

ERCIM NEWS

European Research Consortium for Informatics and Mathematics

Number 15

December 1993

EDITORIAL



After five years ERCIM is firmly established and is ready for further challenges, says Cor Baayen, President of ERCIM (Photo: CWI)

SPECIAL :

Activities in
Large-scale Scientific Computing

C O N T E N T S

Editorial	1
Joint ERCIM Actions	2
The European Scene	4
Research Activities:	5
Large-scale Scientific Computing	
Other Activities	14
Technology Transfer	22
International Relations	26
Events	27
In Brief	31

Next Issue:

Databases

Five years ago ERCIM was established. Now it is a good moment to reflect briefly on its original objectives, achievements and future development. Focal point was, and still is, to create a truly European research community in information technology (IT) and related mathematics. Whereas for example research in physics and in astronomy is already for quite some time strongly organized in Europe, the European map of IT research still presents a scattered picture. A particularly weak point – repeatedly noticed, but not yet essentially improved – is the transfer of basic research results to practical applications in industry. ERCIM members are all national institutes with a mission in technology transfer. Thus ERCIM, as an European consortium, is in an excellent position to meet this challenge on the European scale.

During the period of rapid expansion – ERCIM membership rose from three in 1990 to ten in 1993, the 11th member joining before the end of the year – the partners became well acquainted with each other and several joint actions were initiated, including a series of symposia on technology transfer, the first of which, "Affordable Parallel Systems", will take place in the UK this fall. By now ERCIM, representing some 5000 researchers, has become a significant force in Europe, recognized as such by the EC, as was shown recently by, e.g., its substantial support of ERCIM's rapidly growing fellowship programme.

Today Europe is facing serious political as well as economic problems: instability in Eastern and Central Europe, stagnation of the European unification process and a world-wide recession. ERCIM's objectives remain unaffected by these developments – in fact, their realization is more than ever necessary. Only the means by which to achieve them need adjustment. In particular emphasis will now be put on large-scale projects, to be carried out jointly by ERCIM partners and European industries. A first example is EDGE – a European Distributed Generic Environment, providing state-of-the-art tools for IT engineering design on the European level, which ERCIM will submit as an EC Priority Technology Project (PTP). Thus the ERCIM-generated European integration in the IT research field will be utilized to the utmost and brings ERCIM a step closer to its goal to help create a strong and competitive IT research community in Europe.

Cor Baayen

ERCIM Workshops

4th ERCIM Database Research Group Workshop: Repositories, Methods and Tools for System Engineering

by Eleni Petra

Over 30 participants from nearly all ERCIM sites attended the 4th ERCIM EDRG Workshop at Ierapetra, Crete, 3-5 May 1993, which was organised and hosted by FORTH.

The need for this workshop arose during last year's ERCIM Workshop on Methods and Tools for Software Reuse (Heraklion, Crete, 29-30 October 1992). During this meeting, several interesting similarities on research issues in Software Engineering and in Databases were observed and discussed. As a result, the synergy between the two areas was thought to be of mutual benefit, and a decision was made to conduct an in-depth workshop in the framework of the established EDRG group workshop series.

Continuing the strong tradition, the workshop was well attended and was conducted with lively discussions following the formal presentations which identified the research issues. The main topic was the software repository.

Requirements of software repositories, design and architectural issues, choices for its implementation (Object-oriented databases, traditional databases, knowledge bases, etc) roles assumed by repositories for the integration or interoperability of software development tools, were among the issues analysed during the two days intensive meeting.

The presentations and discussions revealed that a number of relevant projects and activities already exist or are to be

formed in ERCIM sites. Links were established between the research groups and decisions were made for joint work among researchers from different organisations.

The proceedings were distributed to the participants and an edited final version is to appear as an ERCIM Report.

Please contact: Eleni Petra - ICS-FORTH
Tel: +30 81 221171
E-mail: petra@ics.forth.gr

ERRATUM

ERCIM Electronic Mailing-lists

The names of the two mailing-lists set up at ERCIM's central office, announced in the last issue of ERCIM News, will now be called:

1. Ercim-activities@inria.fr

(instead of: Ercim.activities@inria.fr)

This corresponds to the public unmoderated mailing-list for the exchange of information on activities organised by ERCIM.

In order to be added to the above mailing-list, send a message to:

Ercim-request@inria.fr

(instead of: Ercim.requests@inria.fr)

2. Ercim-fellows@inria.fr

(instead of: Ercim.fellows@inria.fr)

This mailing-list is intended for scientists who wish to receive information about the ERCIM Fellowship programme.

Please contact: Alain Michard - INRIA
Tel: +33 1 39 63 5472
E-mail: alain.michard@inria.fr

ERCIM Workshops

The ERCIM fall workshops, which took place in November this year, were hosted by INRIA and located at the different INRIA sites throughout France.

For more information on the four workshops please contact:

• Multimodal Human-Computer Interaction

Nancy, 2-4 November 1993

Please contact: Noelle carbonell

Tel: +33 83 59 3055

E-mail: Noelle.Carbonell@loria.fr

• Development and Transformation of Programs

Nancy, France, 2-4 November 1993

Please contact: Jeannine Souquieres

Tel: +33 83 59 3033

E-mail: Jeanine.Souquieres@loria.fr

• Parallelism and Non-determinism in Active Database Application

Rocquencourt, 2-4 November 1993

Please contact: Stéphane Grumbach

Tel: +33 1 39 63 5446

E-mail: Stephane.Grumbach@inria.fr

• Stochastic Numerical Methods for Solving PDE's, and Applications

Sophia-Antipolis,

8-9 November 1993

Please contact: Denis Talay

Tel: +33 93 65 78 98

E-mail: Denis.talay@sophia.inria.fr

The 1993/94 ERCIM Fellowship Programme

by Annick Theis -Viémont

As announced in the previous issue of ERCIM News, the accompanying tables show the names of the new ERCIM fellows and the institutes they will visit.

The placements of the fellows at the different ERCIM member institutes are given in Table 1. The ten members of ERCIM will host fellows with the following allocation:

AEDIMA (Spain)	1 fellow
CNR (Italy)	1 fellow
CWI (The Netherlands)	3 fellows
FORTH (Greece)	1 fellow
GMD (Germany)	3 fellows
INESC (Portugal)	2 fellows
INRIA (France)	3 fellows
RAL (United Kingdom)	4 fellows
SICS (Sweden)	1 fellow
SINTEF (Norway)	3 fellows

The ERCIM fellows from previous rounds are completing their fellowships according to their schedule. They are presently at CNR (D. Breslauer), CWI (A. Bouali, R. Boucherie, C. Wuthrich), GMD (F. Penz), INESC (E. Siegel).

Concurrently, ERCIM is running another fellowship programme under the Human Capital and Mobility Programme. The different placements are given in Table 2.

Please contact:

Annick Theis-Viémont - INRIA

Tel: +33 1 39 63 53 78

E-mail: Annick.Theis-Viemont@inria.fr

Name ERCIM Institute

Michael Dracopoulos (Greece)	RAL 16/08/93	CWI 16/02/94	GMD 16/08/94
---------------------------------	-----------------	-----------------	-----------------

George Drettakis (Greece)	INRIA 01/01/94	AEDIMA 01/12/94	GMD 01/03/95
------------------------------	-------------------	--------------------	-----------------

Ludovicus Florack (The Netherlands)	INRIA 01/01/94	INESC 01/10/94
--	-------------------	-------------------

César Galindolegaria (Mexico)	CWI 01/09/93	SINTEF 01/03/94	CNR 01/09/94
----------------------------------	-----------------	--------------------	-----------------

Gerd Hillebrand (Germany)	INRIA 04/01/94	FORTH 01/07/94	RAL 01/01/95
------------------------------	-------------------	-------------------	-----------------

Olivier Maffeis (France)	RAL 15/09/93	GMD 15/03/94
-----------------------------	-----------------	-----------------

Danilo Montesi (Italy)	INESC 01/09/93	RAL 01/03/94	SINTEF 01/09/94
---------------------------	-------------------	-----------------	--------------------

Alessandro Panconesi (Italy)	CWI 15/09/93	SINTEF 16/05/94	SICS 16/10/94
---------------------------------	-----------------	--------------------	------------------

Table 1: Placement of the 1993/94 ERCIM fellows

Name ERCIM Institute

Alan Craig (United Kingdom)	INRIA 01/09/93	CWI 01/03/94	SINTEF 01/09/94
--------------------------------	-------------------	-----------------	--------------------

Philippe Nehlig (France)	RAL 01/09/93	CNR 01/09/94
-----------------------------	-----------------	-----------------

Mukesh Patel (United Kingdom)	CNR 01/10/93	FORTH 01/10/94
----------------------------------	-----------------	-------------------

Alan Stewart (United Kingdom)	INRIA 01/01/93	CWI 01/05/93
----------------------------------	-------------------	-----------------

Dimitri Theodoratos (Greece)	RAL 10/09/93	INRIA 10/07/94
---------------------------------	-----------------	-------------------

Gerry Tretmans (The Netherlands)	SINTEF 01/02/93	FORTH 01/10/93	GMD 01/06/94
-------------------------------------	--------------------	-------------------	-----------------

Table 2: Placement of the fellows in the ERCIM/Human Capital and Mobility Programme

ERCIM Comments on the Fourth EC Framework Programme

by Alain Michard

At the end of June, the Commission made public its official proposal for the fourth framework programme for research and technological development. The general organisation of the programme is still the same as what was proposed in the two "working-documents" which were published in November 1992 and March 93. The programme is organised along four major "actions".

- The first one gathers all the R&D activities. Seven broad domains are described: Information and Communication Technologies (ICT), Industrial Technologies, Environment, Life Sciences and Technologies, Energy, Research for a European Policy of Transports, Finalised Socio-Economic Research.

ICTs are spread over four sub-domains: telematic applications of general interest, technologies for integrated information and communication systems, technologies for advanced communication services, information technologies. An important difference with the third framework programme in our domains is that the Computer Integrated Manufacturing projects will now be undertaken by the Industrial Technologies area (formerly BRITE-EURAM).

- The second main action aims at developing R&D cooperation with non-EC countries. Several "targets" are mentioned: cooperation with non-EC programmes or entities within Europe (eg: EUREKA, COST, ESA, CERN, etc.), cooperation with Central European countries and with the Independent States of the former Soviet-Union, cooperation

with non-European industrialised countries, and with developing countries.

- The third action deals with the spreading and transfer of results arising from the EC R&D programmes. Follow-up of well-known programmes such as VALUE and SPRINT are proposed, as well as new "financial measures" aiming at supporting emerging high-tech companies.
- The fourth action is entitled "Stimulation of Training and Mobility of Researchers in the Community". It is in fact a follow-up of the well-known Human Capital and Mobility Programme.

We shall not present here the proposed financial breakdown among actions and domains, as it is likely to change substantially during the negotiations between the Commission, the Council and the European Parliament.

The ERCIM Executive Committee has prepared a short report giving some comments on the CEC proposal. Our main point is that although the general thrust towards a more focused programme is welcome, a strong effort towards a better vertical integration of the projects aiming at designing basic components and generic technologies, integration of these components in higher-level platforms or technologies, and infrastructure and application development, is still needed. ERCIM proposes that the first action (R&D activities) should be organised around a set of well-identified Strategic Goals.

Examples of possible Strategic Goals in the IT domain are:

- the European Design Generic Environment,

- the World-wide Library,
- the Interactive Numerical Television,
- etc.

Whatever the final choices could be, Strategic Goals should lead to significant increases in the scientific and technical knowledge available to European organisations. They should address real needs of society in a credible way, and should require resources exceeding the capabilities of a single organisation or country. Last but not least, they should have a significant chance of success, with the ability to demonstrate results at intermediate stages and produce results of wide applicability. Strategic Goals should be intellectually ambitious, preparing future technologies and methodologies for the long-term. They should therefore be addressed through well-coordinated projects, lasting for several years (six to ten, for instance).

The ERCIM report also comments briefly on the necessity of lighter and more effective management procedures for the R&D specific programmes, and propose deep changes in the mechanisms aiming at technology transfer and exploitation of results.

Negotiations on the fourth framework programme will take place during this summer and are likely to continue at least until the end of this year. ERCIM will continue to make efforts to convince the EC institutions of the importance of strategic research for the long-term economic development of Europe.

Please contact: Alain Michard - INRIA
Tel: +33 1 39 63 54 72
E-mail: Alain.Michard@inria.fr

Management of Distributed Work – ESPRIT Project EuroCoOp Concluded

by Thomas Kreifelts

The GMD Institute for Applied Information Technology has been actively involved in the ESPRIT project "EuroCoOp" for the last two years. Within the scope of this project – devoted to using information technology to provide support to distributed cooperative work – GMD has developed an "activity assistant" for coordinating distributed activities. The project has now been successfully concluded by the German, Danish and British partners at the final review in Denmark.

The need for computer support in the management of distributed work is increasing constantly – a fact illustrated, for example, by consortia operating across organisational and national boundaries for the purpose of implementing large-scale projects or by governmental and parliamentary organisational struc-

tures which are distributed over various different locations and sites. The Danish state-owned enterprise supervising the construction of the bridge across the "Great Belt" provided the ESPRIT project "EuroCoOp – information technology support for distributed cooperative work" with precisely the right user organisation.

The activity assistant enables users to define, enhance or modify cooperative activities, supervise work progress, make joint use of documents and services in the activity context and exchange information. The activity assistant distributes activity descriptions, appended documents and information to the participating users in a consistent manner. The primary user interface of the activity assistant is a hierarchically structured "To do" list which contains all activities which a user is participating in and which permits direct access to the relevant information relating to an activity.

The activity assistant has been subjected to initial assessment in conjunction with the Danish state-owned enterprise. This evaluation was performed within the framework of user workshops at which managers, engineers and secretaries had the opportunity to work with the activity assistant.

At the final review held in Nyborg on the Great Belt, the activity assistant and the other components of the EuroCoOp system – an organisation information system, a distributed hypermedia system and a synchronous desktop conferencing system – were presented in integrated form.

The work on the theme "Management of distributed work" will be continued over coming years within the framework of the new ESPRIT project "EuroCODE" at the GMD Institute for Applied Information Technology. At the heart of the work is the creation of an open development environment for CSCW systems (Computer Supported Cooperative Work) which in particular also provides the necessary basis for mobile forms of cooperative work.

Please contact: Thomas Kreifelts - GMD
Tel.: +49 2241 14 2643
E-mail: thomas.kreifelts@gmd.de

RESEARCH ACTIVITIES - Large-Scale Scientific Computing

Large-Scale Scientific Computing: Activities within the ERCIM Institutes

by Karstein Sørli

For most people, I think, Scientific Computing is quite well defined and has more or less the same meaning. It is concerned with the use of computers for the numerical solution of scientific problems, formulated by scientists and engineers in a diversity of fields. Examples are found in chemistry, physics, mechanics, electronics and civil engineering. As a matter of fact, solving problems in science and engineering, is one of the earliest, and still one of the largest, uses of computers. The common denominator is

to find solutions of mathematical models that represent some physical, chemical or biological problem or process. The techniques that are used to obtain such solutions are part of a general area called Scientific Computing, and the use of these techniques to get better insight into scientific and engineering problems is often denoted by Computational Science or Computational Engineering.

The term "large scale", in front, is perhaps a bit vague to some people. However, the meaning is quite simple. Its focus is on large and complex

models which very often demands tailor-made solution techniques and a tremendous amount of computing power, only obtained from high-end supercomputers.

The most frequently used mathematical modelling technique in this area of science is based on the formulation of given problems by systems of differential equations and some additional conditions. These systems are then discretized by some numerical approximation method, and the resulting system of algebraic equations are solved by an appropriate numerical procedure, implemented on a computer. Techniques for the interpretation of the large scale numerical solution has become an important issue during the last ten years. This is due to the vast amount of numerical data created by such computer calculations. As a matter of fact, it has grown to be a particu-

lar research field on its own, called Scientific Visualization. It is now an important part of Computational Science.

The evolution of computing and visualization equipment technology, environments for software development as well as algorithmic design, will ultimately make it possible to build *interactive* "platforms" for large-scale modelling and simulation of scientific and engineering problems. Up to now the concept of interactive simulation has been limited to relatively simple models. With the rapid development of high-performance computers, including parallel and distributed computing facilities and specialized graphics hardware as well as software, these limitations may not exist in the future. Then there is a new and challenging area of research for computational scientists, and new possibilities for engineering design. Computationally based engineering design and analysis could easily be revolutionized in the future, by means of the interactive and "simultaneous nature" of such simulation environments. This may result in better products in shorter times and safer systems and environments for people.

The ERCIM Workshop "Interactive Modelling, Simulation and Visualisation in Large-Scale Scientific Computing" at the Røros meeting, May 26-28 this year covered several of the above topics. The main goal of the workshop was to bring together scientists and engineers to exchange experience, ideas and thoughts about existing and future technology and techniques, which are believed to be important in the context of interactive simulation of processes and phenomena in science and engineering. The main themes of the workshop were:

- High-Performance Computing with Advanced Computer Architectures;
- Environments and Tools for Interactive Large-Scale Numerical Simulation and Scientific Visualization.

High-Performance Computing

The discussion group on "High Performance Computing" gave the following recommendations. Emphasis should be placed on the parallel solution of differential equations and associated numerical problems, in particular systems of algebraic equations (linear and

nonlinear). Application examples should include problems in physics, engineering, elasticity, CFD, chemistry, environmental problems, weather prediction problems and other similar applications. Model problems involving the solution of two or three dimensional PDE's should be considered as there is a lot to be learnt from solving such problems.

Two directions are considered important for these applications: the development of new parallel algorithms, and the parallelization of existing algorithms.

Research and development in both directions should examine the issues of scalability and portability and should emphasize techniques involving domain decomposition, partition and allocation, grid partitioning and data transposition. Local refinement techniques and adaptivity are particular new and important fields to be examined and developed in parallel computing.

In order to study the behaviour and performance of parallel algorithms it is necessary to improve on paradigms or models of parallel computation and combine the theoretical study with extensive experimentation.

Powerful techniques in interactive visualization are expected to play a key role in several aspects of the development, the study and the efficiency of parallel algorithms. Examples of interest include interactive visualization for the geometric aspects of the computation (e.g. grid generation and mesh adaptation), the postprocessing and presentation of results, and the visualization of the behaviour of algorithms based on monitoring tools and measurements.

Environments, Tools and Scientific Visualization

The working group on "Environments, Tools and Scientific Visualization" brought together a number researchers to discuss and exchange a number of ongoing issues on visualization tools for interactive modelling, simulation and visualization in large-scale scientific computing. The participants of the working group came from a diversity of application areas, which was viewed as very positive and resulted in some stimulating discussions.

The following observations were made:

- All participants noted that the computing resources required for simulation of their models are steadily increasing. This has a number implications for visualization:
 - a) Model Complexity – the requirements on underlying models are becoming more stringent, resulting in larger and more realistic models. The computing resources needed to simulate and verify these models grow exponentially.
 - b) Computation Accuracy – the demands placed on the underlying computational accuracy is ever increasing. This places constraints in the interactive nature of the visualization. Some noted that adaptive techniques might provide interesting solutions to this problem.
- Most participants used visualization for verification purposes only. The most common working practice was simulation of existing models with some experimental data. Interactive techniques are then used to allow the perturbation of physical, geometrical and numerical parameters.

A number of common visualization techniques are being demanded:

- Intelligent probing: current probing techniques are not satisfactory as they can not provide sufficient information of the ongoing simulation. This is due to the decoupling of the visualization data types with the application data types. New techniques need to be provided which allow for information extraction and visualization of the "physical" meaning of underlying phenomena.
- History matching: new techniques must be found which allow matching of visualization on perturbed simulation parameters.
- Computational steering: current visualization techniques are based on post-processing results of a simulation. If direct interaction with an ongoing simulation is to be achieved, then novel techniques for navigation through spatial domains and monitoring of various simulation processes need to be studied.

Please contact: Karstein Sorli - SINTEF
Tel: 46-73596854
E-mail: Karstein.Sorli@sima.sintef.no

Scientific Computing for Large Systems at CNR

by Laura Moltedo

The 5-year CNR Project for "Information Systems and Parallel Computation" began in 1989. The project is organized in eight sub-projects (see ERCIM News14 for a description of the areas covered by each one). This last year is officially oriented towards product transfer activities.

The main objectives of the first sub-project: "Scientific Computing for Large Systems" are the following:

- to facilitate the use of supercomputers to solve high complexity problems;
- to design high level software products;
- to disseminate know-how on parallel computing.

In order to achieve these goals, the sub-project has been divided into three specific areas:

1. Modelling and Simulation
2. General Purpose Software in Scientific Parallel Computing
3. Networks for Access to Supercomputers.

For the system architecture, there has been a remarkable evolution: from an initial interest primarily oriented towards the vectorial adaptation of originally sequenced codes, we have moved to a new formulation of both global and distributed memory programs for supercomputers. The project has been very active in this field: during 1993, twenty-three operating units (4 CNR Institutes, 13 university departments, 6 industries) have been involved producing 406 papers, 175 technical reports and 31 prototype software libraries.

The products derived from this activity have been grouped into two classes: products with industrial scopes (computational codes oriented to the solution of specific application problems) and products interesting the scientific com-

munity (development of methodologies, know-how and general interest programs).

In the first class, software codes have been studied and implemented by industrial operating units for specific application fields, in particular:

- code to solve three-dimensional filtering problems, developed on Alliant and Cray systems;
- a multilevel parallel system for image processing (binocular stereovision), implemented on a transputer based environment;
- parallel software on workstation clusters for data processing in the geophysical area;
- code for a real-time helicopter flight simulator.

The second class regards software which has been developed by academic operating units in and for the following topics: mathematics, physics, chemistry, scientific visualization. Centres for the collection and diffusion of this software have been established within the Sub-project.

The centre for mathematical software, TecnoPolis, Bari, aims at efficient solutions of computational mathematical problems on shared and distributed architectures and concerns linear algebra, the solution of linear equation systems, non-linear optimization of large scale problems, multidimensional quadrature, O.D.E. solution, image reconstruction, Monte Carlo methods.

The centre coordinated by the Physics Department of Trento University is oriented towards theoretical and experimental aspects of solid state physics.

The centre for software regarding chemistry, Cineca, Bologna, aims at the study of molecular properties, spectroscopy, gaseous kinetics, chemiabsorption, phase transitions in liquid crystal and solutions.

The centre for scientific visualization software, again coordinated by Cineca, is dedicated to the tracking and interactive steering of numerical simulations, the representation of multidimensional fields, and animation.

The most important results of this sub-project were presented and demonstrated at the second national congress of the project for "Information Systems and Parallel Computation", held in Naples, 9-11 June 1993. The product catalogue, which contains a brief description of all the products developed during the project's life-time, including those developed within the Large Scale Scientific Computing area, will be available in English very shortly.

Please contact: Laura Moltedo - IAC-CNR
Tel: +39 6 88470210
E-mail: moltedo@iac.rm.cnr.it

Industrial Molecular Mechanics Code Ported to Parallel Platforms

by Horst Schwichtenberg and Gerd Winter

In 1992, GMD and a group at Hoechst AG started cooperation for parallelizing the Hoechst Computer Assisted Molecular Design software (CAMD). The molecular mechanics simulation tool MOLMEC is in regular use at Hoechst for computerized drug design. The CPU time is typically a function of the square of the number of atoms. The size of the molecular system which can be solved is therefore often limited by the CPU time available. GMD has ported the software to parallel platforms, thereby enabling more powerful machines to be used. Further research is planned to change the numerical algorithms, so that the CPU time is a linear function of the number of atoms, and to parallelize them for local memory machines.

The first step with our partners at Hoechst was the analysis of the MOLMEC code for parallelization. One of the main requirements was the need to run the par-

allel version efficiently on various platforms ranging from workstation clusters with Ethernet to monolithic parallel machines. Another condition set by Hoechst was to limit changes to the code so that it could still be integrated into their preprocessor. This preprocessor is able to produce source versions for various different applications and platforms.

The calculation of the force field was the main operation to be parallelized. Here the computer time is used mainly for calculation of the unbounded forces, that is, for finding the nearest neighbour atoms and subsequently evaluating the interaction terms for the atom pairs obtained. To determine the energies, the program also calculates the bonded forces (from the bond length, valence angles, torsion angles). All these different forces make it difficult to segment the computations with less communication.

Instead of a space partitioning, where the atoms are leaving the subdomains, we have chosen a splitting over the atoms. This is optimal for calculating unbounded forces in parallel, but problems arise for the other bonded components of the force field, resulting in large communication requirements. So it was decided to develop a special parallel version for a workstation cluster, replicating all parameters, coordinates and charges of all atoms on each node.

A careful analysis has shown that all components of the force field and the energies could then be calculated without any communication. Only for the final result a parallel summation along a binary tree is done.

This results in a speed-up of nearly 90% for a problem with 7000 atoms on a cluster of 2 to 6 IBM RS6000/350 workstations. The communication time in every iteration step is negligible compared to the calculation on a node.

The Ewald summation technique is another feature of MOLMEC. In this case typically only few atoms are used. So we searched for another splitting and found the images instead of the atoms.

A project is now planned for introducing a hierarchical algorithm, such as the fast multipole method (FMM) or multi-

level technique, for reducing drastically the CPU time requirements of the parallel simulation package.

Please contact:

Horst Schwichtenberg - GMD
Tel: +49 2241 14 2577
E-mail: schwichtenberg@gmd.de
or: Gerd Winter - GMD
Tel: +49 2241 14 2752
E-mail: gerd.winter@gmd.de

Continuous Speech Recognition as a Large-Scale Scientific Computing Problem

by João Neto and Luis Almeida

Continuous speech recognition is a hard task. The approach we use to this problem is based on a hybrid system that combines the statistical nature of hidden Markov models and the power, as static pattern recognizers, of connectionist models. The models that we use are typically large: 250,000 weights and 1,360,000 training patterns. We use a parallel processor to enable their training in a manageable time.

State-of-the-art continuous speech recognition systems are based on statistical approximation using hidden Markov models (HMM). Connectionist models have been used to improve continuous speech recognition systems through the exploitation of some of their features. The kind of connectionist model that we use is the multi-layer perceptron (MLP) which is used as an estimator of the state probabilities for the HMM. This hybrid system yields state-of-the-art recognition accuracy, competitive with the world's most advanced systems.

The MLPs used as probability estimators are typically large. In our baseline sys-

tem we have 182 input units, representing a feature analysis made over seven consecutive frames of the speech signal, a hidden layer of 1000 units and an output layer of 68 units, each unit representing a different phone. Each layer is fully connected to the next layer, giving a total of about 250,000 weights. The training corpus currently in use is the DARPA Resource Management database, which for our problem results in approximately 1,360,000 input patterns, corresponding to a total of 3690 sentences from 109 speakers.

Training this kind of network is computationally intensive. A training in a conventional RISC workstation would take about one week of computer time. For this reason there is a strong necessity of dedicated hardware that takes advantage of some of the features of this kind of network:

- large number of units
- repetitive multiply-add operations
- local information.

Presently we are using, for MLP training, a Ring Array Processor (RAP) machine from International Computer Science Institute (ICSI). The RAP is a system composed of several DSPs (Digital Signal Processors) operating in parallel and connected in a ring arrangement. They work under the control of a host computer running the UNIX operating system. Each DSP has local memory and is called a node. Our system has 12 nodes. The MLP has a highly regular connectivity between the processing units, with easily partitionable data sets for distribution among the nodes. The RAP provides 0.5 GigaFLOPS of processing power and was designed specifically for connectionist processing. Although the RAP is potentially a general purpose machine, its software is mainly tailored for neural network applications. In our specific case, the RAP enabled us to reduce the MLP training time to about one day.

The speech recognition research described above is being performed in the framework of the ESPRIT-BR Project 6487 – WERNICKE.

Please contact: Joao P. Neto - INESC
Tel: +351 1 3100315
E-mail: jpn@inesc.pt

Computational Mathematics for Differential Modelling (CMFDM) at IAC-CNR

by Daniela Mansutti

Researchers at the "Istituto per le Applicazioni del Calcolo" (IAC-CNR) are involved in a number of projects covering many aspects of CMFDM, a broad field ranging from ad hoc model rearrangement for numerical treatment to computer visualization, from qualitative description of solutions to numerical analysis and parallel computing.

In the context of CMFDM, several applications in continuum mechanics, engineering and industry are being tackled from different viewpoints: within fluid dynamics, crystal growth, non-linear elasticity and general conservation laws.

A study on the solution of Navier-Stokes equations, which is partially supported by the CNR Project "Information Systems and Parallel Computation: SP1 Scientific Computing for Large Systems", has recently led to a new finite difference scheme, the Modified Discrete Vector Potential method, for the formulation in vector potential with boundary conditions in velocity/pressure. As local and global mass conservation is met, a domain decomposition technique has been easily built for solving the huge system of discrete equations. Preconditioning is still under development.

The study and experimentation of Essentially Non-Oscillatory schemes for conservation laws and for 2D Euler equations for gases with reacting source terms is also funded by this Project. Multidimensional cases are solved by an A.D.I. like method, which is implemented for parallel computing. Wave propagation in crystals is a possible future application of these numerical methods. A

working group is also involved in the analysis of general conservation laws, the qualitative features of solutions and Hamilton-Jacobi type algorithms are being investigated.

The Bridgman technique for artificial crystal growth is an application of the project on the "Stability of Free Convection in Microgravity Conditions" and is funded by the Italian Space Agency. Here the Hopf bifurcation pattern of the Newtonian isochoric flow in a shallow 3D cavity with a horizontal differential of temperature is detected by a direct method for Hopf points and a spectral discretization. In crystal growth modelling, oscillatory instabilities are the main cause of dishomogeneities in the crystal.

Elementary flows of rate type non-Newtonian fluids which are relevant to fluid dynamics are numerically approached by B-spline collocation and spectral methods. Mechanical stability is often numerically studied.

The solution of solidification processes is being given increasing attention, either for modelling or for the numerical treatment and free-boundary description.

Some problems in semi-linear elasticity and exact controllability are also being studied; for example, the equilibrium of spherical cups. Modelling and qualitative description of solutions are the main contributions in this area. Spectral methods and finite differences are used where numerical tests are required.

General purpose techniques and software relevant to differential model solution are also studied: i.e. multigrid methods with interactive boundary fitting for multiple connected domains, computer representation of 3D vector fields, ray tracing techniques and isosurface identification. Codes are being developed for dynamic interfacing with numerical solvers.

Please contact:
Daniela Mansutti - IAC-CNR
Tel: +39 6 88470264
E-mail: mansutti@obelix.iac.rm.cnr.it

More Reliable Weather Forecasts by Parallel Computing

by Ute Gärtel, Wolfgang Joppich and Anton Schüller

The large economic benefit from a faster and more reliable weather prediction is obvious. Finer resolutions than currently used in weather models would allow more reliable and more precise forecasts, but they would also demand for an essentially larger computer power, which can be provided only by parallel computers. In the Institute for Foundations of Information Technology of the GMD a first breakthrough in the field of parallelization has been achieved under the leadership of Ulrich Trottenberg: For the first time, a meteorological production code has been run on a variety of parallel systems. The cooperation partners GMD and European Centre for Medium-Range Weather Forecasts (ECMWF) met at GMD on May 25th in order to take stock of the current situation and to discuss the future developments.

Up to now parallel computers have been used mainly in universities and large research institutions. New concepts ("parallelization") for the complex computer programs and associated mathematical developments ("parallel algorithms") are necessary to utilize these systems. For large applications and permanent use in industry and in central service institutions – such as weather centers – this kind of software is still missing.

In the cooperation between GMD and ECMWF in Reading (UK), the 2D version of the "Integrated Forecast System (IFS)" has been parallelized as a first step. The IFS is a complex program package, that is run daily for the European weather forecast.

Three employees of the GMD institute, Dr. Ute Gärtel, Dr. Wolfgang Joppich and Dr. Anton Schüller, have already finished the adaption of the integrated 2D model after half a year. This proves that an efficient parallelization of large production programs can be achieved with a justifiable amount of work – a result that has been called impossible by many experts until recently and that has been possible only due to the large know-how and preparatory work in the institute. The full 3D model is under development currently.

The parallelization is performed in a portable way. This means that the parallel program can be run on a variety of parallel computers from the current and future market (and also on workstation nets). On parallel systems such as Parsytec GCel, Meiko i860, CM5 of Thinking Machines, Intel iPSC, nCUBE/2 excellent parallel efficiencies have been achieved. The corresponding program package has been delivered recently.

Although this program does not yet treat the full 3D model with all physical calculations, the corresponding performance measurements and implementation experiences are of great interest to both the applications and the hardware vendors side. For example, at the RAPS-Workshop (Real Applications on Parallel Systems) at the end of May in Geneva, the parallel IFS-2D program has been suggested to be selected as the first application example to be used for a systematic benchmarking of large production codes on all available parallel computers.

During the meeting on May 25th (David Burridge, director of the ECMWF, and Geert Hoffmann, head of the Computer Division of the ECMWF, have been among the participants) the project partners agreed to continue the fruitful cooperation. A new challenge is "Parallel I/O". The efficient and portable Input and Output of a large amount of data still constitutes an unsolved problem in meteorology as well as in other application fields.

Please contact: Ute Gärtel - GMD
Tel.: +49 2241 142746
E-mail: ute.gaertel@gmd.de

Attacking Sources of Overhead in Parallel Applications

by Evangelos Markatos


There are three major sources of overhead in the execution of parallel applications: synchronization overhead which is the result of contention on synchronization variable load imbalance which is due to an uneven distribution of load, and communication overhead which results from non-local memory accesses, or cache misses. Thus, for applications to execute efficiently, their load has to be evenly balanced among all processors, and interprocess communication has to be reduced to a minimal, resulting to a high degree of locality within each processor.

These three overhead dimensions are usually at odds with each other: i.e. load balancing encourages frequent process

migration, which increases communication overhead, leading to a reduction in locality. One task of the parallelizing compiler, the run-time system, and the operating system is to maintain a delicate balance among these overheads, so as to achieve the best performance.

Traditional systems usually address only one, or two of the above overhead dimensions. For example, several shared-memory systems usually ignore communication overhead, as it is often considered to be small. Distributed-memory systems, on the other hand, usually ignore load balancing, because the mechanisms required for it (e.g. migration) are not implemented, or are very expensive to use.

Our work focuses on run-time system support for parallelizing compilers, and thread libraries for parallel programming environments and addresses all three overhead dimensions. One of our recent contributions is a scheduling policy called affinity scheduling, which deals with communication overhead first, and then addresses load balancing and synchronization. We have implemented our policy in the run-time environment of several shared-memory multiprocessors, ranging from the large-scale BNN Butterfly I, to the recent KSR-1 multi-



Gaussian elimination (1024 by 1024) on the KSR-1 multiprocessor.

processor. We have seen up to a factor of three improvement over traditional scheduling policies in the completion time of applications. In the figure, the completion time of the Gaussian elimination application under affinity scheduling is compared with the best traditional scheduling policy on the KSR-1 multiprocessor.

Our work so far, has focused on scientific applications. We plan to extend the evaluation of the above overhead dimensions in engineering, and data management (transaction processing) applications as well. We would like to experiment in environments where locality is dynamic and locality-preserving decisions are not straightforward. In those environments locality can be preserved by either moving data close to processes, or by moving processes close to their data, or by using some combination. Taking the right decision is not easy, since the system may lack lots of information, such as the location of the data, or the reference patterns of processes. To overcome these problems we plan to use tracing information from previous runs.

Part of this work was done jointly with Thomas J. LeBlanc, Computer Science Department, University of Rochester, Rochester, NY 14627, USA. This work is performed as part of the research conducted by the Parallel and Distributed Systems group at ICS FORTH.

Please contact:

Evangelos P. Markatos, ICS-FORTH

Tel: +30 81 234316

E-mail: markatos@csi.forth.gr

or:

Prof. C. Nikolaou, ICS-FORTH

E-mail: nikolau@csi.forth.gr

Interactive Use of CRAY Super-Computers at RAL

by Roger Evans

The idea of interactive supercomputing is still strange to many since the supercomputer is often associated with long running jobs leading to large files with numerical data, at the end of the process.


Based on new applications software, new opportunities for distributed graphics applications and major improvements in the bandwidth of wide area networks following the introduction of the SuperJANET network in the UK, we are expecting a much increased interactive and pseudo-interactive workload over the next two years.

Application Packages

Supercomputer users increasingly expect the comfort of interactive computing when analysing results or setting up computing experiments. Distributed application packages such as Cray's UNICHEM for computational chemistry is an example of how this can be done.

UNICHEM runs partly on a Silicon Graphics workstation and partly on a Cray under Unicos, a graphical front end enables crude molecular design and definition of model parameters before a job is initiated on the Cray to produce a detailed structure of the molecule. Numerical data is then returned automatically to the workstation for display of geometry and electron density maps and associated parameters.

UNICHEM runs efficiently over wide area networks of modest bandwidth and typically deals with no more than 20 - 30 atoms. With the new generation of 100 Mbit/sec networks it is feasible to transmit much more data during the real time calculation on the Cray and in collaboration with Birkbeck College we are constructing an application for molecular dynamics which uses simpler physical mod-



A molecule structure visualised in UNICHEM.

els in order to deal with many more atoms simultaneously. Real time graphics display will be implemented using the PVM (Parallel Virtual Machine) package.

Distributed Graphics Packages

The new generation of graphics applications builders such as AVS and Explorer have the capability to run numerically intensive modules on a more powerful remote machine.

For full 3D vector fields the amount of information that must be sent to give good interactive display response is large and one of these applications (from Imperial College) is running currently as a demonstrator project for the pilot SuperJANET network.

Data Browsers

In major environmental science projects the accumulation of data represents many months or even years of research followed by attempts to extract the maximum understanding from the tens of Gigabytes of numerical data.

Another SuperJANET pilot application is being developed for the study of climate modelling data in the Universities Global Atmospheric Modelling Project (UGAMP). Animation back and fore in time, vertical and horizontal sections and rapid switching between displayed fields gives a new insight into the atmospheric dynamics and energy flows.

Enhancements to the Atlas Cray

In order to ensure that our Cray Y-MP8I/8128 is efficient for interactive as well as batch work we have ordered a 256 MWord (2 GByte) SSD for improved swapping and IO performance and we plan to have a direct connection to the 155 Mbit/sec SuperJANET ATM network by the end of 1993. Later this year our Cray courses aimed at getting the best out of Cray hardware will also include the options for more productive distributed supercomputing.

Please contact: Roger Evans - RAL
Tel: +44 235 44 5656
E-mail: rge@uk.ac.rutherford.unixfe

LANCELOT: a Package for Large-Scale Function Minimization

by Nick Gould

According to legend, Sir Lancelot and his companion Knights of the Round Table were charged by God to return the Holy Grail to Camelot. The Dark Ages having given way to Enlightenment, LANCELOT can still help you to locate the grail and is now the acronym for "Large And Nonlinearly Constrained Extended Lagrangian Optimization Techniques".

The LANCELOT project is an international collaboration between Nick Gould (RAL), Philippe Toint (FUNDP, Namur, Belgium) and Andy Conn (IBM, Yorktown Heights, NY, USA) to develop widely applicable software for solving such problems.

A guiding aim of the project has always been that the algorithms incorporated in the software are based on solid theoretical foundations. The underlying theoretical behaviour of the methods has been, and is continuing to be, published in international research journals.

The basic optimization model studied in the project is that of finding a (local) solution to the problem:

$$\text{minimize } f(x)$$

where $x \in \mathcal{R}^n$ subject to the general (possibly nonlinear) constraints:

$$c_j(x) = 0; j \in E$$

$$c_j(x) \geq 0; j \in I$$

and the simple bounds on the variables:

$$l_i \leq x_i \leq u_i; 1 \leq i \leq n$$

Here, each of the functions f and c_j is differentiable, the index sets E and I are finite and any of the bounds l_i and u_i may be infinite. The aim has been to be able to solve genuinely nonlinear problems with up to, say, 10,000 variables on a reasonable workstation, and this goal is now being realized.

The key to being able to solve such problems is to recognize that, in most real problems, the problem functions may be expressed in the form:

$$\sum_{j \in J} f_j(x) + a^T x - b$$

where each nonlinear function f_j involves just a few (say 2 or 3) of the variables x and the vector a is sparse. Such functions are known as partially separable. This nonlinear form of sparsity may be exploited when solving the linear systems of equations which naturally arise at each iteration of algorithms for the model problem. Both direct and iterative methods may be used to solve the resulting linear systems.

The objective function and general constraints are combined in an augmented Lagrangian function. This composite function is (approximately) minimized within the region defined by the simple bounds, and the parameters defining the combination taken are adjusted until a solution to the original problem is identified. The method used for minimizing the composite function is specifically designed to be efficient in identifying which of the variables lie on bounds at the solution to the problem.

Release A of the software was available in May 1992. The package works in two phases. Firstly, users must describe their minimization problems in a standardized input language. This description is translated into data and Fortran programs which are then compiled and linked to the main LANCELOT load module. Secondly, users select from a number of options which control the minimization method used. The minimization program then starts to solve the problem and, with luck, the desired solution rapidly follows.

There are versions of LANCELOT for IBM mainframes, CRAY, and DECVAX, and also for most UNIX systems (UNIX is a registered trademark of AT&T). The package and its use are described in the book by Conn A.R., Gould N.I.M., and Toint Ph.L.: *LANCELOT: A Fortran Package for Large-Scale Nonlinear Optimization* (Release A), Springer-Verlag, Berlin, 1992.

Please contact: Nick Gould - RAL
Tel: +44 235 44 5790
E-mail: N.Gould@letterbox.rl.ac.uk

High Performance Computing Activities at SINTEF

by Leif Reidar Roekkum and Roar Skaalin

Supercomputer usage at Norway's four universities is varied, but mainly within fundamental science, such as computational chemistry, astrophysics, geophysics and solid state physics. In addition, there is a certain amount of petroleum related research needing supercomputing resources. The Norwegian Meteorological Institute (DNMI) absorbs a lot of computing power for weather and ocean forecasts as well as for environmental research. Supercomputer usage in industry is dominated by the needs of the exploration and production of the offshore oil and gas resources.

The first initiative to coordinate the supercomputing activities in Norway was taken by SINTEF and The Norwegian Institute of Technology (NTH) in 1987. Financing was obtained from the Norwegian research councils and Norwegian industry for a five year project involving the establishing of a research and support group in supercomputing, and installation of a CRAY X-MP/2. Today the high performance computing and related activities at SINTEF/NTH is based on three major projects:

- Production Supercomputing
- Parallel Technology in Research and Education at SINTEF/NTH
- Scientific Visualisation

The production supercomputing project represents a prolongation of the first supercomputing project. The primary objective of this national project is to supply a state-of-the-art supercomputing service to the research community at universities and research institutes, the Norwegian Meteorological Institute and

industry. The service is based on a central installation in Trondheim, at the campus of NTH, operated and supported by SINTEF. The main computer resource is now a CRAY Y-MP/464 with a Silicon Graphics 4D/320S as fileserver. The internal network is FDDI-based. A 34 Mbit/s network is established between all the four universities.

Researchers at the Norwegian universities can apply for free of charge access to the CRAY and the support services. Applications for resources are processed by the research council every 1/2 year. On the average, about 100 projects are honoured. Table 1 shows the CPU usage by application area for the first half year of 1993.

Application Area	CPU hrs	% of total
Civil engineering	1506	15.6
Physics	1633	16.9
Chemistry	2716	28.1
Maths & Informatics	536	5.6
Environment	1548	16.1
Petroleum techn.	814	8.4
Fluid flow	589	6.1

Table 1: Usage by application area first half year 1993

The CRAY is also used by The Norwegian Meteorological Institute, for production of their weather forecasts, as well as for research and development. Some industrial companies also use the service for compute-intensive applications. Currently, most of the industrial usage is in the areas of structural analysis of oil offshore platforms and seismology. Previously, also reservoir simulation was a heavy application, now taken over by workstations. SINTEF's own use of the system is mainly in the areas of seismology and heat and mass transfer in fluids. Table 2 shows use by major user groups in 1992.

User group	CPU hrs	% of total
Universities	5916	61.3
Norwegian Meteorological Institute	930	9.6
Industry	1557	16.1
Research institutes	938	9.7

Table 2: Usage by major user groups first half year of 1993

The parallel technology project was established by SINTEF and NTH in 1992, in co-operation with the Norwegian oil company STATOIL and the research institute IKU Petroleum Research. The goal is to promote use of parallel computers in research and education at NTH/SINTEF as well as in Norwegian industry. The major areas of research on parallel algorithms are weather forecast and climate modelling, computational fluid dynamics, molecular dynamics, seismology, structural engineering, databases and financial analysis. Especially promising results have so far been obtained in the areas of atmospheric models and supersonic fluid flow.

As a part of the parallel technology project NTH and SINTEF entered into a long-term agreement with Intel Supercomputer Systems Division. An Intel Paragon parallel supercomputer was installed at the supercomputing centre in December 1992, as one of the first three in the world, and Intel has named NTH/SINTEF a centre of excellence for parallel computing in Norway.

The visualisation project was started in 1990 in order to meet the demand for efficient data analysis methods from scientists and engineers using supercomputers to solve their problems. Several projects have been carried out in co-operation with industrial partners and scientists from the universities in order to visualise their complex data sets. The project uses equipment from Silicon Graphics, with a Onyx Reality Engine 2 as the current high-end visualisation system. Several visualisation codes have been implemented based on the Silicon Graphics GL-library. Among them is GLview, a high performance 3D general purpose scientific visualisation tool which is extensively used in Norwegian industry and research and is available on standard licence terms from SINTEF. Another major code developed within the visualisation project is NIMBUS, a public weather forecasting presentation system to be used by one of the commercial television channels in Norway.

Please contact:
Leif Reidar Roekkum - SINTEF
Tel: +47 7 592033
E-mail: lreidar@simasintef.no
or: Roar Skaalin - SINTEF
Tel: +47 7 593045
E-mail: roar@simasintef.no

Multimedia Systems

ELO: The Elusive Office

by Ken Robinson

This project, funded by the European Commission under its ESPRIT II programme, has developed a system to support mobile professionals for the mid-1990s and beyond; it has recently had a very favourable final review. The current system is based on a portable Unix workstation, with cellular and conventional communications providing email, fax, telephone, and answerphone services. The organisations involved were empirica (requirements and evaluation), Computer Lern Systeme (help and tutorial system), SEL (communications hardware and system software), the ISI Fraunhofer Institute (requirements

and evaluation), and the ÖVA insurance company (user organisation). RAL provided most of the user interface design and system architecture/implementation roles.

The ELO system adds other features to mobility. These include communications power, a graphical, task-based user interface, and a high level of integration. Communications devices are connected to the workstation via a communications concentrator, a processor which controls the various devices directly. This serves to offload real-time requirements for communications from the Unix host, and also isolate the host from any changes to the communications devices, easing upgrade paths for new device connection.

Mobile workers frequently live an interrupt-driven existence. The ELO system must therefore cope well with this modus operandi. It does this by enabling the ELO user to attain major goals via a series of small, atomic tasks. This approach has a further advantage: the task is completed

via a Task Form, which enables a highly consistent user interface, as very different tasks can be presented in a similar way. These ELO Tasks can also generate other tasks, with information being passed from task to task. Thus data need not be re-keyed, avoiding errors and loss of time. Control of the many tasks which a typical user generates is done via a diary-based task planning and control system known as the TaskOrganiser (see figure for a view of a typical day). Further features showing ELO system integration aspects are: a help and tutorial system, available at any point in the interaction; and a common interface to the various communications devices, independent of the device supplier.

Possible applications of ELO technologies are very wide, and include both vertical and horizontal market segments.

Please contact: Kenneth Robinson - RAL
Tel: +44 235 44 6491
Fax: +44 235 44 5893
Email: kr@uk.ac.rl.inf

The screenshot displays the Task Organiser interface for Monday 13 January 1992 (Week 3). The main calendar area shows a grid with time slots from 9:00 to 20:00. Key appointments include 'Korte Business' at 10:00 and 'Hilton Hotel' at 12:00. The 'To Do' section lists tasks such as 'Rickman Video', 'Jane Visit', 'Conway Bird', 'Lewis Sports', and 'Keil Upgrade'. The 'Phone calls expected' section lists 'McDermott Burglary' and 'Sutcliffe Life'. The 'Pending' section lists 'Mac Rental Holiday', 'Shrimpton ELOdemo', and 'Christie Evaluation'. Navigation buttons on the right include 'Help', 'Next language', 'Previous', 'Next', 'Go to today', 'Views', 'Daily view', 'Weekly view', 'Yearly view', 'Task Browse', 'Update', 'Print page', and 'Close'.

Typical daily view of the Task Organiser, showing appointments, outstanding actions, etc.

Multimedia Archives in Public Networks

by Thomas Rakow

Globally Accessible MultiMedia Archives (GAMMA) is a joint project carried out by GMD – Institute for Integrated Publication and Information Systems and IBM-ENC, Heidelberg, on behalf of DeTeBERKOM, Berlin, a subsidiary of German TELEKOM. The project started in January 1993.

GAMMA is part of a general project intended to provide concepts for the development of multimedia telecommunication services. These services enable numerous users to remotely communicate and cooperate in public networks, where both conventional data like texts, numbers, tables and multimedia data like pictures, graphics, audio and video are transmitted. The aim is to hold these data available after transmission as multimedia documents for further modification. Unlike conventional telecommunication services the nature, quantity and complexity of multimedia documents require new concepts and developments for their management. Archives in multimedia services manage the documents exchanged in these services.

An archive e.g. can manage a multimedia calendar of events in which an information provider may offer details on regional events to information purchasers. The description of events like theater plays, concerts or such, can include texts, maps, samples of vocal performances and video clips. Besides providing this information the archive may also manage the purchasers' requests for tickets entered as text or spoken language.

The access to the archives is accomplished by standardized protocols. Access via the ISO-Standard "Document Filing and Retrieval" (DFR) enables users to work synchronously, while access by electronic mail enables them to work

asynchronously on the basis of a multimedia extension of X.400.

The archives and their management software are realized with the object-oriented database management system VODAK, which has been developed in GMD – Institute for Integrated Publication and Information Systems during the past five years. At present VODAK is being expanded, the department AMOS was formed especially for the support of multimedia applications.

Please contact: Thomas C. Rakow - GMD
Tel: +49 6151 869 944
E-mail: rakow@darmstadt.gmd.de

The MOVIE Project: Mapping Objects on Video by Interactive Editing

by Victoria Burrill

A multimedia application is, by definition, a (probably large) collection of visual information (in the form of text, photographs, graphics, animation and video sequences), and audio information (in the form of voices, noises and music). All these pieces of information are indexable by physical attributes such as word number, frame number or time-stamp, but without some form of keyword categorisation then (in general) only text and computer-generated images are indexable by content. When one considers the vast amount of information that already exists in the form of photographs, videos and audio tapes, then this inability to access that information seems a great opportunity gone to waste.

The MOVIE project – Mapping Objects on Video by interactive editing – is a first attempt at redressing this imbalance.

The intention behind the MOVIE project is to provide a multimedia author with a suite of tools that can be used to delineate significant objects within a video picture and then provide a mapping between that visual area and the contextual object it represents. Once this visual-semantic link has been made, the system can be attached to any hyper-information base the author chooses. In this way, the user can point to an object or character on the video screen and then use that object to access the information base behind it – the user is using the picture to drive the application, rather than adopting the more conventional text-plus-video-window-driven approach.

The delineation is made within a cinematic context, essentially reverse-engineering the screenplay of the film. Interpolation functions are used to determine the visual extents and trajectories of the objects. The data thus generated are stored in a file structure which may subsequently be read and used by any application incorporating that video.

This first stage of the MOVIE project addresses only visual information; audio and other, more exotic sense media such as olfactory and tactile may be addressed in the future.

MOVIE is a satellite project of ShareME (Shared Multimedia Environments) and HyperPicture, and is the result of a 6-month research fellowship between ZGDV and the Informatics Department of RAL.

Please contact: Victoria Burrill - RAL
Tel: +44 235 44 6799
E-mail: vab@inf.rl.ac.uk

IDA: User Interface Design Assistance

by Harald Reiterer

The IDA (User Interface Design Assistance) project has been started in the GMD – Institute for Applied Information Technology. The goal of this project is to develop computer based design aid tools for a human factor based user interface design.

The increasing importance of human factors based user interface design, the lack of human factors knowledge of the designers and the availability of new user interface design tools were starting points for the IDA project. The objective of this project is to develop design aid tools for a human factor based user interface design. These design aid tools are integrated in a UIMS to support the designers of innovative user interfaces (e.g. object-oriented graphical interfaces; 3-D user interfaces; hypermedia user interfaces) during the development process. The aim is to impart human factors knowledge for innovative user interfaces to the designers. Therefore the designers should have direct access to the human factors knowledge from their development tool. The presentation of the design knowledge is based on object-oriented, multimedia and knowledge-based techniques.

An important pre-condition for developing computer based design aid tools is to formalise human factors knowledge to allow a computer based presentation. Especially in the area of graphical user interfaces the GMD Human-Computer Interaction Research Division has a deep understanding of the necessary design principles and design guidelines, the relevant international standards, the available style guides and the construction and use of evaluation methods for usability testing. In the area of 3-D and hypermedia user interfaces the Human-Computer Interaction Research Division is doing much R&D activities. The results of these R&D activities and the available experi-

ences in the field of Human-Computer Interaction are the basis for constructing new design principles and guidelines for innovative user interfaces. The human factors knowledge is summarised in a "Human Factors Style Guide", which is the foundation for the computer supported design aid tools.

The following design aid tools are under development in the IDA project:

- A *composition tool* to construct innovative user interfaces; based on a library of interaction objects and dialogue scripts.
- An *advice tool* to present the human factors knowledge; based on hypermedia documents, a library of interaction objects and an expert system.
- A *quality assurance tool* to evaluate the conformance of the user interface with the human factors knowledge; based on an expert system.

The GMD is starting cooperation with some companies to construct the design aid tools for specific application domains. For this purpose some workshops with members of companies, which are developing software applications and members of companies, which are developing UIMS were arranged. In the realistic context of such application domains the usefulness and usability of the design aid tools will be evaluated. It's clear that the user interface design has to be embedded in the software development life-cycle. Today a lot of methods and tools for the application development are available. There are also some special methods for the user interface development. Till now little work has been done to integrate methods of application development and user interface development. We think that the new object-oriented paradigm offers a good chance for such an integration. This assumption is based on the fact, that modern graphical user interfaces (GUIs) also use the object-oriented paradigm (e.g. X-Toolkit). The aim should be to come to one general method for the whole development process. Therefore we are planning to integrate the use of the design aid tools in an object-oriented development life-cycle.

In the future a software development environment (e.g. CASE Tool) should contain both, a user interface development

environment (UIDE) and an application development environment (ADE). The IDA project concentrates on tools and facilities that are unique to the UIDE. Today commercial UIDE normally includes only tools for specifying user interfaces and in some innovative UIDE you can also find libraries with limited use. In IDA we present some ideas how we can reach a real UIDE including sophisticated libraries of reusable software (composition tool), guidelines and advisers (advice and explanation tool) and tools for evaluating user interfaces (quality assurance tool). We called our tools design aid tools, because their primary focus is not to support the design process (like editors) but to give the designers some advice and evaluation during the design process. We are sure that with the help of such design aid tools the aim "to construct more usable innovative user interfaces" could be reached.

Please contact: Harald Reiterer - GMD
Tel: +49 2241 14-2729
E-mail: reiterer@gmd.de

Parallel Architectures

PARMACS: Portable Message- Passing Programming for Parallel Computers

by Rolf Hempel

The parallel programming interface PARMACS makes application programming independent of the peculiarities of any parallel computer, thus providing portability among the whole range of machines where PARMACS have been implemented. The current work at GMD concentrates on two topics: a new library interface which will be available for both Fortran and C,

and the contribution of PARMACS to the international standardization process of a message passing interface.

The work on PARMACS (PARAllel MACroS) started at Argonne National Laboratory. The idea is that machine-specific constructs for operations like sending a message or creating a process are replaced by macros. These macros are then expanded into the code specific for a particular machine. Thus, the application program becomes machine-independent and can be run on a whole range of parallel computers.

GMD took over the PARMACS development in 1988, at first by developing a Fortran version of the macros. Today the most important difference of PARMACS compared with other interfaces like Express or PVM is the handling of virtual process topologies. The programmer has only to deal with the logical process structure. Automatic mapping tools are provided for an efficient mapping of the processes onto the hardware processors.

PARMACS are a commercial product. GMD cooperates with Pallas GmbH on the distribution and on new implementations. PARMACS are in wide-spread use, mainly in Europe. For example, they were chosen as the common programming platform by the Esprit projects Genesis and PPPE, and the benchmarking project RAPS.

Recently a new PARMACS interface has been defined. The user can now either use macros or function calls, and a C version has been added. The new definitions are fully compatible with the current macro version 5.1. An automatic tool is being implemented which replaces the macros with the corresponding library calls. The future PARMACS development will be based on the library interface.

In 1992 an international initiative was set up for the standardization of a message passing interface. All major vendors of parallel systems are represented as well as the developers of portability packages like Express, PARMACS, and PVM. The project is open to everybody, and it follows the same organizational guidelines as HPFF. Working groups for various topics have been set up. A representative of GMD has been elected to chair the group

on virtual topologies. The goal is to define a standard interface this year.

Please contact: Rolf Hempel - GMD
Tel: +49 2241 14 2575
E-mail: hempel@gmd.de

PEPS: Performance Evaluation of Parallel Systems

by Henk Nieland

Parallel systems have as yet hardly penetrated the market, a main obstacle being the difficulty of assessing their cost effectiveness because no acceptable performance metric is available. In a three year ESPRIT project Performance Evaluation of Parallel Systems, efficient techniques and tools are developed for such an evaluation.

The outcome of this project will provide guidance to

- manufacturers (performance comparison)
- managers (purchasing decisions)
- designers (selecting the best solution)
- application developers (program behaviour analysis)
- system administrators (effective tuning).

PEPS will follow four complementary approaches:

- create a set of internationally accepted benchmarks
- develop portable monitoring tools, using visualisation
- develop reusable system modelling techniques
- identify architecture characteristics on behalf of developers.

The PEPS consortium members are: Thomson Sintra ASM (project leader), NPL, Intecs Sistemi, Simulog, Universi-

ty of Warwick (CS department) and SOSIP.

SOSIP (Standards Organisation of Special Interest Parties) is a group of government agencies and impartial bodies, involved in defining and promoting a methodology for benchmarks. The group is led by NPL (UK) and currently includes AFNOR (France), PTB (Germany), CNR (Italy) and CWI (The Netherlands).

Please contact: Ed Brocklehurst - NPL
Teddington
Tel: +44 81 943 6974
E-mail: ppg@newton.npl.co.uk
or: Jean-Louis Giordana - Thomson-Sintra ASM
Fax: +33 92 963 984
or: Herman te Riele - CWI
Tel: +31 20 592 4106
E-mail: herman@cw.nl

GRACIA: A Software Environment for Graphical Specification, Automatic Configuration and Animation of Parallel Programs

by Ottmar Krämer-Fuhrmann

GRACIA is a graphical programming environment for parallel computer systems being developed at the GMD Department of Supercomputing in co-operation with the Daimler Benz Research group Berlin.

Starting point of the project is the insight that "It appears to be easier to build parallel computers than to use them". Various problems, like the specification of distributed programs, their mapping to

the hardware and the observation of running systems are the main reasons for this fact. The purpose of GRACIA is to support the programming of parallel systems to overcome these difficulties. GRACIA comprises different graphical tools for all phases of the software development cycle.

The design-tool contains a graphical editor which allows the hardware independent specification of process graphs. Each process is an independent, sequential task having access to local memory only. Processes can communicate with each other by message passing constructs. Furthermore, the design-tool allows the user to specify the configuration of the hardware system.

The configuration-tool computes a mapping of the process graph onto the parallel architecture and generates the configuration languages. These languages are different for each parallel computer. GRACIA supports a large variety of target systems, for example seven configuration languages for Transputer based systems or PVM and the PARMACS on all parallel platforms of the laboratory for Parallel Computing at the GMD.

Different tools allow the visualization of program execution. It is possible to observe the hardware, i.e. load characteristics and scheduling information to improve the performance of the application. Another tool allows program animation to display execution states of the software, variable contents and other application dependent data. All tools allow not only on-line animation but also off-line animation, with stepping and backtracking facilities. Furthermore, they support the snapshot concept, which displays actual system states, as well as the visualization of dependencies between events on a time scale.

The development of GRACIA will end in the mid of 1993, but a prototype is still utilized in different GMD research projects and the Daimler Benz project TRAF-FONIC, where a parallel processor is under development which can be embedded into vehicles.

Please contact:
Ottmar Krämer-Fuhrmann - GMD
Tel: +49 2241 14 2202
E-mail: kraemer-fuhrmann@gmd.de

A Dynamically Reconfigurable and Adaptable Parallel Processor

by Olav Landsverk

Major research areas of the group for computer architecture and design at the Norwegian Institute of Technology are dynamically reconfigurable hardware and parallel processing. As a pilot project for the application of these technologies a dynamically reconfigurable parallel processor for neural network applications has been developed. Neural networks lend themselves to parallel processing as they are inherently massively parallel, and to reconfigurability as there are a variety of neural network architectures.

Reconfigurable hardware gives the possibility to combine the advantages of a computer architecture adapted to a specific task with the flexibility of a programmed solution. In addition to loading a program file in a general purpose computer when changing task, a configuration file is first adapting the comput-

er architecture to the new task. A pure hardware solution, e. g. application specific integrated circuit will have the advantages of speed and compactness. On the other hand a pure software solution will have its advantages in flexibility and ability to be adapted to specific applications. The dynamically reconfigurable hardware will to a certain extent be a compromise between these two extremes. It will in many situations be a valuable addition to the tools available in the design of computer systems.

Adaptation of a computer architecture to an application becomes more challenging when capacity or other requirements necessitates the use of parallel processing. The basis for the design of a parallel processing system for a particular application is to adapt the structure of the system to the structure of the application. This area seen in connection with the dynamically reconfigurable hardware offers very interesting possibilities for dynamic adaptability of parallel processors.

Neural networks has been selected in order to link the two above activities to an advanced application. Neural networks lend themselves to parallel processing as they are inherently massively parallel, and to reconfigurability as there are a variety of neural network architectures. It is a pilot project and is used as a target application for the activities in these areas. The implementation of artificial neural networks may be based on




Figure 1: Module

simulation on a conventional computer, often speeded up by an accelerator, by a neural computer or by a dedicated integrated circuit. For systems of some size, new hardware solutions are required in order to achieve realistic computational speed, and research is going on based on various approaches. The approach used in this group is to develop a neural computer suitable for use in a research environment for experimentations with neural networks. This requires high computational speed in order to obtain short turn around times, and reconfigurability to adapt it to the various neural network architectures.

The implementation is based on a fairly general emulator using both parallel architecture and dynamically reconfigurable hardware. Its general form is thus a dynamically reconfigurable parallel architecture. It consists of a set of modules comprising processor, memory and communication (figure 1). The modules are designed to meet the requirements of the neural computer, but are otherwise general. Each module consists of a 32 bits processor with 32 MFLOPS and 16 MB memory. For experimental purposes a small version of the emulator with 16 nodes, giving 500 MFLOPS, is first implemented. The communication subsystem has four independent communication channel controllers and eight unidirectional byte-wide data streams. Each controller handles two data streams, which can be incoming or outgoing. In

addition, limited interchannel communication is possible.

The communication network between the modules is reconfigurable at three levels and can be changed dynamically between and during simulations.

The data-path level handles the communication network topology between the modules. A programmable communication switch makes it possible to connect the data streams in several ways, so that different network architectures can be formed.

The logic level is the control of data streams in the communication module. The controllers use reprogrammable logic, which can be reprogrammed dynamically by loading a configuration file whenever a change in functionality is needed. In this way, communication networks using circuit switching, packet switching or a combination, can be implemented.

The command level is the programming of command-registers normally included in the channel controllers. The number of registers and their function is defined by the current state of the logic level. By programming these registers, different operation patterns of the data streams can be implemented. Operation patterns include receiving, sending, receiving and sending, and bypass.

A simple configuration has all modules connected in one ring. An extension of this is to use several rings to mimic the inherent layer structure of many neural network models. Rings can also be connected to form a hierarchy of intermodule communication. Other possible configurations are to use the communication channels to build a mesh network or a four dimensional n-ary cube, both with one ring in each dimension. The rings are scalable and each module can participate in 4 rings. Figure 2 illustrates one possible configuration with up to 3 rings through a module.

In order to use this general emulator as a neural computer, its architecture will first be adapted to the current neural network architecture and it will then process the neural network in accordance with the description in a specification language. The general emulator adapted to the neural network application has been named RENNS (REconfigurable Neural Network Server). This indicates the use of RENNS as a server for workstations where neural network applications are developed.

Based on the experiences with the first experimental version, expansions and enhancements will be implemented. RENNS will also be adapted to, and applied for various neural network applications. Experience with dynamic reconfigurability will be gained through this work and will establish a basis for further development of this concept.

The general emulator will also be used as a dynamically reconfigurable parallel processor for other applications than neural networks. This will first be done for applications where the computations resembles those found in neural networks. Computations on arrays and vectors, as in some simulations, is one example. This opens the possibility of adapting a parallel processor to various heavy computational tasks.

Please contact
Olav Landsverk - NTH-SINTEF
Tel: + 47 7 59 44 55
E-mail: olavl@idt.unit.no




Figure 2: A possible configuration

Systems and Architecture

High Performance Arithmetic Devices Based on Residue Number Systems

by Enrico Martinelli, Maria Cristina Pinotti

One of the research topics studied in the Systems and Architecture Division at IRI-CNR concerns the design of non-conventional arithmetic units, based on Residue Number Systems. The interest stems from the modular and parallel nature of this arithmetic, particularly suited for high-speed applications and VLSI implementations.

Residue Number Systems (RNSs) have always been considered attractive both theoretically and for practical applications. Their arithmetic is easy and fast, at least as far as addition and multiplication are concerned, due to the inherent absence of carries to be circulated among digits. More precisely, given a set of pairwise prime positive integers $\{m_1, \dots, m_n\}$, an integer X in the range $(0, \prod_{i=1}^n m_i)$ is represented in the associated residue system by means of the n-tuple of "digits" $\{x_1, \dots, x_n\}$ where x_i is the remainder of the integer division of X by m_i . It can be proved that this representation is unique and that the following arithmetic properties hold:

$$|X \cdot Y|_{m_i} = |X|_{m_i} \cdot |Y|_{m_i}$$

$$= |x_i \cdot y_i|_{m_i}, \quad \forall i$$

with $\cdot = +, -$ or $*$

In other words, arithmetic operations are carried out independently, digit by digit, and then every digit is elaborated in parallel. As a consequence, arithmetic units based on RNSs are regular, and modular. They are thus particularly appealing for special applications such as digital signal processing (DSP). However there is a significant drawback: the absence of explicit information on number magnitude means that intermodular operations are necessary when dealing with sign, overflow detection and number comparison. Unfortunately, such operations are sequential in nature and therefore slow. Much care must be given to their implementation to avoid the intrinsic speed of RNS arithmetic being overridden when they are required.

Our activity in this area is focussed on the following aspects:

- the development of efficient algorithms to speed up intermodular operations such as binary-to RNS and RNS-to-binary conversions and base extension;
- the design of devices implementing arithmetic operations in RNS without using table look-up techniques that are expensive and time consuming;
- the definition of new number systems retaining the advantages of positional (i.e. magnitude information) and of residue number systems.

The approach is oriented towards possible VLSI implementations of the proposed devices, exploiting the regularity and the modularity featured by RNS arithmetic. In particular, the logic design of such devices is developed and evaluated according to VLSI asymptotic complexity theory. The complexity figure generally assumed is a function of the product silicon area by response time and, in this respect, some optimal designs have been introduced. Our main attention, however, has been given to speed enhancement.

Please contact:
Enrico Martinelli - IRI-CNR
Tel: +39 50 593446
E-mail: martinelli@iei.pi.cnr.it
or:
Maria Cristina Pinotti - IRI-CNR
Tel: +39 50 593445
E-mail: pinotti@iei.pi.cnr.it

Networks

International Line from Greece to Internet

by Stelios Sartzetakis

As of 22nd of June 1993, FORTH (Crete, Greece) and Ebone (Paris, France) have put in operation the first digital 64Kbps international line from Greece to Internet.

FORTHnet RBS is now connected to Paris EBS via this link. At the same time operation of the two old analog links (to INRIA Sophia-Antipolis, and CNUSC Montpellier) has stopped, while more new international digital circuits are planned to be in operation later. FORTHnet is the first, largest and most advanced computer network in Greece, which serves and links the Technology Parks of FORTH (in Crete, Thessaloniki and Patras) amongst themselves and with the rest of Greece, the rest of Europe, and the rest of the world.

The majority of the links upto today were analog telephone lines. This already started to change with the introduction of the first 128Kbps lines that the Greek PTT delivered to FORTH in May connecting Heraklion to the other three major Greek cities. FORTH-ICS today provides all design, development, and management of the network. It handles all operation and services management, does much of the exploitation of the research project results, as well as building pilot projects and designing the expansion of the network by connecting new sites, providing new services, and broadening the links bandwidth. FORTHnet closely follows the international developments (especially those in Europe), and maintains the best and close relations with all large computer user communities in Greece (both in academia and companies R&D).

Please Contact:
Stelios Sartzetakis - FORTH - ICS
Tel: +30 81 221171
E-mail: stelios@ics.forth.gr

Software Environments

CSCW Shell in Preparation – Natural Cooperation Serves as Model

by Thomas Kreifelts and Wolfgang Prinz

The need for computer support in the management of distributed work is increasing constantly – a fact illustrated, for example, by consortia operating across organisational and national boundaries for the purpose of implementing large-scale projects or by governmental and parliamentary organisational structures which are distributed over various different locations and sites. Through its work on two new projects, the GMD Institute for Applied Information Technology is looking to significantly enhance the basis for developing new systems aimed at supporting cooperation (CSCW systems - Computer Supported Cooperative Work). The first goal is the development of a CSCW shell which will substantially reduce the development work needed for new CSCW systems. The second is the development of new CSCW concepts which are oriented entirely on the model of natural cooperation.

The ESPRIT project "EuroCODE" centres around the creation of an open development environment for CSCW systems which, above all, also provides the necessary basis for mobile forms of cooperative work. Project leader Dr. Thomas Kreifelts outlines the new requirements as follows: "CSCW systems today are being built from scratch as what are essentially self-contained systems. This is a laborious process and will not meet the needs of the user for openness and easy transition between different forms of cooperation support. As a consistently ob-

ject-oriented development environment, the CSCW shell of EuroCODE is intended to remedy this situation. Its performance features include the ease of implementing very different CSCW systems which are capable of interacting, easy adaptation to the peculiarities of the given organisation and the reusability of the shell modules."

The work on the CSCW shell is based on work conducted in the ESPRIT project "EuroCoOp – information technology support for distributed cooperative work" which has seen the further development of a key component in the form of an activity assistant. The activity assistant enables users to define, enhance or modify cooperative activities, supervise work progress, make joint use of documents and services in the activity context and exchange information. The activity assistant distributes activity descriptions, appended documents and information to the participating users in a consistent manner. The primary user interface of the activity assistant is a hierarchically structured "To do" list which contains all activities which a user is participating in and which permits direct access to the relevant information relating to an activity.

The experience which the GMD Institute for Applied Information Technology has gathered in the field of supporting cooperative work is also being invested in a new project known as COMIC (Computer-Based Mechanisms of Interaction in Cooperative Work) which is likewise being funded by the EC within the scope of ESPRIT Basic Research and is being conducted in conjunction with researchers from Denmark, the UK, Italy, the Netherlands, Sweden and Spain. Project leader Wolfgang Prinz outlines its goals: "In order for CSCW systems to enjoy greater acceptance with users than is currently the case and to become a commercial success, the natural conditions of cooperative work must be taken into account much greater than at present". While it is a known fact today that cooperation-supporting systems ought to image a whole array of natural features of human cooperation, the way in which the technical simulation of many obvious features of human cooperation is to be achieved is unresolved.

The new GMD project concentrates on two key conditions which are needed for human cooperation to function properly, namely an awareness of the status of the work task and the experiencing of changes. In natural work groups, every member essentially experiences developments in his sphere indirectly. Similar impressions must be possible in a work group linked via technical tools.

The COMIC project examines these new requirements in four work packages. The first deals with the question of which of the often complex perceptions are relevant in an organisational context and how they can be modelled if required. These spatial, time-related and social conditions of cooperative work are viewed today as "foreign societies" and are examined with the aid of methods from the field of ethnography. The second involves the further development and interdisciplinary application of methods which have already been tried out by members of the project in the Air Traffic Control study. The third deals with the question of which description languages are suitable for the notation of the interaction mechanisms. The fourth is to use virtual reality to create new user interfaces which allow the greatest possible level of "experiencing". In addition, special metaphors (shared objects, common objects) are to provide information on the system and its capabilities in a simple manner.

Please contact: Thomas Kreifelts - GMD
Tel: +49 2241 14 2643
E-mail: thomas.kreifelts@gmd.de
Wolfgang Prinz - GMD
Tel: +49 2241 14 2730
E-mail: wolfgang.prinz@gmd.de

Ultrasound Material Inspection

by Edoardo Bozzi, Salvatore Minutoli and Andrea Marchetti

An integrated system for the non-destructive testing of manufactured goods has been developed at IEI-CNR in collaboration with Alenia GAT of Pomigliano, Naples, as part of the CNR Special Project for Robotics.

The system includes several dedicated work stations for the acquisition of signals using different techniques, e.g. thermography, shearography, microwave tomography and ultrasound.


The Ultrasound Inspection Station uses a 6 axes mechanical arm produced by Asea Brown Boveri, commanded by an S3 control unit. The robot system can be programmed for both on and off-line operations. For off-line operations, the program is written in ARLA, transported to the S3 through a serial port and activated locally. For on-line operations, S3 dialogs with the host computer using an RS-232 line and the ABB Communication Tool. The robot is commanded by the host computer and by programs resident on S3. The samples to be examined are put into the immersion tank and an ultrasound impulse generator synchronized by S3 supplies the echo signal, which is sampled using an A/D converter. Currently, a LeCroy 9450 oscilloscope, connected by a PC through a GPIB IEEE 488 port, is employed.

Programs have been developed to inspect samples of 400 x 1000 mm or less and with surfaces defined with respect to the robot reference system $\Omega(x,y,z)$ by a generating function $f(y,z)$. The programs are activated by defining the parameters specifying the scanning dimensions and step ($>0,1$ mm), and the acquisition type, i.e. the peak or full wave. In the first case, the value of the echo signal peak within a time gate is estimated; for each scanning coordinate i,j the value $L_{i,j}$ is obtained and a C-scan map representing the reflection of the echo signal in corre-

spondence with a constant pre-established depth for the sample under exam is constructed. Using the full wave, the entire echo signal $V_{i,j}(t)$ is acquired for each i,j coordinate and signal $V_{i,j}(z)$, which depends on the spatial coordinate z measured in the signal propagation direction is then obtained; $V_{i,j}(z)$ can be represented by A-scan or B-scan. The scanning rows are parallel to the X axis of the robot reference system. A row scanning subprogram translates the transducer mounted on the mechanical arm, whose direction is controlled by the host computer. If the sample has a flat surface, the probe remains in the same position during scanning; if the surface is defined by a profile generator, commands containing the values of the spatial coordinates x,y,z and the quaternions ($q1, q2, q3, q4$) for each row scanned are sent to S3 to keep the probe orthogonal to and at a constant distance from the surface.

The data acquired are transferred to a processing system equipped with AVS software. A procedure for the 3D representation of volumetric data derived from full wave inspection is now being developed; tomographic images on sections of the data can be derived in order to identify defects and delineate the zones of interest from which useful information for the classification and characterization of the defects can be extracted.

Please contact: Edoardo Bozzi - IEI-CNR
Tel: +39 50 593496
E-mail: salva@rix6000.iei.pi.cnr.it



An ultrasound inspection station for the non-destructive testing of manufactured goods consisting of a robot arm, an immersion tank, a control unit, a PC, an ultrasound impulse generator and a system for the acquisition of the echo signal acquired by the transducer. (Photo: CNR)

Official Visit to INESC by Portuguese Minister of Industry

by João Bilhim

The premises of INESC's new technological infrastructures, which were co-financed under PEDIP (Specific Programme for the development of Portuguese Industry), were officially opened by the Minister of Industry in June last.

This inauguration represents INESC's intention to provide a new type of support and a new attitude to the requirements for industrial modernisation in its field of activity - Telecommunications and information technologies. INESC will pro-

vide development, demonstration and advanced technology transfer to Portuguese industry, as part of a coherent strategy whose objectives have been classified as priority under PEDIP (a Community Programme for the development of Portuguese industry).

The objectives of these new infrastructures, which represent a total investment of around PTE 6 billion and which have been 80% financed under PEDIP, are to endow Portuguese industry with the technological capacity and consequently, strategic and competitive advantages in the international marketplace, in addition to providing incentives for the creation of technology based small and medium sized companies.

The infrastructures have been set up in their own two buildings in Lisbon and Oporto and, in 1994, will house around 200 specialists working in the field of technological development. They will provide support services to industry through 14 Technology Transfer and Demonstration Centres organised into two divisions: Technology and Systems.


A Business incubator called AITEC has also been set up. Its objective is to create conditions for encouraging the incorporation of an economic network of technology based small and medium size companies through the provision of facilities and support for the analysis, development and identification of fresh ideas with business potential in the Information Technology field in an attempt to increase the rate of success.

INESC will be renewing and expanding its pioneering activity in the field of company incubation which was launched with the AITEC – Tecnologias de Informacao, S.A. incubator through the Company Incubation Centre which will, for the first time, enjoy the support of the Portuguese Government. This company will now take on the role of a corporate investor. It is, at the present time, active in 22 technology based small and medium size companies.

Please contact: João Bilhim - INESC
Tel: +351 1 572895
E-mail: jafb@eniatic.inesc.pt

Parallel Codes for Circuit Analysis and Control Engineering

by Henk Nieland




In many control engineering and circuit analysis problems real-time response is required. This implies the necessity of high-speed processing, which can be achieved by exploiting parallelism in new computer architectures. In a four year project CWI will design and implement new algorithms for solving the underlying systems of differential equations.

Mathematical modelling of many problems in circuit analysis and in the dynamics of control systems involves the numerical solution of initial-value problems for large systems of nonlinear ordinary differential (-algebraic) equations. Examples of applications where these problems arise include the modelling of VLSI and MOS (metal-on-oxide) digital devices, control of navigation systems, of thermal processes, of manipulators (robotics), etc. Existing codes for solving such problems are designed for one-processor computers and have little scope for parallelization. Therefore, instead of trying to parallelize existing initial-value-

Left: 6 degrees-of-freedom anthropomorphic robot manipulator used at the University of Amsterdam. (Photo: CWI)

Right: circuit diagram of the Philips TDA8760 monolithic bipolar 10-bit AD converter, containing over 2600 transistors, 400 resistors and 6 external capacitors (courtesy Philips Research Laboratories).



In both cases high-speed (parallel) processing is required to solve the large systems of equations involved.

problem solvers, completely new solvers especially for use on parallel computers will be designed. In such an approach, much larger speed-up factors can be obtained.

In general, we encounter here so-called stiff and non-stiff problems. In non-stiff problems the time scales associated with the various physical components do not differ greatly. This situation is frequently encountered in, e.g., robotics and celestial mechanics. In stiff problems the time scales differ by several orders of magnitude. In chemical reactions and in circuit analysis, for example, the different reaction speeds and sizes of the electrical components, respectively, lead to stiff equations. These two different families of problems require essentially different solution techniques. In general, for non-stiff problems explicit methods are applied, whereas for problems implicit methods are preferable. Explicit methods are cheap, but may easily become unstable. Implicit methods are much harder to handle, but can be given much better stability properties.

The project's main objectives are:

- design and implementation of a variable step, variable order algorithm based on parallel, explicit Runge-Kutta methods of high order for solving non-stiff differential equations
- design and comparison of variable step, variable order algorithms based on either parallel block methods or on diagonally-implicit iteration of implicit R-K methods for solving stiff differential equations
- building in facilities to allow algebraic side conditions in order to solve differential-algebraic systems from circuit analysis and control engineering.

Within a problem there can be a transition from non-stiff to stiff and vice-versa. For example, in orbit equations describing the motion of a number of planets, the stiffness is inversely proportional to the (smallest) mutual distance. Therefore the resulting code will be equipped with an automatic stiffness detector and a facility to change from the step-expensive stiff solver to the step-inexpensive non-stiff solver.

Participants in this project, financed by the Dutch Technology Foundation (STW), are: Philips Research Laboratory Eindhoven (PRLE), University of Twente (UT, department of Modelling and Simulation), University of Amsterdam (department of Computer Systems) and CWI. The developed codes will be implemented in the software used at PRLE (circuit analysis) and UT (control engineering).

Please contact:
Piet van der Houwen - CWI
Tel: +31 20 592 4083
E-mail: senna@cw.nl

GMD TechnoPark Open for Tomorrow's Technologies

by Karlheinz Schunk

In the course of increasing complexity of research and development activities in the field of information and communication technology national and international cooperation more and more gains in importance. To create a suitable environment in order to strengthen this cooperation between GMD, universities and other science institutions and industry GMD is going to establish a TechnoPark on the campus of its headquarters at Schloß Birlinghoven near Bonn.

Due to its various fields of research and development and on the basis of an excellent technical infrastructure GMD invites other national and international research institutions as well as technology oriented manufacturer and application industry to join the new TechnoPark. For this reason modern new buildings especially prepared for all kinds of research and development in a stimulating and challenging scientific environment and next door to the groups of researcher in

our four institutes on the campus will be available. These four institutes are:

- the Institute for Foundations of Information Technology with its fields of research: Parallel numerics and applications, discrete algorithms and their application in science and technology, scientific visualization and a laboratory for parallel computing,
- the Institute for Applied Information Technology with its fields of research: Artificial intelligence, cooperation systems and human computer interaction with a laboratory for man machine communication,
- the Institute for System Design Technology with the fields Specification of complex systems, responsive systems, framework integration technology, automated design of digital circuits, design of innovative applications and with a laboratory for VLSI-Design
- the Institute for Application-Oriented Software- and Systems Technology with its fields Software technology for applications, application of new technologies, network centre for science and its technical services for all kinds of computers.

Schloß Birlinghoven near Bonn is a central location in the heart of Europe and a unique site for a TechnoPark. The decision of the German Parliament to relocate the parliament and the government from Bonn to Berlin implies that compensatory measures are taken to maintain the economic and cultural strength of the Bonn region by pushing an extensive concept of a Scientific Region Bonn. In this context the establishment of a TechnoPark GMD in the high tech field of information and communication technology will be an essential contribution for this region to attract attention both in scientific and economical sense.

Please contact: Karlheinz Schunk - GMD
Tel.: + 49 22 41 14 32 15
E-mail: schunk@gmd.de

Technology and Cultural Heritage: An Image Spectroscopy System for Non-Destructive Testing of Paintings

by Franco Lotti
and Andrea Casini

We describe an experimental technique which can be used to investigate the composition and state of conservation of the first layers of painted surfaces (paintings on plates, canvas and frescoes). Sequences of multispectral images are acquired, and the spectral information can be reconstructed at any point in the visible and near infrared regions.

Reflectography is one of the most powerful non-destructive techniques used to investigate works of art and to evidence the presence of drawings below the surface ("underdrawings") and "pentimenti". Nowadays, the reflectance spectrum analysis of paintings is both possible and easy, employing spectrophotometers equipped with fiber optics probes, transparent to visible and near infrared light. They can give accurate spectral information (typically from 400 to 2200 nm), even though only over very small areas (few square millimeters).

We propose a technique for visible and near infrared imaging spectroscopy (VNIR-IS), which furthers the "local" methods by offering multispectral information on the whole image. Knowledge of the spectral profiles can help to characterize the pigment composition, with the help of reference spectra, and can also give information about layers below the painted surface by revealing underdrawings, retouches and pentimenti. The set of multispectral images can be stored for later comparison, in order to check variations of the state of conservation over time of the work of art.

This technique consists in the acquisition of multiband image sequences using a visible-IR camera (sensitive from the visible wavelengths up to 2200 nm) equipped with a set of narrow-band filters, a frame grabber and a couple of light sources, with variable intensity tungsten-




Figure 1:
Detail of the image spectroscopy equipment: visible and near infrared camera, motorized lens, 8-piece interference filter wheel.
(Photos: CNR)




Figure 2:
Representation of a narrow band filtered image sequence for image spectroscopy. Each frame includes black and white reflectance references and chess board pattern for automatic focussing and geometrical alignment.

halogen lamps, placed symmetrically, to illuminate the painting in a uniform way (see Figure 1). All the equipment is interfaced and controlled by a computer. The procedure includes automatic control and synchronization of all the measurement steps. The acquired narrow band filtered images constitute a three-dimensional array of multispectral data, in which the third dimension represents the central wavelength of each band (see Figure 2).

In order to make this operation physically meaningful, that is to obtain reliable measurements for subsequent repetitions and comparisons, a series of experimental and computational problems must be considered. Time averaging filtering is used to improve the low signal-to-noise ratio due to the narrow band of filters; other image processing techniques are then applied to the sequences acquired, to account for the environmental conditions that can affect the measure-

ments. They include: calibration of sources and sensors, correction of distortions introduced by changes of focus at different wavelengths and intensity calibration for non linearity and spatial inhomogeneities of the sensor.

Once the images have been made consistent by these corrections and calibrations, the reconstruction of spectral profiles on any point of the image is possible and easy, by reading the array along the axis of wavelength.

Preliminary results, which agree with local measurements obtained with fiber optics spectrophotometers, have been obtained by a prototype VNIR-IS system, applied to Florentine Renaissance paintings.

Please contact: Franco Lotti - IROE-CNR
Tel: +39 55 42351
E-mail: lotti@iroe.fi.cnr.it
or **Andrea Casini - IROE-CNR**
Tel: +39 55 42351
E-mail: casini@iroe.fi.cnr.it

Starting shot for Real World Computing Partnership

by Thomas Hagemann

The official starting shot for the Japanese Real World Computing Partnership (RWCP) was sounded on 8 March 1993 in Tokyo. GMD is the research partnership's first foreign member. The Chairman of the Board of GMD, Prof. Dennis Tschritzis, was appointed a member of the RWCP's Board of Directors on the same day. The Chairman of GMD's Supervisory Board, Dr. Werner Gries, also welcomed the participants to the festivities via a direct link between GMD and Tokyo.

More than 350 persons attended the opening ceremony in Tokyo's Prince Hotel. They included Minister Yoshio Mori from Japan's Ministry of International Trade and Industry (MITI), parliamentary deputy Tadashi Kuranari and eight further parliamentarians, representatives of universities, research institutes, participating industrial enterprises, embassies and the press.

The RWC programme extends over ten years. During this time, more than

¥60,000 million (almost DM 1,000 million at the current rate of exchange) are to be spent on research into massively parallel and optical systems.

GMD is involved in two projects within the scope of the RWC research work. One group led by Prof. Wolfgang K. Giloi is working in the field of "massively parallel and neuronal systems" on the development, implementation and evaluation of a programming model for massively parallel systems. The future user of such a system will no longer have to involve himself with the special architecture of the parallel computer, but instead will be able to concentrate fully on his application problem. Dr. Heinz Mühlenbein is studying the theme "Statistical interference as a theoretical base for genetic algorithms" in the field of "Theory and novel functions". Genetic algorithms constitute efficient searching methods for problems involving numerous variables.

On 24 May 1993, Stichting Neurale Netwerken (SNN) from the Netherlands, the Swedish Institute of Computer Science (SICS) from Stockholm and the Institute of Systems Science (ISS) of the National University of Singapore were admitted as new members to the RWCP.

The festivities on 8 March opened with a speech by the President of the NEC Corporation, Tadahiro Sekimoto, who is also President of the RWCP. Tadahiro Sekimoto commented that the Fifth Generation Computer Project has dealt


with logic which is probably concentrated in the left hemisphere of the human brain. The task now, he added, was to invent computers which function like the right hemisphere. This was a task in which Japan was not alone and in which international cooperation would play an important role in reaching a solution. In this context, he explained, the involvement of GMD in the RWC project was to be greatly welcomed.

MITI Minister Yoshio Mori summarised his impression of the Real World project in a single sentence: "The task ahead is to construct a new brain". Future computers, he explained, must be able to express sensations, perform translations, process visual information and thus perform functions which up until now have been the domain of man.

In the direct link which was then established to GMD in Sankt Augustin, Dr. Werner Gries, chief of section at the Federal Ministry for Research welcomed the participants to the opening ceremony. Dr. Gries stated that the particularly innovative areas which were to be addressed in the RWC programme called for new forms of international and interdisciplinary cooperation. The participation of GMD had opened up new territory in the research sector. He also hoped that industry in both countries could be involved in an appropriate form. The Federal Ministry for Research had demonstrated great commitment to this RWC partnership and would continue to do so in the future.

Yuji Tanahashi, Administrative Vice Minister of MITI, stated that Japan's goal with the RWC programme was to contribute towards innovation into the 21st century and to the solution of topical problems such as environmental pollution. He expressed the hope that the large-scale project, which was heavily oriented to international cooperation, could create a model for new forms of cooperation.

Please contact: Eckart Bierd}mpel - GMD
Tel.: +49 2241 14 2293
E-mail: bierduempel@gmd.de
or: Thomas Hagemann - GMD
Tel.: +81 3 3586 7104
E-mail: hagemann@gmd.co.jp



The chairman of GMD's Supervisory Board, Dr. Werner Gries, welcomes the participants to the festivities via a direct link between GMD, St. Augustin and Tokyo. (Photo: Münch, GMD)

COMETT Programme: Postgraduate Course on Geographic Information Systems

Pisa, Italy, 1994

The first COMETT postgraduate course on Geographic Information Systems (GIS) will be organized in 1994 at San Miniato, Pisa, by the Technical University of Vienna with the support of CNUCE-CNR.

The course will provide participants with a common background on GIS architectures, their functionalities and their integration into various administrative processes. Students will be expected to have a degree in any GIS-related subject area including Geography, Surveying, Geology, Economics or Computer Science. Due to the diversity of backgrounds, basic topics will also be taught such as elements of statistics, computer science and geography.

The course will consist of a sequence of three intensive units of two weeks, each taught in the classroom, together with a practical project unit, completed by each student at home. This set of four units will be presented over a period of one year. Students should be able to continue their regular work duties with minimal interruption while at the same time receiving over 200 hours of classroom and laboratory instruction plus more than 70 hours of practical experience. Upon successful fulfillment of all the course requirements, the students will receive an internationally recognized diploma.

Schedule:

The schedule for the Course is:

- Part 1: Spatial Information for GIS, 7-13 February 1994
- Part 2: Information Systems for GIS, 5-15 April 1994

- Part 3: Practical Project (at home), May - August 1994
- Part 4: Using GIS in Organization, 12-23 September 1994

Deadline:

15 January 1994. Applications

Please contact: Rebecca Winn -
Technical University of Vienna
Fax: +43 1 504 3535
E-mail: winn@geoinfo.tuwien.ac.at
or for Italian candidates:
Roberto Scopigno - CNUCE-CNR
Tel: +39 50 593304
Fax: +39 50 904052
E-mail: scop@vm.cnuce.pi.cnr.it

- Future Applications and New Services
- Multimedia Servers
- Public Directory Services
- Migration and Interoperability
- Security and privacy
- Legal and Economic Issues.

Please contact: Radu Popescu-Zeletin
Tel: +49 30 25499 206
E-mail: zeletin@fokus.gmd.d400.de

First IEEE International Symposium on Global Data Networking

Cairo, Egypt,
13-15 December 1993

The need for open global connectivity is removing the distinction between private and public networks in the wide area. Technologies such as Frame Relay, SMDS and B-ISDN were developed to support this merger. The Global Data Network provides the suppliers with the opportunity to offer various Value Added Services. The first IEEE Conference on Global Data Networking will focus on new ideas for the necessary infrastructure for such networks and services.

Topics:

Papers in the following areas will especially be discussed:

- Public Network Architecture
- Network Access technologies
- ATM Cell-Based Switching
- Mobile Data Communications
- Addressing and Routing
- Transport Services and Protocols
- Trials and Standards
- Network Operation and Management

16th International Conference on Software Engineering

Sorrento, Italy, 16-21 May 1994

The objective of ICSE is to foster the development of the software engineering field by:

- providing a forum for the introduction and discussion of new software engineering research results;
- providing the practicing engineer an evaluation of evolving research;
- providing the research community exposure to the problems of practical applications of software engineering;
- encouraging the exchange of advanced software technology and experience within the international community.

To address these goals ICSE-16 will integrate a variety of events: presentations of technical papers, panels, tutorials, research workshops and tools demonstrations.

Three pre-Conference workshops will explore research issues at the intersection between Software Engineering and other disciplines. The subject areas and contact points for these workshops are:

- Software Engineering and Databases (Roger King: roger@cs.colorado.edu),

- Software Engineering and Artificial Intelligence
(Steve Fickas: fickas@cs.uoregon.edu),
- Software Engineering and Human-Computer Interaction
(Dick Taylor: taylor@ics.uci.edu).

Each workshop will develop a research agenda, which will be presented at a special session of the conference. There will also be a post-Conference workshop on Software Engineering Education (Anthony Finkelstein: acwf@doc.ic.ac.uk).

Deadlines:

6 September 1993: Submission of papers, panels and tutorials

15 December 1993: Acceptance notification

31 January 1994: Final papers

Please contact:

Bruno Fadini - University of Naples

Tel: +39-81- 768 3193

Fax: +39-81- 768 3186

E-mail: fadini@vm.cised.unina.it

CALL FOR PAPERS

Fourth Eurographics Workshop on Object-Oriented Graphics

Sintra, Portugal, 9-11 May 1994

The general aim of this series of workshops is to investigate the applicability of object-oriented methods to computer graphics – computer graphics in a very broad sense, including human computer interaction, image synthesis, dynamic graphics, and computer aided design.

At the past workshop major issues have been identified as critical for further investigation to form a better consensus. What are the solutions to graphics problems provided by object-oriented design and programming? How are entire ar-

chitectures, e.g., OMG, suited to support graphics? How can a graphics kernel be defined and developed to support different object-oriented language paradigms? How can – besides of multiple inheritance and delegation – other mechanisms such as composition, protocols, polymorphism, multiple dispatching be used to design flexible and re-usable systems?

Topics:

Besides the above mentioned topics, contributions to the following areas are encouraged:

- support for extensibility in graphics systems,
- object-oriented architectures to integrate new modelling paradigms into existing systems,
- introduction of new types of output primitives by providing extensible renderers,
- support for concurrency and distribution in graphics applications,
- specific graphics object types (interactors, plugging objects, synchronisation objects),
- specific object-oriented issues for 3-D graphics, animation, and multimedia,
- generic object-oriented frameworks for co-operative graphics applications,
- object-oriented graphics standards.

Workshop Format:

The workshop will be limited to 60 participants. To encourage discussions, different types of presentations are foreseen: full paper presentations, presentations of position papers, presentation of results from ad-hoc working groups, demonstrations of systems.

Deadlines:

15 January 1994: Submission of papers
15 March 1994: Notifications of acceptance

1 September 1994: Revised papers.

Please contact: Peter Wisskirchen - GMD

Tel: 49 2241 14 -0

E-mail: wisskirchen@gmd.de

CALL FOR PAPERS

Eurographics Workshop on Design, Specification, Verification of Interactive Systems

Carrara, Italy, 8-10 June 1994

Making systems easier to use implies an ever growing complexity in managing communication between users and applications; an increasing part of the application-code is being devoted to the user interface. In order to manage this complexity, it is very important to have tools, notations, and methodologies which support the designer's work during the refinement process from specification to implementation.

The purpose of the workshop, organized by Eurographics, is to review the state of the art in this area. The different existing approaches will be compared in order to identify the principal requirements and most suitable notations, and to indicate significant results.

Suggested Topics:

- Foundations and Reference Models for Interactive Systems.
- Verification of user interfaces, application of theorem-provers.
- Methodologies for abstract design, comparative studies of methods/description techniques.
- Specification of human-computer interaction, multi-modal user interfaces and virtual realities.
- Design of Graphics Systems and Window Systems.
- Formal description of users' related properties.
- Model-based user interface software tools.

Participation:

Participation will be limited to encourage discussion. Selection will take place on the basis of full papers (up to 15 pages) reviewed by the Programme Committee. Papers and conclusions will appear in the Workshop Proceedings. Invitations to submit revised versions for a post-workshop book will depend on the quality of the contributions. Participation without the submission of a full paper may be possible in a few cases, on the presentation of position paper.

Deadlines:

31 December 1993: Expression of interest
 20 February 1994: Full papers/position papers
 10 April 1994: Notification of participation/acceptance

Please contact:

Fabio Paternó - CNUCE-CNR
 Tel: +39 50 593289
 Fax: +39 50 589354.
 Email: paterno@vm.cnuce.cnr.it

CALL FOR PAPERS

2nd International Workshop on Massive Parallelism: Hardware, Software and Applications

Naples, Italy, 3 - 7 October 1994

The 2nd International Workshop on Massive Parallelism: Hardware, Software, and Application is sponsored by the {em Progetto Finalizzato Calcolo Parallelo e Sistemi Informativi}, a special project promoted by CNR to advance knowledge in all areas of parallel processing and related technologies. The Workshop will be organized by the Istituto di Cibernetica, Naples, in cooperation with the Department of Computer Architecture,

Barcelona, Spain, the Department of Computer Science, Patras, Greece, and the Center for Supercomputing Research & Development, Urbana Champaign, U.S.A.

In addition to technical sessions, MP '94 will offer tutorials, a parallel systems fair, and commercial exhibits.

Topics:

Topics include but are not limited to:

- Parallel Algorithms
- Parallel Architectures
- Parallel Languages
- Programming Environments
- Parallelizing Compilers
- Performance Modeling/Evaluation
- Signal & Image Processing Systems
- Other Application Areas

Tutorials:

Proposals are solicited for organizing full or half-day tutorials to be held on the first day of the Symposium. and should be presented by 15 January 1994 to the Tutorials Chair:

Prof. C. Polychronopoulos, CSRD, University of Illinois
 Tel: +1 217 244 4144
 E-mail: cdp@csrd.uiuc.edu

Parallel Systems Fair:

This all-day event will include presentations by researchers who have parallel machines under development, as well as by representatives of companies with products of interest to the Massively Parallel Processing community. A presentation summary should be submitted by 15 January 1994 to the Parallel Systems Fair Chair: Prof. A. Massarotti Istituto di Cibernetica, Naples
 Tel: +39 81 853 4126

E-mail: massarotti@cib.na.cnr.it

Deadlines:

15 January 1994: Full papers
 15 January 1994: Proposals for Parallel Systems Fair
 15 January 1994: Proposals for Tutorials
 30 April 1994: Acceptance letter sent
 31 May 1994: Camera ready copies

Please contact: Mario Mango Furnari - Istituto di Cibernetica, Naples
 Tel: +39-81-853-4229
 Fax: +39-81-526-7654
 E-mail: furnari@cib.na.cnr.it

CALL FOR PAPERS

Structure in Complexity Theory Ninth Annual IEEE Conference

**Amsterdam, The Netherlands,
 28 June - 1 July 1994**

The conference seeks original research papers or technical expositions in all areas of complexity theory.

Topics:

- Structure of complexity classes
- Resource-bounded reducibilities
- Circuit complexity
- Interactive proof systems
- Structural aspects of distributed and parallel computing
- Properties of complete sets
- Theory of relativizations
- Descriptive complexity
- Kolmogorov complexity
- Cryptographic complexity

Contributions:

To submit a paper for consideration, send 10 copies of an extended abstract to the program committee chair:

Uwe Schoening, Abt. Theoretische Informatik, Univ. Ulm, Oberer Eselsberg, 7900 Ulm, Germany;

E-mail: schoenin@informatik.uni-ulm.de

Other conference activities include open rump sessions, research abstracts, and program committee talks. Conference proceedings will be published by the IEEE Computer Society.

Deadline:

1 December 1993: Papers received

Please contact: James Royer - School of Computer & Inf. Science, Syracuse University, NY 13244, USA
 E-mail: royer@top.cis.syr.edu

CALL FOR PAPERS

ECAI '94: 11th European Conference on Artificial Intelligence

Amsterdam, The Netherlands,
8-12 August 1994

The European Conference on Artificial Intelligence (ECAI) is the European forum for scientific exchange and presentation of AI research. The aim of the conference is to cover all aspects of AI research and to bring together basic research and applied research. The Technical Programme will include paper presentations, invited talks, panels, workshops, and tutorials. The conference is designed to cover all subfields of AI, including non-symbolic methods. An industrial and academic exhibition will be organized from August 9-11, 1994. The conference will take place at the Amsterdam RAI, International Exhibition and Congress Centre.

Topics:

Original research papers are sought that represents a significant contribution to any aspect of AI, including:

- the principles underlying cognition, perception, and action in humans and machines;
- the design, application, and evaluation of AI algorithms and intelligent systems;
- and the analysis of tasks and domains in which intelligent systems perform.

Theoretical and experimental results are equally welcome. Of special interest this year are papers which address applied AI. Two kinds of papers are sought.

- AI applications that address significant real-world problems and which are used outside the AI community itself;
- novel AI techniques and principles that may enable more ambitious real-world applications.

Deadlines:

8 January 1994: Submission of papers
12 March 1994: Acceptance
19 April 1994: Camera-ready papers

Please contact: Tony Cohn - University of Leeds, UK
Tel.: +44 532 33.54.82
E-mail: ecai94@scs.leeds.ac.uk

CALL FOR PAPERS

5th Eurographics Workshop on Rendering

Darmstadt, Germany,
13-15 June 1994

Following four successful workshops (Rennes-1990, Barcelona-1991, Bristol-1992, Paris-1993) we announce the fifth workshop on rendering techniques. In the recent years the workshop has been well established as a major international forum in exchanging experience and knowledge between people from universities, research and industry interested in the different aspects of rendering techniques.

Topics:

- Radiosity
- Ray tracing
- Illumination models
- Colour, texture
- Sampling, filtering, anti-aliasing
- Parallel solutions for rendering

Two special themes of this workshop are:

- Illumination & rendering of participating media (volume objects, clouds, ...)
- Rendering of architectural & CAD models (illumination simulation, real-time rendering, walkthroughs, handling of large datasets, ...)

Deadlines:

5 April 1994: Submission of papers
5 May 1994: Notification of acceptance
30 May 1994: Full-paper deadline

Please contact: Stefan Haas - Fraunhofer-IGD, Germany
Tel: +49 6151 155 133
E-mail: haas@igd.fhg.de

CALL FOR PAPERS

5th Eurographics Workshop on Visualization in Scientific Computing

Rostock, Germany,
30 May - 1 June 1994

This is the fifth in an annual series of workshops organized by the Eurographics Working Group on Visualization in Scientific Computing. Contributions are invited with respect to all aspects of scientific visualization. The emphasis should lie on challenging problems and innovative solutions, involving methods of for instance design of algorithms and data structures, computational geometry, applied mathematics, hardware and software systems.

Topics:

- visualization and modeling of large data set
- interface techniques: control and steering, sonification, virtual reality
- context visualization, hybrid visualization, comparison of simulation results and measured data
- volume visualization, non regular grids, vector and tensor fields
- parallel and distributed architectures and algorithms for visualization
- human factors and perception
- applications, as far as new approaches of visualization are involved.

Deadlines:

31 January 1994: Submission of contributions
15 March 1994: Notification of acceptance
1 May 1994: Full papers


Please contact: B. Urban - FhG -IGD, Germany
Tel: +49 381 4545 110
E-mail: urban@egd.igd.fhg.de

GMD – The GMD Institute for Foundations of Information Technology under Ulrich Trottenberg was very successful in the international Mannheim SuParCup'93 competition. The second and the third prizes were awarded the institute. The cup is dedicated to an outstanding contribution in the field of parallel computing. An international committee of 10 experts selected 4 out of 24 contributions from numerous countries. The second prize was awarded to Ute Gärtel, Wolfgang Joppich and Anton Schüller for their contribution "Portable Parallelization of the ECMWF's Weather Prediction Code: The 2D Case" (see also article "More reliable weather forecasts by parallel computing"). A third prize was won by Ute Gärtel, Horst Schwichtenberg, Ulrich Trottenberg and Gerd Winter together with Michael Griebel, Walter Huber, Christian Zenger and Thomas Störckuhl from the Institute for Computer Science of Technical University Munich for their common contribution "The parallel ASMG algorithm for 3D Poisson-like equations" (see also ERCIM News No.13, pp.18-19).

CWI – CWI announces the birth of a Dynamical Systems Laboratory (DSL). The goal of DSL is to create a research environment for workers in dynamical systems theory and applications using parallel computers and graphical workstations. With support from the Netherlands Organization for Scientific Research (NWO), computers and software will be installed at CWI as well as at other sites in The Netherlands, and the possibility will be created for researchers, including those from abroad, to spend some time at DSL to use the environment and the available expertise, or to use CWI's CRAY S510A parallel computer over the network. Initially DSL is staffed with three part-timers and expects to have a similar number of visitors annually.

GMD – IBM Germany and GMD have agreed to intense cooperation on parallel computing. A contract was signed on 7th June 1993 in GMD, Sankt Augustin. The cooperation will concentrate on the main problem in this area, namely software. All software will be developed jointly, particularly at application level. Dr. Joachim Redmer, head of IBM training and research,

emphasised the vital importance of exchanging knowledge and experience between research and industry, illustrating this with examples from aerospace and chemistry. In GMD the cooperation is the responsibility of the Institute for Foundations of Information Technology, led by Prof. Ulrich Trottenberg. His institute had already produced first results: the ASMG program for solving Poisson-like differential equations had been implemented on a cluster of 8 IBM workstations. A problem with over 109 unknowns was solved in about 30 seconds!



Prof. Stefan Jähnichen
(Photo:
Münch, GMD)

GMD – Prof. Stefan Jähnichen, lecturer in software engineering at the Technical University of Berlin and one of the three heads of the GMD Institute for Research Architecture and Software Engineering, has been appointed Chairman of the European Activity Committee 1993 by the Board of the IEEE Computer Society (Institute of Electrical and Electronics Engineers). The IEEE is the world's largest professional organisation for electrical engineers and computer scientists.

INRIA – The Science et Defense Prize, awarded by the French Ministry of Defense, was won by Marie-Odile Bristeau (INRIA) and Jacques Periaux (Dassault-Aviation) for their work on variational and control methods for numerical simulation in fluid mechanics and electromagnetism.

RAL/FORTH – A meeting about common interests in research into intelligent user interfaces, and about collaboration possibilities in this field between RAL and FORTH took place at RAL on 30th June 1993. Constantine Stephanidis, group leader at ICS-FORTH, had initiated the meeting. David Duce welcomed the visitors and gave them an introduction to RAL and the Informatics Department. During their visit, they talked to several group leaders at RAL as well as project staff. Some of the dis-

cussions were enhanced by demos of current work going on at RAL. Both sides benefitted from the exchange of views, and gained a better understanding of how the partner organisations work. We would encourage further contacts of this sort, to complement the larger-scale ERCIM meetings which are currently taking place twice a year. Such smaller-scale meetings could lead to valuable input to decisions concerning the formation of new ERCIM research groups, or concerning the definition of ERCIM research priorities.

INRIA – IBM-France will distribute two products, "Grif Editeur SGML" and "Grif Builder", marketed by a start-up of INRIA called GRIF S.A., on its RS/6000 stations. These products (editors of structured texts) were initially developed by INRIA in collaboration with the CNRS.




Christoph Rudolph
(Photo:
Münch, GMD)

GMD – Christoph Rudolph, head of the GMD Staff Group for International Coordination retired on 31 March 1993. Joining GMD in March 1969, he played a leading role in shaping GMD's international affairs. As the Board's representative he was GMD's intermediary to numerous foreign government bodies and GMD partner organisations. In 1972 he was GMD's signatory to the cooperation agreement in the fields of mathematics and informatics which was concluded with Brazil and promoted by the Federal Ministry for Research and had since been in charge of GMD's International Office. He was also responsible for the agreement between GMD and the Academia Sinica. From the start of the 1980s, GMD's participation in the information technology promotion programmes of the European Communities formed a further key area of activity for Christoph Rudolph. He was also one of the fore-runners and co-organisers of the European Research Consortium for Informatics and Mathematics (ERCIM) which was founded in 1989.

**European Research
Consortium
for Informatics
and Mathematics**

ERCIM



The European Research Consortium for Informatics and Mathematics (ERCIM) is an organisation dedicated to the advancement of European research and development, in the areas of information technology and applied mathematics. Through the definition of common scientific goals and strategies, its national member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry. To further these objectives, ERCIM organises joint technical Workshops and Advanced Courses, sponsors a Fellowship Programme for talented young researchers, undertakes joint strategic projects, and publishes workshop, research and strategic reports, as well as a newsletter.

ERCIM News is the in-house 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.

Central Office:
ERCIM
Domaine de Voluceau
Rocquencourt
B.P. 105
F-78153 Le Chesnay Cedex
FRANCE

Lay-out and Design:
Helena du Toit

Central Editor:
Helena du Toit

Local Editors:
Wernher Behrendt (RAL)
Lars Bergman (SISU)
João Bilhim (INESC)
Truls Gjestland (SINTEF)
Siegfried Münch (GMD)
Pierre Népomiaschty (INRIA)
Henk Nieland (CWI)
Carol Peters (CNR)
Constantine Stephanidis (FORTH)

E-mail:
helena.dutoit@inria.fr

wb@inf.rl.ac.uk
lars@sisu.se
jafb@eniacc.inesc.pt
truls.gjestland@delab.sintef.no
siegfried.muench@gmd.de
pierre.nepomiaschty@inria.fr
henkn@cwi.nl
carol@vm.iei.pi.cnr.it
cs@csi.forth.gr


Telephone:
+33 1 39 63 5401

+44 235 44 5281
+46 8 752 1613
+351 1 572 895
+47 73 59 26 45
+49 2241 14 2303
+33 1 39 63 5646
+31 20 592 4092
+39 50 593 429
+30 81 22 64 87

AEDIMA



FORTH



**Asociacion Española de
Informatica y Matematica
Aplicada**

UPC
Av. Gregorio Marañon, s/n
E-08028 Barcelona

Tel: +34 3 401 72 89
Fax: +34 3 401 62 10

**Centrum voor Wiskunde
en Informatica**

Kruislaan 413
NL-1098 SJ Amsterdam

Tel: +31 20 592 9333

**Consiglio Nazionale
delle Ricerche**

IEI-CNR
Via S. Maria, 46
I-56126 Pisa

Tel: +39 50 593 400
Fax: +39 50 554 342

**Foundation of Research
and Technology – Hellas**

Institute of Computer Science
P.O. Box 1385
GR-71110 Heraklion, Crete

Tel: +30 81 22 93 02
Fax: +30 81 22 93 42

**Gesellschaft für
Mathematik und
Datenverarbeitung mbH**

Schloß Birlinghoven
P.O. Box 1316
D-53731 Sankt Augustin

Tel: +49 2241 14 -0
Fax: +49 2241 14 2889



**Institut National de
Recherche en Informatique
et en Automatique**

Domaine de Voluceau,
Rocquencourt, B.P. 105
F-78153 Le Chesnay Cedex


Tel: +33 1 39 63 55 11
Fax: +33 1 39 63 53 30



**Instituto de Engenharia
de Sistemas e Computadores**

Rua Alves Redol 9
Apartado 10105
P-1017 Lisboa Codex

Tel: +351 1 310 0217
Fax: +351 1 572 897



**Rutherford Appleton
Laboratory**

Chilton
DIDCOT
GB-Oxon OX11 0QX

Tel: +44 235 82 1900
Fax: +44 235 44 5385



**Swedish Institute
of Computer Science**

Box 1263
S-164 28 Kista

Tel: +46 8 752 1500
Fax: +46 8 752 7230



**Stiftelsen for Industriell
og Teknisk Forskning ved
Norges Tekniske Høgskole**

SINTEF DELAB
N-7034 Trondheim

Tel: +47 73 59 26 00
Fax: +47 73 59 43 02