sites. The themes are clear:

Understanding: the most basic objective is the understanding of system state and estimation of future state of individual nodes, complete clusters, and heterogeneous compositions of these to form a computational grid. This will be explored using control theory via the analagous problem of real-time grid input-output, which is a very difficult concept due to the non-determinism of the grid.

Evolution: traditional models of grid computation (control-driven models) encourage programming practises that are polarised between the shared memory model common to symmetric multiprocessors and the message passing model common to ensembles of independent nodes. There is, however, a middle ground where lateral thinking and systematic exploitation may bear fruit; this hybrid model is essentially evolutionary.

Revolution: in contrast, data-flow and demand-driven execution models take a revolutionary approach, viewing an execution as a set of computational nodes, connected by paths along which data values move in a manner determined by the topology of the

Figure 1: 16-node cluster at Trinity College Dublin.

Figure 2: left: tracer/analyzer for the Scalable Coherent Interconnect (see http://www.cs.tcd.ie/coghlan/scieuro/), and right: in a heavily instrumented environment.

graph and by the rules underlying that model. Their superset, condensed graphs, offers a generalisation of both data-flow and demand-driven configurations.

Algorithms: the grid is different, ergo the algorithms must be different? We will explore suitable applications algorithms for the grid, particularly the emerging techniques for adaptive grid simulations. The programme will involve both the development of parallel Monte-Carlo algorithms and testing of these in a clinical environment both locally and nationally.

The grid is very interesting in terms of behaviour because it is extreme - tightly coupled processors will be interconnected via a high latency low bandwidth network. Moreover it is diverse. A great deal of basic research is needed into the conceptual models for almost every aspect of the grid environment, and these efforts will need to be driven with hard applications in diverse fields, more than is typical at present.

Understanding

A major aim of those involved in grid research is to enable applications to utilise resources in an optimal way, by recognizing current utilisation of the grid, by exploiting those resources that are under-utilised, and by avoiding or moving away from those that are over-utilised. This requires monitoring, analysis, and control mechanisms. Here we fail at the first hurdle - even current monitoring mechanisms are extremely minimal.

Moreover, there is little or no understanding of how the grid behaves. The computational grid is recognized to be a very dynamic system. Monitoring and analysis must be conducted on-line - offline approaches are not viable. Furthermore, attempting to control a poorly understood system is a recipe for catastrophe. There have been a number of famous failures of the power grid.

The most essential thing is to understand the dynamics of the grid, with all its non-determinism. Here we propose to investigate concepts from process control theory. This very fledgling work is being conducted in the context of research activities related to the global state estimation and optimisation of compute-clusters. Measurement is crucial. This is a subject in which Trinity College Dublin has long had an interest, particularly at the instrument and analysis level.

Evolution

For traditional (control-driven) execution models, maximally-optimized message-passing provides the upper bound for efficiency. The optimizations may be complex and counter-intuitive - maximally optimizing a message-passing program requires significantly greater effort and expertise than for shared memory. The reasons for this are clear: message passing requires explicit data placement and communication, whereas for shared memory this is done implicitly. The grid will exacerbate this situation. The motivation for shared memory is to eliminate the need to expend so much effort.

From the programmer's point of view, it is easier if at all times all processes have a consistent view of the contents of shared data structures. However, parallel programs require synchronisation for correct (race-free) execution. Many shared memory optimizations exploit this fact and weaken consistency. Generally the weaker the consistency, the less the message traffic, and on the grid this will assume greater importance. The unfortunate reality is that, despite this, message-passing is the most efficient solution, at least in critical places in a program. In an attempt to overcome this, we propose to explore the realm between these extremes. There are a number of very interesting possibilities. This is very speculative research.

Please contact:
Brian Coghlan – Trinity College Dublin
Tel: +353 1 6081 766
E-mail: coghlan@cs.tcd.ie
http://www.cs.tcd.ie/coghlan/

Michael Manzke – Trinity College Dublin
Tel: +353 1 6081 797
E-mail: michael.manzke@cs.tcd.ie
http://www.cs.tcd.ie/Michael.Manzke/

grid-ireland – New Execution Models and Different Algorithms

by John Morrison and Andy Shearer

grid-ireland has been established to link compute clusters at three partner sites in Ireland: the Departments of Computer Science in Trinity College Dublin, University College Cork and NUI Galway.

What if grids were better served by a totally different execution models and substantially different algorithms than heretofore? University College Cork and NUI Galway intend to explore these possibilities, both on grid-ireland.

Revolution

Like both dataflow and demand-driven execution models, condensed graphs (CGs) view an execution as a set of computational nodes, connected by paths along which data-values move, in a manner determined by the exact graph, and the rules underlying that model. CGs however, offer a generalisation of both data-flow and demand-driven configurations: each can be straightforwardly transformed into a CG equivalent, and furthermore CGs can be constructed which contain elements exhibiting both behaviours. The same model is inherently parallel, but can equally well express sequential, imperative-style computations.

An important property of CGs is their multi-level, hierarchical, structure. Replacing a cluster by a single node representing the same computation may condense structured graphs; this is analogous to a fold operation in a source-level program transformation. The dual operation, evaporation, is comparable to a program transformation unfold. By this device, it is possible to alter how much detail of a computation it is desired to expose at any given time, thereby making the representation more or less abstract.

Condensed Graphs can act equally well as both an abstract or physically realized machine, and as an intermediate representation: compiling the C intermediate code now requires, not an actual translation into a different representation, but rather is a matter of transforming the code into graphs which are less abstract, and which have more specificity to given hardware. Many of the techniques of source-to-source program transformation can be carried over into the CG world, allowing code optimisations to be expressed as graph transformations. Having done this we can then either execute the 'intermediate' representation directly, using the Condensed Graph Abstract Machine (CGAM) model; or it can be used as a true intermediate representation, compiling the CG code down to native code as a final step.

The Centre for Unified Computing in University College Cork is implementing CGs on a rich variety of platforms, and this work is already well under way. The grid offers a valuable opportunity to test the process of obtaining different machine-level CG realizations.

Algorithms

It is clear to the emerging grid community that the grid is so different that existing application-level algorithms are often quite unsuited. NUI Galway intends to explore and advance the algorithmic base via real medical applications. NUI Galway is currently developing algorithms for determining the passage of low energy optical photons through the body, and intends to apply them to the higher energy regime where different nuclear cross sections and scattering behaviour is expected.

These algorithms will be in the area of simulations for radiotherapy purposes. Patient treatment planning is a crucial aspect of any course of radiotherapy. In essence this consists of some scans that determine the best machine configuration to give the optimum dose to the patient. The best method is via Monte-Carlo simulations where the track of a large number of particles is followed through the scattering material. Although analytical techniques have been used they are not as accurate. Monte-Carlo simulations, however, require the use of many particles to gauge accurately the degree of scatter and attenuation, see Figure 1. Such calculations can also help determine when a particular round of treatment has not given the tumour the

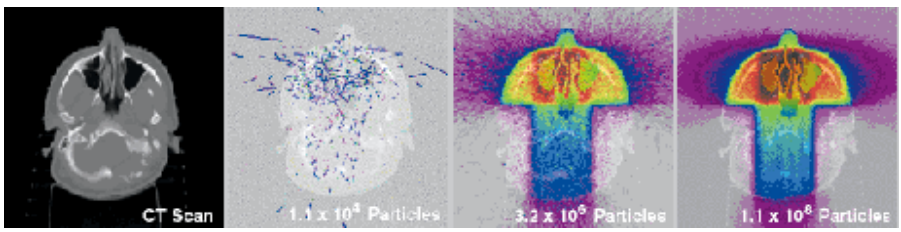

Figure 1: A scan followed by simulations involving 10,000 to 10 million particles (courtesy Lawrence Livermore Labs).

Figure 2: The tumour is shown in the first figure, the planned treatment in the second (analytical approach). In the third a Monte Carlo simulation shows that not all the tumour has been irradiated.

desired dose. Figure 2 illustrates this problem.

NUI Galway intends to develop similar Monte-Carlo code but without the requirement for a large computer to be on-site within each hospital. This can be done by taking the code developed for optical imaging and photo-dynamic therapy purposes, and modifying the underlying physics for X-rays rather than photons. The results could be used for setting radiotherapy machines, and so are expected to be of interest to the major local hospitals employing radiotherapy at each of the grid sites, that is, in Dublin, Cork and Galway.

Finally we intend to investigate the emerging technique of adaptive mesh simulation. In a typical analytical simulation a fixed grid is used to describe a physical phenomena and the forces at each node are calculated for a series of time steps. A better solution is to vary the grid spacing and structure dynamically through the calculation. The grid with its different architectures is an ideal platform for both calculating the node forces and the subsequent grid topology.

Links:
http://www.cuc.ucc.ie/
http://www.it.nuigalway.ie/

Please contact:
John Morrison – University College Cork
Tel: +353 21 4902 914
E-mail: j.morrsion@cs.ucc.ie

Andy Shearer – NUI Galway
Tel: +353 91 524 411
E-mail: andy.shearer@nuigalway.ie

Grid-related Activities in the Laboratory of Parallel and Distributed Systems of SZTAKI

by Péter Kacsuk and Ferenc Vajda

The Laboratory of Parallel and Distributed Systems (LPDS) of SZTAKI has a long and successful history in the research and development of distributed systems, tools and applications. A number of these activities and results are closely related to the current grid-oriented projects of the laboratory. The largest computer cluster in Hungary runs under the supervision of LPDS.

In the past the laboratory developed different graphical, parallel programming environments. One of them is the P-GRADE that is an integrated set of programming tools for general message passing applications to be run in heterogeneous computing environments or supercomputers. It was developed in the frame of several EU and Hungarian projects, and supported by the Hungarian National Committee for Technological Development. The most interesting tools in the P-GRADE package are:

- GRAPNEL graphical programming language which supports parallel programming activities by graphics while the other part of a program can be written in a textural language (eg C). GRAPNEL programs can be compiled to existing message passing systems (PVM or MPI).
- GRED, a graphical editor to write parallel applications. It can be used to construct program graphs both in the Application and the Process Windows supported by GRAPNEL.
- GRM, a semi-online monitoring tool to generate and collect information (trace data or statistical information) about an application. Host and network monitoring sensors can be connected to GRM. It delivers the collected information to PROVE.
- PROVE, a visualization tool to analyze and interpret the trace file information and to present it graphically to the user during execution. GRM and PROVE provided solid bases to the current grid monitoring activities of the laboratory.

LPDS is the coordinator of a large-scale research/development project supported by a grant of the Hungarian Ministry of Education, Research and Development Division. The main goals of the project are:

- to develop a virtual supercomputer based on the computers of the institutions connected to academic computer network
- to develop and use a metacomputing supervision system to supervise the heterogeneous systems
- to support pilot-applications of the system
- as side effect of the project, to establish a 'knowledge center' to provide professional support to the proliferation of these new technologies.

The Hungarian Scientific Research Fund (OTKA) provided a grant to the laboratory for the period of 2000-2003. The title of the project is 'A graphical supervisory system for geographically distributed heterogeneous metacomputing environment'. The goal is to extend P-GRADE to grid applications.

Participation in the DataGrid Project
Represented by our laboratory, SZTAKI is an associate member of the DataGrid Project supported by the European Commission in the 5th frame of the IST program. The objective of the project is to enable next generation scientific exploration, which requires intensive computation and analysis of shared large-scale databases (the size from hundreds of TeraBytes to PetaBytes) across widely distributed scientific communities. It is based on emerging computational grid technologies (mainly on the Globus Project). LPDS takes part in the Monitoring Services Workpackage.

The ability to monitor and manage distributed components is critical for enabling high performance distributed computing. Monitoring data is needed to determine the source of performance problems, to support tuning systems and applications for better performance, for fault detection and recovery and for performance prediction services.

The system consists of the following main components:

- event consumers, which could be any program that requests event data for the purpose of real-time monitoring (performance analysis), overview monitoring (combining information to make some decision), process monitoring (triggering an action on an event from a server process) and archiving (providing ability to do historical analysis of system performance)
- event producers, which are gathering event data from host, network, process, application, storage, I/O, etc.,
- directory service, which locates names and describes the structural characteristics of any data available to the grid.

The main tasks of the laboratory in the project are defining and developing the producers' activities and their communication principles and protocols with the consumers and event directories.

The researchers of LPDS took an active part in the European Grid Forum and now in the united international Global Grid Forum. The laboratory actively participated in the Cactus Testbed demonstration of the European Grid Forum at the Supercomputing Conference (4-10 November 2000) in Dallas. Cactus is an open source problem-solving environment designed for scientists and engineers. The demonstration simulated the collision of black holes and was executed on the machines of the six participating European institutions including the computer cluster of SZTAKI.

On the occasion of ERCIM's 10th anniversary, a Prospective Report entitled 'Network-based Distributed Computing (Metacomputing)' by Péter Kacsuk and Ferenc Vajda was supported by an ERCIM grant. The document is available at the ERCIM website under http://www.ercim.org/publication/prosp/.

Link:
http://www.lpds.sztaki.hu

Please contact:
Péter Kacsuk – SZTAKI
Tel: +36 1 329-7864
E-mail: kacsuk@sztaki.hu

Ferenc Vajda – SZTAKI
Tel: + 36 1 329-7864
E-mail: vajda@sztaki.hu

MetaCenter – Building a Virtual Supercomputer

by Ludek Matyska, Michal Vocú and Ales Krenek

The MetaCenter project builds a nation-wide Czech computing and data storage GRID. Started in 1996 with an initiative from Masaryk University, its main goal is development and deployment of middleware products to create homogeneous environment on top of heteregeneous geographically distributed computing resources.

Three academic high performance computing centers, located at Masaryk University in Brno, Charles University in Prague and West Bohemia University in Pilsen, currently under the umbrella of CESNET, are connected in projects whose long term goal is a creation of an academic GRID for the Czech Republic, the environment which supports large scale distributed and parallel applications and in the same time leads to more efficient use of available computing resources. To fulfill this goal, ie to create MetaCenter GRID environment which hides the details of individual computing resources and their distance from the end users, set of appropriate middleware components are build. The work is focused on the following areas:

• **Information services**, both for end users, via a web based interface, and for program devellopers, via appropriate APIs. Information about users, data sets and application programs is stored in Oracle database within the perun system, which was developed as part of the MetaCenter project. The relevant data from the database are regularly exported to LDAP based directory service, and to the Kerberos authentication system.

• **Security**, including uniform access to all computing resources. A single sign-on system based on Kerberos 5 protocol was created, which allows once per access authentication. The Kerberos 5 implementation is based on the Heimdal system, which is being extended to suits MetaCenter requirements. Currently, GRID wide login name is associated with each user, but research towards virtual accounts mapping is under way. This will allow a seamless collaboration with other GRIDs. The perun system based information services are built to allow on-demand creation of user accounts on individual machines and are therefore well prepared to the incorporation of virtual accounts and their mapping to actual physical persons.

• **Shared data space**, with an illusion of location independency of users (and application) programs and data sets. This is achieved via use of AFS, distributed file system with global unique file naming, and also via simple file transfer protocols using the single sign-on mechanism (curently scp a ftp protocols

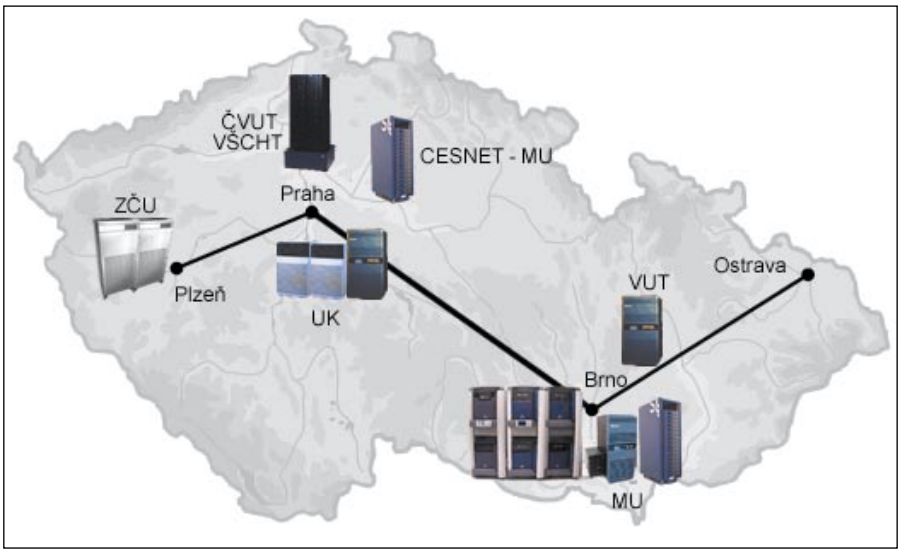

Czech computing and data storage GRID centres.

are supported, with the goal to support their GRID-aware extensions like Grid-FTP). High capacity backup storage is also available and serves uniformly all GRID nodes. The AFS is fully integrated within the single sign-on mechanism, the location independency is supported through replica servers of read-only data. All information about physical placement of individual data volumes is stored in perun and hidden from end users, making administration of the whole GRID rather easy. Applications are also installed in AFS and are accessible using system of modules. A module is a virtual entity representing access to a particular application. Instead of remembering different locations of applications, scratch space etc., users willing to access a particular application issue a single 'add <application>' command which among other things creates appropriate shell environment necessary to run the apllication. Global module name space (also stored in perun) secures uniformity in access to applications regardless of their GRID physical location (and also support easy use of floating licenses).

• **Batch systems** are used to control all non-interactive use of the MetaCenter GRID (in fact, ideas to support even interactive jobs via batch queue system are currently discussed and tested). The original system of choise was LSF (Load Sharing Facility), but despite it advantages it is currently replaced by OpenPBS batch system (the main reasons are the cost – while LSF is very costly, OpenPBS is available under open licensing terms – and ability to repair errors and create own extensions, which are required to support more advanced and experimental scheduling policies – LSF is available in binary form which cannot be directly modified nor extended). Set of global batch queues is created and users, when submitting their jobs, can left the decision where the programs will be run completely on the batch system, increasing thus the total system throughput.

• **High speed network infrastructure** is provided by the national-wide academic backbone CESNET 2 (2.5 Gb/s between Prague and Brno, and 34 Mb/s between Pilsen and Prague to be this year upgraded to at least 1 Gb/s) and metropolitan area networks with 155 Mb/s connections. This network provides high bandwith and low latency necessary for distributed and parallel applications. Recently purchased PC clusters located in Prague and Brno are connected directly to the backbone via 1 Gb/s uplinks, providing environment allowing to study influence of latency on distributed applications.

Since 2000 the MetaCenter project became involved in two important European wide Grid activities – the European Grid Forum (http://www.egrid.org) and the Datagrid project (http://grid.web.cern.ch). Under the framework of the former one we took an active part in a successful demonstration of the EGrid functionality during the SC2000 conference in November 2000. The demonstration presented a dynamic, migrating scientific computation (a 'worm') built on the Cactus and Globus metacomputing toolkits. Nine supercomputing centres of seven European countries were involved in the experiment. Under the Datagrid project scheduling, related security problems and information services are our primary are of interest.

Link:
http://meta.zcu.cz

Please contact:
Ludek Matyska – CRCIM / Masaryk University
Tel: +420 5 41 512 310
E-mail: Ludek.Matyska@muni.cz

Metacomputing Applications in All-optical Networks

by Thomas Eickermann, Helmut Grund, Wolfgang Ziegler and Lothar Zier

Today's commodity networking technology is one limiting factor for the performance of distributed applications with high bandwidth communication. Another obstacle is the still rudimentary middleware layer for convenient set up, operation and steering of metacomputing applications. The proposed project RePhoNet is aimed to provide a testbed bringing together optical network technologies, advanced middleware for metacomputing and real world applications. It will be funded by the German government and the DFN (German Research Network).

Following two other successful advanced networking projects of the last years (RTB-NRW and Gigabit Testbed West) RePhoNet is a new joint R&D project of 15 partners from industry, governmental research centres and universities: German Telekom, Siemens, GMD, FZJ, DLR, Caesar Bonn, Universities of Bonn and Cologne, FH Rhein-Bonn-Sieg and others. It consists of three major building blocks:

- OptiNet, provision of the underlying photonic network technology of the testbed
- MetaComp, developing and implementing the essential technology to run metacomputing applications in the testbed
- a number of real world applications ranging from simulation of molecular dynamics to distributed virtual reality systems.

Most applications will use PC-clusters with shared memory multi-processor nodes as compute resource. There are six PC-clusters available in the testbed today, distributed over five sites, having between 14 and 144 CPUs. Each cluster has a fast Myrinet internal network. It is planned interconnect the local Myrinet infrastructures to the optical network through special high speed links.

OptiNet will establish an optical network in the Bonn/Cologne region. Later on an extension to Aachen and Darmstadt is planned. The network is based on optical switches of several vendors. Signalling will allow the dynamic, on demand allocation of multiple lightpaths. In collaboration with MetaComp the Meta-scheduler will be extended with features for optical signalling. Each of the optical connections will provide several Gigabit capacity with QoS guarantees. Various networking technologies in the end systems (eg ATM, Gigabit Ethernet, 10GigE, Myrinet) can be connected with the all-optical equipment in the network core.

MetaComp is active in two main areas: a metacomputing enabled communication library (MetaMPI) and the Meta-scheduler. Additionally MetaComp covers a survey of existing communication software and interconnect hardware, system management, and end-user advising.

MetaMPI will be ported to Myrinet/Intel clusters to allow an efficient MPI-communication between PC-clusters and other parallel computers. Apart from this work it is necessary to generalize the communication model of MetaMPI to make an optimal use of the different architectures of supercomputers and clusters: while supercomputers still communicate through dedicated router-nodes the planned Myrinet infrastructure allows a direct point-to-point communication between all nodes

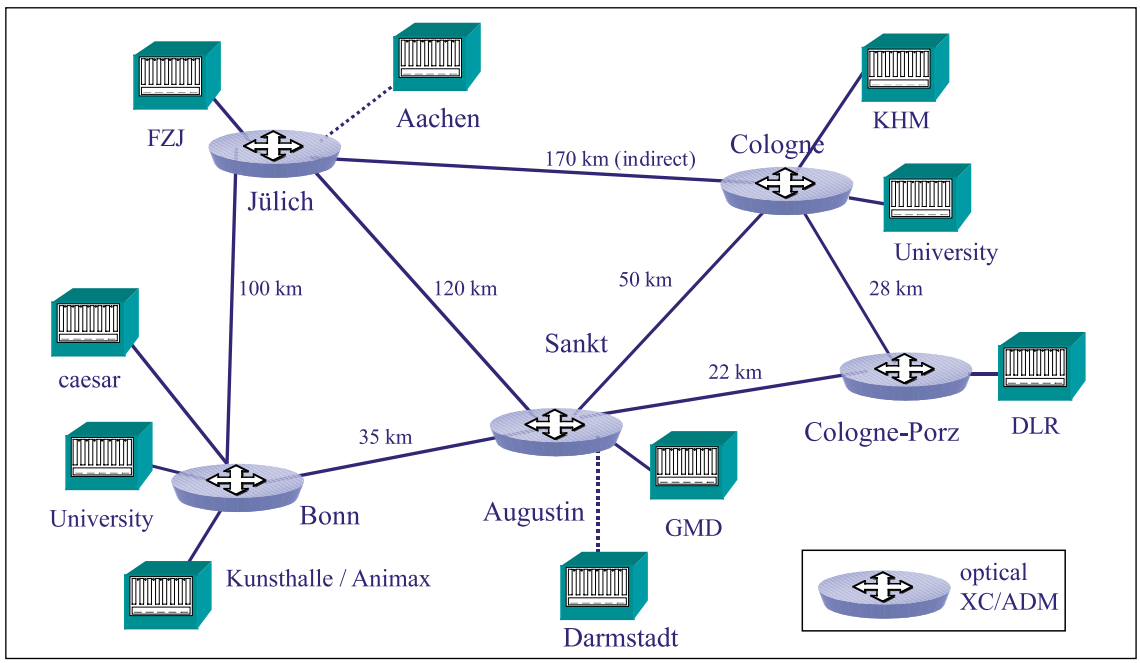

Topology of the all-optical network testbed.

of a distributed PC-cluster. In order to support MetaMPI's dynamic features (MPI-2 process spawning) an interface to the Meta-scheduler has to be implemented.

Using the Meta-scheduler requires a local job-scheduler running on all PC-clusters supporting the RAA (Resource Allocation Protocol). This might be done by either porting and implementing GMD's EASY scheduler or through modification and extension of a local existing and suitable scheduler. Another important task is the integration of a system that allows resource allocation on the networking level, i.e. scheduling of wave-lengths, or binding a dedicated QoS to a job in the RePhoNet network.

During the first phase of the three years project the applications will focus on simulation and virtual reality systems. The goal of the simulation subproject (Molecular dynamics with proteins) is to run the first simulation of protein folding in an ion solution while breaking the nano-seconds barrier which limits protein folding simulations until today. The RemoteCAVE subproject aims at building an infrastructure that interconnects museums and powerful scientific institutions allowing museums and exhibitions to get graphics power on demand over the network. The approach is similar to allocating compute power over the net but has never been done up to now. The third subproject is a virtual multi-user environment with a large number of (distributed) participants which could not be realized due to lack of both compute power and network bandwidth until today. The RePhoNet project is expected to start April 2001.

Link:
http://rephonet.gmd.de

Please contact:
Thomas Eickermann, Helmut Grund, Lothar Zier – GMD
Tel: +49 2461 616596, +49 2241 14 2298, +49 2241 14 2943
E-mail: Th.Eickermann@fz-juelich.de, Helmut.Grund@gmd.de, Lothar.Zier@gmd.de

The Distributed ASCI Supercomputer Project

by Henri Bal

The Distributed ASCI Supercomputer (DAS) is an experimental testbed for research on wide-area distributed and parallel applications. The system was built for the Advanced School for Computing and Imaging (ASCI), a Dutch research school in which several universities participate.

The goal of DAS is to provide a common computational infrastructure for researchers within ASCI, who work on various aspects of parallel and distributed systems, including communication substrates, programming environments, and applications. Like a metacomputer or computational grid, DAS is a physically distributed system that appears to its users as a single, coherent system. Unlike metacomputers, we designed DAS as a homogeneous system.

The DAS system consists of four cluster computers, located at four different universities in ASCI, linked together through wide area networks. Three clusters (at the Universities of Amsterdam, Leiden and Delft) have 24 nodes; the cluster at the Vrije Universiteit has 128 nodes. All four clusters use the same processors (Pentium Pro) and local network (Myrinet) and run the same operating system (RedHat Linux). Each university has fast access to its own local cluster. In addition, a single application can use the entire wide-area system, for example for remote collaboration or distributed supercomputing. DAS can be seen as a prototype computational grid, but its homogeneous structure makes it easier to avoid the engineering problems of heterogeneous systems. DAS can also be seen as a cluster computer, except that it is physically distributed.

A snapshot of the cluster at the Vrije Universiteit, clearly showing the Myrinet cables (courtesy Henri Bal, Vrije Universiteit, Amsterdam).

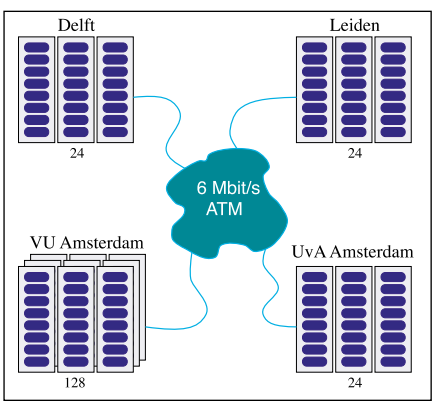

The wide-area DAS system.

The DAS project started in June 1997. The participants are the four above-mentioned universities and the university of Utrecht. Research has been done on low-level communication protocols, languages, schedulers, parallel applications, wide-area (distributed supercomputing) applications, Web-based applications, interactive applications, and worldwide distributed applications. Recently, funding has been obtained to build a successor system, DAS-2, which will consist of five SMP-based clusters.

Link:
http://www.cs.vu.nl/das/

Please contact:
Henri Bal – Vrije Universiteit Amsterdam
Tel: +31 20 444 7733
E-mail: bal@cs.vu.nl

Metaψ: A Web-based Metacomputing Problem-Solving Environment for building Complex Applications

by Ranieri Baraglia and Domenico Laforenza

The increasing complexity of large distributed scientific applications raises the problem of the coordination of diverse computational resources (computers, data bases, etc.). Multi-disciplinary applications often make use of coupled computational resources that cannot be replicated at a single site. There is the need for smart and user-friendly Problem-Solving Environments (PSE) that free scientists from concerns related to the location and complexity of the computing platform being used.

In this article we describe the main features of Metaψ, a software tool developed at CNUCE-CNR to build PSEs for the execution of complex applications on a Web-based metacomputer. This tool is designed to supply a completely transparent support to the user, who thus does not need to be aware of the location and the allocation of computing resources.

Metaψ has a 3-tier architecture with the following layers, as shown in Figure 1:

- Client side: a Web browser (eg Netscape Navigator)
- Middleware: Web Server and LDAP (Lightweight Directory Accesss Protocol) Server
- Back end: the set of computing resources.

Metaψ has been implemented on an appropriately extended web server, mainly using standard tools. In particular, the Java Servlet and Directory Service facilities of LDAP have been used. As can be seen, in order to guarantee easy maintenance and extendibility of the prototype, the design is modular.

The client side consists of a Web browser representing the graphical interface which guides the user when selecting an application, and in the input of data and the retrieval of results. After authorisation (see Figure 2b), the user selects the application (Figure 3a) and then provides the input data (Figure 3b).

The middleware (see Figure 1) layer consists of a Web server and an LDAP server. The Web server takes care of the interaction with the client and executes the Java servlets handling the user requests. The servlets residing on the Web server are: Serv1 which handles the user authentication stage; user profiles establish who can do what and where results of a run should be stored. Serv2 handles the application profiles management, and guides the user in the data input, according to the selected application. Serv3 is the most complex servlet taking care of validation of input data, resource location, allocation and configuration, remote execution of the application, collection and forwarding of results. When responding to user requests, the middleware employs LDAP functions

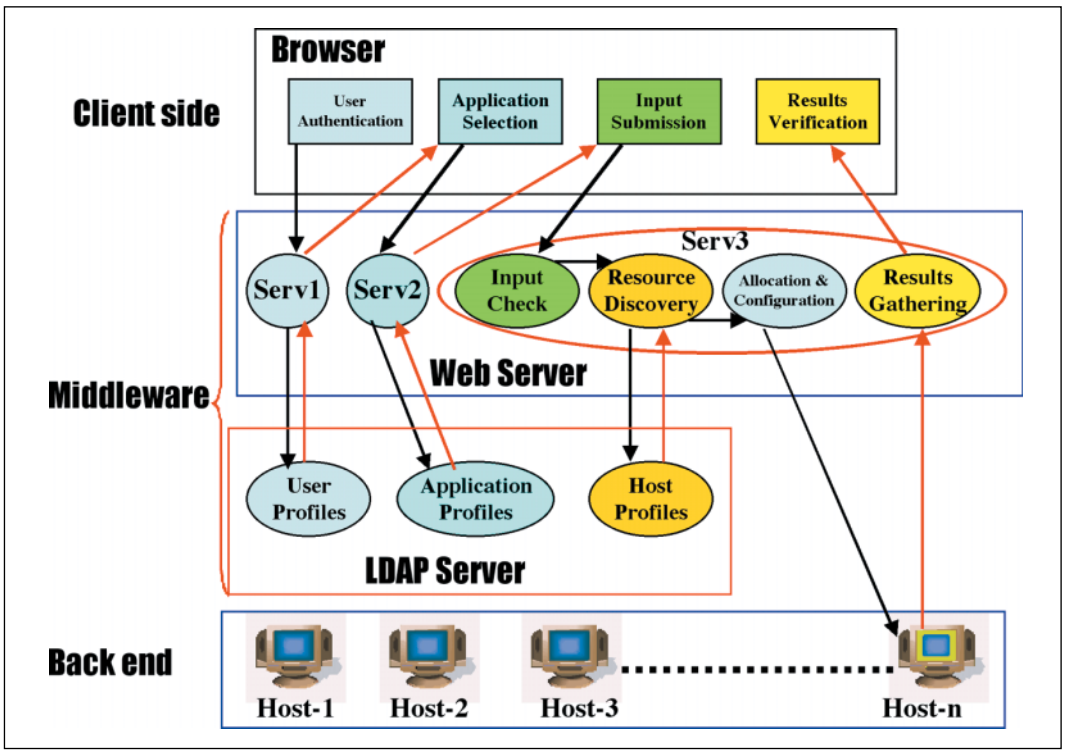

Figure 1: Key interactions among the system components.

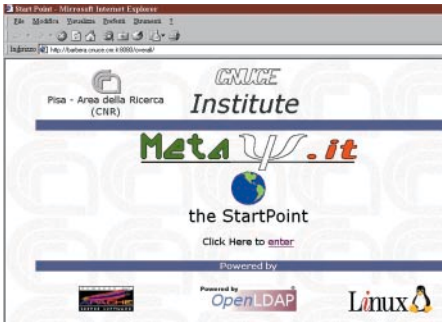

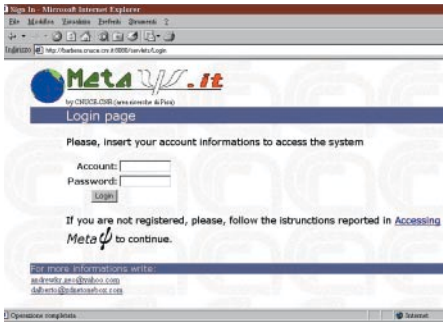

Figure 2a (left-hand side) and 2b (right-hand side): Welcome and User Authentication Windows.

to locate those computing resources that can provide the requested services. LDAP provides information on the computing resources by accessing a Directory Information Tree (DIT). The tree consists of entries representing the computing resources through a group of attributes. The server then activates a remote execution of the selected application. HTML forms are used in the interaction between the user and the Web server. These forms activate the Java servlet that executes the requested action (application selection, data input, etc.). The initial page (see Figure 2b) allows the specification of the userid and password. In order to control access to the system's resources, different user profiles can be defined according to a predefined strategy, implemented using LDAP. After authentication, the user is offered a list of applications that can be run on the machines belonging to the system. An application is selected by clicking on the hyperlink related to the application. Each application has an associated profile describing its computational characteristics/requirements: eg, name, version, documentation available, type of input required, sources of data, etc. The application profile, stored in the LDAP server, is exploited by the Web server to drive the input process. An HTML page guiding the data input according to the characteristics of an application is produced (see Figure 3a). The input data can be submitted (see Figure 3b) to the application in three different ways: by manual data entry, by selecting a file from the client's local disk, or by choosing a link to a remote data source (ie, a file located on a metacomputer).

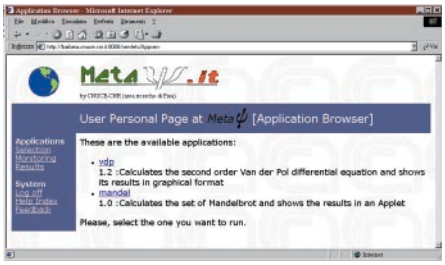

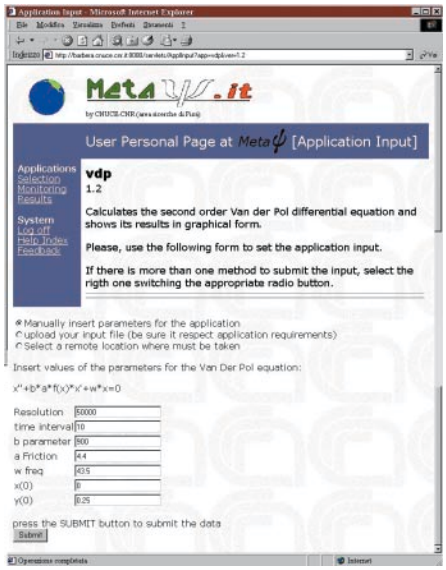

Figure 3a (left-hand side) and 3b (right-hand side): Application Selection and Input Windows.

The back end (see Figure 1) is made up by high performing computing resources, multiprocessor systems and workstation networks which provide computing power to the applications of the metacomputer. The Web server has user accounts on these machines that allow the execution of the applications. In general, the execution of an application on the back end takes a significant amount of time; the user can thus exit the metacomputing session once the application has been submitted. At job completion, the user is notified by an email message where to access the page of the results built by a Java servlet. When execution is completed, the results are passed to the server that forwards them to the Client (Figure 4).

In the near future Metaψ will be used to implement SIMBEX, a simulation of crossed molecular beam experiments. SIMBEX is a computational procedure based on a priori calculations of structures and processes of molecular systems. The procedure runs over several modules; each module consists of alternative or coordinated computer codes which accomplish particular tasks.

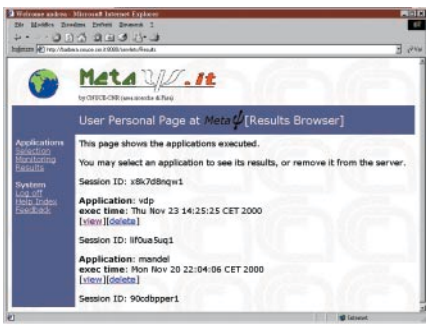

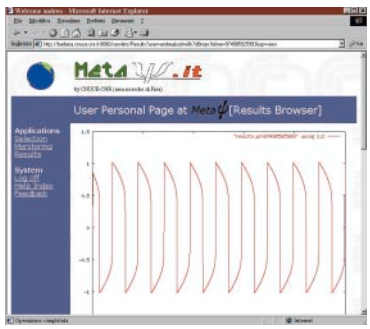

Figure 4a (left-hand side) and 4b (right-hand side): Application Execution Control and Results Verification Windows.

Please contact:
Ranieri Baraglia, Domenico Laforenza – CNUCE-CNR
E-mail: {ranieri.baraglia, domenico.laforenza}@cnuce.cnr.it

DIET: A Distributed Interactive Engineering Toolbox for Client-Server Applications in a Grid Environment

by Frédéric Desprez

Huge problems can now be computed over the Internet thanks to Grid Computing Environments like Globus or Legion. Because most of current applications are numerical, the use of libraries like BLAS, LAPACK, ScaLAPACK or PETSc is mandatory. The integration of such libraries in high level applications using languages like Fortran or C is far from being easy. Moreover, the computational power and memory needs of such applications may of course not be available on every workstation. Thus, the RPC seems to be a good candidate to build Problem Solving Environments on the Grid. Several tools following this approach exist, like Netsolve, NINF, NEOS, or RCS.

In 1998, we have started a project for the parallelization of a Matlab-like environment called Scilab which developed at INRIA. One of the chosen approach was to link the Scilab tool to Netsolve.

Then, in 2000, we started the DIET project (Distributed Interactive Engineering Toolbox), for the development of a hierarchical set of components to build Network Enabled Server (NES) applications. Our target platform is the fast network VTHD connecting several research centers (and their clusters) from INRIA.

This project involves several research teams in CS laboratories accross France: ReMaP at LIP (Lyon), Résédas at LORIA (Nancy), and SDRP at LIFC (Besançon).

DIET architecture

Usually, NES environments have five different components: CLIENTS that submit problems to SERVERS which solve them, a DATABASE that contains information about software and hardware resources, MONITORS that get informations about the status of the computational resources, and finally a SCHEDULER that chooses an appropriate server depending of the problem sent and information contained in the database.

In DIET, a server is built upon Computational Resources Daemons and a Server Daemon. We have a hierarchical set of agents including Leader Agents and Master Agents. A redirector is used to choose a master agent which is close to the client. Requests for computation from a client are sent to the nearest agent. We believe that such a hierarchy is mandatory when building scalable environments for the Grid. We have designed several tools for resource discovery and monitoring:

- SLiM's (Scientific Libraries Metaserver) goal is to make the junction between problems submitted by clients and the implementations available on servers. In most case, there is no one-to-one mapping: a single problem can be solved by many implementations from several libraries, while another problem may need more than one computational step to be solved. All needed information are stored in a LDAP tree. LDAP is a distributed database protocol which was chosen for its read and search optimizations.
- FAST (Fast Agent System Timers) is a tool for dynamic performance forecasting in a Grid environment.

FAST is composed of several layers and relies on low level software. First, it uses a network and CPU monitoring software to handle dynamically changing resources, like workload or bandwidth. FAST uses and enhances the Network Weather Service (NWS), a distributed system that periodically monitors and dynamically forecasts the performance of various network and computational resources. FAST also includes routines to model the time and space needs for each triplet (problem; machine; parameters set). They are based on benchmarking at installation time on each machine for a representative set of parameters and polynomial data fitting. To store these static data, FAST uses the same LDAP-tree as SLiM.

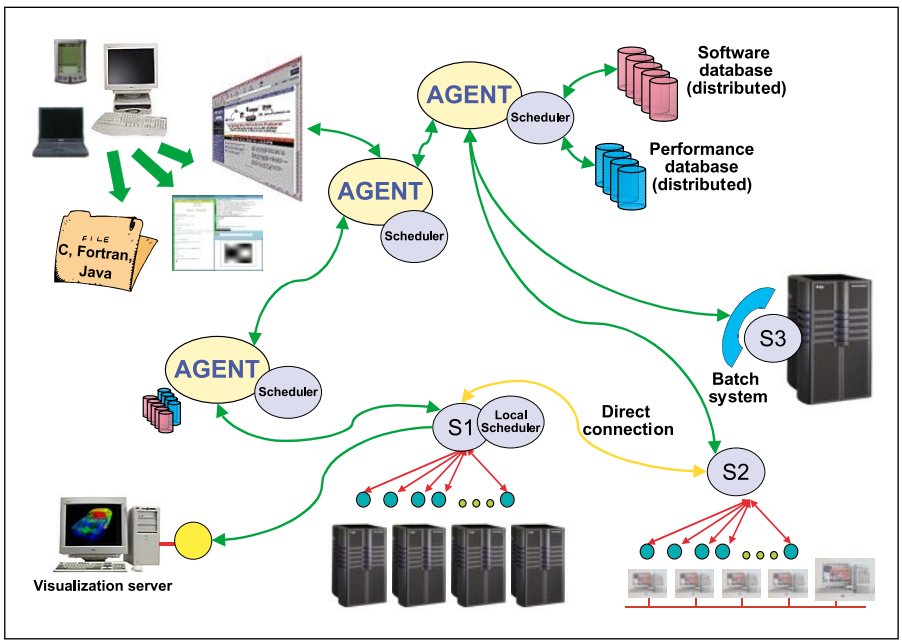

DIET architecture.

In order to implement a network enabled problem solver, one can choose between many communication layers. Low level layers like the socket interface eventually allow the best performance. Higher level layers such has ones complying with the CORBA norm although provide interfaces for a easier and quicker development.

The great variety of metacomputer components (computer, networks, files, software services) and their dynamical behavior raise special problems from the resource management point of view. In such architecture, management will play a major role and it appears very important to be able to offer a unified framework for the management of network, services and applications. The framework implies the definition of an architecture that guaranties the interoperability between applications through the portability of management information wherever they are originated.

We claim that middleware layers based on Java technologies, and more especially JMX (Java Management eXtensions), offer new opportunities to support NES applications in a Grid environment. Our approach is based on WEBM (Web-Based Enterprise Management) which the standardization effort is relayed by the DMTF (Distributed Management Task Force).

Conclusion and Future Work

Our future work is to first test this approach on real applications. As our target platform allows 2.5 Gb/s communications between several INRIA research centers, connecting several clusters of PCs and parallel machines, we think that tightly coupled applications written in a RPC mode could benefit of such an approach. Another problem we would like to address is the optimization of data distributions for parallel library calls using a mixed data and task parallel approach. We also would like to connect our developments to infrastructure toolkits like Globus to benefit from the development of security, accounting, and interoperability services.

A concerted effort is on his way to define the overall architecture and basic functionalities of Grid environments. A working group is dedicated to Advanced Programming Models which include, of course, Network Enabled Solvers. We think that this effort is very important to be able to get efficient software infrastructure soon and we would like to be part of it.

We are of course open to collaborations with other ERCIM members.

Link:
DIET project: http://www.ens-lyon.fr/~desprez/DIET/index.htm

Please contact:
Frédéric Desprez – École normale supérieure de Lyon
Tel: +33 4 72 72 85 69
E-mail: desprez@ens-lyon.fr

Co-allocating Compute Resources in the Grid

by Gerd Quecke and Wolfgang Ziegler

Multidisciplinary simulations – as the simulation of air-plane wings for example – repeatedly require different compute resources being available at the same time to successfully perform such simulations. Co-allocating resources today usually requires a substantial amount of human communication on all levels. MeSch is a solution for the problem of (synchronous) resource allocation and job scheduling in a distributed heterogeneous environment developed in the Institute for Algorithms and Scientific Computing (SCAI) at GMD.

Resource management and job scheduling in the typical Grid environment based on multi-MPP systems or clusters is still one of the challenging problems today. Especially in a geographically distributed and heterogeneous environment, it turns out, that although scheduling tools and policies are available for each subsystem, there is a lack of global resource management and thus, resource allocation is far away from being performed automatically. On the contrary: a substantial amount of human communication on all levels is necessary to partition the application, locate resources, and observe the behavior of distributed modules.

MeSch is a light-weight solution for the problem of resource allocation and job scheduling in a distributed heterogeneous environment. The same way, a Grid application uses the Grid resources as a metacomputing environment allowing the use of more than one MPP system or cluster, MeSch leads to the idea of building a metascheduler, which takes the burden of resource co-allocation for a metajob. The approach here is to build the metascheduler such that it can use schedulers of all subsystems involved for all co-ordination and resource allocation tasks. The MeSch metascheduler prototype allows co-ordination of the whole scheduling process during the application lifetime including resource allocation. The algorithm was especially designed to allow simultaneous access to the requested resources, a requirement typically needed by parallel applications.

Until now, the only solution to overcome these problems is to use scheduling systems that are able to completely handle resource management for all resources involved. However, trying to use heterogeneous environments as they are becomes difficult if such attempts will be based on a single task approach as a regular service, without any need to change local administration rules and policies. Or, for example, to introduce local components like the GRAMs of the Globus system building an additional

Allocation Agreement Protocol.

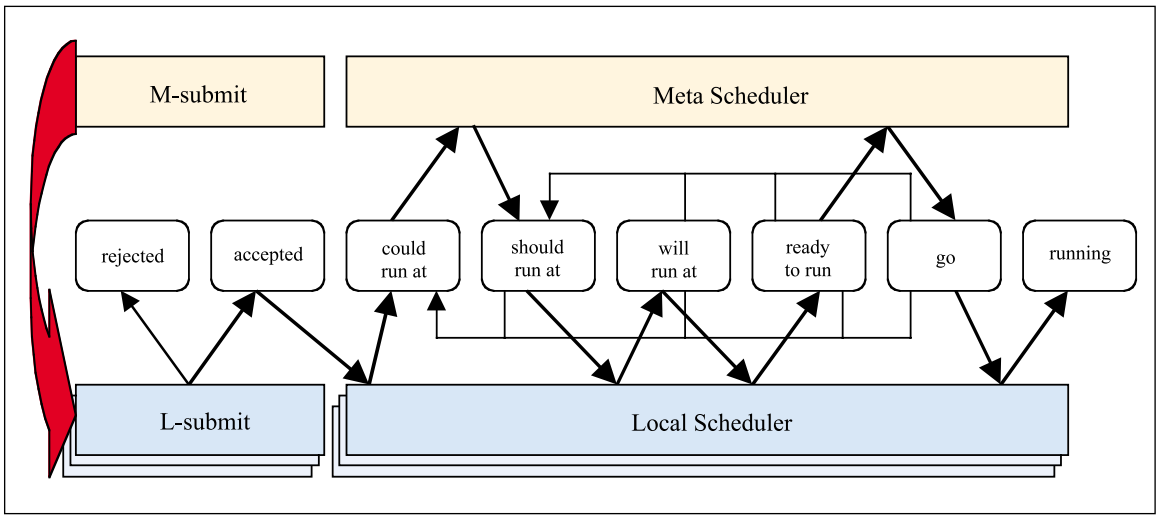

encapsulating layer that interfaces to local resource management systems.

However, this approach implies an 'overhead' which may not be desirable. We are well aware that there are other powerful systems like Globus, Legion or Unicore providing a broader range of integrated tools. As those systems make more and more use of evolving standards it becomes possible to exchange standard components against others that are more suitable and efficient in a certain context, eg a meta-scheduler may replace the standard scheduler to manage distributed parallel applications. In addition MeSch may also be used "stand-alone" providing a simple and efficient way of bundling distributed computing resources for the "bigger" parallel jobs of a user without the need to install one of the systems mentioned above.

The MeSch approach handles resource allocation as a global task which can be divided into subtasks that may be delegated to co-operating schedulers of the subcomponents of a Grid environment. Ideally, we won't discard local schedulers; instead, we build the metascheduler on top of the local ones. This allows us to build a hierarchy of schedulers.

In the same sense as a traditional scheduler maintains the nodes/processors as allocatable resources, the metascheduler does so with systems (or partitions of systems). The advantage is, that all subsystems can act in their usual way with their own policy. Moreover, allocation of processors remains in the responsibility at the local system level and is not explicitly done by the metascheduler. As subsystems remain responsible for allocation, the local use of subsystems is not affected. No restriction is imposed on local scheduling strategies and administrative policies.

The MeSch approach does not impose any restrictions on the type of Grid system: they may be homogeneous or heterogeneous, geographically distributed, any combination of MPP, cluster, and dedicated systems. However, MeSch requires some local scheduler attributes in order to be able to take over the burden of the overall scheduling task's global synchronisation: To provide simultaneous access to required resources, methods of getting reliable information about suitable time slots must be available. This information enables MeSch to determine a common time slot on all Grid components that are required for a Grid application. First subscheduler suggestions about available time slots in general will not lead to a solution for the complete metajob. Thus, we must be able to ask for alternative time slots to have a chance to determine a commonly agreed time slot for simultaneous access.

If a commonly suitable time slot can be determined, the MeSch metascheduler must be able to inform each subscheduler to reserve the time slot and to guarantee that it will allocate required resources at the agreed start time for the agreed time interval.

MeSch synchronization management requires several iterations of interaction with subschedulers to find a solution for a suitable time slot. Obviously, offered time slots must be (pre-)reserved by subschedulers, while they are under consideration for suitability. An allocation agreement protocol eases the synchronisation process by defining a set of states a job may have from an scheduling request to its final execution (see figure).

MeSch is a prototype metajob scheduler approach for a Grid environment. Its main advantage is that local scheduling policies are not affected by Grid jobs. The meta job scheduling can be viewed as using local schedulers as resource managers in a scheduler hierarchy. However, for an easy to implement allocation agreement protocol, local schedulers must provide a run time estimation facility for submitted jobs and accept and guarantee dedicated start time specification. The practicability of the approach has been demonstrated by a prototype implementation based on an enhanced EASY scheduler version.

Currently we are investigating how to implement scheduling for visualization devices such as a workbench for applications with real-time visualisation demand.

Link:
http:// www.gmd.de/SCAI/popcorn/

Please contact:
Gerd Quecke, Wolfgang Ziegler – GMD
Tel: +49 2241 14 2375, +49 2241 14 2258
E-mail: Gerd.Quecke@gmd.de,
Wolfgang.Ziegler@gmd.de

Open Services for Information Sharing among Legacy Systems

by Mark Roantree

The problems associated with defining interoperable protocols for heterogeneous information systems has been the subject of researchers for many years. While numerous solutions have been offered, the problem remains unsolved, mainly due to the distinctive nature of each target environment, and the uniqueness of each solution. This is not helped by the use of proprietary common data models or the use of non-open or non-standard technologies to provide an infrastructure and services. In the OASIS project, services and processors were defined and implemented using a platform of existing standards such as ODMG and CORBA. This provides a more open platform for the GRIDs architecture, and enhances forward interoperability with new systems over time.

The Interoperable Systems Group at Dublin City University conducts research in issues concerned with the sharing of complex objects between information systems. The OASIS project (www.compapp.dcu.ie/~oasis) which used ODMG databases as interfaces to integrate heterogeneous healthcare systems, was started in October 1998 and completed in December 2000. Research was focused on the usage of modern technologies, standards, and common database models rather than the introduction of new proprietary solutions.

This type of GRIDs architecture comprises five layers and numerous processor types for the integration of non-ODMG information systems. Any number of local systems may be integrated to form a federated schema. At the local information system level, two services are required to prepare information systems for participation in the GRIDs system: a wrapper service which binds the information system to a common model representation; and a view service to define an object-oriented subschema of shareable data.

Essentially, the Wrapper Service comprises a specification language for mapping ODMG to non-ODMG entities, and a language parser combined with a storage mechanism. Briefly, a wrapper may contain any number of entities specified in the form of a class; a class has a name and is comprised of attributes and relationships; and a class may have relationships with other classes. In other words, the wrapper language permits an ODMG-style specification. Due to its close association with the ODMG model, the wrapper specification language is known as ODLw, the Object Definition Language for Wrappers.

The View Service comprises a specification language for defining local and global views, and a language parser combined with a storage mechanism. The view language, ODLv (Object Definition Language for Views) is more complex than ODLw due to the nature of the operations involved. The purpose of the view language is to define subschemata which contain as much semantic information as possible, and to provide an operator set which supports the integration of subschemata. In figure 1, the view definition (an ODLv file) is processed by the View Service and stored as a collection of metaclass instances. A standard interface to the schema repository is crucial to both the view and display services. During the design of ODLv, it was necessary to extend the ODMG schema repository specification to include the concept of virtual types. These extensions maintained the structure and naming scheme of the original repository interface for the purpose of interoperability.

Once ODMG views are defined, they are passed to a global information server known as the Federated Kernel. Here they can be integrated using the same ODLv

Figure 1. *Definition of metadata for sharing with external Information Systems*

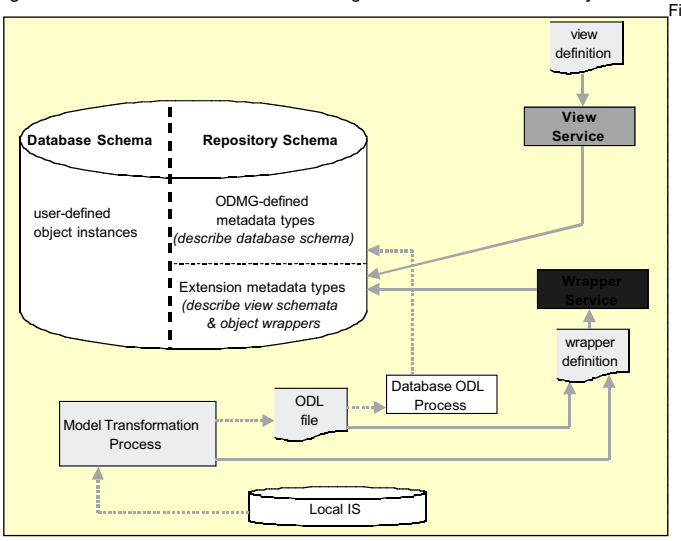

Figure 1: Definition of metadata for sharing with external Information Systems.

Figure 2. *Legacy Systems in the GRIDs Architecture*

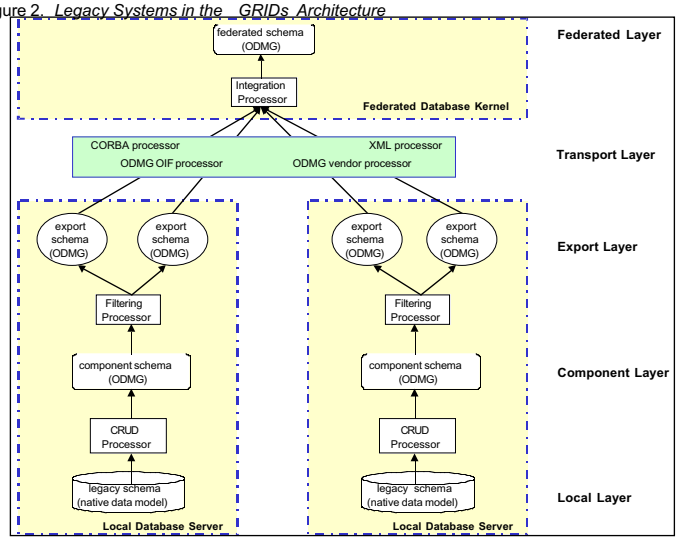

Figure 2: Legacy Systems in the GRIDs Architecture.

language to form global or federated schemata.

The Transport Processor is used to move data between local ISs (at the export layer) and the federated database kernel (at the federated layer). In practical terms metadata is moved when export schemata are extracted from local ISs and transferred to the federated kernel; query data is passed from the federated kernel to local ISs; data is moved to the federated kernel, or to local ISs which may request data. The Transport Layer (displayed in figure 2) contains different implementations of transport processors depending on the required function and the constituent parts which comprise the federation. This facilitates the participation of heterogeneous software and hardware systems in the GRIDs architecture.

The service architecture described in this article has been implemented sufficiently to create and query federated schemata. Several healthcare systems with ODBC interfaces were employed as participating systems, and using the view language, federated schemata were defined. The View and Wrapper Services were implemented for NT platforms using Java and ANTLR to code the BNF production rules for each language, and C++ to code the semantic actions to process and store view and wrapper definitions. The Versant object-oriented database was used to provide ODMG storage for component, export and federated schemata. Finally, at the transport layer, three versions of transport processor were implemented: a Versant (vendor) processor which provides a direct connection between each local and global Versant database for direct migration of meta-objects; an XML processor which converts ODMG objects into XML objects for transfer between any ODMG database in the federation; and a CORBA processor implemented in Orbix. A Display Service has been implemented to display local and global views, including those which have been integrated using any of the join operators earlier. It is implemented using C++ for NT platforms. Details of how to build the prototype system is available at the OASIS web page.

Link:
Oasis web page:
http://www.compapp.dcu.ie/~oasis/

Please contact:
Mark Roantree – Dublin City University
Tel: +353 1 7005 636
E-mail: mark.roantreecompapp.dcu.ie

Grid for Remote Sensing

by Giovanni Aloisio

The continuous monitoring of the earth's surface by remote sensing space systems produces huge quantities of raw data that are stored in the ground station segments of the National Space Agencies. These data must be processed to produce useful images and then delivered to the final user. Remote sensing imaging based both on traditional sensors and on more sophisticated radar technologies (SAR-Synthetic Aperture Radars) can provide a potential source of information for users working in many fields, such as archaeology, geology, drawing of maps, ecology and others.

Several Earth Observation Systems (EOS) have been developed by the National Space Agencies and web technologies are usually exploited to simplify the information access. An Earth Observation System (EOS) is made up of a sensor mounted on a satellite or a space shuttle that gathers images of the Earth's surface, of a distributed database in which these images are stored and of heterogeneous distributed computing resources which process and distribute them via web to the final users (see Figure 1).

Computing resources are needed to transform the raw data gathered by the sensor in an actual image. This preliminary step of processing is what is called pre-processing to distinguish it from the post-processing that may be needed to extract knowledge from the image. The real-time pre-processing of SAR raw data requires the use of high performance computing resources.

The traditional 'static' approaches normally used in the construction of EOS systems allow a user to request a pre-processed image of a place (if the required image is found in the distributed archive!) but they lack the intelligence needed to start a specific post-processing in response to high level user requests. On the contrary, a 'dynamic' EOS (DEOS) should be able to manage high level requests such as "show me the corn fields in the surroundings of that place" or "give me the image of that place and find the computer that can process that image in the fastest way and with the lower cost" (see Figure 2).

To satisfy these requests, a DEOS should have flexible management and control of the distributed resources. Further challenges for a DEOS are represented by the integration of information coming from different sources, eg, the integration of maps and images of the Earth's surface with any information associated with the geographic location; the integration of the functions of a friendly web browser with those of GIS and related technologies. A DEOS must also provide secure access to the distributed information and the steering of remote applications.

A Dynamic Earth Observation System can be implemented using Grid technologies. Grids can be thought of as a mixture of software and hardware infrastructure whose aim is to provide a coherent, unified view of geographically spread computing resources, smart instruments and distributed data archives. This seamless integration will enable users to share their resources and to build new classes of applications based on resource pooling and selection.

SARA-Digital Puglia (Synthetic Aperture Radar Atlas-Digital Puglia Project) is a Grid-based Dynamic Earth Observation

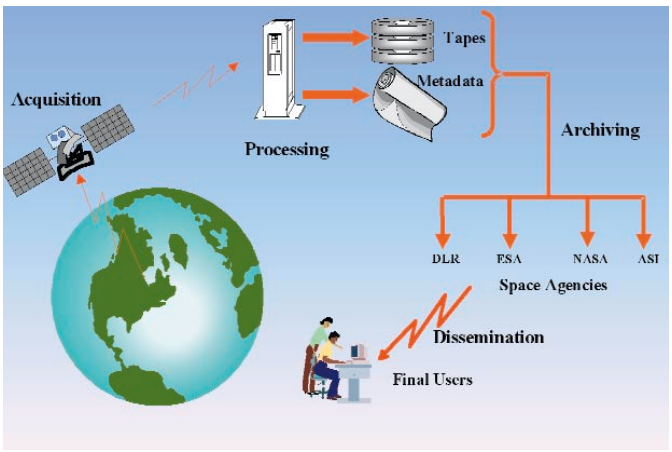

Figure1: Earth Observation System.

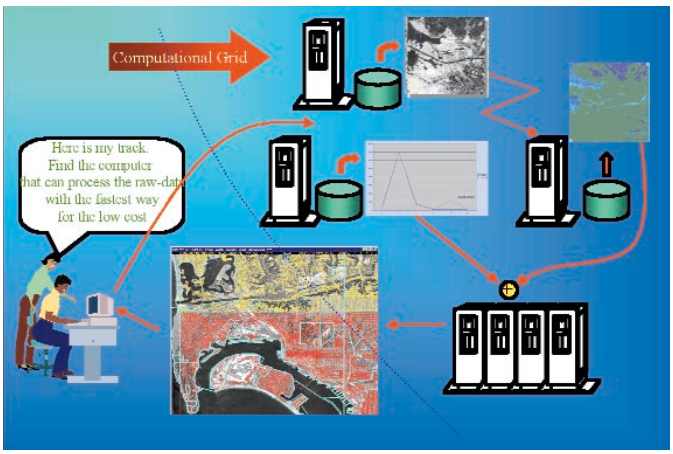

Figure 2: Dynamic Earth Observation System.

System. The system is under development at the ISUFI HPC Laboratory of the University of Lecce, Italy, in a joint project with the Center for Advanced Computing Research of the California Institute of Technology and the Italian Space Agency.

The first prototype implements an active digital library of SAR images of the Puglia region of Italy. The purpose of SARA/Digital Puglia is to allow a user to choose from a web browser a geographic area of interest and to specify, using high level requests, the information to be extracted from the selected images. High level user requests usually require not only the retrieval of the selected image but also further post-processing on it. This must be transparently activated on some of the remote machines belonging to the 'Computational Grid'.

A grid portal, the Grid Resource Broker (GRB) has been also implemented, to allow authorised users to create and handle computational grids on the fly. GRB was designed to be a very general grid tool providing location-transparent and secure access to Globus services using a simple web-based graphical user interface. Using the GRB, users do not need to write specialized code nor to rewrite their existing legacy codes. We plan to add other features, and to enhance the basic functionality provided, implementing several scheduling algorithms. We have also derived a library that can be used to build grid enabled applications from GRB. A demonstration of the GRB capabilities was given in the SuperComputing 2000 conference, held in Dallas, Texas. Another demonstration will be given in the context of the NPACI All Hands Meeting 2001, San Diego, California.

Links:
http://sara.unile.it/sara/
http://sara.unile.it/grb/grb.html

Please contact:
Giovanni Aloisio – University of Lecce
Tel: +39 30 832 320221
E-mail: Giovanni.aloisio@unile.it

Factoring Large Numbers on a Grid of Computers

by Herman te Riele and Walter Lioen

Probably the first grid computing project on Internet was initiated by Arjen Lenstra in 1988. The purpose was to answer the question: how large integers can we factor with our present algorithms? This question was, and is, of interest to cryptographers because any serious attempt to answer it gives crucial information about the security of the RSA public-key cryptosystem. This cryptosystem is used on a world-wide scale in secure communication. CWI has contributed actively to the state-of-the-art in factoring: as provider of algorithmic improvements, efficient implementations and a large computing grid, and as coordinator of the grid of people, computers and data storage resources which has led to the present world record: the factorization of a difficult 512-bit RSA modulus.

Factoring large numbers is a classical mathematical problem. The best methods available are complicated and extremely time-consuming, but the larger part of the computations can be split up trivially into smaller parts and distributed among as many processors as we like. Data transport is negligible compared with computing time. The most time-consuming part is a sieving step: essentially, on a large two-dimensional grid G of points (a,b), where both a and b run through a given set of consecutive integers, values of a polynomial $P(x,y)$ have to be found consisting only of small prime factors. Such 'smooth' values are extremely scarce, but there is one property of polynomials which can facilitate this search considerably: if we know that some integer q divides some polynomial value $P(a,b)$, we may add/subtract any

Figure 1: Arjen Lenstra (Citibank, New York), one of the pioneers of the Grid on Internet.

Figure 2: Breaking RSA-512 made international headlines.

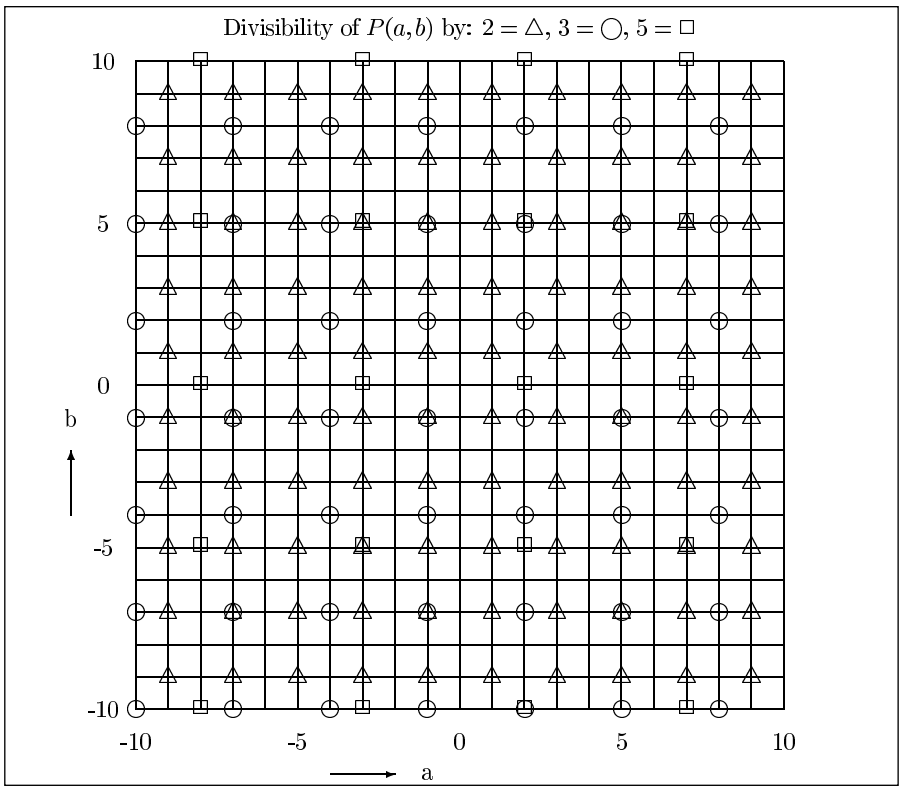

Figure 3:
2D grid for the polynomial $P(x,y) = x^2 - y^2 + 2xy - 3y + 1$
Δ means divisible by 2, grid points (1+2k, 1+2l)
O means divisible by 3, grid points (2+3k, 2+3l)
□ means divisible by 5, grid points (2+5k, 0+5l)

multiple of q to/from a and b, and still obtain a polynomial value which has q as divisor. For example, for the quadratic polynomial $P(x,y) = x^2 - y^2 + 2xy - 3y + 1$, we have $P(2,2) = 3$, so that, eg, $P(2,5)$, $P(5,2)$, $P(5,5)$, $P(-1,2)$, $P(-1,5)$ all have a divisor 3. Now, the sieving consists of marking all grid points $(2+3k, 2+3l)$ in G as having the prime divisor 3 *without* doing any trial division by 3. The larger G, the more efficient this sieving. This is carried out for all primes in a given set of small primes. Figure 3 illustrates this for the three primes 2, 3, and 5. Grid points marked by many small primes are selected as candidates for giving a smooth polynomial value. For the largest numbers which have been factored, the grid G contains between 10^{14} and 10^{15} grid points. This grid size is much too large to be handled by one processor of a modern computer, so it is split up into subgrids, which are distributed among different processors. Each processor receives the characteristics of a subgrid, sieves it in search of smooth polynomial values, and stores the values found. Processors are allowed to drop out, because what counts is the *number* of smooth values found, not which ones. The size of the grid G is chosen such that a drop-out of some computers in the network is taken into account.

Since 1994, CWI is running a 'factory': during night time and weekends the idle time of servers, workstations, and PCs is being used for carrying out the sieving step, needed to factor large integers. In our first approach, the factory was running continuously and we monitored the keyboard and mouse activity to stop the factorization processes. However, the user had to wait a couple of seconds for his machine to stop the factorization process, swap it out, and swap in his own process. After user complaints we ran the factory only during non-office hours. However, we still used the UNIX stop signal, and a large amount of swap space. So we still had complaints from people running out of swap space. Finally, we restarted jobs from the point where they stopped by looking at the already available output (as a sort of checkpoint). Workstation owners are able to stop the jobs by running a setuid program. By running our factorization jobs on UNIX machines at priority *nice* -9, we were able to secure sufficient CPU for them, and our current setup is running stable for years now. This factory has contributed to several factoring world records which have been established in 1999 and 2000 by an international grid of about ten sieving sites called 'The Cabal', coordinated by CWI. A second ERCIM participant in this project is INRIA Lorraine and Loria (Paul Zimmermann). The most appealing record was the factorization of a 512-bit RSA modulus: a number of 155 decimal digits. Since moduli of this size are used as actual RSA keys in E-commerce and communication protocols, this record attracted much publicity. One future aim in this project is to collect so much grid computing power, that it becomes feasible to factor a 640-bit RSA key within one year.

Please contact:
Herman te Riele or Walter Lioen – CWI
Tel: +31 20 592 4106/4101
E-mail: {hermanlwalter}@cwi.nl

The DataGrid Project

by Robin Middleton and David Boyd

The DataGrid project is a large European collaboration, supported by the EU, to develop a pan-European Grid infrastructure linking the various science Grids of the participants and to demonstrate its utility through demanding scientific applications. With 21 partners and a challenging technical agenda, the project is calling on the expertise of several ERCIM partners to help achieve its goals.

The LHC Challenge

The initial driver behind development of the DataGrid project was the recognition that the scale of computing and data management required by the particle physics experiments on the Large Hadron Collider (LHC), which will become operational at CERN near Geneva in 2005, far exceeded the capability or capacity of existing resources. With data volumes of several petabytes per year from all the experiments on the LHC feeding into the global particle physics community, who will then progressively reconstruct, filter and analyse the data, the aggregate computational and data throughput required is massive.

The DataGrid solution

To address this challenge, and similar requirements in other sciences, the DataGrid project was conceived and funded to the extent of approximately 10Meuro through the EU 5th Framework Information Society Technologies programme. In addition to trying to solve the problems of Europe's scientists, this 3 year project will have a wider remit to develop and prove a technological infrastructure which could potentially revolutionise commercial and social activities throughout Europe.

Project Structure

The DataGrid project logically has 4 major components:

- the underlying fabric of computational, data and communications resources required
- middleware to make these components accessible and controllable, initially based on the Globus toolkit
- management tools to monitor and control this infrastructure
- three application areas which will exercise the resulting Grid environment, namely particle physics, earth observation and bioscience.

The initial task, now underway, is to define the overall architectural vision of the project and to establish a detailed technical framework within which the project can progress.

WorkPackages

The above components have been further broken down into workpackages and these are currently defining their detailed work programmes and assigning tasks to the partners. The figure shows these workpackages and indicates the relationships between them.

- WP1 – Workload Management – will address distributed scheduling and resource management
- WP2 – Data Management – will develop and demonstrate the necessary middleware to ensure remote access to petabyte databases and the replication and caching of data in a secure environment
- WP3 – Monitoring – will produce the means for users and managers to monitor and optimise performance
- WP4 – Fabric Management – will develop new automated system management techniques to support the deployment and operation of tens of thousands of commodity processors
- WP 5 – Mass Storage Management – will agree and implement interfaces to mass storage systems in use within the partners
- WP6 – Integration Testbed – will evaluate effectiveness of the integrated DataGrid architecture for production use across European networks and provide a platform for computation by the applications
- WP7 – Networking Services – will oversee the networking aspects of the project
- WP8 (High Energy Physics), WP9 (Earth Observation) and WP10 (Biology) will build on the framework created by the other workpackages to demonstrate use of the DataGrid environment.

Project Partners

There are 6 main partners: CERN, ESA, PPARC (UK), CNRS (France), INFN (Italy) and NIKHEF (The Netherlands) with CERN as the co-ordinating partner and 15 associated partners from 10 countries across Europe. The industrial partners are IBM, Datamat and Compagnie des

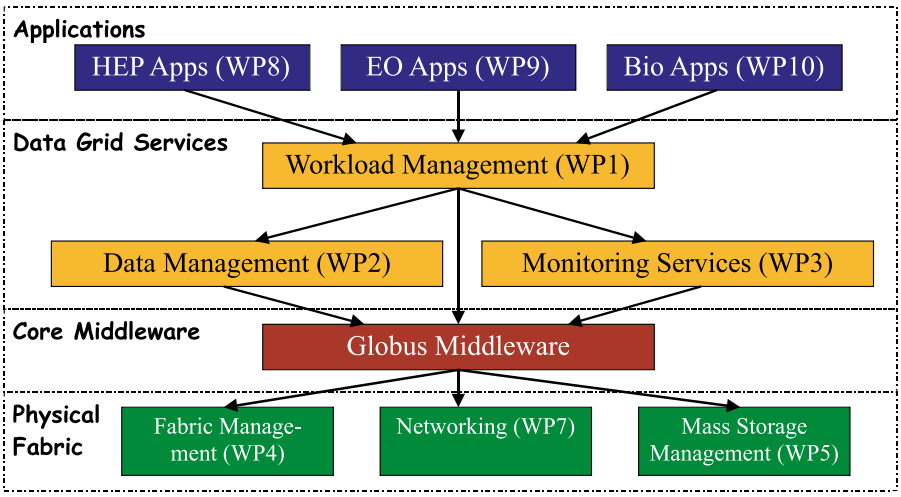

Workpackage relationships.

Signaux. They will be contributing their technical expertise and commercial experience and addressing the issue of how to effectively disseminate the new technology developed by the project into the marketplace so that European society and business can also benefit from these advances initially driven by science.

Collaboration

The project will be working very closely with several groups in the US including the Globus team, the Grid Physics Network (GriPhyN) and the Particle Physics Data Group (PPDG). By sharing technology and agreeing on joint development programmes, the resources of all these partners can be brought together most effectively to tackle what must be the largest global computing challenge ever undertaken.

Link:
DataGrid web site:
http://www.datagrid.cnr.it/

Please contact:
Robin Middleton – CLRC
Tel: +44 1235 446348
E-mail: R.P.Middleton@rl.ac.uk

The NL Datagrid Initiative

by Kors Bos

In the framework of initiatives on the European and world level to realize the Grid concept, the NL Datagrid Initiative proposes to set up an infrastructure in The Netherlands as a test bed for the grid software to be developed by the Dutch participants in the European Datagrid project, and for the software developed in that project as a whole. The national institute for nuclear and high-energy physics NIKHEF plays a key role in this local Grid initiative.

The European Commission has accepted in its IST Programme a proposal to develop a computational and data grid last fall. This follows a number of Grid projects initiated and proposed in North America, in particular the Globus project. A number of European companies and organizations have started the E-Grid Forum (http://www.egrid.org) , paralleling the American organization (http://www.gridforum.org). Recently the two have combined into the Global Grid Forum of which the first meeting took place in Amsterdam on 4-7 March 2001.

Whereas the US Globus project is primarily computational, the EU project focuses on datasets, described in databases where bulk data storage is widely distributed. Areas like particle physics, astronomy, bio-informatics, telecom, and E-commerce increasingly face this situation. The datasets to be processed are not big themselves, but they are many, and very distributed. Three areas serve as a test bed for the software to be developed in the European Datagrid Initiative: particle physics, earth observation, and bio-informatics. The main contractors are: Centre National de la Recherche Scientifique CNRS (France), European Space Agency ESA/ESRIN (Italy), Istituto Nazionale de Fisica Nucleare INFN (Italy), Particle Physics and Astronomy Research Council PPARC (UK), European Organization for Nuclear Research CERN (Geneva, coordinator), and NIKHEF (The Netherlands). Dutch assistant contractors are: Academic Computer Centre in Amsterdam SARA, and the Royal Dutch Meteorological Institute KNMI.

Providing an adequate computing infrastructure is left to the partners in the EU project. In The Netherlands it is estimated that several million guilders are required for this purpose. The work in the EU project is distributed over several Work Packages (WPs). NIKHEF and SARA will develop tools for the Grid middleware WPs Fabric Management and Mass Storage Management, and for the WP Network Services. Not surprisingly, NIKHEF and KNMI participate in the application WPs for High Energy Physics and Earth Observation, respectively.

NIKHEF's interests in Grid development are derived primarily from the computing needs in the LHC (Large Hadron Collider) era, in view of its involvement in the Atlas, LHCb and Alice experiments. KNMI is, apart from research connected with its operational tasks, a global change research centre. It will use the Datagrid for the exploitation of products from different (satellite) sources and the calculation of, eg, ozone profiles. SARA's interest is based on its long term strategy, where services will no longer be limited to its own resources. Here complex problems related to data distribution, brokerage, scheduling and resource optimisation, automatic disaster recovery, etc., have to be solved.

The Dutch participants in the European Datagrid project - NIKHEF, SARA, and KNMI - have started the national NL Datagrid initiative to set up an infrastructure in The Netherlands as a test bed for the grid software to be developed by them, and for the software developed in that project as a whole. The infrastructure includes groups of interested people, existing hardware and network facilities. Applications which will serve as a test, are real projects taken from particle physics and earth observation, which are carried out in any case, with or without a grid. SARA plays a special role as a service provider for any scientific enterprise in The Netherlands, and for its data storage expertise. The NL Datagrid initiative envisages a large compute farm at NIKHEF serving all ongoing experiments, as well as local farms at Nijmegen University (D0/Atlas), the Free University in Amsterdam (LHCb) and the University of Utrecht (Alice). The data will be stored at either the tape robot at SARA, at Fermilab, or at CERN. NIKHEF can use resources at SARA such as the new Cray/SGI

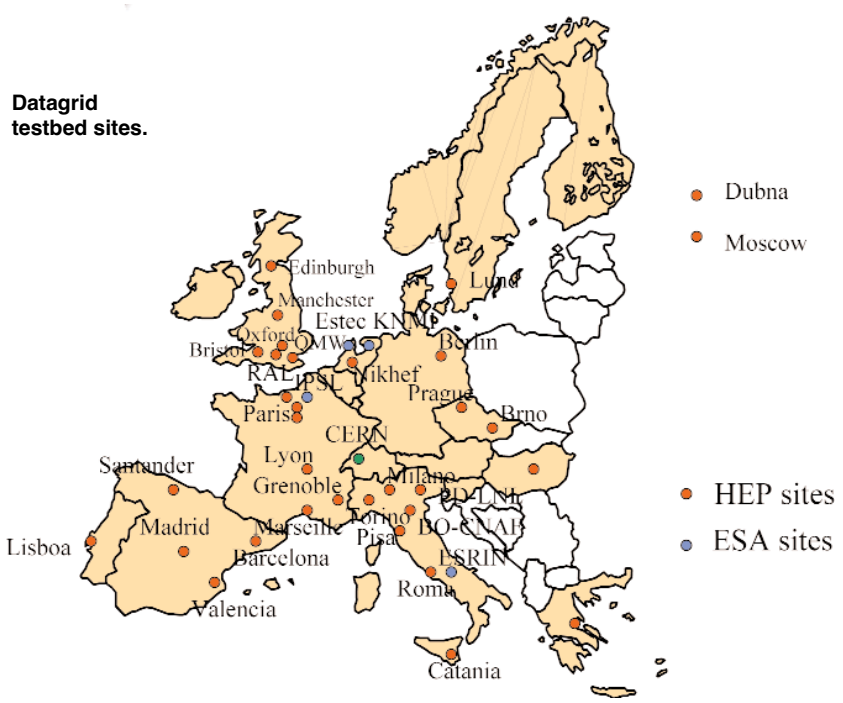

Datagrid testbed sites.

multiprocessor computer, and data storage in the SARA tape robot. KNMI should similarly be 'hooked up' to the high bandwidth infrastructure to profit from facilities elsewhere and make its own resources available to the other partners.

Several local parties have expressed interest in the NL Datagrid initiative, all bringing in their own often very relevant expertise. SURFnet, the National Academic Computing Service provider, considers the Datagrid project as an application candidate to challenge the limit of its network capacity. The SURFnet-5 backbone with a bandwidth of about 20Gb/s is expected to be operational in 2001, thus providing the necessary high bandwidth connection. The Computer Science group of the University of Amsterdam has experience with the Globus tools and contributes to the NL Datagrid initiative through the government sponsored Virtual Laboratory project, in which the Grid will be used for data transfer, processing and storage. The Department of Mathematics and Computer Science of Delft University of Technology is interested in some software aspects of the Datagrid. Five Dutch universities cooperate in the Distributed ASCI Supercomputer project (see elsewhere in this issue). The Datagrid initiative could very well profit from the computer science expertise built up here. CWI has experience in distributed multimedia databases and network management, and is interested in using the Grid for cell biology imaging by virtual reality engines. Finally, the national astronomy foundation ASTRON, hosting the Joint Institute for Very Long Baseline Interferometry in Europe JIVE, is interested in the common use of a broadband fibre network infrastructure (of at least 16Gb/s) for data transport from many radio telescopes in different countries to the data processing centre in Dwingeloo, The Netherlands.

Links:
DataGrid website: http://www.cern.ch/grid/
NIKHEF: http://www.nikhef.nl/
NL Datagrid Initiative: www.datagrid.nl
First Global Grid Forum: http://www.ggf1.nl/

Please contact:
Kors Bos – NIKHEF, Dutch National Institute for Nuclear Physics and High Energy Physics
Tel: +31 20 592 5083
E-mail: kors.bos@nikhef.nl

The Particle Physics Grid Programme at CLRC

by John Gordon and David Boyd

Particle physics will present some of the most challenging data handling and computing problems over the next few years, culminating in operation of the Large Hadron Collider (LHC) at CERN near Geneva in 2005. The volume of data which will be produced is staggering, several petabytes per year from all the LHC experiments combined. Meanwhile experiments at other particle physics laboratories are producing data which is driving prototype Grid developments.

LHC Tier1 Centre

The deluge of real data will start in 2005 but its forerunner, the Monte Carlo simulations needed to develop the experiments and prepare the data analysis system, is already here. The proposed solution is to devolve much of the data processing to about 10 regional centres, known as Tier1 centres (CERN itself is the Tier0 centre). Each centre will be equipped with petabyte-scale data storage facilities and computing farms containing thousands of commodity processors to process the incoming data from CERN and pass it on to the hundreds of research institutes and groups which will analyse it further to extract new physics results.

CLRC's Rutherford Appleton Laboratory (RAL) will act as the Tier1 centre for the UK and is now working closely with other UK particle physics groups to build a prototype testbed system using the Globus toolkit. This testbed is already carrying out simple computing tasks using these tools such as remote job submission and data retrieval. The diagram shows schematically the multi-tier hierarchy of centres.

With such a large distributed system, controlling access becomes a serious problem. The current solution to this is to establish a trusted Certification Authority (CA) which issues digital certificates to members of the community. These certificates are the passports to entry to the system and access to its growing computing resources. RAL now operates a CA on behalf of the UK particle physics community.

CDF

CDF is an experiment on the Tevatron at Fermilab. The CDF institutions in the UK (Glasgow, Liverpool, Oxford, University College/London) and RAL are planning to use Grid distributed computing ideas partly to help with CDF data analysis and partly as a rehearsal for computing for the LHC experiments at CERN. They are currently developing a demonstration project.

The plan is to implement three user services over five sites coupled using components of the Globus software toolkit. The three services are:

- The presentation of datafile catalogue metadata about datasets from the local catalogues at each participating site. The Globus capability used for this service is an LDAP server acting as a front end to the existing local metadata catalogue which has been implemented using an RDBMS. The LDAP server is used currently to supply information about local computing and disk resources over the Grid network.
- Remote job entry between sites. These Globus tools exist and interface to various batch systems for running the jobs. In the initial version, the user chooses where to send the job but a later development will choose the best location.
- Copying data between sites when bandwidth is available. A user interface accessing the information from the various catalogues will present the user with the information he needs to copy data and metadata to a convenient location and to submit his job. The replication would initially be driven by the user. This will be developed into an automatic policy-driven system later.

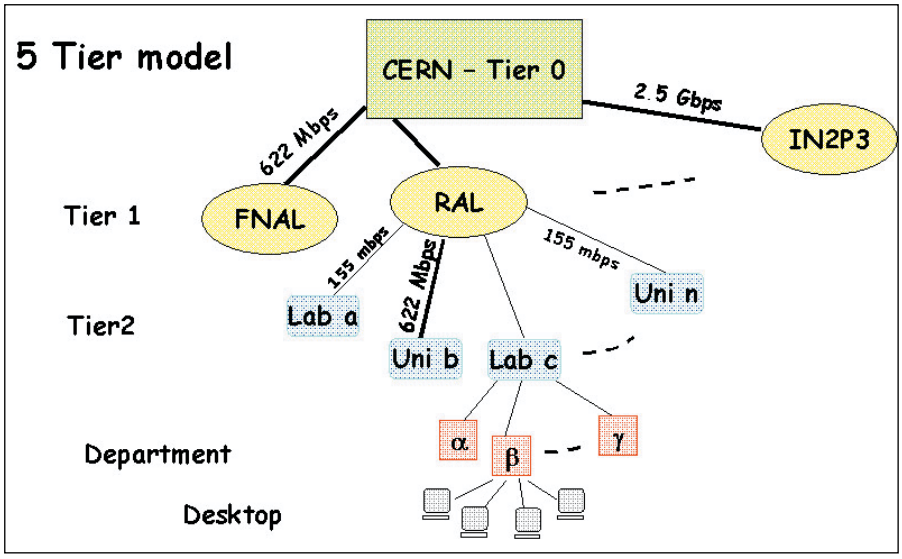

Schematic of the 5 Tier model for LHC data processing.

After successful completion of this demonstration project, the CDF-UK groups plan to ask the CDF team at Fermilab to install Grid software so that their central analysis cluster could participate in these three services.

BaBar

BaBar is an experiment on CP violation at the Stanford Linear Accelerator Centre (SLAC). The BaBar collaboration already operates a distributed database cache in the UK which is used to hold experimental data for analysis at 9 collaborating universities. This cache system is currently being augmented by a 600 processor, distributed, simulation and analysis facility. The resources located at each site primarily support local researchers, but developments have started to implement a Grid infrastructure based on the distributed data catalogue and to give all collaborators access to all sites to take advantage of data cached elsewhere and to use spare CPU capacity for simulation studies.

The main data store for the UK BaBar collaboration is at RAL and plans are now well developed to significantly enhance this centre. Current plans call for disk and tape storage of order 100TB by the end of 2002. This storage at RAL will be augmented by processing facilities to support bulk data queries/processing and the system will be fully integrated into the existing distributed database cache. These BaBar developments are complimentary to the prototype LHC Tier1 centre which will be developed at RAL.

Link:
UK High Energy Physics Grid web site:
http://www.hep.grid.ac.uk/

Please contact:
John Gordon – CLRC
Tel: +44 1235 446574
E-mail: j.c.gordon@rl.ac.uk

e-Science and Grids at CLRC

by David Boyd and Paul Jeffreys

e-Science is enabling increasingly complex and challenging scientific problems to be addressed through the use of advanced IT. Current e-Science activities at CLRC are putting in place the groundwork for a new multi-million pound three year programme starting in April 2001. The pilot projects described here are exploring how Grid techniques can enhance the multi-disciplinary science programme supported by CLRC.

StarGrid

The StarGrid project, involving CLRC's Space Science and Technology and IT Departments, is integrating the Globus toolkit into the KAPPA package which is part of the Starlink astronomical data processing software suite. This enables image files to be retrieved from remote data archives using Grid techniques, processed locally and, if necessary, reloaded into the archive. The figure shows KAPPA accessing an archived image. The final goal of the project will be to integrate these facilities into the GAIA graphical user interface. This project will provide useful experience for the Astro-Grid project when it starts.

Earth Observation Grid

The British Atmospheric Data Centre (BADC) is located at CLRC's Rutherford Appleton Laboratory (RAL) and makes a variety of data resources available to the UK environmental research community. This project is investigating how Grid techniques can improve the services which the BADC offers its users including being able to remotely identify and retrieve data from a range of sources, if necessary with subsequent processing to reduce the volume of data which needs to be transferred over the network. The project will also review ways in which the Grid can improve UK access to data from instruments on the ENVISAT satellite to be launched by ESA in mid-2001.

SMART Experiments

Many university-based researchers who carry out experiments at CLRC's two accelerator-based facilities, the SRS at Daresbury Laboratory (DL) and ISIS at RAL, travel to the facility in order to be able to control the experiments as their data is collected. This project, which is a collaboration between teams at DL and the crystallography group at Birkbeck College in London, is studying ways in which the Grid can facilitate remote control of instruments and rapid recovery of data back to the user's university so it's quality can quickly be checked. These studies are paving the way to increase the use of automated beamlines on these facilities in the future.

HPC in the Grid

This project aims to develop a computational Grid which will enhance access to high performance computers and enable novel combinations of computer simulation, remote measurement and data analysis. Computing groups at DL and RAL are co-operating with the universities of Edinburgh and Manchester to form the UK High End Computing Consortium. This consortium has initially established a working Grid infrastructure using the Globus toolkit. This is now being exercised with a variety of applications including remote medical imaging to help surgeons during operations, combining MRI scanned data of a tooth with the results of a finite element analysis in a VR facility and real time flood warning analysis to assist emergency services.

Grid Reference

With any new development such as the Grid, it is necessary to get many people up to speed quickly and to define best practice so experience gained is quickly shared. This project is achieving this through establishing a reference Grid implementation platform within CLRC's IT department which can then be cloned by teams in other departments developing Grid applications. The reference platform is a Linux PC running Grid middleware based initially on the Globus toolkit. The project is also operating a Grid Certification Authority, currently to the UK particle physics community. An internal Globus technical forum is held regularly to disseminate newly acquired technical knowledge.

Gigabit Networking

With many scientific programmes rapidly increasing the volume of data which they

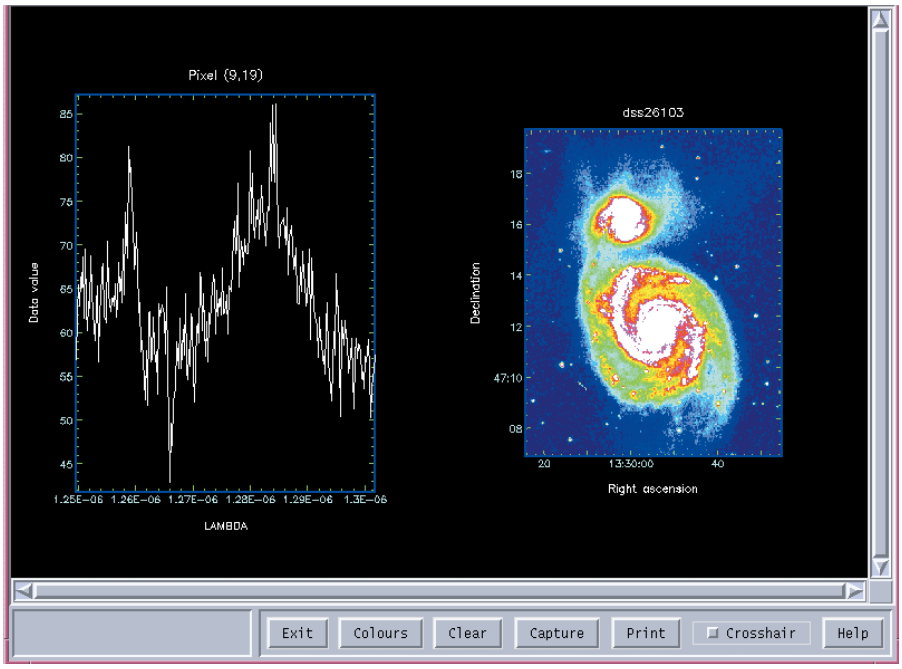

The Starlink application KAPPA accessing image data remotely using Globus software.

generate, for example by using more complex computational models or higher resolution experimental detectors, the bandwidth of networks through which the data must be moved becomes a limitation. The development of Grid technology will increase this trend through making it easier to move data around as part of a distributed data analysis process. Within CLRC, the internal networks must carry data between experimental facilities, computers, disk and tape stores, and external locations. To meet the growing data traffic requirements, local area networks at RAL and DL are being increased to a gigabit capacity now with further increases to 10 gigabits planned within 2 years.

Petabtye Data Storage

The Atlas Data Store (ADS) at RAL provides secure and affordable long term storage for experimental and computational data from many of the scientific facilities in CLRC and elsewhere. Data curation will be increasingly important as the cost of facilities, and therefore of the data they produce, continues to rise. The future demands of the particle physics community in particular will require major upgrades to the ADS capacity over the coming years. A major increase towards petabyte capacity is currently underway with further increases planned over the next 3-5 years. As the Grid provides easier access to large scale data storage facilities, intelligent data access will become an increasingly important issue. Metadata-based data location tools are currently under development in CLRC to meet this need.

Link:
CLRC e-Science web site:
http://www.e-science.clrc.ac.uk/

Please contact:
David Boyd – CLRC
Tel: +44 1235 446167
E-mail: d.r.s.boyd@rl.ac.uk

Grid Computing with Off-the-Shelves Middleware Technologies

by Christian Perez and Thierry Priol

JACO3 is a Grid environment that supports the execution of coupled numerical simulation. It is being developed jointly by EADS, INRIA, INTECS and KTH. It is a set of CORBA services that manage remote computing resources through the Internet.

The PARIS research group at INRIA-Rennes is conducting research activities in Grid computing. A Computational Grid acts as a high-performance virtual computer to users to perform various applications such as for scientific computing or for data management. It is made of several computing resources interconnected together by various networks. Although Grid Computing is still an emerging academic research field, it is foresee that the industry will soon make a stronger demand on Grid environments. Indeed, today in the industry the trend is to replace prototyping as much as possible by simulation to reduce costs and time to market. This new process is highly demanding in computing power and often involves partnerships with world-distributed sub-contractors. In such a distributed and competitive environment, antagonist requirements arise such as the need to preserve each partner's intellectual property, while permitting collaborative visualization and analysis of the results. A Grid environment should give a competitive advantage to carry out such numerical simulation in a distributed way.

The PARIS research group is involved in a research and development project aiming at developing a distributed simulation environment allowing industrial partners to work together while preserving their intellectual property. The development of this environment is carried out in the context of a European Union funded R&D Project named JACO3. JACO3 stands for Java and CORBA based Collaborative Environment for Coupled Simulations. The implementation relies heavily on off-the-shelves middleware technologies which are well accepted by the industry. More precisely, the JACO3 environment is based on the CORBA middleware from the OMG on top of which several services are introduced to solve specific problems such as managing user sessions, security, data movement or computing resource allocation. A code coupling tool allows chaining of simulation codes and steering of the whole computing process. A visualization toolkit as well as custom visualization tools ease the online monitoring and final interpretation of the results by the specialized scientists. These tools are based on Java to achieve portability. The environment is being tested with two distinct applications for distributed numerical simulations. One application is to simulate the same physics on different objects, each of them being simulated on a separate computing resource to avoid the exchange of confidential data (ie the model of the object to be simulated). In the JACO3 project, Saab-Ericsson Avionics and Allgon are collaborating to study the electromagnetic coupling of an antenna on an aircraft. The second application is the simulation of different physical behaviour that encompass multi-disciplinary aspects: structural mechanics, computational fluid dynamics, electromagnetism, noise analysis, etc... This is typically what is being done by Alcatel Space Industries, who is studying simultaneously dynamics, structural mechanics and thermal behaviour of a satellite.

Besides the development of several of the JACO3 CORBA services, the PARIS project is studying how to combine parallel and distributed programming. More precisely, our objective is to allow the introduction of a parallel SPMD (Single Program Multiple Data)

programming style in the design of CORBA objects. In this matter, we aim at encapsulating parallel SPMD codes (such as MPI-based ones) into CORBA objects in such a way that simplicity and efficiency are achieved. With such approach, it would be possible to couple simulation codes together using CORBA. However, one of the key problem to solve is let simulations codes exchange data as efficiently as possible.

To reach this goal, we are investigating several directions such as the design of fast Object Request Broker (ORB) and the design of a new kind of CORBA object we called parallel CORBA object. The first direction is motivated by allowing an ORB to best use the underlying communication networks that are available within clusters of computing resources. Most of the current ORB takes only benefit from an Ethernet network but few of them from System Area Network (SAN) such as Myrinet or SCI. We are studying how to design an ORB on a low level communication layer, such as Madeleine (http://www.pm2.org), in order to efficiently transport CORBA requests on such high speed networks.

The second direction aims at allowing a scalable connection between CORBA objects when they encapsulate SPMD codes. We introduced the concept of parallel CORBA object as a collection of identical CORBA objects. Each object of the collection is connected to the ORB since it is standard CORBA object. Therefore, communication between two parallel CORBA objects can be performed with several communication links allowing a better usage of networking resources if available. The calling of an operation to a parallel CORBA object by a client will result in the execution of the associated method by all objects belonging to the collection at the server side. This parallel activation is done transparently by the system. Data distribution between the objects belonging to a collection is entirely handled by the system. However, to let the system carry out parallel execution and data distribution between the objects of the collection, some specifications have been added to the object interface. A parallel object interface is thus described by an extended version of IDL, called Extended-IDL. It is a set of new keywords, added to the IDL syntax to specify the number of objects in the collection, the shape of the virtual node array where objects of the collection will be mapped on, the data distribution modes associated with parameters and the collective operations applied to parameters of scalar types. The result of this work formed the basis of a response to the Aggregated Computing RFI (Request For Information) issued by the OMG (Object Management Group) that is the organization in charge of the CORBA standard.

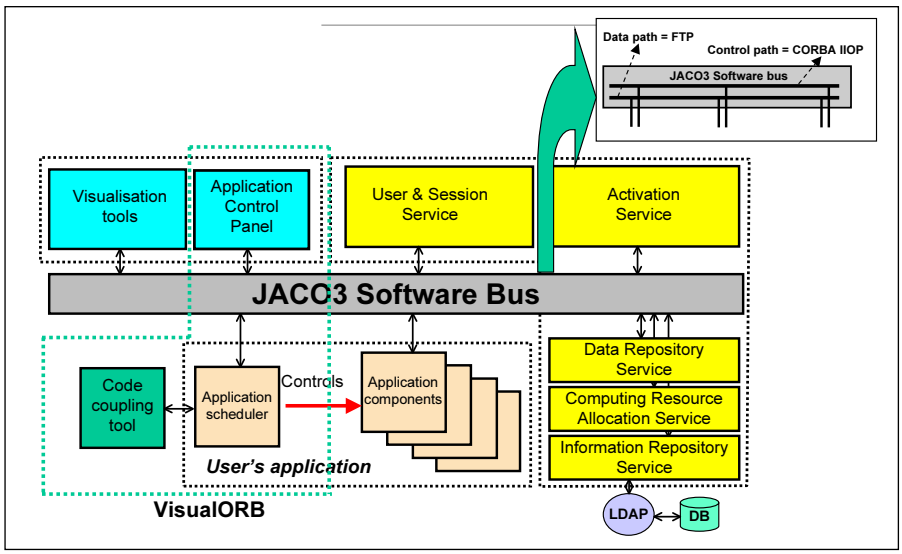

Figure 1: The software architecture of the JACO3 environment is based on several CORBA services connected to the JACO3 software bus.

Our current work is focusing on a portable version of the concept of parallel CORBA object. We are investigating how to implement such concept without modifying the CORBA standard. We are also starting an activity related to introduction of a SPMD programming model within the CORBA Component Model. The objective is perform the coupling of simulation codes using a component model instead of a distributed object one.

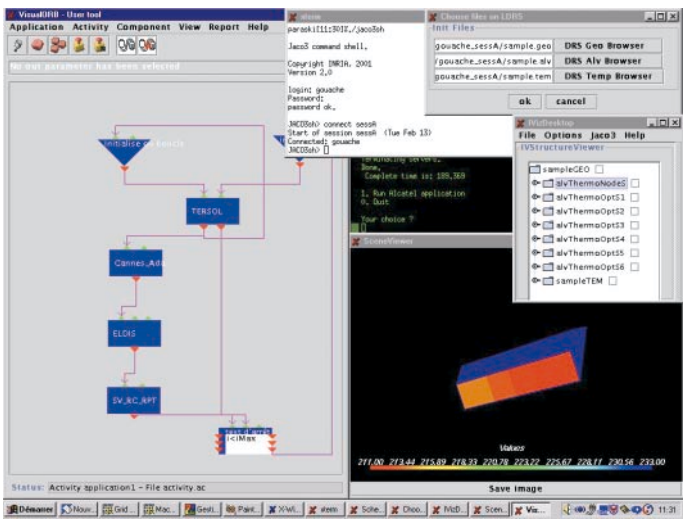

Figure 2: Execution of the ALCATEL application with the JACO3 environment. VisualORB, from EADS, is used to model the interaction between simulations codes. Visualization tools, from INTECS, have access to the data repository to perform real-time visualization during the simulation.

Link:
JACO3 Project:
http://www.arttic.com/projects/JACO3

Please contact:
Thierry Priol – INRIA
Tel: +33 2 99 84 72 10
E-mail: Thierry.Priol@inria.fr

Christian Perez – INRIA
Tel: +33 2 99 84 72 05
E-mail: Christian.Perez@inria.fr

ServerGrids: Available and Scalable Global Computing Platforms

by Evangelos Markatos

The widespread use of computing and communication devices in personal computers, (mobile) telephones, vehicles, and home appliances, has given rise to a new computing paradigm that is mainly based on co-operating personal computers and thin clients. These clients usually have low-computing capabilities (compared to current desktop computers), limited (battery-based) power supply, and a slow communication connection. Scientists at ICS-FORTH propose the use of powerful ServerGrids to make the cooperation of such thin clients effective.

ServerGrids are networks of computers that provide an available and scalable service to their clients. Limited versions of ServerGrids have already started to appear. For example, existing ServerGrids include the Domain Name Service (DNS) system which translates Internet domain names to IP addresses, several existing networks of cooperating web proxies, and the content-delivery networks that replicate content of popular web servers all over the world and transparently serve clients from the closest replica. Examples of future ServerGrids include video-on-demand servers, Internet-based distribution networks of television channels from anywhere in the world, digital libraries, medical information systems, multinational department stores, etc. Clients of current and future ServerGrids include traditional Personal Computers, PDAs, Mobile phones, and various computing devices embedded in home appliances, vehicles, etc.

ServerGrids, and in general similar approaches to the creation of networks of scalable and available servers, will transform the way we live, the way we work, and the way we entertain ourselves. For example, capitalizing on the existence of ServerGrids, mobile phones have started to become indispensable communication and entertainment devices, much like a telephone, a radio, TV, and VCR combined together. Besides calling other people, mobile telephones will soon be able to download and play songs, movies, and receive TV Channels. Mobile phones will be used more and more to surf the web (or its descendant), locate people, alert their owners of friends and family that happen to be nearby, and notify about that long-wanted tie that went on sale at the Marks & Spencer down the road. Besides mobile phones and PDAs, a wide variety of other devices will soon be equipped with a computing and communication capabilities. For example, cars, which already have several embedded microcontrollers, will consult ServerGrids about maps of the area, nearby hotels, possible obstructions ahead, detours, and service stations. Based on information downloadable from ServerGrids, cars will help drivers find the shortest possible way to their destination bypassing heavy traffic and avoiding adverse weather conditions. Cars (through the use of ServerGrids) will be able to guide drivers through their journey and make sure that they are always on the right road. In return, cars will send information about current road conditions to ServerGrids.

ServerGrids are an excellent platform to execute (mobile) agents that perform a wide range of tasks on behalf of their owners. For example, shopping agents will continually search on-line stores to find (new or used) products that match a user's preferences, scientific information agents will discover papers that match a researcher's interests, stock market agents will continually search for (world-wide) stock market opportunities, television agents will locate channels (anywhere in the world) that play a user's favourite cartoon. But, besides making our life easier, ServerGrids can make our life safer as well. For example, car accidents will be reported the second they happen: the moment a car crashes, a microcontoller will broadcast the location and conditions of the accident; paramedics will be on their way to help the driver and passengers immediately. By interacting with ServerGrids, cars will alert their drivers about severe weather conditions that may be up ahead, suggesting an alternate route. Home appliances will also report adverse operating conditions that according the ServerGrid's databases may lead to damage or even fire.

One could parallel future ServerGrids with the current (mobile) telephone network. Currently, several people have lightweight (mobile) telephones and can talk to each other through local telephone centers that wire calls through a maze of telephone centers and switches. If a switch goes down, calls can (to some extent) be re-routed through other switches, trying to keep the telephone service as available as possible. Similar in spirit, ServerGrids will provide more advanced communication services with higher availability.

To make ServerGrids a reality, several challenging research problems need to be addressed. Two of them are availability and scalability.

Availability is the first and foremost requirement of ServerGrids. ServerGrids must provide their useful and vital services around the clock. When a service is not available, profits may be lost, customers may switch to other providers, cars may crash, and lives may be in danger. To put it simply, a ServerGrids' down time should be less than a few seconds per year, a goal that is 3-5 orders of magnitude away from the availability even the most advanced current web servers provide.

The next most important property of ServerGrids is scalability. ServerGrids provide services to various clients spread over several different countries in different continents. ServerGrids will be required to serve millions of clients per second, a goal that is two orders of magnitude away from the number of clients even the most advanced current content delivery networks serve.

Finally, ServerGrids must be reliable and easy to maintain. Experience with power networks and telephone networks suggest that human error is the single most important factor of failure. Therefore ServerGrids should be able to maintain themselves, so that human intervention is kept to a minimal.

Therefore, the major research challenges that ServerGirds face are availability and scalability. In addition to them several other research areas need to be explored. Such areas include security, authentication, resource management, division of work between fat servers and thin clients in order to improve performance and conserve energy (on the client side), etc.

In order to achieve these requirements, the computers of a ServerGrid will probably be hierarchically build, clustered in several centers which will be physically distributed all over the Globe. The client population will be dynamically distributed among those centers according to the current load, and network conditions. In case of failure or overload, clients will be transparently redirected to other centers, making every possible effort to keep the service available at all times.

Summarizing, ServerGrids are networks of computers that provide a scalable and available service to (usually thin) clients. Although small-scale versions of ServerGrids have already started to appear, major research challenges need to be met, before they can be widely deployed to a variety of applications.

Please contact:
Evangelos P. Markatos – ICS-FORTH
Tel: +30 81 391655
E-mail: markatos@csi.forth.gr

The M3-Project - Living in Simulated Worlds

by Helge Rosé and Matthias L. Jugel

Complex real world systems are currently becoming a decisive instrument for IT-supported problem solving for a great number of problems in science, economy and society. This enables us for the first time to find new solutions, visions, and strategies for sustainable development on the basis of complex, real world simulations.

The main objective of the M3-project is the realization of a distributed, interactive real world simulation system. The idea of embedding multidisciplinary simulation models in a virtual reality environment involving real human actors in the simulation is fundamental. The distributed design of the M3-system aims at an integrated simulation of complex processes belonging to heterogeneous levels of reality, paying specific attention to real human modes of conduct. It represents a generic implementation of the GRID concept in the field of real world simulation.

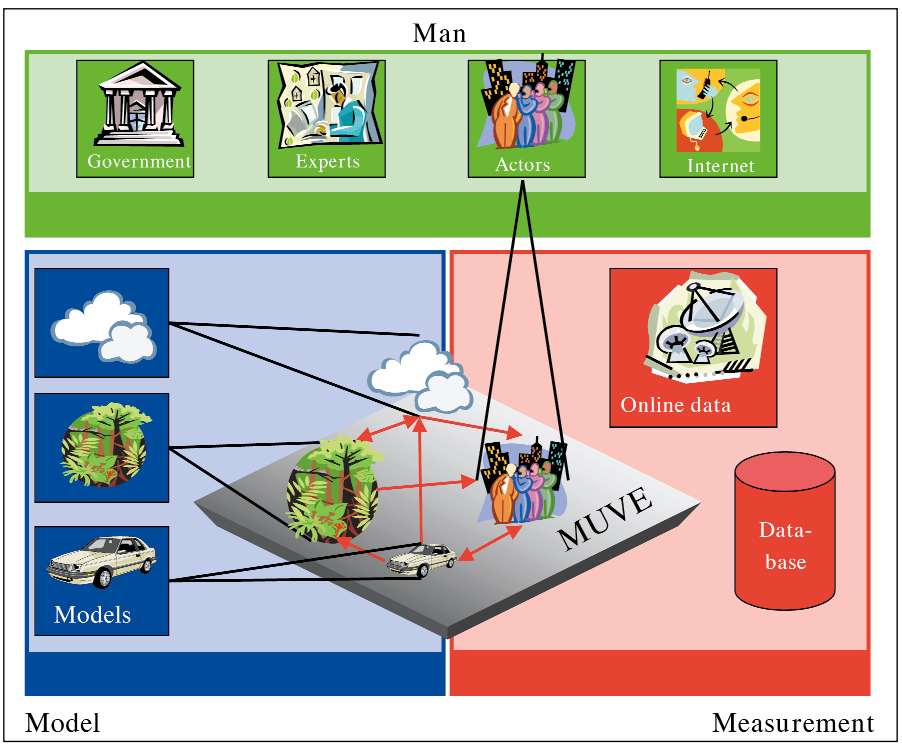

General structure of the M3-System.

The social behavior of man doubtlessly plays a major role in the investigation, simulation and sustainable management of ecological, economical and political systems. A simulation system that attempts a realistic consideration can not only simulate human behavior - but has to involve real actors directly into the simulation. They participate in the simulation by their actions and can develop strategies securing long-term development. On the other hand, they can experience possible consequences of their behavior and, thus, improve their strategic decisions in a process of evolutionary learning. A simulation based on scientifically-founded simulation models, using current data of measurement and actively involving man in this way, we refer to as a 'Man-Model-Measurement simulation'.

M3-simulation combines the scientific approach of conventional simulation

systems with multimedia capabilities from the technology of virtual reality in a distributed way. The four main features of the system are: Modeling and Data management of real systems by scientifically-founded models and generalized databases, Visualization and Information by intuitive representation using a VR-environment, Communication and Interaction by active involvement of people in the simulation, Evolution and Enhancement by an open software API. The general structure of the M3-system is unified by a Multi User Virtual Environment (MUVE) representing the logical structure of the simulated world and coordinating the communication between the three main components of the framework: Multi-Purpose Graphical User Interface for visualization of the simulated world in specialized views adapted for the different user groups (actors, experts, decision makers, publicity), Model Server providing a distributed network of scientific simulation models connected via generic control and data mapping interfaces (RMI, XML), Measurement Database providing current measuring data, reference configurations and parameters of the real environment in question.

A crucial advantage of modeling a real world system by means of a MUVE is given by the possibility to map the real world into the simulated world in a consistent object-oriented manner. A state-process-relation model of the objects which constitute the MUVE is not only essential for building a virtual image of the considered reality, it is also a promising approach to overcome the 'coupling problem' of modeling multidisciplinary systems. The creation of the M3-world in orders of magnitude is more complex than most current available MUVEs. To investigate new software architectures to handle such complexity will be one of the major tasks. The control and data communication between the models will be realized by specialized markup languages which employ the XML technology for easy extension and validation.

The realization of the M3-concept is an encouraging and very challenging enterprise. Successful realization necessitates the far-reaching cooperation of all involved scientific fields. In fact, the interactive character of the M3-world itself provides an efficient environment for development, fast communication and cooperation. Because the M3-framework is not only a 'software' but a simulated world with real humans, it has the ability to evolve itself. For this it is necessary to conceive the system as an open software project with a controlled programming interface. In this way, the M3-concept will not only simulate the real world - it will also employ its natural process of creation and advancement to reach the ambitious aim of opening up a new kind of simulation initiating manifold possibilities of application for the investigation, simulation and imparting of knowledge, as well as for management of real world systems in the fields of:

- science: as integration instrument
- economy: as planning and marketing instrument
- policy: as instrument for decision support and strategic representation
- education: as instrument for comprehending complex context
- entertainment: as basis of a new generation of reality-based simulation games.

Links:
The M3 Project: http://mmm.first.gmd.de

Please contact:
Helge Rosé – GMD
Tel: +49 30 6392 1835
E-mail: rose@first.gmd.de

The CLRC Data Portal

by John Ashby, Juan Bicarregui, David Boyd, Kerstin Kleese-van Dam, Simon Lambert, Brian Matthews and Kevin O'Neill

A web-based data portal is currently under construction and based on a new metadata model of scientific data, for exploring and accessing the content of the data resources held within CLRC's main laboratories. This system comprises a web-based user interface incorporating access control and a metadata catalogue which interfaces to distributed heterogeneous data resources.

CLRC operates several large scale scientific facilities for the UK research community including accelerators, lasers, telescopes, satellites and supercomputers which all create copious quantities of data. These data resources are stored in many file systems and databases physically distributed throughout the organisation with, at present, no common way of accessing or searching them to find what data is available. It is often necessary to open and read the actual data files to find out what information they contain. There is little consistency in the information which is recorded for each dataset held and sometimes this information may not even be available on-line, only in experimenters' log books. This situation could potentially lead to serious under-utilisation of these data resources or to the wasteful re-generation of data. It could also hinder the development of cross-discipline research as this requires good facilities for locating and combining relevant data across traditional disciplinary boundaries.

To address these problems, a web-based data portal is being developed with the aim of offering a single method of browsing and searching the contents of all the CLRC data resources through use of a central catalogue holding metadata about all of these resources. The structure and contents of this catalogue are based on a metadata model for representing scientific data which is being developed by the project. Extensive use is being made of XML and related W3C standards for representing, transferring and manipulating the metadata.

The objective of this demonstrator project is to prototype these ideas by developing a pilot implementation of the proposed system which will enable researchers to access and search metadata about data resources held at the ISIS and SRS accelerator facilities in CLRC.

System Architecture

The system being developed has 3 main components:

- a web-based user interface
- a metadata catalogue
- data resource interfaces.

In the pilot system these are integrated using standard Web protocols. Eventually it is anticipated that the system will exploit the emerging Grid infrastructure to offer a distributed interface to scientific data resources both inside and outside CLRC. The figure shows the architecture of the Data Portal system.

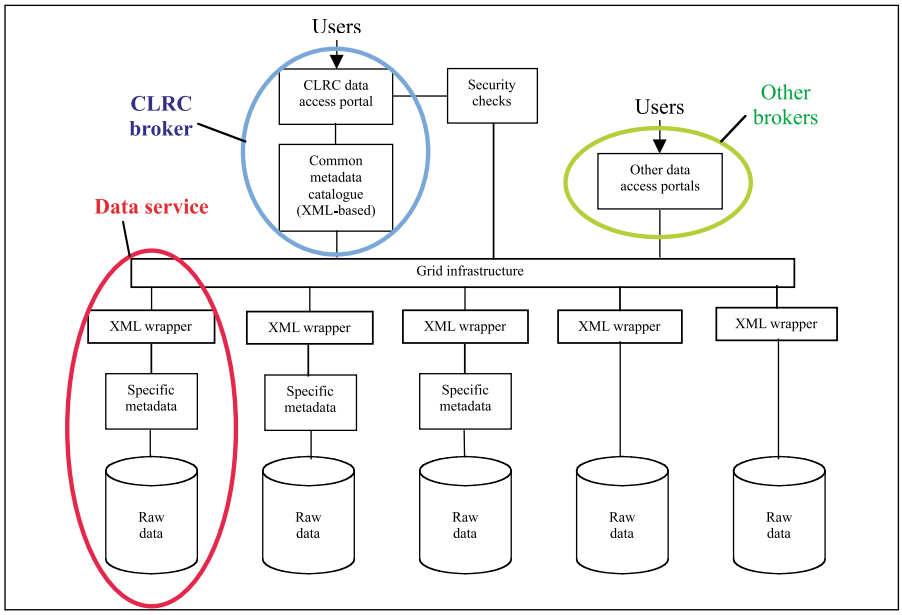

Architecture of the Data Portal system.

User Requirement Capture

To help us understand and capture the requirements of potential users of the data portal, we interviewed several experimental scientists, instrument scientists and data resource providers at each facility and also some external users. Based on their input, we constructed a series of user scenarios. These scenarios have been validated with those we interviewed and, based on them, we are developing use cases which are being used to design the user interface dialogue and to determine the required functionality of the database system underlying the metadata catalogue.

User Interface

The user interface client for the data portal is being implemented using a standard web browser which will accept XML. At each step in the user dialogue, users are presented with a page generated from the metadata available at that stage of their search. They are prompted to refine or broaden their enquiry. User responses are interpreted by a user interface server process and converted into queries which are sent to a database holding the metadata. The results of each query are generated by the database as XML which is sent to the server process, parsed, converted into a displayable page using an XSL script and presented to the user.

Metadata Catalogue

The logical structure of the metadata in the catalogue is based on the scientific metadata model being developed in the project. This model exploits our experience gained in developing metadata models for other domains. It is defined in XML using a DTD schema. The metadata catalogue is implemented using a relational database with a schema derived from the metadata model schema. This offers views of the data supporting the expected user enquiry patterns. Once the specific datasets required by the user have been identified using the available metadata, the catalogue will provide links to the files holding the actual data. Users can then use these links to access the data with their own applications for analysis as required.

Data Resource Interfaces

The data resources accessible through the data portal system may be located on any one of a number of data servers throughout the organisation. Interfaces between these existing data resources and the metadata catalogue are being implemented as wrappers which will present the relevant metadata about each resource to the catalogue so it appears to the user to be part of the central catalogue. These wrappers will be implemented as XML encodings of the specific metadata relating to that resource using the metadata model schema.

Project Status and Future Plans

This pilot project is planned to complete at the end of March 2001 with the operation of a working prototype data portal system. At the time of writing, the metadata model has been developed and validated, the user scenarios have been defined and the user interface, metadata catalogue and data resource wrappers are under development. The longer term goal of this work is to extend the data portal system to provide a common user interface to metadata for all the scientific data resources held in CLRC. We also envisage it being useful for locating and accessing data held in other laboratories. Where catalogues already exist in specific areas, we will not attempt to duplicate these but instead we will provide the user with a smooth connection into these domain-specific systems. It is anticipated that the resulting system will have wide applicability across many scientific disciplines. We are keen to develop its potential in partnership with others, particularly within a European context.

Link:
CLRC e-Science web site:
http://www.e-science.clrc.ac.uk/

Please contact:
David Boyd – CLRC
Tel: +44 1235 446167
E-mail: d.r.s.boyd@rl.ac.uk

BioOpera: Grid Computing in Virtual Laboratories

by Gustavo Alonso

Many scientific disciplines are shifting from in vitro to in silico research as more physical processes and natural phenomena are simulated in a computer (in silico) instead of being directly observed (in vitro). In these virtual laboratories, the computations involved are often a major bottleneck and a significant source of inefficiencies. To spare scientists such limitations, at ETH Zürich we have developed BioOpera, an extensible process support management system for virtual laboratories.

In existing virtual laboratories, storing, manipulating, and keeping track of complex computations is done manually through ad-hoc pieces of code. The data processing logic is typically written using conventional programming languages (eg, C or Fortran) for the basic algorithms and collections of operating system scripts (mainly Perl scripts) as the glue between the different components. Such an approach leads to logic that is extremely difficult to modify and rather primitive, unsystematic methods for driving and monitoring computations. Given the increasing relevance of the work done in virtual laboratories, better software tools are becoming critical to the success of any virtual experiment. As an example of the functionality such tools should provide, consider dependability. An adequate software environment for a virtual laboratory should dependably run computations for months at a time with minimal user intervention. This requires to automatically and transparently handle issues such as efficient scheduling of jobs, load balancing, tracking of progress and results of the computation, recovery from system errors and machine crashes, access to intermediate results as they are computed, automatic accounting of statistics concerning computing time, and a systematic method for storing all necessary meta-data.

The first step towards providing this functionality involves finding an appropriate representation for the computation. We have chosen the notion of process, similar to that used in workflow management systems (although the final implementation is rather different since workflow tools are not entirely adequate for virtual laboratories). A process is an annotated directed graph where the nodes represent tasks and the arcs represent the control/data flow between these tasks. The notion of process allows one to capture sequences of invocations of computer programs in a distributed and heterogeneous environment and the corresponding data exchanges between these programs. From here, the process can be encoded in such a way so as to allow its efficient storage in a database. Once in a database, this information is persistent, allowing us to both automatically manage the computation and increase its dependability.

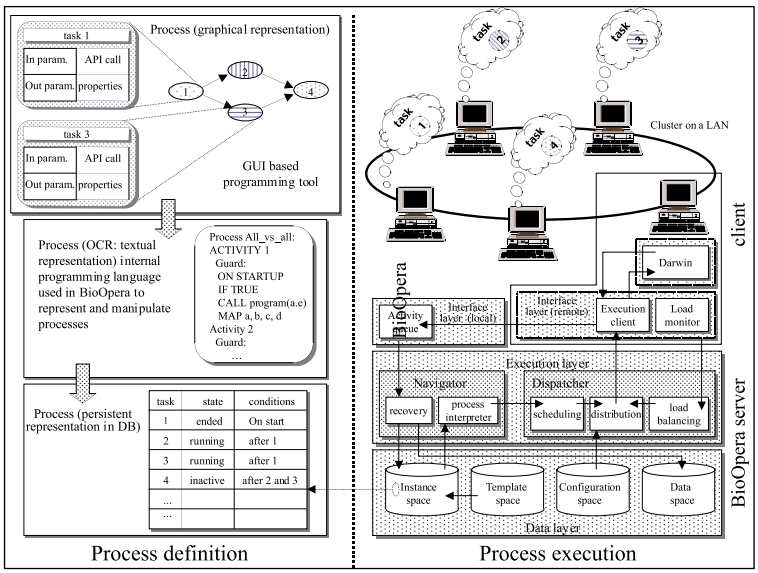

Architecture of BioOpera.

We have implemented these and many other novel ideas in BioOpera, a process support system for virtual laboratories in bioinformatics. BioOpera is based on Opera, a workflow-like middleware tool that has evolved into a programming and runtime environment for cluster computing with the capability to define, execute, monitor and manage a broad range of large-scale, complex scientific computations. BioOpera is being developed at the Information and Communication Systems Research Group of ETH Zürich in collaboration with groups in bioinformatics at ETH Zürich and McGill University in Canada.

The best way to describe BioOpera is as a high-level distributed operating system managing the resources of a computer cluster. Using BioOpera, we have already achieved interesting results both in computer science and in bioinformatics. From the point of view of research in grid computing, we have shown how BioOpera can effectively sustain complex distributed computations during long periods of time. The system is capable of surviving total failures, software and hardware upgrades, and node crashes without manual intervention being required for the computation to resume after the failures have been repaired. We have also proven how BioOpera can duplicate the processing capacity of a grid cluster by making more efficient use of the available resources. In bioinformatics, several search algorithms have been considerably optimized using BioOpera, in some cases reducing the cost to less than 30% of the original cost.

Link:
http://www.inf.ethz.ch/personal/bausch/bioopera/main.html

Please contact:
Gustavo Alonso – SARIT/ETH Zurich
Tel: +41 1 632 7306
E-mail: alonso@inf.ethz.ch

A University Ph.D. Course on Cluster and Grid Computing

by Ferenc Vajda

Since the fall semester of 1999, a postgraduate course has been offered in the frame of the doctorate (Ph.D.) program of the Faculty of Electrical Engineering and Informatics of the Budapest University of Technology and Economics (BME). This is a unique course of this new and fast developing topic and as far as we know there is no similar course available in Europe and only a very few in the US.

The initial idea of computer clusters (a collection of interconnected computers working together as a single system) was developed in the 1960s. However, cluster computing did not gain momentum until three trends came around in the 1980s: high performance microprocessors, high-speed networks and standard tools for high performance distributed computing. A possible fourth trend is the increased need of computing power for scientific and commercial applications coupled with the high cost and low accessibility of traditional supercomputers. This process has led to grid computing in the mid 1990s to provide geographically distributed, heterogeneous computing infrastructure for advanced science and engineering. The course is motivated by the fact that these two computing technologies of networked computers have increased attention recently and a noticeable progress toward usable systems has been made.

The purpose of the course is to provide basic knowledge on the most important principles, methods, tools, systems, standards, etc. behind these two evolving basic technologies. Detailed description of the topic of the course is beyond the scope of this article.

Course Outline:

- Introduction to distributed and high-performance computing. Basic terms: distributed computing, HPC, HPCC, network computing, Internet computing, cluster, grid, meta-computing, middleware, etc; milestones of the history, some representative applications
- Classification: Taxonomies, MPP, SMP, CC-NUMA, cluster: dedicated high performance (HP), high availability (HA), CoPs, PoPs, CoWs; distributed, on-demand, high-throughput, collaborative, data-intensive computing
- Basics of communication media and protocols: TCP/IP, Internet2, QoS, ATM, Fast Ethernet, etc.
- Programming models: Message passing, client-server, peer-to-peer, broker computing, code shipping, proxy computing, mobile agents. Toolkit and OO systems
- Higher level communication: Light-weight communication, sockets, standard APIs, active messages
- Storage and file problems: Network RAM, RAID and software RAID. Distributed File systems: NFS, AFS, OSF-DSF, RSF
- Message passing standards: PVM (Parallel Virtual Machine), MPI (Message Passing Interface)
- Object-oriented de facto standards CORBA and DCOM
- Java-based methods: JVM, RMI, Bytecode, Applet and Servlet, JavaBean and JavaSpaces, Jini
- Grid toolkit approach: Globus Hourglass concept, communication, resource and process management, data access, security
- Object-oriented approach: Legion Language support, component wrapping, program support, resource management
- Security: Confidentiality, integrity and availability. Authentication, authorization, assurance, auditing, accounting
- Scheduling: Algorithms, policies and techniques, high performance and high throughput schedulers, resource scheduling
- Grid monitoring: Tasks, types, architecture, components.

The course requires a student presentation on a selected topic by each participant. Students are provided by a reading list and copies of slides presented at the lectures in PDF or PS format. The course is aimed at bridging the gap between the distributed and high performance topics in computer science and engineering university education and the current hot research topics and activities of the field.

The course has a strongly applied outlook with numerous case studies and examples and therefore provides a solid basis for Ph.D. students to take part in European and international research projects and programmes.

Please contact:
Ferenc Vajda – SZTAKI
Tel: + 36 1 329-7864
E-mail: vajda@sztaki.hu or vajda@mit.bme.hu

Collection and Management of EMG Examination Cases

by Mária Bieliková, Pavol Návrat and Mária Smolárová

As a part of EMG-net project, researchers at the Slovak University of Technology developed a software tool for collection and management of electromyography (EMG) examination cases called EMMA. They concentrated on data representation, which is a crucial task in such data intensive area as the EMG examination. The proposed data model serves as a basis for a multifunctional platform for EMG studies aiming at assisting EMG practitioners in developing standard examination procedures as well as in analysing and evaluating of existing EMG procedures.

The European Research Network for Intelligent Support of Electromyographic Studies (EMG-net) is an INCO Copernicus research project which aims at establishing a research network combining electromyography and information technology (IT). EMG-net comprises teams of physicians and IT experts from nine Western and Central/ Eastern European countries. EMG-net is a successor of ESTEEM (European Standardised Telemetric tool to Evaluate knowledge-based EMG systems and Methods) project that introduced a set of standards for EMG studies.

EMG is a standard method for monitoring neuromuscular activity at the level of bioelectrical signals. A physician has to deal with large amount of data. Computer based information processing supporting diagnostic process in neurophysiology may enhance the physician's ability to make appropriate decisions and find the right diagnosis. A software tool for collecting and managing EMG examination cases is useful because of its capability of:

- storing EMG examinations, which enables retrieving previous examinations and comparing or evaluating them
- exchanging and distributing EMG examinations, which serves as a communication between physicians
- tracking condition of a patient by effective retrieving of EMG examinations for a particular patient at different times.

EMG data can be collected at several levels: the local level, the national level, and the European level.

The local level is represented by a particular EMG workplace. The national and the European level are intended mainly for exchange of the EMG expertise accumulated during the years of practice. This would improve the quality of an early diagnosis and prevention of neuromuscular diseases in each particular country, or a whole region.

The national level of EMG data collection becomes also important when the aim of data collection is development of normative EMG data.

Data Representation

During the ESTEEM project, the software tool for EMG data collection (called CASETOOL) was developed together with Communication Protocol EPC/ECCO 3.2 for exchange of EMG data. Binary ECCO format has been designed as an efficient format from the point of view of storage. Several existing tools (including the CASETOOL) use it. Moreover, EMG laboratories, which form EMG-net consortium, have collected in the last years together more than 1000 EMG cases stored in this format. There exists so called golden EMG data collection, which stores more than 200 cases resulted from the consensus of partners. The problem with ECCO format lies mainly in its inflexibility with respect to extensions. The problem arose during the EMG data collection and evaluation.

Binary format is also not very suitable for exchange using the Internet for obvious reasons. In order to preserve continuity in the EMG-net project as well as to support future developments, we use several data representation formats:

- binary format (Communication Protocol EPC/ECCO 3.2)
- relational format (represented by a relational database)
- text format (represented in the eXtensible Markup Language).

New software tool supports all the three formats. The ECCO binary format is retained in the form of export/ import capabilities.

Relational format is used for storing EMG examination cases. Design of the relational format is based on the ECCO format, which defines all the necessary data and data types stored in a database. Moreover, we proposed additions, which are inevitable for efficient data management and extensions proposed by EMG experts – users of the CASETOOL. Proposed relational data model is further capable to store normative EMG data.

The relational format is advantageous because of its flexibility and extensibility. Created database can be easily used by different tools through standard interfaces to the relational database and can be shared between different applications. This increases the portability of the final solution. Relational format presents also a base for data mining. What we should pay for this flexibility is less effective storing than the one that is achieved by binary format (such as ECCO).

We identified some parts of the data model, which were designed with the likely future changes of requirements in mind. The identified entities contain their symbolic values in separate tables. Each of such tables contains attribute, which determines version of the given code table (at the present the values correspond to ECCO, v. 3.2 values). This allows changes and support of different symbolic value assignment versions without a change of the software tool. As an example we can mention flexible representation of clinical diagnosis names and their numeric representations, or EMG conclusion names and their numeric representations.

The described approach complicates the data model, so we had to carefully decide between simple data model and its flexibility. A complete report of the proposed relational model is available at the Slovak EMG-net local webpage. Finally, the XML format is suitable for distribution of EMG expertise between the users of the system. Information exchange is important in the process of consensus exercises, where physicians discuss particular EMG cases in order to reach consensus regarding an EMG diagnose as well as a clinical diagnose.

The proposed markup language for EMG data exchange (EMG-ML) is based on the proposed relational model. EMG-ML document contains information about all the patient's examinations. Using XML files as the exchange format has several advantages:

- the content of the document is dynamic, depending on the actual anatomical structures that are examined and used examination techniques
- the content of the document is self-describing because of the text-based nature of XML files
- the EMG data transmission is highly error-free because its non-binary character
- the transmitted data are 'open' for third party software processing.

The main disadvantage of such a format is the length of the document and inefficient manipulation of the data. However, its combination with a relational database solves the mentioned problems.

Design and Implementation of EMMA

A software tool for collection and management of EMG cases, EMMA, is developed in Visual C language in the Windows operating system. The user interface comprises several windows and a set of forms for filling and maintenance of EMG data. Relational database is built in Microsoft Access. The design of EMMA follows the process of EMG examination, in which a physician works with three groups of data:

1. General data: patient information, examination information, clinical information
2. Examination data: test conditions, parameters
3. Inferred data: symbolic parameter values, pathophysiological test conclusions, pathophysiological structure conclusions, EMG diagnoses, clinical diagnoses.

We also incorporated into the design EMG normative data, which serve for accurate interpretation of an examination. EMMA supports inputting and modification of normative data in several formats (eg, in table format, functional format) and calculating symbolic parameter values according these data in an actual test.

EMMA is designed in such a way that a patient is considered a central entity in the system. A physician can create several EMG examinations for the patient, which enables observing changes in subsequent examination of a particular patient.

The designed data representation allows future extensions of EMG data management features: support of tests and techniques selection during EMG examination; support of conclusion statements determination (on all levels, ie structure conclusion, EMG diagnose conclusion and clinical diagnose conclusion); support of retrieving cases based on their similarity; support of monitoring patient's course of illness.

Acknowledgement

We would like to thank Birger Johnsen and Anders Fuglsang-Frederiksen from Gentofte Hospital, Copenhagen, who provided expert consultations during design of EMMA and to students from Slovak University of Technology, Bratislava Matej Makula, Tomás Milicka, Ivan Noris, Boris Vasilovcík and Karol Vlasko, who contributed to implementation of EMMA software. The work reported here was partially supported by project INCO Copernicus, No. 977069 (EMG-net).

Links:
EMG-net project web page:
http://www.inriaalpes.fr/sherpa/emgnet/emg_index.html
Slovak node page:
http://www.dcs.elf.stuba.sk/emg/

Please contact:
Pavol Návrat – Slovak University of Technology
Tel: +421 7 602 91 548
E-mail: navrat@elf.stuba.sk

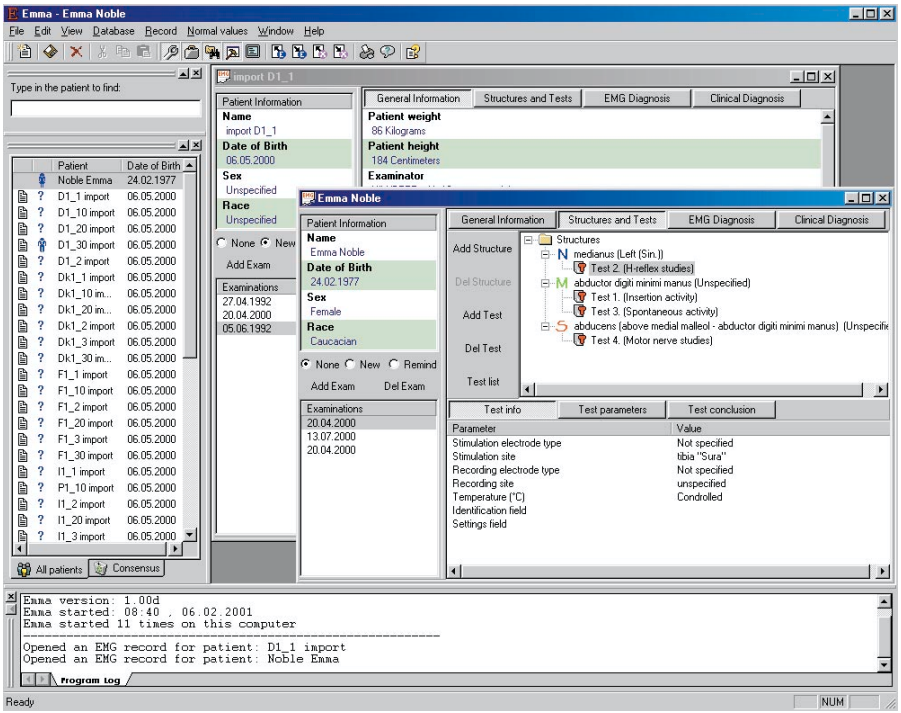

Example of the EMMA user interface.

Similarity Hashing for Metric Data

by Claudio Gennaro, Pasquale Savino and Pavel Zezula

Similarity searching has become a fundamental computational task in a variety of application areas, including multimedia information retrieval, data mining, pattern recognition, machine learning, computer vision, genome databases, data compression, and statistical data analysis. This problem, which originally was mostly studied within the theoretical area of computational geometry, has recently been attracting more attention in the database community, because of the increasing need to deal with large, often distributed, volumes of data. Consequently, high performance has become an important feature of cutting edge designs. Scientists at IEI-CNR, Pisa, and Masaryk University, Brno, propose a novel access structure for similarity search in metric data, called Similarity Hashing (SH).

Similarity Hashing has a multi-level hash structure which consists of separable buckets on each level, supporting easy insertion and bounded search costs, because at most one bucket needs to be accessed at each level for range queries up to a pre-defined value of search radius. The number of distance computations is significantly reduced by use of pre-computed distances obtained at insertion time. Buckets of static files can be arranged in such a way that the I/O costs never exceed the costs of scanning a compressed sequential file. Experimental results demonstrate that the performance of SH is superior to available tree-based structures. Unlike tree organizations, the SH structure is suitable for distributed and parallel implementations.

We consider the problem from a broad perspective and assume the data to be objects from a metric space where only pair-wise distances between objects are known. Contrary to traditional storage structure designs where the I/O costs form the dominant component of the insert/delete and retrieval functions, the CPU costs in generic metric structures can also be quite high, as in some applications the computation of distances is time consuming. Our major objective is thus to develop a similarity search structure that would minimize both the I/O and the CPU costs. This has not been a major goal in previous studies which have tended to concentrate on main memory structures, not considering paged disk environments. The idea of reusing pre-computed distances at later stages, ie in the retrieval phase, seems to be promising. In this way, the number of necessary distance computations to evaluate a query can significantly be reduced and the search time decreased.

To the best of our knowledge, all metric data designs are trees, and the reported node utilization is typically poor (often much less than 50%). This implies high space occupancy and random access to read nodes. Note that sequential files can be allocated with the minimum of necessary disk memory, and the sequential scan of continuous disk areas is very fast. There are two additional reasons for arguing against developing metric structures as trees. First, insertion costs are high and node splitting strategies, both top-down and bottom-up, require a lot of distance computations - published articles do not usually report on this issue. Second, trees are not convenient for parallel and/or distributed implementations. On the other hand, parallel (distributed) memories are available and proper exploitation of their potentiality can significantly reduce the search costs.

The main challenge of our work is to build a similarity search organization based on

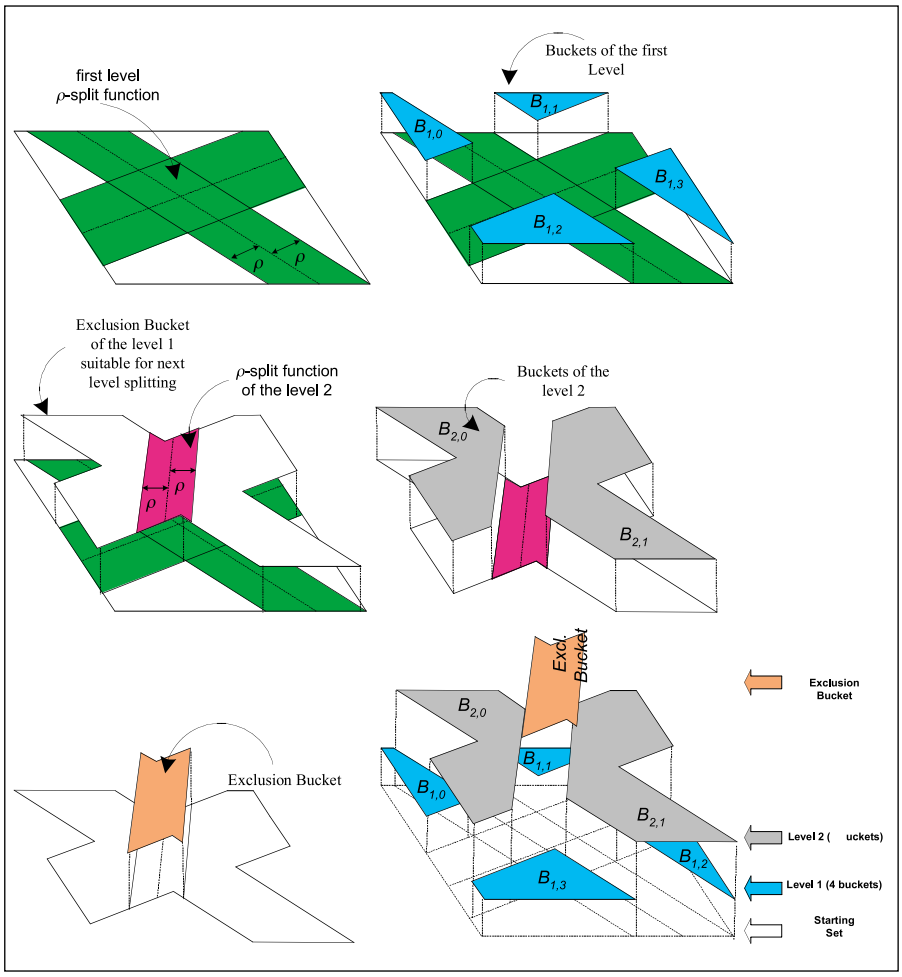

An example of creation of a two level SH structure.

hashed partitioning, that is the Similarity Hashing (SH) technique. This is a multi-level hash structure that takes advantage of the excluded middle partitioning. In fact, all buckets on a level are separable so that maximally one bucket must be accessed for any query up to a specific value of a query radius. Objects that at one level do not conform to such arrangements are excluded from storage on this level and become candidates for storing on the next level. Depending on the number of levels and the data file, some objects can be excluded, and are thus stored in a separate exclusion partition that must always be accessed. Once computed distances are memorised at practically zero storage costs - a distance between objects is a number. At query time, a computation of simple functions determines for each level maximally one partition to access. However, distance computation between the query and accessed objects is not always required. In fact, the knowledge about pre-computed distances can infer that the object does not belong to the query response set. To summarize, the main features of our approach can be characterized as follows:

- each object is stored through hashing in one bucket from the multilevel structure of separable buckets
- queries need to access maximally one bucket per level, plus the exclusion partition, and the number of distance computations is significantly reduced through pre-computed distances
- an upper bound of the number of accessed partitions is the number of hash levels plus the exclusion partition
- by storing the partitions in a sequential file and using the hashed structure as the main memory directory of its parts, the I/O costs are upper-bound by the costs of optimized sequential scan, but are typically much lower
- contrary to tree organizations, the SH structure is suitable for parallel and distributed implementations.

As can be seen on the figure, the SH structure is defined recursively. First, we define the ρ-split function of the first level. This function allows the objects of the data set to be spread among the buckets. The objects that are not accommodated in the buckets of the first level are reorganized in the second level, by means of second level ρ-split functions, and so on. Eventually, the remaining objects of the last level form the exclusion bucket.

We evaluated a prototype SH system and measured its performance in terms of the bucket reads and distance computations to solve range queries. We conducted our experiments on two synthetic sets of Euclidean vectors. The first set of experiments concerned 50,000 uniformly distributed vectors in 10 dimensional Euclidean space. The second set of experiments was conducted on 20 dimensional Euclidean vectors generated in clusters of equal sizes. We compared the search performance of the mvp-trees with the SH structure for the two different vector sets. The performance of the SH technique results from 5 to 25 times higher than that obtained when using the standard tree-based search structures.

Please Contact:
Claudio Gennaro – IEI-CNR
Tel: +39 050 315 2888
E-mail: gennaro@iei.pi.cnr.it

Quantum Computing Pioneers in Amsterdam

by Henk Nieland

Quantum computing has become a hot topic in the past five years. Worldwide interest rapidly increased, for example from IBM, Caltech, Lucent, AT&T, and NSA, since existing techniques will soon reach their limits, and several computations could be done much faster on a quantum computer. The Fourth Workshop on Quantum Information Processing was organized by CWI 9-12 January 2001 in Amsterdam.

Remember that a cat is said to have nine lives? A Quantum Cat does even better: it can be dead and alive at the same time. This wizardry is at the heart of quantum computing, a novel way of computing based on certain characteristics of quantum mechanics. It emerged in the 1980s as a theoretical alternative to traditional computing, which will be faced with its physical limitations soon. The possibilities are very promising, but the field is still in its infancy, and realization of a working quantum computer is still an enormous challenge.

Quantum mechanics offers a novel way of computing, enabling computations out of reach for traditional computers, even if these could be miniaturized to the same level. A crucial notion is 'superposition'. Quantum mechanics describes matter as a superposition of all of its possible states, each with a certain amplitude – a complex number whose modulus squared is interpreted as a probability. These amplitudes thus have addition properties different from ordinary probabilities – a feature that becomes apparent in 'interference'. Precisely this feature together with the superposition principle gives quantum computing its power. The evolution of a system is described by a unitary operation on the superposition which preserves the probability interpretation of the amplitudes. In traditional computing the smallest unit is a 'bit' which can only take the values 0 or 1. Quantum computing is based on 'qubits' which consist of a superposition of the two classical states 0 and 1 each with its own amplitude. Several physical realizations of qubits have been proposed, for example an atom in the ground state (0) or excited state (1). A computation starts with a number of qubits in a well-determined state, on which a series of unitary operations is performed (the

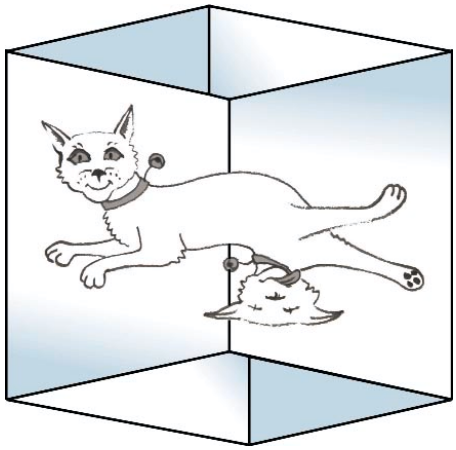

Schrödinger's Paradox

Nobel Prize winner Erwin Schrödinger invented in the 1930s a thought experiment to elucidate a seemingly absurd consequence of quantum mechanics. A cat sits in a closed box together with one radio-active atom. This atom will decay with a certain probability, according to the laws of quantum mechanics. Upon decaying, the emitted nuclear particle crushes a thin glass tube filled with cyanide and the cat dies instantly. Being outside the box, we don't know whether the radio-active particle has decayed and thus whether the cat is dead or alive. Quantum mechanically the animal is in a superposition of both states with a certain probability. If we look inside the box, however, we see only one of the two states: the cat is either dead or alive. An observation of the 'superposition' of states makes it collapse to one of the states with a certain probability. (Drawing by Tobias Baanders, CWI.)

algorithm). Because of interference certain superpositions are intensified, whereas others cancel each other out. After a number of steps the final state (the result) is observed. During the intermediate evolution all possible computational paths are followed simultaneously (quantum parallelism), but they remain hidden in a box. In certain cases this form of computation may lead to a tremendous speed-up compared to traditional methods, but at the same time it poses equally tremendous problems: the smallest disturbance from the environment may ruin the delicate superposition and may render the computation meaningless.

An important notion in quantum mechanics and quantum computing is 'entanglement': two (or more) qubits can be prepared in such a way that, although they are separated in space – one could be on Mars and the other here on earth – they have correlations that can not be explained by classical probability theory, for example two atomic nuclei having unknown but opposite spins. As soon as one qubit is measured, the content of the other is also known, no matter how far they are apart. This property can be used for error correction during the computation, as well as for more efficient transmission of information and for certain forms of distributed computations.

Quantum computing gained momentum after P.W. Shor showed in 1994 how to construct an efficient algorithm for factoring large numbers, which is of crucial importance for, eg, internet security, followed by an algorithm by L.K. Grover (1996) to search a database quadratically faster than any classical algorithm.

The Workshop in Amsterdam (http://www.cwi.nl/~qip) drew 150 participants from 24 countries, including several pioneers of the field, such as Charles Bennett and David DiVincenzo (IBM Yorktown Heights), Richard Jozsa (Bristol), Gilles Brassard (Montreal), Umesh Vazirani (Berkeley), as well as Nobel laureate Gerard 't Hooft (Utrecht).

CWI started the first quantum computing research group in The Netherlands (and one of the first in Europe) in 1995, and has contributed significant discoveries to the field. CWI has applied quantum computing notions to communication complexity, and found general limitations of quantum computers, as well as some new speed-ups. CWI also studies quantum information theory. The European Union has recognized the importance of this research and has given the group substantial support.

Link:
http://www.cwi.nl/ins4/

Please contact:
Harry Buhrman – CWI
Tel: +31 20 592 4076 / 4078
E-mail: Harry.Buhrman@cwi.nl

Ecological Modelling meets Interactive Geovisualization at GMD

by Oleg Chertov, Alexander Komarov, Natalia Andrienko, Gennady Andrienko and Peter Gatalsky

In December 2000, two forestry experts and specialists in ecological modeling, Prof. O. Chertov and Dr. A. Komarov from the European Forest Institute, Finland, visited the Knowledge Discovery Team of the Autonomous Intelligent Systems Institute of the GMD. They discussed the prospects and possibilities for future scientific co-operation and for the development of an effective tool for decision-making in modern forestry based on new computer technology. Two weeks of intensive work lead to impressive results.

The concept of Sustainable Forest Management (SFM) brings together different approaches and different branches of science. The need to integrate our knowledge of forest development with broader ecological ideas is evident in Europe, where forests are used for the maintenance of biodiversity, aesthetic values, forest products and various other things. There are many different tasks related to forest ecology, economics, politics etc. but uniting these should be done with the help of mathematical modeling and information technology. Many tasks are linked to the spatial

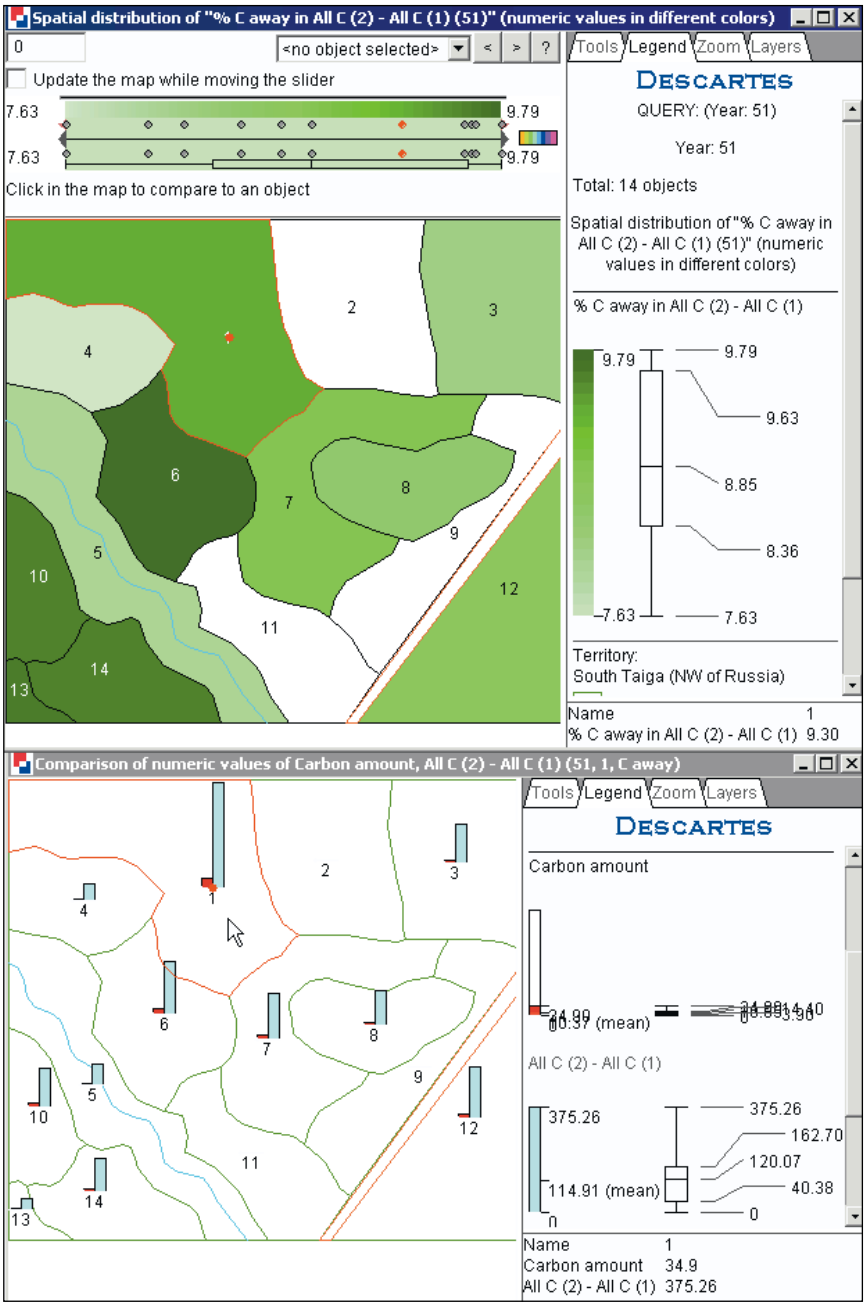

A comparison of total ecosystem carbon loss (as the difference between total carbon in the first and second scenarios) and carbon removed from the ecosystems as cut wood (lumber) at the end of the simulation. The top figure shows the percentage of the total carbon loss due to removed wood. The lower figure shows removed carbon and total ecosystem loss (tonnes per stand) after cuttings in the first scenario. The figure shows that ecosystem loss of carbon due to clear cutting is about ten times the amount of used wood carbon. The reaction is altering in various forest stands.

distribution of forest territories among surrounding settlements, roads, and agricultural land with different land-use. If we thoroughly analyze the Criteria & Indicators of SFM then we find that all tasks concerned with forest productivity can be analyzed by simulation dynamics models. Tasks concerned with spatial distribution of forest areas such as ecological corridors and the establishment of protected or recreation zones are currently analyzed only by experts. The main task is therefore how to unite the simulation models with spatial visualization of results to create a Decision Support System (DSS) for SFM.

We attempted to integrate long-term forest ecosystem modelling with modern techniques of exploratory data analysis. The purpose of the integration was to build a prototype system for 'decision making' in sustainable forestry at landscape level. The system should support the spatially oriented tasks arising from the Criteria & Indicators (C&I) of Sustainable Forest Management. We suggest that analysis of the landscape mosaic be done on the basis of a combination of single stand simulation models and software for visual exploration of spatial-temporal data.

We performed a model test on a small forest plot consisting of various stands with two silvicultural regime scenarios; the first is a clear cut system with some ecological restrictions dampening the effect of wood harvesting; the second scenario is a full protection of the forest for fifty years. A combined, spatially explicit forest simulation model EFIMOD II, and the DESCARTES software system designed to support visual exploration of spatially referenced data, were used in the experiment. The results of the experiment were impressive. The test shows the high potential of the integration of forest ecosystem modelling with the visual exploration of data on maps. The visualization allows direct representation of time series and spatial patterns of forest dynamics in graphical form, facilitating analysis of the dynamic trends and of the effectiveness of various silvicultural regimes.

We note the usefulness of representation of spatially distributed time series on maps and the informative value of the spatial combination of different attributes. The prototype system enables us to explore the diversity of ecosystem reactions in various forest compartments and to test various spatial combinations of strategies and zoning of the forest area.

There is a strong demand for a new user-friendly modeling system integrating various types of forest models with exploratory data visualization for methodologically easier and more expressive decision-making based on a long-term simulation at forest enterprise or landscape level. Currently we are preparing a project proposal incorporating these topics to be submitted for EU funding.

Please contact:
Gennady Andrienko – GMD
Tel: +49 2241 142486
E-mail: gennady.andrienko@gmd.de

CIPRESS – Cryptographic Intellectual Property Rights Enforcement SyStem

by Christoph Busch

The ReEncryption system resulting from the CIPRESS project was developed by the Fraunhofer Institute for Computer Graphics based in Darmstadt, Germany on behalf of Mitsubishi Corporation, Tokyo, Japan. It offers a high security environment for mission-critical applications based on commercial operating systems and standard applications. Data is protected by the patented ReEncryption™ providing mechanisms for access and use control including detailed auditing of usage trails. Even illegal analog audio and video recordings as well as eg printouts of images can be traced back due to the use of patented digital watermarking technologies.

The system consists of independent components, namely Client software extending the operating system with ReEncryption™ and digital watermarking for both file system and network communication channels to be installed on all protected systems; a Key Center maintaining a database of access permissions and keys and their associations with protected data objects along with audit trails for all relevant operations, and a number of data archives (Content Servers). All communication is performed over an integrated VPN layer. The system is integrated with directory services and standard public key infrastructure services.

These servers maintain storage of arbitrary digital data. They supply the content either directly and in encrypted form over a CIPRESS protocol to the client systems or use a standard WWW server for distribution of the materials and act as archives for protected data. All data registered with a content server will be marked with invisible (or inaudible) digital watermarks providing for both the identification of the copyright holder or owner of the data, permitting auditing and recovery of original digital data even based on analog copies of data.

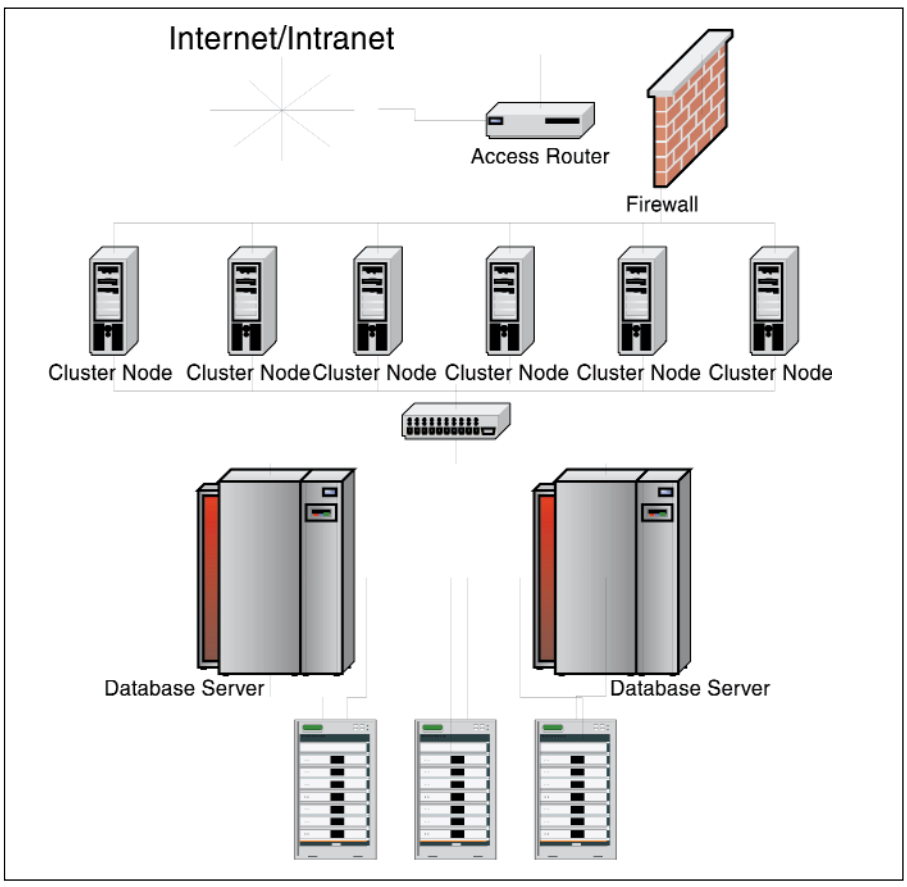

Cipress architecture.

The client system consists of a layer of operating system extensions providing security enhancement through the addition of mandatory encryption and access control using ReEncryption™ technology. Both network and all file system accesses are protected transparently without applications or users becoming aware of the process as long as only legitimate operations are performed. In addition, all multimedia data are marked with a digital watermark containing the identity of the user accessing the data object. Access to the archives contained on Content Servers is possible using CIPRESS applications, a WWW front-end, or a published application programming interface for use in value-added applications.

The Key Center is a trusted system providing access control and auditing services for all objects under its control. For this purpose, databases containing users, groups, access permissions, and Content Servers are maintained. The Key Center is also the repository for all ReEncryption™ keys associated with data objects. Access to controlled objects is verified by the client system every time a data object is referenced, so any change in access rights or classification is effective immediate regardless of where a copy of the data object resides. The Key Center also can audit the user and document identity along with the identity of the client computer used to access the object and the precise timestamp. This permits precise auditing of who used which data object when and where and even where a data object was obtained from.

CIPRESS provides a VPN mechanism transparently embedded into the operating system. This mechanism is invisible to application programs and, since it is implemented at the transport layer, does not interfere with network management and is suitable for communicating over wide area networks employing network address translation.

On the basis of the CIPRESS system attractive applications can be realized. For instance Data Linkage providing effective mechanism for protection of secondary copyright and integrity of interlinked documents. CIPRESS answers questions such as: "who accessed which document when and where through compound documents". With the Trusted Mail users have not only a reliable email service but also convenient new features as proof of receipt and 'unsend' functionality of misdirected messages.

First field trials were conducted in 1999, the ReEncryption system is available commercially at the beginning of 2QCY01. Based on the core CIPRESS system, additional applications such as a secure and convenient messaging/email system and document management systems will be available. CIPRESS provides a open framework for third party application integration and can be extended efficiently to meet elaborate customization requirements.

Please contact:
Christoph Busch – Fraunhofer-Gesellschaft
Tel: +49 61 51/1 55-1 47
E-mail: busch@igd.fhg.de

Efficient Computation of Ship Waves

by Piet Hemker

A new solution method for the efficient computation of water waves around a ship was developed at CWI, in close cooperation with the Dutch maritime research institute, MARIN.

The ability to predict the behaviour of flows is essential in many engineering applications, such as ship hydrodynamics. Although the physical laws describing fluid flow and their mathematical formulation (the full Navier-Stokes equations) are well understood for more than a century, due to its tremendous complexity the most accurate mathematical model is still unmanageable for engineering purposes. Therefore, those concerned with practical fluid-mechanics problems have always considered more tractable models. For practical computations by means of a computer, the unknown functions are approximated by tabular values, and the fluid-flow problem then translates into the problem of solving a very large system of equations in which the tabular values are the unknowns. This technique is called 'discrete approximation'. The approximation becomes increasingly accurate as the tabulation becomes finer.

One challenge in the hydrodynamic design of ships is to find the optimal shape of the ship hull satisfying prior requirements of use, economy, safety, etc. CFD computations have led to improvements of the hull shape, so that the high amplitude of the long waves made by the ship could be strongly reduced. (Photo courtesy MARIN.)

On the other hand, however, the system of equations grows correspondingly and becomes more difficult to handle.The large system of equations is solved to arbitrary accuracy by a computer in an iterative (ie, cyclic) process. Although the rapid development of computers has allowed the treatment of increasingly larger systems, it still remains a challenge (and a practical necessity) to obtain, within a given computational effort, better accuracy for more complex flow problems.

One of the flow problems that have recently become tractable is: free-surface viscous flow in three dimensions, eg, waving water flow around a ship hull. Here the problem is to determine both the (viscous) flow and the shape of the wavy water surface. Previously, only simpler models could be applied, and the surface and the flow problem were treated separately: either the wavy surface of an inviscid fluid-flow, or a viscous fluid-flow below a flat water surface was determined. However, simultaneous treatment is required to obtain more reliable results. The essential problem was the inefficiency of the available iterative methods for the viscous free-surface flow problem. As a consequence, in practice the accuracy of the approximation was limited, and the more sophisticated model gave less accurate results. Hence, an essentially more efficient iterative method had to be developed. The new method was shown to be efficient for a test problem. Since October, computations for actual ship flows are in progress.

Please contact:
Piet Hemker – CWI
Tel: +31 20 592 4108
E-mail: pieth@cwi.nl

Collaboration between IEI-CNR and Ericsson Lab Italy

by Antonia Bertolino

A long term agreement for collaboration between IEI-CNR and Ericsson Lab Italy (ERI) was signed in December 2000. The collaboration also involves research groups from the University of Pisa and from Scuola S. Anna.

Ericsson Lab Italy is Ericsson's Italian product company; it was created in July 2000 from the R&D division of Ericsson Telecomunicazioni S.P.A. It is a strategic supplier for the Ericsson world-wide market with the mandate to identify, develop and manage globally successful products using the best Italian talents. Ericsson Lab Italy (ERI) develops products for mobile and fixed networks using "state of art technology" and the most advanced methods and tools. When participating in research activities with Italian universities and research centres, the aim of ERI is both to collaborate on the refinement and further development of the latest technologies and also to promote the creation of innovative solutions that will lead to the design of new products.

Within this framework, the objectives of the IEI-CNR/ERI collaboration agreement are multiple: to promote a tight, hands-on interaction between a research institute and an industrial laboratory, that share common research interests and similar working attitudes; to establish a privileged, institutional communication channel between industry and research; and last, but not least, to train young researchers on strategic themes, identified on an annual basis.

These ambitious goals will be pursued through the joint coordination of a set of research projects. Starting in January 2001, teams from IEI, Pisa University, Scuola S. Anna and ERI will conduct joint applied research on the topics of Software Engineering, and its application to the Telecommunications Sector. The hope is that, on the one hand, these projects will make evident to the research community the complexity of the activities and the real world problems and constraints faced by the industrial partner while, on the other, they will provide an effective benchmark for the validation and refinement of the latest research results from the academic labs.

Within the scope of this agreement, a new Software Laboratory is now being established in IEI, at the new CNR Research Area in Pisa. A number of young researchers will work full time in the laboratory, and will act as the connecting agents between senior researchers from IEI and ERI. For the first year, ERI has funded three doctoral grants, and two annual fellowships. These have already been assigned, after a public selection process. Staff from the ERI headquarters in Rome will spend short periods in the Laboratory in Pisa for face-to-face, effective interaction and exchanges of ideas with the these young researchers and the IEI team.

Four research projects have been launched for the first year of activity:

- Quality and Validation of Software Architectures
- Programming Techniques and Operating Systems for Multi-processor Embedded Real-time Systems
- Protocols and Quality of Service for the connection between Wireless Networks and Internet
- Traffic Control in 3rd Generation Mobile Core Networks.

These projects share a common formulation: they will consider real world case studies and problems provided by ERI, and, although the researchers will be free to experiment innovative ideas and techniques, ie, the projects will be conducted off the production line, they will be monitored and reviewed at fixed, short term intervals.

The initiative is now in the preparatory phase, the organization of both the logistics and the human resources is currently being completed. Great care is being taken to create all the conditions needed to achieve an effective, result-driven research and a valid technology transfer programme. Yearly workshops will be scheduled in which the first results of the projects will be presented, and ideas can be exchanged for the identification of new directions for future projects. We hope to return in these columns within a year's time to report on the highlights of the first workshop and the status of ongoing projects.

Link:
http://www.iei.pi.cnr.it/ERI/

Please contact:
Antonia Bertolino – IEI-CNR
Tel. +39 050 315 2914
E-mail: bertolino@iei.pi.cnr.it

SPONSORED BY ERCIM

SOFSEM 2000 Conference Report

by Václav Hlavác and Gabriela Andrejková

The SOFSEM (SOFtware SEMinar) is an annual international computer science and computer engineering conference of generalistic and multidisciplinary nature. SOFSEM roots are in the former Czechoslovakia. SOFSEM 2000, the 27th in a row, was held in Milovy, Czech Republic from 25 November to 2 December 2000. SOFSEM 2000 received support from ERCIM and several other institutions from the IT industry.

SOFSEM traditionally builds on invited talks and the 2000 conference saw 15 invited speakers. The programme was divided into three main streams: Trends in algorithmics, Information technologies in Practice, and Computational Perception. The programme was, as usual, accompanied by 18 submitted refereed papers (called contributed talks at SOFSEM). Of the 159 attendees 49 were foreign participants, many of these from ERCIM institutes.

The programme committee chair Václav Hlavác, co-chairs Jirí Wiedermann, Keith Jeffery, as well as the organising committee chair Jan Staudek were approved by the SOFSEM Steering Committee for SOFSEM 2000. Serving as a meeting ground for professionals from both the theory and the practice of computing has always been an important ingredient of the SOFSEM mission. The busy time schedule of SOFSEM 2000, in which invited and contributed talks were supplemented by commercial presentations by sponsors and a cultural programme. The concentration of all events in a single hotel complex allowed a longer conference duration and left time for lively discussions, which often focused on sketching research projects with newly discovered partners. There were many students and young researchers among the participants. For the sixth time, the SOFSEM 2001 proceedings were published by the Springer-Verlag in the Lecture Notes in Computer Science (LNCS) Series (Volume 1963) and contain papers from the invited and contributed talks.

SOFSEM moves forward, of course. The new chair of the Steering Committee was elected on SOFSEM 2000, which is Branislav Rovan from the Comenius University in Bratislava, Slovakia. The previous Steering Committee chair Jirí Wiedermann from the Institute of Computer Science, Prague, Academy of Sciences of the Czech Republic, was given a big thanks for his five-years job. SOFSEM 2001 will be held in the Piestany spa town in Slovakia. The new PC chair is Peter Ruzicka from the Comenius University in Bratislava, Slovakia.

Workshop on Softcomputing

This year, a very successful two-days workshop on Softcomputing was joined to SOFSEM. The workshop was held in November 27-28, 2000.

The area of Soft Computing can be characterised using words of Prof. L. Zadeh (see also his homepage at http://www.cs.berkeley.edu/ ~zadeh) who describes the field of Soft Computing as follows: "Soft computing is an association of computing methodologies centering on fuzzy logic (FL), neurocomputing (NC), genetic computing (GC), and probabilistic computing (PC). The methodologies comprising soft computing are for the most part complementary and synergetic rather than competitive.

The guiding principle of soft computing is: exploit the tolerance for imprecision, uncertainty, partial truth, and approximation to achieve tractability, robustness, low solution cost and better rapport with reality. One of the principal aims of soft computing is to provide a foundation for the conception, design and application of intelligent systems employing its member methodologies symbiotically rather than in isolation".

The organizer and chairman of the workshop was Petr Hájek from the Institute of Computer Science of the Academy of Sciences of Czech Republic. The institute is one of the research centres having rather good credentials for research in SC: these follow from the experience given by their long term investigation in neural networks, fuzzy logic as well as in probabilistic computing. The program committee decided to invite Prof. Kruse for a plenary invited lecture of SOFSEM with the title 'Information mining with fuzzy methods'.

The domain of SC was understood rather broadly and openly - from strictly mathematical foundations to practical applications. The following contributions were accepted and published in the Special Issue on SOFSEM 2000 - Neural Network World, Vol. 10, No. 5, 2000 (International Journal on Neural and Mass-Parallel Computing and Information Systems).

Petr Hájek characterised the workshop as follows: "The presented works appears to be a reasonable selection of views, approaches, methods and results which, on the one hand, are interesting contributions to the field of SC, and on the other hand, invite to stronger cooperation and cross-fertilization - a necessary condition for the pursuit of the main idea of SC. Needless to say, many problems are still unsolved and some research tasks are only sketched. The future will show how far reaching and deep the idea of soft computing is. Contributing to this understanding is certainly an exciting programme."

The participants of the workshop proposed the creation of an ERCIM Working Group on Soft Computing.

Links:
SOFSEM 2000:
http://www.sofsem.cz/sofsem00
http://www.sofsem.sk/

Please contact:
Václav Hlavác – CTU
Tel: +420 2 2435 7465
E-mail: hlavac@fel.cvut.cz

SPONSORED BY ERCIM

Two Days in Honor of Alain Bensoussan

by Agnès Sulem

A two days event has been organized in honor of Professor Alain Bensoussan, on the occasion of his 60th anniversary in Paris on 4 and 5 December 2000: a conference on 'Optimal Control and Partial Differential Equations' and panels on 'Innovations and Applications of Space and New Technology'.

The first day of the conference was dedicated to fundamental research: Invited lectures were delivered by international wellknown scientists: Professors Jens Frehse, Hitashi Ishii, Jacques-Louis Lions, Sanjoy Mitter, Umberto Mosco, Bernt Øksendal, George Papanicolaou and Albert Shiryaev.

The second day has been devoted to the tight links between research and applications and the impact of research on the 'real world' especially through valorization of research results. This theme has been enlighted through round tables on the following topics: space, spatial applications, science and technology for information and communication. A last panel discussed the role of Europe as a priviledged example of the collaboration between research and industry.

Moreover, a special volume, entitled 'Optimal Control and Partial Differential Equations' edited by José Luis Menaldi, Edmundo Rofman and Agnès Sulem has been published by IOS Press on this occasion. It includes the proceedings of the lectures given on December 4th together with articles written in honor of Alain Bensoussan by his colleagues, collaborators and former students. Their topics include domains where Alain Bensoussan has offered leading contributions such as functional and numerical analysis, stochastic partial differential equations, nonlinear filtering and identification, dynamic programming and stochastic control, mathematical finance and operations management.

Please contact:
Agnès Sulem – INRIA
Tel +33 1 3963 5569
E-mail: Agnes.Sulem@inria.fr

Left to right: Edmundo Rofman, Agnès Sulem, Alain Bensoussan and José Luis Menaldi.

CALL FOR PARTICIPATION

2nd International Workshop on User Interfaces to Data Intensive Systems

Zurich, Switzerland,
31 May - 1 June 2001

UIDIS 2001 continues the tradition of previous workshops in Interfaces to Database Systems (IDS 1992, IDS 1994, IDS 1996) and the UIDIS 1999, and seeks to bring together researchers in information management systems, human-computer interaction and knowledge based methodologies. It provides a platform for the presentation of current results and encourages the exchange of ideas on how user interfaces to data intensive systems should be improved from a conceptual as well as from a practical viewpoint. UIDIS 2001 will provide a forum to present original research contributions in an informal atmosphere allowing sufficient time for discussions and debate as well as presentations and demonstrations covering the following topics:

- Information retrieval techniques
- Interactive systems design
- Information visualization.

Further information:
http://www.uidis01.ethz.ch/

CALL FOR PARTICIPATION

e-Smart 2001

Cannes, France,
19-21 September 2001

e-Smart 2001 is an international conference on research in smart cards jointly organized by Java Card Forum, INRIA and Eurosmart. The conference will provide a forum for discussion and exchange of results on smart cards development, security and applications.

Scopes and Topics

- Smart card cryptography and electronic signatures
- Smart card security solutions (hardware and software technology, contact and contactless)
- Formal Methods for Smart card Evaluation and Validation
- Advanced architectures for multi applications and secure open platforms
- Object-oriented methodologies for Smart cards
- Smart card programming environments
- Middleware for Smart cards
- Novel applications of Smart cards.

More information:
http://www-sop.inria.fr/eSmart2001/

ERCIM NEWS

ERCIM News is the magazine of ERCIM. Published quarterly, the newsletter reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. ERCIM News has a circulation of 7000 copies.

ERCIM News online edition is available at http://www.ercim.org/publication/ERCIM_News/

ERCIM News is published by ERCIM EEIG, BP 93, F-06902 Sophia-Antipolis Cedex
Tel: +33 4 9238 5010, E-mail: office@ercim.org
ISSN 0926-4981
Director: Bernard Larrouturou
Central Editor:
Peter Kunz
E-mail: peter.kunz@ercim.org
Tel: +33 1 3963 5040
Local Editors:
CLRC: Martin Prime
E-mail: M.J.Prime@rl.ac.uk
Tel: +44 1235 44 6555
CRCIM: Michal Haindl
E-mail: haindl@utia.cas.cz
Tel: +420 2 6605 2350
CWI: Henk Nieland
E-mail: henkn@cwi.nl
Tel: +31 20 592 4092
CNR: Carol Peters
E-mail: carol@iei.pi.cnr.it
Tel: +39 050 315 2897
FORTH: Constantine Stephanidis
E-mail: cs@csi.forth.gr
Tel: +30 81 39 17 41
GMD: Dietrich Stobik
E-mail: stobik@gmd.de
Tel: +49 2241 14 2509
INRIA: Bernard Hidoine
E-mail: bernard.hidoine@inria.fr
Tel: +33 1 3963 5484
SICS: Kersti Hedman
E-mail: kersti@sics.se
Tel: +46 8633 1508
SRCIM: Gabriela Andrejkova
E-mail: andrejk@kosice.upjs.sk
Tel: +421 95 622 1128
SZTAKI: Erzsébet Csuhaj-Varjú
E-mail: csuhaj@sztaki.hu
Tel: +36 1 209 6990
TCD: Carol O'Sullivan
E-mail: Carol.OSullivan@cs.tcd.ie
Tel: +353 1 60 812 20
VTT: Pia-Maria Linden-Linna
E-mail: pia-maria.linden-linna@vtt.fi
Tel: +358 0 456 4501

Free subscription
You can subscribe to ERCIM News free of charge by:
- sending e-mail to your local editor
- contacting the ERCIM office (see address above)
- filling out the form at the ERCIM website at http://www.ercim.org/

CALL FOR PARTICIPATION

ISDL 2001 – First DELOS Summer School on Digital Libraries Technologies

Pisa, Italy, 9-13 July 2001

The first 'DELOS International Summer School on Digital Library Technologies' (ISDL) will take place in Pisa, Italy on July 9-13, 2001, organized by the DELOS Network of Excellence on Digital Libraries (an initiative funded by the 5th Framework Programme of the European Commission). The School is directed by Professor Maristella Agosti, University of Padua.

The main purpose of the school is to foster research and understanding in the fundamental technologies underlying the Digital Libraries (DL) field. It is directed towards members of the research community in the wide sense, that is, graduate students and young researchers and professionals involved in R&D in DL-related areas, possibly representing both the information technology scientist, the industrial communities (electronic publishing, broadcasting, software industry, etc.) and the user communities interested in DL technologies (libraries, archives, museums). The one-week intensive course will consist of 8 to 10 half-day lectures, including the following ones:

- Introduction to DL of Multimedia Digital Collections (Carl Lagoze, Cornell Univ., USA)
- DLs: New Models for Scholarly Dissemination (Robert Wilensky, Univ. of California, Berkeley, USA)
- DL architectures and Open Access to DLs (William Arms, Cornell Univ., USA)
- DL and IR: IR models and methods, metadata, and evaluation (Norbert Fuhr, Univ. of Dortmund, Germany)
- Online Information Access from Handheld Devices (Andreas Paepcke, Stanford Univ., USA)
- Cross-language Retrieval (Carol Peters, CNR-IEI, Italy)
- Text Categorization and Information Filtering (Fabrizio Sebastiani, CNR-IEI, Italy)
- Video DLs (Howard Wactlar, Carnegie Mellon Univ., USA).

Further information:
http://www.ercim.org/delos/

CALL FOR PARTICIPATION

First IFIP Conference on e-commerce, e-business, e-government

Zurich, 4-5 October 2001

This conference provides a forum for users, engineers, and scientists in academia, industry, and government to present their latest findings in e-commerce, e-business, or e-government applications and the underlying technology to support those applications.

Further information:
http://www.ifi.unizh.ch/I3E-conference/

CALL FOR PARTICIPATION

Cross-Language Evaluation Forum

Registration is now open for the CLEF 2001 evaluation campaign.

There will be three main evaluation tracks in CLEF 2001, testing multilingual, bilingual and monolingual (non-English) information retrieval systems. There will also be a special sub-task for domain-specific cross-language evaluation and an experimental track testing interactive cross-language systems.

The CLEF test collection for 2001 consists of a multilingual corpus of newspaper and newswire documents for English, French, German, Italian, Spanish and Dutch from the same time period. Query sets will also be provided in additional languages including Russian, Japanese and Greek.

Important Dates

- Data Release: 1 March 2001
- Topic Release: 9 April 2001
- Receipt of results from participants: 15 May 2001
- Release of relevance assessments and individual results: 15 July 2001
- Submission of paper for Working Notes: 7 August 2001
- Workshop: 3-4 September 2001 (in conjunction with ECDL 2001)

Further information:
http://www.clef-campaign.org/

Please contact:
Carol Peters – IEI-CNR
Tel: +39 050 315 2987
E-mail: carol@iei.pi.cnr.it.

CALL FOR PARTICIPATION

SIM2001 – Trade Fair for Simulation and Visualisation: The Virtual Product: from Automobile to Zeppelin

Freiburg i.Br., Germany, 17-21 June 2001

GMD Institute for Algorithms and Scientific Computing (SCAI) will demonstrate innovative products and services in Simulation, Optimisation and Visualisation. The eye-catcher of the GMD stand is the new 10sqm huge virtual reality environment 'Powerwall' of its partner T.A.N. Projektions-technologie. Using this new high-brightness DLP-based panoramic system, featuring TAN's new INFITEC stereo-viewing 3D technology, current developments of the institute SCAI can be presented in 3D with an unprecedented size and accuracy. Important applications providing innovative solutions in Computational Fluid Dynamics, Crash simulation and – as a premiere – Coupling based on the MpCCI software will be shown.

To compete in the market today, companies need to shorten production times, increase quality, and minimise development and production costs. Cross-company logistics processes, the digitalisation of products (digital mock-up), application service portals and simulation technology will be a key to the global success of companies in the future. For the first time, a trade fair devoted specifically to the wide variety of industrial applications for this new technology is to be held under the same roof. In addition, a programme of accompanying events will be offered.

SIM2001 will provide an overview of the current key issues in industrial product development. Apart from innovative hardware and software solutions, new sector-specific applications will also be on show. The importance being attached to this new trade fair is clearly underlined by the involvement of companies such as Compaq, CD adapco, Crossair, DaimlerChrysler, Fluent, Hewlett-Packard, MSC Software, NEC, Silicon Graphics and SUN Microsystems. Further, the German Arbeitsgemeinschaft für Simulation (ASIM) and GMD's institute SCAI are participating in the SIM2001 conception.

SIM2001 will also be accompanied by two conferences, the Simulation and Visualisation 2001 Conference and World Fluid Dynamics Days 2001 with a special Aviation Session, aimed at users and experts in the field of simulation. The trade fair will be rounded off by SIM2001 Recruiting Day for Engineers aimed at the newly qualified engineers and SIM2001 Partnering Day to encourage new cooperative ventures.

Further Information:
http://www.sim2001.com/
http://www.gmd.de/SCAI

Please contact:
Stephan Springstubbe – GMD
Tel: +49 2241 14 2337
E-mail: springstubbe@gmd.de

CALL FOR PAPERS

TestCom 2002 – Testing Internet Technologies and Services – The IFIP 14th International Conference on Testing of Communicating Systems

Berlin 19-22 March 2002

TestCom denotes a series of international working conferences on testing communicating systems in the datacom and telecom domain. The conference will provide a forum for researchers, vendors and users to review, discuss and learn about new approaches, concepts and experiences in the field of testing of communicating systems. The conference will consist of tutorial sessions on the first day, followed by presentations of reviewed and invited papers, demonstrations of tools, and panel and discussion sessions. This conference will, in particular, encourage exchange and discussion of concepts and practices in testing Internet related systems, services and protocols. The conference will be held at Berlin-Brandenburg Academy of Sciences and Humanities at the beautiful historic place in Berlin, the Gendarmenmarkt.

Important Dates
- 8 September 2001: Submission deadline for regular papers and tutorials
- 26 October: 2001 Notification of acceptance
- 23 November 2001: Camera-ready copies
- 15 January 2002: Submission deadline forposition statements and tool demonstrations

Topics of interest include, but are not restricted to

- Types of testing: Conformance testing, Interoperability testing, Performance testing, Scalability testing, Robustness testing, Regression testing
- Phases of testing: Test case generation, Test selection and management, Test setup and execution, Test coverage and results analysis
- Classes of systems to be tested: Internet protocols and services, Internet and IN, Internet and UMTS, Active networks, Ad-hoc networks, Middleware platforms, Multimedia systems, E-commerce systems
- Theory and Practise of testing: Theoretical frameworks, Test specification, Test derivation, validation and analysis, Applications of testing theory, Industrial testing experience, Test standards, Test tools.

Several types of contributions are solicited:
- Full papers - research or practice results (no more than 16 pages)
- Tutorial proposals
- Position statements - work in progress/experience reports (no more than 5 pages)
- Tool demonstration proposals.

All accepted full papers will be published by Kluwer. TestCom'2002 is organized by GMD FOKUS, the German Research Institute for Open Communication Systems, and by BTU Cottbus, the Brandenburg University of Technology at Cottbus.

Further Information
http://www.fokus.gmd.de/events/testcom2002
http://www.ifip.or.at/
E-mail: testcom2002-org@fokus.gmd.de

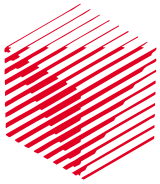

ERCIM – The European Research Consortium for Informatics and Mathematics is an organisation dedicated to the advancement of European research and development, in information technology and applied mathematics. Its national member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry.

Central Laboratory of the Research Councils
Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire OX11 0QX, United Kingdom
Tel: +44 1235 82 1900, Fax: +44 1235 44 5385
http://www.cclrc.ac.uk/

Consiglio Nazionale delle Ricerche, IEI-CNR
Area della Ricerca CNR di Pisa, Via G. Moruzzi, 1; 56124 Pisa, Italy
Tel: +39 050 315 2878, Fax: +39 050 315 2810
http://www.iei.pi.cnr.it/

Czech Research Consortium for Informatics and Mathematics
FI MU, Botanicka 68a, CZ-602 00 Brno, Czech Republic
Tel: +420 2 688 4669, Fax: +420 2 688 4903
http://www.utia.cas.cz/CRCIM/home.html

Centrum voor Wiskunde en Informatica
Kruislaan 413, NL-1098 SJ Amsterdam, The Netherlands
Tel: +31 20 592 9333, Fax: +31 20 592 4199
http://www.cwi.nl/

Foundation for Research and Technology – Hellas
Institute of Computer Science, P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
Tel: +30 81 39 16 00, Fax: +30 81 39 16 01
http://www.ics.forth.gr/

GMD – Forschungszentrum Informationstechnik GmbH
Schloß Birlinghoven, D-53754 Sankt Augustin, Germany
Tel: +49 2241 14 0, Fax: +49 2241 14 2889
http://www.gmd.de/

Institut National de Recherche en Informatique et en Automatique
B.P. 105, F-78153 Le Chesnay, France
Tel: +33 1 3963 5511, Fax: +33 1 3963 5330
http://www.inria.fr/

Swedish Institute of Computer Science
Box 1263, SE-164 29 Kista, Sweden
Tel: +46 8 633 1500, Fax: +46 8 751 72 30
http://www.sics.se/

Swiss Association for Research in Information Technology
Dept. Informatik, ETH-Zentrum, CH-8092 Zürich, Switzerland
Tel: +41 1 632 72 41, Fax: +41 1 632 11 72
http://www.sarit.ch/

Slovak Research Consortium for Informatics and Mathematics
Dept.of Computer Science, Comenius University, Mlynska Dolina M SK-84248 Bratislava, Slovakia
Tel: +421 7 726635, Fax: +421 7 727041
http://www.srcim.sk

Magyar Tudományos Akadémia – Számítástechnikai és Automatizálási Kutató Intézete
P.O. Box 63, H-1518 Budapest, Hungary
Tel: +36 1 4665644, Fax: + 36 1 466 7503
http://www.sztaki.hu/

Trinity College, Department of Computer Science,
Dublin 2, Ireland
Tel: +353 1 608 1765, Fax: 353 1 677 2204
http://www.cs.tcd.ie/ERCIM

Technical Research Centre of Finland
VTT Information Technology, P.O. Box 1200, FIN-02044 VTT, Finland
Tel:+358 9 456 6041, Fax :+358 9 456 6027
http://www.vtt.fi/