ERCIM NEWS

European Research Consortium for Informatics and Mathematics Number 13 April 1993

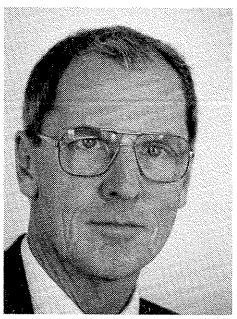

Siwert Sundström, Managing Director of the Swedish Institute of Computer Science, considers the role of computers in our lives. (Photo: SICS)

SPECIAL :

Activities in Parallel Architectures 8

CONTENTS

Next Issue:

Software Reuse

EDITORIAL

A few weeks ago my first grandchild was born. A nice little baby named Clara. Her brain which is smaller than a tennis ball contains about 100 billions of neurons, and her brain power exceeds that of a supercomputer. Her memory is today completely empty, and it will take many years before its contents are useful. The process of storing information just started, and even after a full life time, only a fraction of the available memory space will be used. Without any training in knowledge acquisition, little Clara is already using extremely efficient methodology to acquire and store knowledge, to train her brain to interpret pictures and sound, and to control the millions of muscle fibres she needs in order to speak and move. The way she represents knowledge in her brain and uses that knowledge to think and act is still to a great extent unknown.

The goal of Swedish Institute of Computer Science (SICS) is neither to find out how the human brain is functioning, nor to develop computers which copy its abilities. The goal is rather to make the computers efficient parts of an infrastructure that can serve Clara in as many situations as possible in her private as well as her professional life. We like to see computer systems as aids to complement rather than compete with human abilities. The processing elements in this computer based infrastructure should be as invisible as possible. Clara should be able to benefit fully from it without having to be a computer specialist.

This goal puts heavy requirements on the functionality of computer systems. This in turn requires drastically increased processing power. High performance computing systems can be achieved in different ways. Multiprocessing, which is the theme of this issue is one way. The challenge here is to interconnect and to program the processing elements in such a way that they can share the work load with a minimum of loss in efficiency. A number of concepts are available, and some European initiatives have successfully reached the market. The SICS Data Diffusion Machine Project which is presented in this issue is a promising approach to scalable shared memory multiprocessing systems. Another article describes the MUSE system, developed by SICS. This system is marketed by BIM as the first commercially available parallel Prolog system on the market under the name of BIM-MUSE.

The collaboration with European researchers in the multiprocessing field as well as in other areas has given valuable input to the research at SICS, and we believe that the benefit has been mutual. We are convinced that our membership in ERCIM will further improve such collaboration to the benefit of ourselves as well as European IT research as a whole.

Siwert Sundström

The Swedish Institute for Computer Science (SICS) becomes ERCIM's Ninth Member Institute

by Janusz Launberg and Lars Bergman

SICS is one of some thirty industrial research institutes in Sweden. Three of these institutes are IT-oriented with altogether more than 200 research staff. The three institutes, SICS, SISU (Swedish Institute for Information Systems Development) and IM (Swedish Institute of Microelectronics) are located in the Electrum building in Kista just north of Stockholm. Electrum also contains the department of Information and System Science of the University of Stockholm and several departments of the Royal Institute of Technology (RIT).

The SICS environment is a quite interesting combination of research and higher education and industrial activities of both big companies such as Ericsson Radio Systems, Ericson Components and IBM Sweden, and a great number of SMEs active in electronics, software, and application segments of the IT market. SICS also has well developed contacts with all major research universities in Sweden. In the ERCIM setting the role of SICS is to function as an active node to the Swedish IT research and industrial organisations and vice versa.

SICS is to ensure effective transfer of knowledge and technology to promote practical use of research results. To this end SICS has established several joint projects with industry. Some have been part of the SICS Basic Research Program and carried out by guest researchers from industry; others have taken the form of research assignments. An important effect of such cooperation is the establishment of a network of personal contacts between SICS and the companies while creating opportunities to evaluate research results in industrial applications. Guest researchers from industry play the important role of bringing new knowledge back to the companies and bringing industrial problems to SICS. Contacts are reinforced by SICS employees moving to industry.

Distributed Systems (DSLAB)

Distributed systems are central to a large section of Swedish information technology. The goal of the laboratory is to develop methods and techniques for the design of embedded real-time systems for communication and control of data networks while meeting the challenges of new technical opportunities, such as high-speed networks and multimedia interaction technology.

Research in the area of formal design methods has concentrated on the further development of process algebra, real-time logic, synthesis technology, and type theory, with the goal of developing formal methods that support (semi-) automatic synthesis and reasoning about the correctness of protocol specifications.

Research in the area of distributed system architectures falls entirely under the MultiG research program and has been conducted jointly with various groups at the Royal Institute of Technology (RIT), industry and SISU. The goals of the program are (1) to demonstrate how the high-speed network can be utilised in applications requiring extensive bandwidth, (2) to produce methods and techniques for the development of such, and (3) to study the man-machine interaction of systems in which different communication forms (data, speech, pictures) cooperate.

The areas of application are computer-supported cooperation (the distributed office), correspondence schooling (the distributed university), and various types of command and control centres (air traffic control, highway traffic control, network management, rescue control, and C3I. Two experimental interactive environments have been produced: "Collaborative Desktop", which is an extension

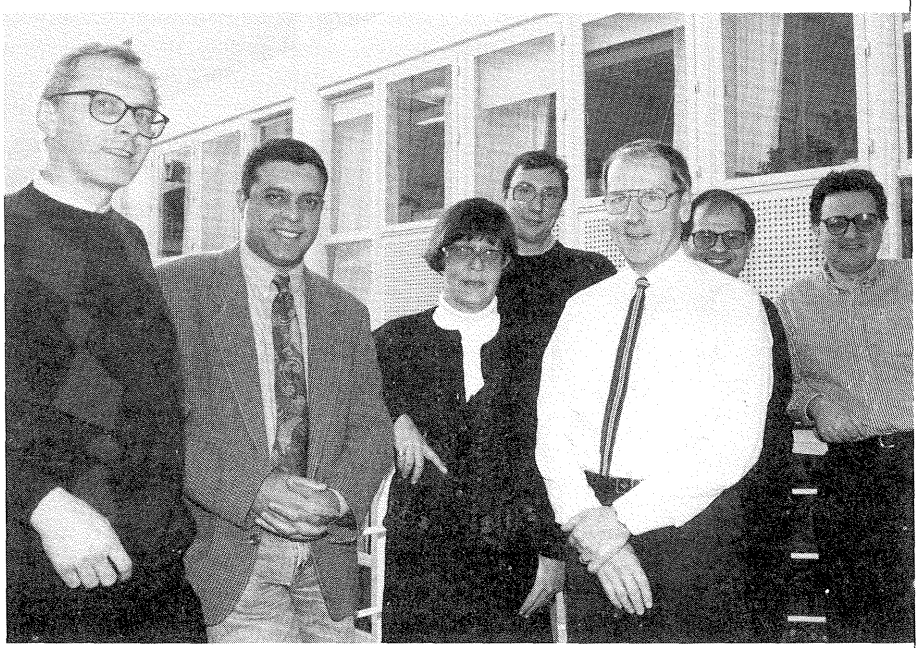

The managing group of SICS gathered in the mall of the Electrum building where SICS, SISU and IM, the three IT oriented institutes, are located. (Photo: SICS)

ERCIM Research and Workshop Reports

If you wish to receive any of the following ERCIM Research and Workshop Reports, please circle the report number(s) of interest, fill in your name and address at the bottom of the page, and return the form to:

ERCIM, B.P. 105, Rocquencourt, F-78153 Le Chesnay Cedex, France
Fax: +33 1 39 63 53 30

ERCIM-92-R001: Numerical Solutions of Some Semiconductor Devices by a Domain Decomposition Method, Choi-Hong Lai and Chris Greenough

ERCIM-92-R002: Dynamic Gesture Machine, Monica Bordegoni

ERCIM-92-R003: An Iteration Scheme for Non-Symmetric Interface Operator, Choi-Hong Lai

ERCIM-93-R004: An Imperative Language for Task-Level Programming: Definition in Temporal Logic, Eric Rutten and Lionel Marcé

ERCIM-93-R005: An Imperative Planning Language: From Temporal Representation to Real-Time Execution, Eric Rutten

ERCIM-93-R006: The qwertz Document Types- Version 1.1 Reference Manual, Thomas Gordon

ERCIM-93-R007: Iterative Refinement of Linear Least Squares Solutions, Alexander Malyshev

ERCIM-93-R008: Parallel Aspects of Some Spectral Problems in Linear Algebra, Alexander Malyshev

ERCIM-93-R009: A Comparison of OBJ3 and ASF+SDF, Steven Eker

ERCIM-93-R010: Texture Synthesis, Michal Haindl

ERCIM-93-R011: Verification of Parameterised Synchronous Concurrent Algorithms with OBJ3: The Pixel Planes Architecture Revisited, Steven Eker

ERCIM-92-R012: Associative Matching for Linear Terms, Steven Eker

ERCIM-93-R013: Formal Specification of Manifold: a Preliminary Study, Eric Rutten, F. Arbab and Ivan Herman

ERCIM-92-W001: Theoretical and Experimental Aspects of Knowledge Representation, Pisa, May 21-22, 1992

ERCIM-92-W002: Numerical Linear Algebra, Pisa, May 21-22, 1992

ERCIM-92-W003: Software Quality Principles and Techniques, Pisa, May 21-22, 1992

ERCIM-92-W004: Parallel Architectures for Computer Vision, Heraklion, October 29-30, 1992

ERCIM-92-W005: Network Management, Heraklion, October 29-30, 1992

ERCIM-92-W006 : Methods for Software Reuse, Heraklion, October 29-30, 1992

ERCIM-92-W007: Numerical Methods for Linear and Nonlinear Problems in Wave Propagation, Heraklion, October 29-30, 1992

Name: ..... *Address:* .....
.....
.....
.....

of the desktop metaphor that includes support for computer-aided cooperation (CSCW), and DIVE (distributed interactive virtual environment), which extends the desktop metaphor toward a 3-D "workroom metaphor".

To disseminate results from the MultiG program, SGN, Stockholm Gigabit Network, a fibre-optics network between several research and industrial sites in the Stockholm area is currently being established. Extensions throughout the country are planned for the next three-year period. SGN is currently being used for an FDDI loop around Stockholm and an ATM test facility.

Logic Programming and Parallel Computer Systems (LPSLAB)

Programming methodology is being studied from the viewpoint of logic and concurrent constraint programming. Research in computer systems focuses on architecture and operating systems for parallel computer systems.

In the programming methodology area, the laboratory is working in two directions. The first is based on PROLOG technology and efforts are directed toward making it available as a tool for problem-solving. The results, after applied-case studies, are integrated into the widely distributed SICStus PROLOG system. The laboratory has also shown that it's solution to automatic parallelisation and scheduling of Prolog programs, MUSE, is successful on parallel computers with either common or distributed memory. The second direction is represented by work with the development of Andorra Kernel Language (AKL), the goal of which is to create a programming methodology that systematically combines the expressivity of logic languages, the structuring properties of object-oriented methods, the ability of the process-oriented language to handle events in the surrounding world, and, mechanisms for solving optimization problems. Another important goal is that AKL shall support the utilisation of sequential, parallel, and distributed systems.

LPSLAB's other area of research, computer systems, is conducted as an interplay between architecture and operating systems for parallel computer systems. One goal is the development of performance-modular parallel systems, i.e., systems whose capacity can be raised step by step and guarantees an attractive price-performance relationship. Another goal is to make it possible for parallel systems to be programmed based on a well-developed and broadly applicable programming model based on shared memory. The operating system research is based on UNIX and Mach.

The laboratory has developed the concept Cache Only Memory Architecture (COMA). A bus-based prototype of the COMA architecture, called the Data Diffusion Machine (DDM), is currently being developed.

Knowledge-Based Systems (KBSLAB)

Application areas for research at KBSLAB are design, planning, and computer-based understanding of natural languages. Prototype and tool development generally take place together with partners in universities and industry.

Research in man-machine communication has been conducted in the area of understanding natural languages, Computer Supported Cooperative Work (CSCW), and user modelling. In the area of computer-based understanding of natural languages, an advanced platform for experiments (BCI) has been developed in co-operation with SRI Cambridge forming the basis for experiments with applications such as machine translation and advanced interfaces with data bases. Techniques in automatic learning have been used both in BCI and to define language models for speech recognition in real time.

In CSCW, research has mainly been conducted in the area of models for negotiation and dialogue. Recently work has been initiated on so-called virtual worlds and mechanisms for communication. Applied studies have been carried out, for instance in the area of computer support at the construction site. User modelling, as a tool for generating acceptable user-system dialogue, has been studied under the Prometheus project called Interactive Route Guidance.

Research in the area of environments for software development is being carried out in part under the European KADS II project. Common KADS has been established as a reference standard for other projects in the area. An initial prototype for a workbench has been developed. KBSLAB is concentrating on technologies and methods for the construction and reuse of knowledge bases in the application areas of technical design, diagnosis, and planning.

Reasoning models research is done in connection with applications domains such as planning/re-planning, and electronics design. Several prototype applications have been developed, mainly in commissioned projects. The laboratory has also developed a logical language based on partial inductive definitions, GCLA II, to formalise and implement certain types of reasoning in design and planning applications.

International co-operation

SICS participates in six ESPRIT-projects: General Purpose MIMD (GPMIMD), KADS II, CONCUR2, COMIC, CONFER and PARFORCE, ACCLAIM in Prometheus under Eureka, as well as in two DRIVE program projects: CITIES and SOCRATES KERNAL. Formal co-operation projects have been carried out with Stanford Research Institute, Cambridge and the Institute for New Generation Computer Technology (ICOT) in Tokyo.

Short facts

The annual budget is about 7 MECU. The staff consists of 80 persons. The basic research program is funded 60% by companies and 40% by the Swedish Technology Development Board (NUTEK). ■

Please contact: Janusz Launberg - SICS
Tel: +46 8 752 1500
E-mail: janusz@sics.se
or: Lars Bergman - SISU
Tel: +46 8752 7230
E-mail: lars@sisu.se

ERCIM Fall Workshops in France

by Helena du Toit

The next set of ERCIM workshops, taking place in November later this year, will be hosted by INRIA and located at different sites throughout France. There will be four workshops dealing with aspects in the following research areas: human-computer interaction, software development, databases and numerical methods. A short description of each workshop is given below.

Multimodal Human-Computer Interaction

Nancy, France,
2-4 November 1993

This workshop will focus on the following issues relating to multimodality in the context of human-computer communication:

- human-computer interaction through natural language and graphics;
- theoretical aspects and applications;
- multimodal communication acts (see M.T. Maybury's taxonomy of communicative acts for multimedia and multimodal dialogue);
- architectures for multimodal user interfaces including voice besides standard input/output (keyboard, mouse, graphics);
- ergonomic experimental studies and guidelines for the design of multimodal human-computer dialogue.

Please contact: Noelle Carbonell - INRIA
Tel: +33 83 59 3055
E-mail: Noelle.Carbonell@loria.fr

Development and Transformation of Programs

Nancy, France,
2-4 November 1993

This workshop is intended to address the scientific, technical and pragmatic issues involved in the development of specifications and programs using transformations. Topics include, but are not limited to:

- formal transformation rules;
- development strategies;
- description of the program transformation process;
- reuse and adaptation of programs;
- tools and environments to support program transformations.

Please contact:
Jeannine Souquieres - INRIA
Tel: +33 83 59 3033
E-mail: Jeanine.Souquieres@loria.fr

Parallelism and Non-determinism in Active Database Application

Rocquencourt, France,
2-4 November 1993

The ERCIM Database Research Group organises periodically workshops on different themes of interest. The first workshop was held at CWI in Amsterdam, the second at RAL in Abingdon on scientific databases, and the third one at CNR in Pisa on updates and constraints handling. This workshop will be on parallelism and non-determinism in databases. Parallelism and non-determinism are two fundamental issues both from a theoretical and a practical point of view. The workshop will focus on both orientations.

Following the desire of the community, tutorials will be organised. Patrick Valduriez (INRIA) will speak about parallel database systems and Victor Vianu (UCSD) about non-determinism and parallel complexity of queries.

Abstracts should be sent to: Stéphane Grumbach, EDRG'93, INRIA BP 105, 78153 Rocquencourt Cedex, France, before August 31, 1993. The abstract should consist of a one page description of the work to be presented with a few references.

Please contact:
Stéphane Grumbach - INRIA
Tel: +33 1 39 63 5446
E-mail: stephane.grumbach@inria.fr

Stochastic Numerical Methods for Solving PDE's, and Applications

Sophia-Antipolis, France,
8-9 November 1993

The objective of the session is to present recent developments concerning stochastic numerical algorithms for solving deterministic partial differential equations, in particular non-linear, with applications in Fluid Mechanics or Statistical Physics e.g: Monte-Carlo methods, stochastic particle methods.

A particular emphasis will be brought to theoretical results on the rates of convergence, and to the numerical implementations on parallel architectures.

Contributions studying the probabilistic interpretation of non linear PDE's (as Mc-Kean Vlasov limits of large systems of interacting particles, or descriptions of statistics of laws of branching processes, for example) are also welcome.

Please contact: Denis Talay - Sophia Antipolis-INRIA
Tel: (33) 93 65 78 98
Fax: (33) 93 65 76 02
Email: Denis.Talay@sophia.inria.fr

■

EDRG Workshop: Repositories, Methods and Tools for Systems Engineering

Heraklion, Greece, 3-5 May 1993

Recent research in the Information Repositories area is aimed at developing a uniform approach to the management of software artifacts produced during the software development process, such as requirements, designs, programs and test data which traditionally were managed in ad-hoc and arbritrary manner.

The 4th ERCIM Database Research Group (EDRG) Workshop aims to provide a forum to exchange ideas and experiences on the research issues concerning the development and application of methodologies, tools and techniques in the area of Information Repository Systems. A strong connection with evolving database techniques exists and will be a major focus of the workshop.

Topics:

- Repositories and the Software Engineering/Development Process
- Repositories and Software Reuse
- Repositories and Database Technolo
- Present and future in Information Repositories

■

Please contact: Janis Bubenko - SISU
Tel: +46 8 752 1600
E-mail: janis@sisu.se
or: Rogerio Carapuca - INESC
Tel: +351 1 310 0237
E-mail: rsc@eniac.inesc.pt
or: Stephane Grumbach - INRIA
Tel: +33 1 39 63 54 46
E-mail: grumbach@seti.inria.fr
or: Keith Jeffery - RAL
Tel: +44 235 44 6103
E-mail: kgj@ib.rl.ac.uk
or: Wolfgang Klas - GMD
Tel: +49 6151 869 943
E-mail: klas@darmstadt.gmd.de
or: Arno Siebes - CWI
Tel: +31 20 592 9333
E-mail: arno@cwi.nl
or: Arne Solvberg - SINTEF
Tel: +47 759 3438
E-mail: asolvber@idt.unit.no
or: Yannis Vassiliou - ICS-FORTH
Tel: +30 81 22 11 71
E-mail: yannis@csi.forth.gr

ERCIM Workshop on Theory and Practice in Verification

by Stefania Gnesi and Alessandro Fantechi

More than thirty researchers from 4 ERCIM Partners (CNR, FORTH, INRIA and RAL) and from other European research institutes attended the first ERCIM workshop on "Theory and Practice in Verification", organised by the Istituto di Elaborazione della Informazione (IEI-CNR) in Pisa, 9-11 December 1992.

The aim of the workshop was to provide ERCIM researchers with an opportunity to share theoretical results and practical experiences in the field of formal verification; both technical presentations and demonstrations were thus given. Professor Rance Cleaveland from North Carolina State University gave an invited paper on "Model Checking for Modal μ-calculus".

Fourteen papers were presented in the three days, together with some tool demonstrations, before a lively audience, revealing that a lot of attention is now beng paid to this area.

The main topics included:

- Verification Theories for Infinite State Systems
- Verification Tools and Techniques for Distributed Concurrent Systems
- Automatic Verification of Real Time Systems
- Model Checking for Temporal Logics

There was particular interest in verification techniques for non finite-state systems, case studies, and attempts at integrating different verification techniques and tools.

Stimulating discussions were raised by the presentations and it is hoped that this workshop is just the beginning of a fruitful cooperation between ERCIM members in this area of research. The overall feeling was that the scope of meetings of this type should be mostly to permit the exchange of information and to encourage a general awareness of the activities of others in the field, thus favouring the creation of smaller groups closely collaborating on specific arguments, and a better interaction between them.

■

Please contact: Stefania Gnesi - IEI-CNR
Tel: +39 50 593486
E-mail: gnesi@vm.iei.pi.cnr.it

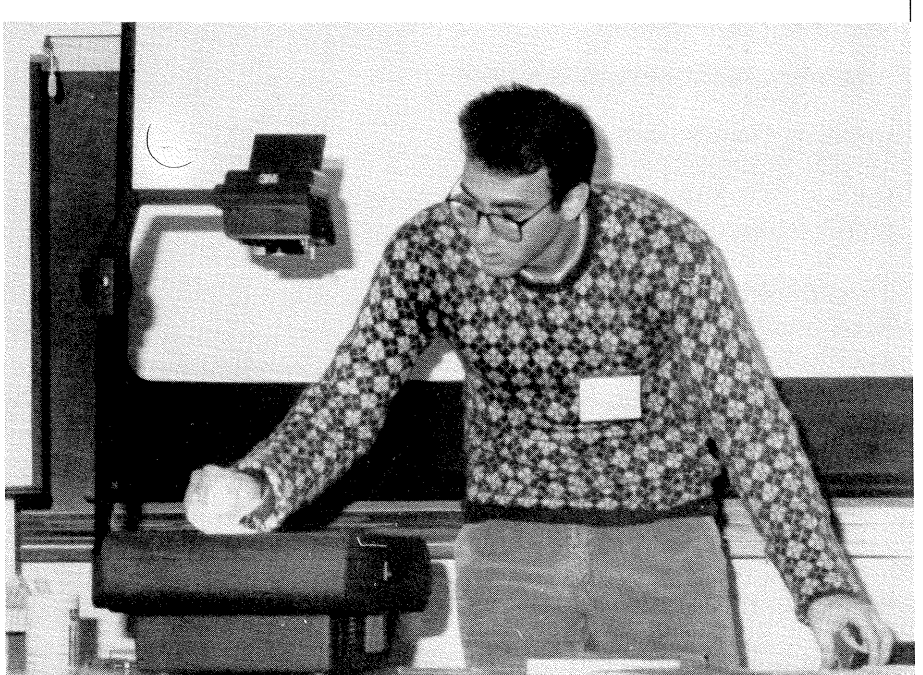

Daniel Yankelevich, H.P. Labs and Pisa University, presenting a parametric verification tool for distributed concurrent systems (Photo: CNR)

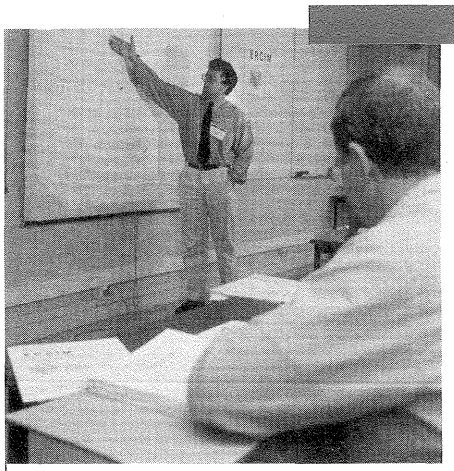

(Photo: Eidelman, INRIA)

ERCIM Advanced Training Programme '93

by Annick Theis-Viémont

The following Advanced ERCIM Courses are scheduled for the remainder of the first half of 1993:

Compiler for Parallel Machines

Rocquencourt, France, 25-27 May 1993

Organiser: François Tapissier - INRIA
Tel: +33 1 39 63 56 00
Fax: +33 1 39 63 53 30
E-mail: Francois.tapissier@inria.fr

Programming Techniques for Parallel Architectures

Amsterdam, The Netherlands, May 1993

Organiser: Frans Snijders - CWI
Tel: +31 20 592 4171
Fax: +31 20 592 4199
E-mail: franss@cwi.nl

Three Dimensional Vision: Application to Mobile Robots and Medical Data

Sophia-Antipolis, France,
28 June - 2 July 1993

Organiser: Catherine Juncker - INRIA
Tel: +33 93 65 77 78
Fax: +33 93 65 77 65
E-mail:Catherine.Juncker@sophia.inria.fr

Please contact:
Annick Theis-Viémont - INRIA
Tel: +33 1 39 63 53 78
E-mail: Annick.Theis-Viemont@inria.fr

The 1993/94 ERCIM/HCM Fellowship Programme

by Annick Theis -Viémont

After applying within the European programme, Human Capital and Mobility (HCM), ERCIM was granted additional funding for its fellowship programme. Among the applicants best ranked within the ERCIM Fellowship Programme, nine applicants were selected and asked to file for HCM.

We are pleased to take this opportunity to announce the nine fellows and their allocations to the different ERCIM member institutes in the accompanying table.

Please contact:
Annick Theis-Viémont - INRIA
Tel: +33 1 39 63 53 78
E-mail: Annick.Theis-Viemont@inria.fr

The nine successful applicants and their placement at the various ERCIM institutes (at present, consisting of nine member institutes). Where possible, the fellows will spend a period of six months at three institutes, totalling a research period of eighteen months each.

Name	ERCIM Institute		
Alan Craig (United Kingdom)	INRIA 01/09/93	CWI 01/03/94	SINTEF 01/09/94
Wilhem Heinrichs (Germany)	INESC 01/04/93	CNR 01/10/93	
Giulia Iori (Italy)	INESC 01/07/93	INRIA 01/05/94	CWI 01/11/94
Bernard Manderick (Belgium)	GMD 01/07/93		
Gerhart de Meer (Germany)	CWI 01/04/94	CNR 01/10/93	
Philippe Nehlig (France)	RAL 01/09/93	CNR 01/09/94	
Mukesh Patel (United Kingdom)	INRIA 01/04/93	FORTH 01/04/94	
Alan Stewart (United Kingdom)	INRIA 01/01/93	CWI 01/06/93	
Jan Tretmans (The Netherlands)	SINTEF 01/02/93	FORTH 01/10/93	GMD 01/06/94

Next ERCIM Workshops in Norway

The next ERCIM workshops, hosted by SINTEF-DELAB, will take place at Trondheim and Røros from 26-28 May 1993. The emphasis will be on discussions in small groups and formulation of concrete plans for future work. Poster sessions will be organised to get input and contributions from all participating ERCIM members.

The three workshops are:

- **Storage and Retrieval of Multimedia Information**
- **Modelling and Simulation of Industrial Process Systems**
- **Interactive Modelling, Simulation and Visualisation in Large-scale Scientific Computing**

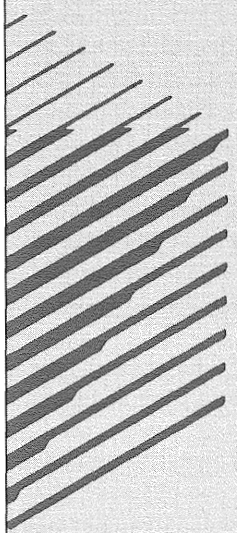

Please contact: SINTEF-DELAB
Tel: +47 7 59 26 00
E-mail: ercim@delab.sintef.no

EDGE: European Distributed Generic Environment

by Bob Hopgood

At the ERCIM Directors Meeting in Crete on October 30, 1992, the Directors had a detailed discussion concerning the mechanisms for strategic research that would be most effective in the Fourth Framework Programme. There was considerable support for the new instrument, the Priority Technology Project (PTP), being proposed by the Commission. A PTP is a major project aimed at improving European industry's competitiveness in world markets by sharpening the focus of R&D onto generic technologies of wide applicability that will have a major impact on industrial competitiveness.

The breadth and depth of R&D activities carried out by the ERCIM partners suggested that it should be feasible for ERCIM to put forward a significant project to the Commission that would be of benefit across a range of industries. ERCIM had already noted the problems that SMEs have in achieving technology transfer from CEC programmes. In consequence, a specific target would be to define a PTP that addressed the needs of SMEs as well as other sectors.

Professor Tribolet from INESC volunteered to oversee the project. Much of the work took place between October and the end of 1992 with the main meeting taking place in Lisbon on December 14th. Further refinement occurred in January with the resulting EDGE proposal outline being completed by the end of January 1993.

The ERCIM partners focused on the needs of Europe in the design of products. Considerable resources have been put into improving production in manufacturing industry but less effort has gone into improving design. The question posed was "What did SMEs need to get a substantial improvement in their design capability?". Many problems were identified ranging from the lack of integration of existing tools, their restricted scope, and the inability to access resources easily. Here, resources ranged from appropriate and affordable processor power to information concerning parts, regulations and standards.

This led to the view that the generic technology required was an integrated distributed environment that could be used by designers across Europe. If the generic environment could support all aspects of the more demanding design tasks, the belief was that the environment would almost certainly be appropriate in many other less demanding areas, including case and data processing associated with commerce and government.

EDGE is, therefore, a proposal to build a modern European Distributed Generic Environment aimed at the manufacturing design sector but having much wider applicability. It would be based on a high bandwidth communications system and would incorporate modern concepts such as object-oriented data and knowledge bases, multi-media and multi-modal interaction, cooperative working etc. The EDGE user would have access to all European product data, an integrated set of simulation and design tools and various processor and storage resources. The widespread availability of cooperative working would allow SMEs to more easily get involved with and contribute to large projects involving many partners from different countries.

The ERCIM partners are aware that this is not an easy task to complete and forecast the need for effort at the level of 400 manyears per year for 4 years. However, such an integrated project would have major deliverables of value across a range of disciplines.

The 15-page EDGE Report is available from ERCIM for anyone interested in obtaining further information.

Please contact: Bob Hopgood - RAL
Tel: +44 235 82 1900
E-mail: frah@ib.rl.ac.uk

Applications and Exploitation of Parallel Architectures: Activities within the ERCIM Institutes

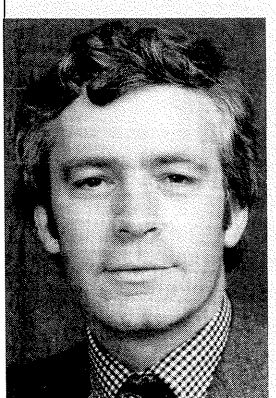

Dr. Chris Wadsworth, Head of Parallel Processing Group, Informatics Department, RAL (Photo: RAL)

by Chris Wadsworth

This series of short articles provides a snapshot of ongoing activities among the ERCIM partners aimed generally at exploiting the potential of parallel processing.

Many parallel computers employing a variety of architectures and with varying performance capabilities have become available commercially over the past few years. At one extreme parallel supercomputers strive to provide the ultimate in absolute performance; other approaches seek to optimise price/performance for affordable parallel computing. One example of the latter is distributed workstations (multi-workstations) interconnected either as a parallel resource or for using spare workstation cycles to useful purpose.

The development of algorithms and applications software remains the biggest challenge in realising the full potential. Typical applications possess much inherent parallelism: in some cases, e.g. ray tracing, the parallelism may be self-evident; in others a new parallel design may need to be devised. The resulting parallelism must then be expressed in terms of an algorithmic formulation and translated into software. Successful exploitation thus relies both on the discovery and development of parallel solutions at the applications level and on the provision of languages, tools and techniques to assist the mapping of conceptual parallel designs ultimately onto real parallel hardware.

The series includes articles on activities addressing both these aspects. The articles are presented approximately in an order ranging from the applications end through to tools and techniques. Two of the articles also address architecture developments (virtual shared memory) that are expected to make a large impact soon on commercial offerings.

One measure of the strength of ERCIM activity in this area is that a number of otherwise acceptable articles have had to be omitted for lack of space. ERCIM, it may be said, is addressing the area in parallel, and in numbers. ■

Please contact: Chris Wadsworth - RAL
Tel: +44 235 44 5101
E-mail: cpw@inf.rl.ac.uk

Medium-Range Weather Forecast Production Code on Parallel Computers

by Ute Gärtel, Wolfgang Joppich and Anton Schüller

GMD and ECMWF (European Centre for Medium-Range Weather Forecasts) joined forces some months ago in order to parallelize ECMWF's production code for medium-range weather forecasts, IFS (Integrated Forecasting System). Recently, the first milestone of this cooperation has been reached: The 2D model of the IFS, which contains already all relevant data structures and algorithmic components of the corresponding 3D models, has been parallelized and run successfully on the Intel iPSC at GMD.

The central goal of this project is the parallelization of the main components of the IFS before summer 1993. The IFS consists of approximately 150,000 lines of Fortran code and supports both a variety of meteorological models and several algorithmic alternatives.

At the beginning of the project the important decision of the basic parallelization strategy had to be made taking into account the narrow time-frame and the complexity of the sequential code. A further side condition was set by the wish of the ECMWF to avoid program modifications in the numerical kernels of the code whenever possible, with the need to introduce to some extent somewhat more sophisticated data structures in the parallel version. A careful analysis showed that the data transposition strategy allowed the fulfilment of all these requirements. In particular, it comprises advantageous features with respect to the parallel efficiency as compared with alternative parallelization approaches.

This strategy re-distributes the complete data to the processes at various stages of the algorithm such that the arithmetic computations between two consecutive transpositions can be performed without any interprocess communication. This approach is feasible since there are only data dependencies within one coordinate direction, this direction being different within the main algorithmic components. Thus, a parallelization with respect to the remaining two dimensions is possible for 3D weather simulations, making the strategy suitable for highly parallel systems.

Problems of the parallelization are caused by the variety of data structures employed in the code (one triangular and two rectangular data spaces are used by the spectral transform method), which requires non-trivial mappings of the data to the processes and special considerations for the efficient parallelization and the data structuring in the parallel program.

Beginning with the 2D shallow water model, the concrete parallelization of the sequential program started in September 1992. In order to guarantee portability among the variety of different parallel computers currently on the market, the portable PARMACS message passing interface developed at GMD was selected for the parallelization. The results on the iPSC confirm the expected parallel efficiencies of up to 90 % and more. ■

Please contact: Anton Schüller - GMD
Tel: +49 2241 14 2572
E-mail: anton.schueller@gmd.de

Adaptive Multilevel Methods on Parallel Computers

by Hubert Ritzdorf and Klaus Stueben

The development of self-adaptive algorithms is the future research area in the numerical solution of complex applications. The realization of such algorithms on massively parallel computers is a particular challenge. A major difficulty in developing adaptive approaches (for grid-oriented problems) on parallel computers arises from the fact that, in general, local refinements (or coarsenings) are generated dynamically during a calculation and thus dynamic strategies for grid-to-processor mapping are required. Despite various investigative approaches, this problem is essentially unsolved.

Complex physical processes, such as, for example, the combustion in an engine, the flow around cars or the development of the weather, are described by large systems of partial differential equations. If a precise resolution of extremly "fine-scaled" physical phenomena is required for a realistic numerical simulation, one has to impose correspondingly high requirements on the fineness of the discretization grid. (Examples: the spreading of shock and flame fronts after ignition by the spark plug, the influence of physical forces inside boundary layers of a flow and the precise spreading of fronts of low-pressure areas and influences of wind effects.) If a static grid of such a degree of refinement was used globally, a numerical solution would be impossible due to the enormous computational work, even with "optimal" (multilevel) methods.

Most applications, however, require extreme grid resolutions only in special (a priori unknown) areas of the computational grid. In such cases, the use of adaptive grids which are automatically generated and dynamically and locally adjusted during the computation according to the requirements of the solution, permits a numerical simulation.

We are working on the development of adaptive multilevel strategies for massively parallel systems of MIMD type with distributed memory. The underlying grid structures are hierarchies of block-structured grids (currently 2D). Implementations are based on the GMD program package LiSS. This package was originally developed for static grids; it offers a flexible environment for the parallel solution of arbitrary (elliptic) systems of nonlinear partial differential equations via multilevel methods. LiSS has a high-level interface to the parallel machine consisting of a self-contained library of communications routines (COMLIB).

This library will be extended by adaptive routines in such a way that the complete communication, including the dynamic grid-processor mapping and global data re-distributions, will be done exclusively by the COMLIB-routines. In this way, the actual application routines are kept free from any machine-specific language constructs and special communications requirements: with the exception of certain calls to COMLIB-routines, the numerical routines look the same as on sequential machines. A first COMLIB version with restricted functionality has been completed. ■

Please contact: Klaus Stueben - GMD
Tel: +49 2241 14 2749
E-mail: stueben@gmd.de

Image Rendering on Massively Parallel Systems

by Wolfgang Krüger and Peter Schröder

The work of the recently founded Scientific Visualization Group of GMD is concentrated on basic research in graphic and visualization algorithms, applications in Computational Physics, Computer Medicine, and in programming of massive parallel machines such as the CM-2. Main research topics include visual control of simulations, real-time walkthrough scenarios for large data sets, development of new visualization models, and integrating multi-media tools (e.g., stereoscopic display and Virtual Reality).

One of the main research projects is the implementation of basic image rendering models such as polygon rendering, radiosity, raytracing on the massive parallel system itself.

Since the early days of Computer Graphics a great deal of attention has been devoted to the design of high performance graphics hardware. This development has led to performance levels of 1 MPolygons/second performance available on some desktop machines today. These performance levels have pushed current hardware technology to its maximum. Current approaches do not appear to scale any further. Thus there have been proposals for new architectures and in particular a great deal of interest in exploiting massive parallelism for rendering tasks.

According to commonly accepted calculations it should be possible to achieve performance on the order of 100 MPolygons/second on current generation parallel machines such as the Intel Paragon and the Thinking Machines CM-5. However, in order to achieve this performance major research is required in algorithm development. We argue that only object parallelism in polygon rendering can achieve these performance levels.

Recently researchers have begun to consider object parallelism, while most previous work has only considered image parallelism. In image parallelism each processing element (PE) is assigned some number of pixels on the screen and is responsible for scanconverting all primitives (e.g. polygons) which cover its respective screen region. In this approach the generated network traffic between processors is a function of the number of primitives in the data base and hence, unbounded.

This implies unfavorable scaling properties of image parallelism which is worsened by the fact that with increasing numbers of processors any overhead has to be amortized over smaller and smaller regions of the screen. The load balancing problems also become more and more severe due to the uneven distribution of primitives across the screen. Object parallelism on the other hand distributes primitives to processors irrespective of their location on the screen.

Each PE contains a local framebuffer into which primitives are rendered. These local framebuffers then need to be merged on a backend network into a global framebuffer. The traffic on the backend network is a function of screen-size and depth complexity and thus, for all intents and purposes, bounded. The main issue in this approach is the question of how to achieve fast pixel merging. Current proposals include buses, trees, and pipelines. This research is conducted on the Intel Paragon and Thinking Machines CM-5 computers, which support various network topologies and programming primitives.

■

Please contact: Wolfgang Krüger - GMD
Tel: +49 2241 14 2366
E-mail: krueger@viswiz.gmd.de

Parallel Image Processing

by Bob Maybury

RAL is a partner in the UK-funded project PRIPS (Parallel Reconfigurable Image Processing Systems). The idea behind the project is that a configuration of processors which is optimal for one phase of a problem may be inefficient at another stage, hence it is desirable to reconfigure the processors as the computation progresses.

Image processing has been identified as an area where parallel processing should have a major impact. However an application may start with a simple operation carried out on all pixels and then proceed to the identification of features in the image in less regular ways. The technique of breaking the image into equal parts and farming these parts out works well for many low-level operations but gives a poor return at the later stages.

It therefore appears that image processing is a good application area within which the value of dynamic reconfiguration can be assessed. The hardware which will be used at RAL is a Parsys SuperNode with 24 T805s. As presently set up 20 of these can be configured into networks and the suitability of different configurations can be tested. The project aims to demonstrate the value of the concept, to identify common patterns for the usage of reconfigurability for imaging applications, and show where gains and losses occur.

■

Contact: Bob Maybury - RAL
Tel: +44 235 44 6292
E-mail: rm@inf.rl.ac.uk

Parallel Volume Visualization

by Claudio Montani, Raffaele Perego and Roberto Scopigno

The need to visualize volume datasets is common to many activities, both in research and applications. Volume visualisation can be performed using a ray tracing approach. Although volumetric RT is computationally less complex than geometric RT, interactive throughputs require the exploitation of the high degree of inherent parallelism common to most visualization techniques.

Recently, we proposed a scalable parallel implementation of a volume ray tracing algorithm for distributed memory multiprocessors, based on a novel hybrid parallelization strategy. The algorithm has been implemented on a hypercube multicomputer and works on both voxel-based and cell-based volume datasets.

The parallelization of a ray tracing algorithm is somewhat trivial if the parallelism is exploited at pixel level with a subset of pixels assigned to each processing node storing the whole dataset. Unfortunately, the huge amount of memory needed to represent high-resolution voxel datasets, even if the datasets have been previously classified and compressed, makes partitioning and distributing the data among the nodes' local memories a must. As a consequence of partitioning, a ray cannot generally be entirely traced by a single node: during the tracing phase, when the ray exits the local partition, it has to be sent for further tracing to the node holding the partition next crossed by the ray. This means a lot of ray transmissions on the communications network, and thus high communication overheads and low scalability if the number of processing nodes has been increased causing the dimension of the partitions assigned to the nodes to decrease.

Following these guidelines and with the aim of achieving good scalability, an original and flexible approach (see figure below) in which the computing nodes are organized into a set of clusters has been designed and evaluated. The image space is partitioned and a subset of pixels is assigned to each cluster, which will compute pixel values independently. Each cluster is composed of a set of nodes that cooperate by mutually exchanging ray descriptors for computing the colours of the assigned pixels. The dataset is therefore replicated on each cluster, but is partitioned over the local memories of the cluster's nodes. In order to optimize global efficiency, the number of clusters and the number of nodes per cluster can be flexibly chosen as a function of dataset resolution.

The results reported in the table below are related to the visualization of a 350x250 image from a 97x97x116 dataset which represents the electron density map of an enzyme. They show that the proposed hybrid parallelization strategy fulfils the initial goal by obtaining good scalability in the visualization of high resolution volume datasets on MIMD distributed memory architectures. Even though the actual run times suffer slightly from shortcomings in the base algorithm used, the efficiency obtained for such a highly communicating algorithm (0.74 on 128 nodes) is higher than those reported elsewhere and validates the correctness of the design choices.

Timings	time	speedup	efficiency
1 node	485.24	1	1
2 nodes	260.35	1.86	0.93
4 nodes	130.19	3.73	0.93
8 nodes	65.82	7.37	0.93
16 nodes	32.96	14.72	0.92
32 nodes	17.26	28.11	0.88
64 nodes	8.90	54.52	0.85
128 nodes	5.14	94.41	0.74

Times (in seconds), speedup and efficiency with increasing number of nodes of an nCUBE 6400 hypercube

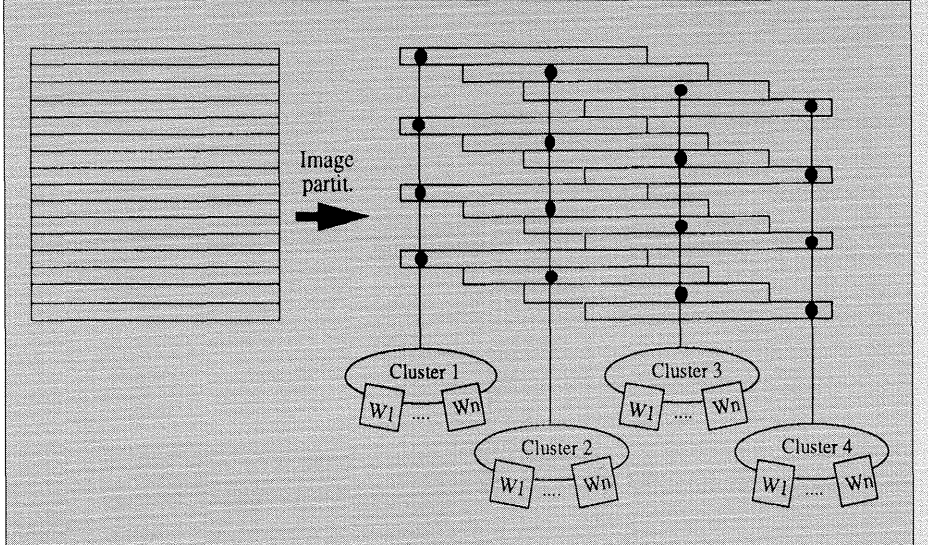

A scalable approach to highly parallel volume visualization.

Our on-going activities in the field of parallel computer graphics are oriented towards Delauney triangulation algorithms, visibility algorithms on digital terrain maps and solid modelling.

Please contact:
Claudio Montani - IEI-CNR
Tel: +39 50 593451
E-mail: montani@iei.pi.cnr.it
or: Raffaele Perego - CNUCE-CNR
Tel: +39 50 593253
E-mail: raffaele@suncube.cnuce.cnr.it
or: Roberto Scopigno - CNUCE-CNR
Tel: +39 50 593304
E-mail: scopi@icnucevm.cnuce.cnr.it

Parallelism in Factoring Large Numbers

by Herman te Riele

In connection with the safety of cryptosystems it is important to find out what are the largest numbers which can be factorized in a reasonable time on the available computer systems. To this end, in a 4-year project at CWI some of the best-known factoring algorithms are studied, implemented and optimized on the fastest parallel and vector computers.

Factoring large numbers and testing primality have received much attention since the publication, in 1978, of the RSA public-key cryptosystem. The safety of this system is based on the difficulty of factoring large numbers. As far as we currently know, cracking this system requires an intrinsically exponential effort (in complexity terminology, an NP-hard problem). However, since the invention of RSA the rapid developments in factoring algorithms, computer power and memory size have learned that the size of the RSA public key has to be chosen very carefully if one wants to be sure that the system is actually secure for a long term of, say, 25 years to come. In 1978 a public key size of 100 decimal digits was expected to be safe for the remainder of this century, whereas nowadays there are at least five research groups in the world (including CWI) where one is capable of factoring difficult 100 decimal digits numbers within a few days on the available computing equipment.

The project at CWI, carried out in co-operation with the University of Leiden, focuses on the Number Field Sieve factorization method (NFS) of J.M. Pollard et al. NFS has proved its tremendous power in the recent accomplishment (a new world record) of the factorization of the ninth Fermat number, $2^{512}+1$. Computers from all over the world (including CWI) have contributed idle computer cycles to this success. The NFS-method is appropriate for numbers of the form a^b+c, for small a and c. Another method, the Multiple Polynomial Quadratic Sieve (MPQS) factorization method, is applicable to general numbers. At CWI this method has been implemented on various supercomputers and (parallel) workstations and presently several recently discovered practical improvements are being investigated, particularly for parallel systems. However, if NFS could be turned into a method suitable for general numbers, then it can be expected to be more powerful than MPQS for numbers with more than 120 –130 decimal digits (the precise cross-over point has to be determined experimentally, requiring a considerable amount of CPU time).

The factoring algorithms studied in this project can be parallelized nearly perfectly, needing only a very modest amount of host-node communication. One of the subprojects is the study of using the so-called Condor system for factoring. Condor is designed to provide convenient access to unutilized workstations and other computing systems in a given network, while preserving the rights of their own users. ■

Please contact: Herman te Riele - CWI
Tel: +31 20 592 4106
E-mail: herman@cwi.nl

Parallelism in Iterative Methods

by Henk van der Vorst

The study of robust iterative schemes, suitable for parallel as well as conventional architectures, is of importance for many large-scale applications where large linear systems arise often quite naturally. The design of reliable and efficient software for iterative linear system solvers is studied at the University of Utrecht, in cooperation with the University of Tennessee (LAPACK) and CWI.

Parallel computers are expected to help us solve very large scientific problems and to lead to insight in phenomena which are poorly understood so far. Usually a major part of the computing time is taken by the solution of linear systems. In problems for which parallel computing is the only way to obtain answers, those linear systems are typically so large that direct solution methods are unfeasible. As alternatives, multigrid and iterative solvers are studied intensively. Our research focuses on the class of so-called Krylov subspace iterative methods, which comprises well-known methods such as Conjugate Gradients, GMRES, CGS, QMR and Bi-CGSTAB. The time consuming computational kernels of these methods are: preconditioning operations, inner products, sparse matrix-vector products, and vector updates.

Inner products, matrix-vector products and vector updates are relatively easy to parallelize. Preconditioning is often the most problematic part in a parallel environment. Hence, a major part of our attention is devoted to this aspect. Incomplete decompositions of the matrix A form a popular class of preconditionings, in the context of solving discretized PDEs. In this case the preconditioner is $K=LU$, where L and U have a sparsity pattern equal or close to the sparsity pattern of the corresponding parts of A (L is lower triangular, U is upper triangular). Solving the system $Kw=r$ leads to solving successively $Lz=r$ and $Uw=z$. These triangular solvers lead to recurrence relations which are not easily parallelized. Although there are a number of approaches to obtain parallelism in the preconditioning part, considerable effort is required to solve this problem, especially for massively parallel machines.

Our approach is based on a combination of the following techniques:

- Reordering the computations:
  Depending on the structure of the matrix a frontal approach may lead to successful parallelism. By inspecting the dependency graph one can select those elements that can be computed in parallel. For coarse-grained parallelism this approach is insufficient.
- Reordering the unknowns:
  One may also use a colouring scheme for reordering the unknowns, so that

unknowns with the same colour are not explicitly coupled. This means that the triangular solvers can be parallelized for each colour. Of course, communication is required for couplings between groups of different colours.

- Forced parallellism:
  Parallelism can also be forced by simply neglecting couplings to unknowns residing in other processors. This is like block Jacobi preconditioning, in which the blocks may be decomposed in incomplete form. This can be improved by making incomplete decompositions on slightly overlapping domains.

The problems with parallelism in the preconditioner have led to searches for other preconditioners as well. Often simple diagonal scaling is an adequate preconditioner and this is trivially parallelizable. Still another approach is to use polynomial preconditioning:

$$w=p_j(A)\,r\,, \text{ i.e. } K^{-1}=p_j(A)\,,$$

for some suitable j-th degree polynomial. This preconditioner can be implemented by forming only matrix-vector products, which, depending on the structure of A, are easier to parallelize. ■

Please contact: Henk van der Vorst - Utrecht University
Tel: +31 30 53 3732
E-mail: vorst@math.ruu.nl

Benchmarking of Concurrent Architectures

by Chris Greenough

BECAUSE is one of a number of European projects considering how best to assess the usefulness of the new generation of concurrent computing architectures. The aim of the project is to develop a process for evaluating these new systems on real applications. Three sample applications areas are represented in the project consortium: computational fluid dynamics, electromagnetic design, and semiconductor device simulation.

Specification and implementation of a BECAUSE Benchmark Set (BBS) have been completed. Work is well advanced on metrics and ways of measuring system performance for the benchmarking phase, the results of which will be used in the final phase, the critical assessment of systems.

The BBS has three major classes, each with several sub-classes, of test programs: Basic Kernels; Application Kernels; and Potential Library Routines. The first implementations have been completed on four machine ranges: the reference machine (IBM RS/6000), the Intel hypercube (iPSC/2 and iPSC/860), the Parsys Supernode SN1000, and the Thinking Machines CM-2. The figure shows results measured for one of the BBS routines, 2D Fast Fourier Transform. A serial version of the BBS has been tested on the reference machine prior to being disseminated to the European community via an information server at INRIA Sophia-Antipolis.

A successful European Workshop was held in October 1992 at INRIA Sophia-Antipolis to which the European community was invited to discuss the test programs and their parallel implementation.

At the conclusion of the project in early 1993 the European community will have a well documented benchmarking suite, tested and critically assessed on a number of currently available parallel systems. ■

Contact: Chris Greenough - RAL
Tel: +44 235 44 5307
E-mail: cg@inf.rl.ac.uk

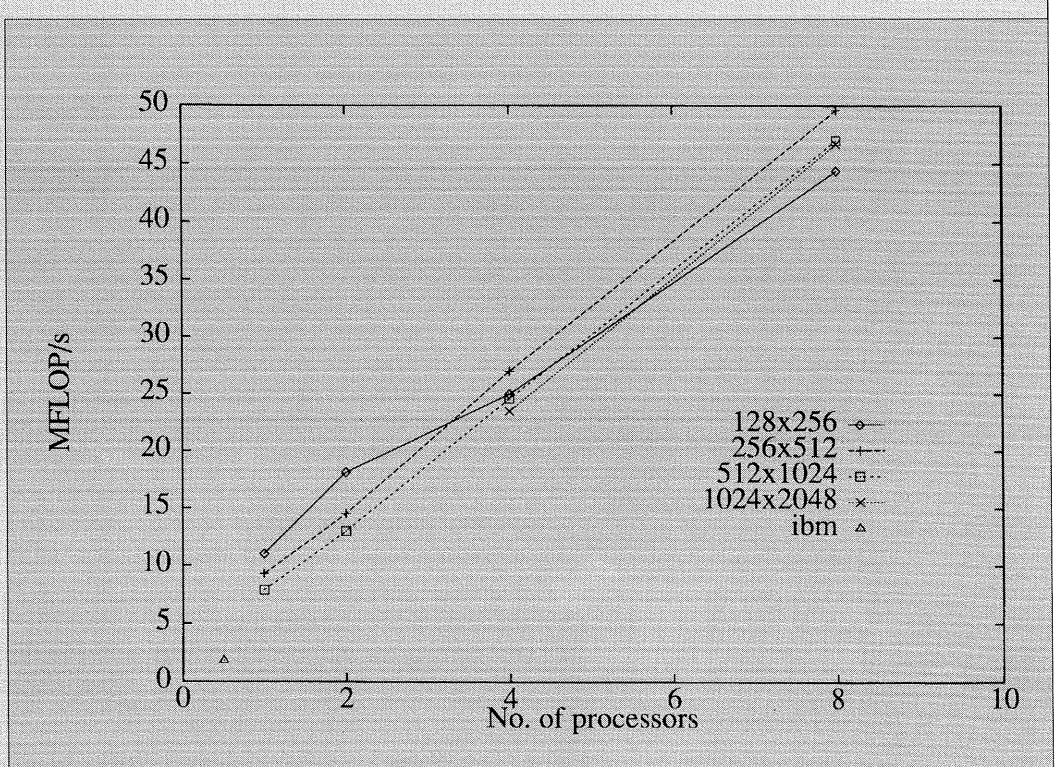

Results of 2D FFT on iPSC/860 showing the computation rate (MFLOP/s) as a function of number of processors for 4 different mesh sizes. The performance of the reference machine is also shown.

Parallel Technology in Education and Research

by Karstein Sørli

The Norwegian Institute of Technology, SINTEF and STATOIL have initiated a project aiming at utilizing massively parallel supercomputers for solving scientific problems.

The project has a timeframe of three years, and includes a collaboration with one of the most experienced vendors of massively parallel computers, Intel. As part of the "parallel initiative" the 3 partners have purchased an Intel Paragon massively parallel computer, the latest product from Intel's Scientific Supercomputer division. Research in porting old sequential codes and developing new parallel algorithms will be persued over the coming three years in key areas like weather forecasting, chemistry, reservoir and seismic processing, computational fluid dynamics and structural mechanics. In this period SINTEF will serve as a "center of excellence" for Intel.

At SINTEF, the Division of Industrial Mathematics has been involved in research in parallel processing since the division was formed in 1990. Areas of activity have been parallel algorithms for the solution of ordinary and partial differential equations, computational fluid dynamics, molecular dynamics and weather forcasting.

As a part of the "parallel initiative" the division has been actively involved in the purchase of the Intel Paragon supercomputer mentioned above. The division will coordinate research utilizing the Paragon computer and will lend its expertise to the other members of the initiative in order to facilitate widespread use of parallel processing within these organizations. ■

Please contact: Karstein Sørli - SINTEF
Tel: +477 59 68 54
E-mail: ks@sima.sintef.no

Parallel Applications at CNUCE-CNR

by Domenico Laforenza

Programming parallel architectures is a difficult task due to the lack of consolidated programming paradigms, environments and general methodologies. This software gap makes parallel programs un-portable, expensive to write and hard to test and evaluate. On the basis of these considerations, some parallel applications have recently been designed at CNUCE both for shared and distributed MIMD computers, with the principal aim of gaining familiarity with the most common parallelism paradigms and being able to evaluate their effectiveness on different architectures.

The CNUCE group has much experience in certain strategic areas such as Computational Chemistry and Computer Graphics which constitute real computing-intensive problems requiring the exploitation of parallelism. These activities are outlined below.

Parallel Computer Graphics

The use of massively parallel methods is highly recommended in Computer Graphics. It has been estimated that current high-quality graphics techniques require a speedup of six orders of magnitude to be used in real time systems. Such a power increase can only be obtained with massively parallel architectures. Therefore the design of new parallel algorithms is a must in this field. Research in graphics algorithms is now under way and regards both basic topics and parallel issues. Stable fields of interest for us are visualization and data representation schemes, paying particular attention to volumetric datasets and addressing both representation issues and visualization techniques for the display of such data. More recent fields of interest are computational geometry and solid modelling.

Parallel Algorithms for Reactive Scattering Codes

Innovative algorithms designed for vector and parallel architectures are of great interest for molecular dynamics calculations and particularly for reactive scattering codes. This type of algorithm makes it feasible to systematically exploit the accuracy of the potential energy surfaces evaluated for atom-diatom reactions by using "ab initio" techniques for a comparison of calculated and measured values of the physical properties of the system. Vector and parallel implementations of reactive scattering codes have been carried out. Comparisons of the performance of shared and distributed memory machines have been made. ■

Please contact:
Domenico Laforenza - CNUCE-CNR
Tel: +39 50 593270
E-mail: free@icnucevm.cnuce.cnr.it

Automatic Partitioning of Data Parallel Programs for Distributed Memory Multiprocessors

by Thomas Brandes

MIMD architectures with distributed memory are the kind of parallel machines that are scalable and can be used for a wide range of scientific applications. Normally these architectures are programmed with explicit message passing between the processes. As the message passing programming model is very error prune and difficult to use, many efforts have been made to offer other programming models that are easier to use.

The High Performance Fortran Forum has defined language extensions and modifications for Fortran to overcome these difficuties by supporting data parallel programming. This kind of programming can be defined as single threaded, global name space, and loosely synchronous parallel computation. The new language allows code tuning for various architectures and should guarantee top performance on MIMD and SIMD-computers with non-uniform memory access costs.

The High Performance Fortran Forum has also a European counterpart that has been organized by Clemens-August Thole, GMD.

The parallelization tool Adaptor (Automatic Data Parallelism Translator) makes it possible to translate data parallel programs to message passing porgrams already now. It transforms data parallel programs written in Fortran 77 with array extensions, parallel loops, and layout directives to parallel programs with explicit message passing.

Adaptor is a big package that consists of:

- the interactive source to source transformation tool XAdaptor
- the distributed array library (DALIB)
- documentation files in Postscript format,
- example programs.

The source language of ADAPTOR can be defined shortly in the following way:

- Fortran 77 with some restrictions,
- with many array extensions of Fortran 90 (inclusive dynamic arrays and intrinsic functions for arrays),
- with some features of Connection Machine Fortran (timing, random numbers),
- with layout directives to specify the data distribution.

ADAPTOR is not a compiler but a source to source transformation that generates Fortran 77 host and node programs with message passing. The new generated source codes have to be compiled by the compiler of the parallel machine. The following parallel machines are supported:

- Alliant FX/2800
- iPSC/860
- Net of SUN 4 Workstations (based on PVM)
- Net of IBM Risc Workstations (based on PVM)
- Parsytec GCel
- Meiko Concerto

The software packages, documentations files in PostScript, and a number of example programs are available via "anonymous ftp" from:

ftp.gmd.de (129.26.8.90)

in subdirectory: gmd/adaptor.

In GMD the tool will be used to develop optimization strategies for High Performance Fortran mappers in cooperation with NA Software, Liverpool, and to develop and test MIMD extensions that could be part of the next version of High Performance Fortran. ■

Please contact: Thomas Brandes - GMD
Tel: +49 2241 14 2492
E-mail: brandes@gmdzi.gmd.de

Architectural Neutral Distribution Format Parallelisation in OMI/GLUE

by Jose Delgado

GLUE is an ESPRIT project started in July 92 in the OMI section (Open Microprocessor systems Initiative). Its main objective is to develop the basic software architecture for OMI, based on the concept of a VBI (Virtual Binary Interface).

The VBI chosen is ANDF (Architectural Neutral Distribution Format), being promoted by OSF, and allows software to be distributed in a code format common to all machines and from which source code can not be retrieved. All that is needed is one compiler (called producer) from each source language to ANDF (Fortran, C, Ada, Lisp, Occam and Cobol will be supported) and one translator (called installer) from ANDF to each microprocessor's machine language (SPARC, MIPS, ARM, T9000, O1 – a new microprocessor being developed in OMI – and Alpha will be contemplated). Among other benefits, this avoids the need to maintain and to purchase different versions of the same software for different machines.

The work of INESC is concerned with automatic parallelisation using ANDF, which will be important to:

- widen the software base for a parallel machine, taking advantage of parallelism for conventional, sequential applications;
- provide a migration path for programmers, making use of familiar, sequential environments;
- increase software portability between parallel architectures (by rearranging the parallelism in the program to suit the architecture).

Automatic parallelisation techniques are usually done in a by language and by architecture basis, with techniques specific for a language and/or an architecture embedded in each compiler. Although in many cases an intermediate form is used, this is a rather non-universal approach, in the sense that the intermediate form differs from case to case, constituting only an implementation technique.

With ANDF, the idea is to decouple the parallelisation techniques into three main components, capitalising on the universality of ANDF and on the fact that it retains essentially all the necessary high level information:

- Language dependent techniques. The parallelisation of Fortran, for instance, relies heavily on DO-loop analysis, whereas in Lisp the interprocedural analysis is of paramount importance. Including parallelising transformations in the producer (source language to ANDF translator) allows them to be optimised to that particular language and minimises the compromises of a completely universal paralleliser.

- Architecture dependent techniques. For instance, a shared memory multiprocessor has different tradeoffs/requirements from those of a distributed memory parallel computer. Again, performing these parallelising transformations at the ANDF to architecture translator (installer) level avoids compromises and maximises efficiency, while retaining the advantages of ANDF from the machine manufacturer point of view: these transformations will apply to all languages.
- System (language and architecture) independent techniques, those that can be efficiently used for a wide range of languages and/or architectures are probably best done at the ANDF level. This means applying parallelising transformations to a sequential ANDF program and obtaining a parallel ANDF program.

Besides researching in this line of work, INESC will also implement an installer of parallel ANDF for a parallel workstation based on SPARCs.

■

Please contact: Jose Delgado - INESC
Tel: +351 1 3100211
E-mail: jcd@inesc.pt

Fine Grained Parallelism in an Object-Oriented Environment

By Paulo Verissimo and Paulo Guedes

Distributed memory systems where nodes cannot access each other's memory scale very well and are readily available but they require a programming model different from the one offered by most programming languages, which makes them difficult to program and requires applications to be rewritten. One interesting approach is to provide a distributed shared memory in software so that applications written for a shared memory model still continue to work in a distributed environment. The big challenge is to make them perform well.

The Distributed Object Systems group at INESC has built a support for fine grained objects in the IK distributed object system developed in Esprit projects Comandos and Harness. Currently it is investigating techniques for efficiently sharing these objects using a distributed object memory where updates of replicas are integrated with the synchronization primitives. This involves compiler support to detect access to distributed objects, a user library to provide synchronization and replica instantiation and system support for updating and maintaining the coherence of replicas using weak consistency algorithms.

Another promising field is coarse grained parallelism on workstations. Distribution of task execution through LAN-based networks of workstations, either for load balancing or functional partition purposes, has been attempted in the past. The advent of group tools in operating system support, such as group communication and membership, group replication and cooperation management, should ease the task of implementing and controlling this kind of applications.

The Navigators group at INESC is developing a new generation of such tools. It will capitalize on past experience in group-oriented programming, such as the xAMp group communication suite developed in the ESPRIT DELTA4 project. Together with group communication and membership services for large-scale systems which remain efficient over LANs (or MANs), replication and cooperation management protocols oriented for specific application domains (such as parallel computing, computer supported cooperative work) will be developed. The task of application building will be supported by an object-oriented programming interface, capable of translating attributes of replication, cooperation, etc., into the necessary distributed interactions using group protocols.

■

Please contact: Paulo Verissimo - INESC
Tel: +351 1 3100228
E-mail: pjv@pandora.inesc.pt
or: Paulo Guedes - INESC
Tel: +351 1 3100287
E-mail: pjg@sabrina.inesc.pt

DDM: A Cache-Only Memory Multiprocessor

by Seif Haridi and Erik Hagersten

The shared-memory multiprocessor model allows the processors to share data uniformly. This approach simplifies dynamical scheduling and load balancing and is often regarded as the most general programming model. One negative attribute though has been its poor performance and poor scalability compared to the alternative message-passing model.

Our approach has been to take today's microprocessors as a starting point and to add the means for successful cooperation via shared memory. There exists, however, no physical shared memory in our system. Instead, all memory is divided among the processors and organized in such a way that data can be duplicated, moved freely, and allowed to reside in any memory. This behavior of data is not visible to the programmer, who sees the popular shared memory abstraction. It is, however, beneficial to performance and adapts well to different application behaviours without bothering the programmer.

We have introduced a new class of architectures based on the above model comprised of caches and processors connected by a network – Cache-Only Memory Architectures (COMA) – and prototype design thereof – the Data Diffusion Machine – (DDM). The large caches are managed by a cache-coherence protocol which makes sure that all copies of a datum have the same value. The protocol also find a datum that is not in the local cache of the requesting processor. The commercial KSRI architecture is an example of a COMA. We have shown that COMAs have superior performance over alternative shared memory architectures in a quantitative analytical performance study, and the implementation proposal for the DDM has been simulated with good performance

results. In 1993 we plan to have the first DDM prototype DDM operational. In parallel with the hardware effort, a scalable operating system for the DDM is being developed based on Mach 3.0.

This research is conducted at the Logic Programming and Parallel Systems Lab (LPS-lab). Other related activities at LPS-lab focus on different latency-hiding techniques, scalability issues for large COMA architectures and their operating systems. In addition LPS-lab is active in the field of parallel software based on logic programming and concurrent constraint programming.

■

Please contact: Seif Haridi - SICS
Tel: +46 8 752 1500
E-mail: seif@sics.se

Parallel Systems and Applications at ICS-FORTH

by Manolis Katevenis, Christos Nikolaou and Stelios Orphanoudakis

Three different groups at ICS-FORTH are currently involved in R&D activities related to parallel systems and applications. These activities are: Computer architecture and VLSI chip design for parallel computers, Systems software for parallel computers and Parallel computer vision applications.

Hardware Architectures

The VLSI Design & Computer Architecture Group of ICS-FORTH, headed by Prof. Manolis Katevenis, conducts R&D work in high performance computers and networks. The group is focusing on very-high-speed networks for connecting together the components of parallel computers – processors and memories.

While the general characteristics of uniprocessor architectures tend to more-or-less become standardized nowadays, the same is not true for the parallel system interconnection network. The research on RISC architectures in the early 80's distinguished the functions and operations of a uniprocessor that it is imporant to implement in hardware from those that it is better to implement in software, and showed ways to implement the hardware operations at high speed. Today, the gross traits of most high performance processors are similar: there is a superscalar RISC CPU, with caches, MMU, and a memory system. On the contrary, the chip components for interconnecting processors together and thus for building parallel computers, are nowhere near a convergence to some particular architecture style.

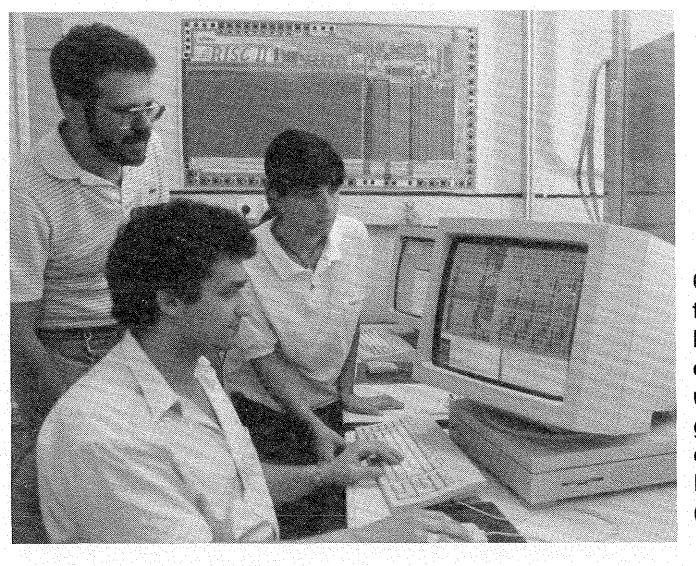

Chip design in the VLSI laboratory of ICS-FORTH under the guidance of Prof. Katevenis. (Photo: FORTH)

At ICS-FORTH, we view our research on generic components for high-speed multiprocessor networks in this context: we aspire to contribute to a convergence analogous to that of uniprocessor architecture. The most central question is: which are the network functions that it is essential to implement in hardware, and how can these be minimized and simplified so that the chips can route and switch traffic with a latency of only a few clock cycles. Immediately related to this are the choices that one makes on the routing algorithm, packet size, buffering and switching architecture, and flow control. At the very end, will the architectures for parallel computer networks, LAN's, MAN's, and WAN's converge?

At the next level up the hierarchy of architectural abstraction, the question is what kind of traffic one uses the network for. What are the communication and sharing mechanisms in the parallel computer? Nowadays, it is clear that "shared memory" is a convenient but expensive illusion, while "message passing" is an efficient hardware mechanism that suffers from software overheads and poor programmability. How can the two be married together, so that we only keep the good features of each?

Please contact:
Manolis Katevenis - ICS-FORTH
Tel: +30 81 229302
E-mail: kateveni@ics.forth.gr

Systems Software

The Parallel and Distributed Systems Group of ICS-FORTH, headed by Prof. Christos Nikolaou, conducts R&D work on the design, implementation, and experimentation with the efficient placement of data objects and threads of control on the processors and memory of multiprocessors. Because of the complexity of the problem, the group is using various interfaces between the application programmer, the parallelizing (or parallel) compiler, and the run-time environment, so as to facilitate the flow and exploitation of information about explicit or hidden parallelism and data locality.

The group works on developing both software tools and algorithms for enhancing the degree of parallelism of parallel or parallelized applications. With good profiling and visualization tools the programmer can give effective hints to the compiler; in turn, these hints can be crucial as input to data alignment and object allocation algorithms, callable by ei-

ther the compiler or the run time environment, for the generation of the parallel code.

The group is currently focusing on scientific and engineering applications, and will later extend its studies to commercial applications like complex query processing.

Please contact:
Christos Nikolaou - ICS-FORTH
Tel: +30 81 229302
E-mail: nikolau@ics.forth.gr

Parallel Computer Vision

The Parallel Computer Vision Group of ICS-FORTH, headed by Prof. Stelios Orphanoudakis, is involved in research which emphasizes the performance of parallel implementations of intermediate level image analysis tasks on distributed memory architectures.

These tasks are characterized by a substantial data reduction (compared to low-level processing of pixel data) occuring between different subtasks, a nonuniform – image content dependent – load distribution, and different data structures used to represent intermediate results. The load balancing requirements of such tasks are being investigated and efficient parallel implementations are being considered. The dependence of the performance of parallel implementations of intermediate level image analysis tasks on architectural features, the characteristics of different algorithms used to implement each task, load distribution strategies and image content is of particular interest.

Todate, many computer vision tasks and related algorithms with different computational and communication requirements have been implemented on the Connection Machine and the iPSC/2 Hypercube. In each case, performance has been evaluated with respect to the tradeoff between computational cost and communication overhead, different load distribution strategies, and different embeddings of interprocessor communication patterns. Similar work is currently carried out on a transputer architecture, as well as on our own implementation of a distributed processing environment based on the Linda model.

This work aims at producing a more general methodology for the efficient implementation of intermediate level image analysis tasks on parallel and distributed architectures.

Please contact:
Stelios Orphanoudakis - ICS-FORTH
Tel: +30 81 229302
E-mail: orphanou@ics.forth.gr

Multi-Workstations as Parallel Systems

by Ulrich Trottenberg

Recently, a 3D sparse multigrid solver for Poisson-like equations has been developed and very successfully used on a large cluster of workstations at the GMD. For a problem with 1 billion grid points on the finest grid, up to 87 processors can be employed, yielding a speed-up of about 80!

This result shows that large "multi-workstations" (large clusters of workstations) can indeed be used like highly parallel systems efficiently - provided that the algorithms are designed in an appropriate way.

What is available, what is needed?

Communication hardware: Multi-workstations are usually connected by a standard network (eg., ethernet), whose communication capabilities (throughput and start-up time) are quite poor in comparison to their computational power: the communication/computation ratio for multi-workstations is usually about a factor of 100 worse than that of closely-coupled ("real") parallel computers. In this regard, the hardware characteristics of multi-workstations and of parallel computers differ essentially.

Parallel software interface: In order to use multi-workstations, the relevant software interfaces must be implemented. PARMACS, which was developed in the GMD, is such a programming interface that is available now on most parallel systems. Application programs written using PARMACS run also on multi-workstations. In this way, software portability between "real" parallel computers and multi-workstations is achieved. PVM and Express are other programming interfaces with features similar to PARMACS.

Low- or "no"-communication algorithms: Most of the well-known parallel algorithms (eg., those in scientific computing) are communication intensive. The limits of the usefullness of multi-workstations (with weak communication hardware) become apparent quickly. In the classical benchmarks such as LINPACK, speedup (compared with one workstation) is obtained only when few workstations (2 to 4) are used. If more are used, the communication overhead dominates: computation time does not decrease, it increases!

Therefore, large multi-workstations can only be used efficiently for:

- applications which can be distributed onto the workstations (without or with very little communication) and than be treated independently,
- applications which allow the use of a low communication parallel algorithm.

The work on low-communication algorithms for communication intensive applications is one of the working fields of the GMD.

Results

A breaktrough has recently been achieved by combining the idea of "sparse grids" with the anisotropic 3D multigrid technique.

At the "Schloßtag" of GMD a program and results were presented, with the following features:

- testprogram: 3D Poisson-like equation on the unit cube
- h^2-discretisation: h = 1/1024, i.e. the full finest grid consists of 10^9 grid points
- accuracy ~ 10^{-5}
- use of up to 87 workstations
- speed-up of more than 80

Perspectives

The underlying idea of sparse anisotropic multigrid is very general. For its realisation only an efficient multigrid solver on flexible grid structures is needed; and certain mathematical conditions (weak extrapolation assumptions) have to be fulfilled. As a next step, the idea will be applied to a more complex, real problem in direct simmulation in turbulence. Apart from sparse anisotropic multigrid, other approaches of low communication algorithms are being investigated in the numerical group at the GMD. ■

Please contact: Ulrich Trottenberg - GMD
Tel: +49 2241 14 2759
E-mail: trottenberg@gmd.de

MUSE: A Parallel Prolog System

by Khayri Ali and Roland Karlsson

MUSE is an OR-parallel implementation of the full Prolog language. It is based on incremental stack-copying. It currently runs on a number of shared memory multiprocessor machines, e.g., TPSS1V from Tadpole Technology, Sequent Symmetry and SUN Galaxy. It also runs on the BBN Butterfly I and II.

The sequential SICStus Prolog, a fast, portable system, has been adapted to OR-parallel implementation. The extra overhead associated with this adaptation is negligible. The speedup factor is very close to the number of processors in the system for a large class of problems. The Prolog company BIM (Belgium Institute of Management) is currently using the MUSE approach to make a commercial product running on SUN Galaxy. ■

Please contact: Khayri Ali - SICS
Tel: +46 8 752 1500
E-mail: khayri@sics.se
or: Roland Karlsson - SICS
Tel: +46 8 752 1500
E-mail: roland@sics.se

General

A presentation of SINTEF Industrial Mathematics

by Marit Odegaard

SINTEF Industrial Mathematics, SIMA for short, is active in the following areas: numerical algorithms for vector and parallel computers and distributed applications, scientific visualisation, concepts related to interactive modelling, simulation and visualisation, modelling and simulation of fluid flow and heat transfer, stochastic signal processing and multivariate analysis applied to oceanographic and remote sensing, and statistical production control.

Most of the researchers have a degree from The Norwegian Institute of Technology, NTH, and have a background in industrial mathematics, statistics and numerical mathematics. The group is significantly oriented towards computational science. The latter is heavily directed by the fact that since 1987, we have been responsible for the Norwegian High Performance Computing Project, funded by Norwegian research council authorities and co-funded by Norwegian industry. The group also contributes significantly to the NTH-SINTEF initiative in "Parallel Computing in Education and Research", taken last year, and to be continued for another three years.

As a consequence of the high performance computing project, activities in scientific visualisation were started 3-4 years ago. These activities are mainly related to large-scale computations in science and engineering. Specialised hardware and software for graphics presentation, including animation, is used and developed.

This year SIMA is starting a new and strategic research program, called "Physical and Mathematical Simulation of Fluid Flow and Processes". It is a broad collaboration between eight different research groups at NTH and SINTEF. The main goal of the program is to "anchor" NTH/SINTEF's base of knowledge in physical and mathematical simulation of fluid flow and processes on a high international level. Applications are to a great extent optimal design of engineering products and processes in the marine, energy, material and environmental sectors of technology.

The remote sensing activities are focused on the use of Synthetic Aperture Radar (SAR) to observe ocean wave fields. The research which is supported by the Norwegian Space Agency is aiming at developing a complete system for obtaining directional wave spectra from the SAR on ERS-1.

Finally, ERCIM News readers may be interested to know that SIMA is responsible for coordination of one of the Workshops at the next ERCIM meeting, 26-28 May, to be held in Trondheim/Røros in Norway. This workshop "Interactive Modelling, Simulation and Visualisation in Large-Scale Scientific Computing" is described in the January issue of ERCIM News. ■

Please contact:
Marit Odegaard - SINTEF-SIMA
Tel: +477 59 68 42
E-mail: marit@sima.sintef.no

Research Activities at IAN-CNR

by Mario Arioli and Franco Brezzi

The main research activities at the Institute for Numerical Analysis, CNR, Pavia, are concentrated on the numerical solution of mathematical models, which are applicable in many scientific areas. Our interests range from the study of computer science techniques for qualitative models in medicine to the solution of partial differential equations.

We can identify six specific areas of interest:

- Basic properties of finite dimensional approximations; in particular, finite element methods (with many variants: mixed, hybrid, etc.), spectral and pseudo-spectral methods, domain decomposition techniques, wavelets, etc;
- Analysis of algorithms in numerical linear algebra; in particular, we study round-off error analysis and stability of direct solvers for large and sparse linear systems, and block iterative methods in parallel environments;
- Modelling and numerical solution of problems which are relevant in industrial applications: e.g. evolution with free boundary (Stefan problems), semiconductor, linear and nonlinear elasticity, and fluid dynamic problems;
- Modelling and simulation of the cardiac electromagnetic field produced by fronts of excitation;
- Analysis and implementation of qualitative models with particular emphasis on applications in expert systems for diagnostic purposes;
- Development of methodologies for using computers in education.

The location of our institute within the Mathematics Department of the University of Pavia greatly favours collaborations with researchers from the department; we also have very good scientific relationships and collaborate actively with the engineering faculty of the university.

Many of our research activities are carried out in collaboration with foreign research institutes or organisations. In particular, we are involved in the EC project AIM-GAMES, participating in the topic "Knowledge representation formalism and methods" (System Models); we collaborate with INRIA (Paris) within the framework of the "MODULEF Club" for the implementation of software for finite elements; our activity in numerical linear algebra for sparse matrices is conducted in collaboration with the Numerical Algorithm Groups at CERFACS (Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique) in France and at RAL in England. Significant research work is also under way on reflection and Steklov-Poincaré operators for fluid dynamics in collaboration with the Department of Numerical Mathematics of the Russian Academy of Sciences in Moscow.

Finally, we maintain close scientific relations with a number of industries: Adina Inc., A.S.I., Enidata, Himont, Marcel Dassault B.A., Olivetti, Philips, S.G.S., SNAM and Whirlpool. ■

Please contact: Mario Arioli - IAN-CNR
Tel: +44 39 382 21185
E-mail: arioli@bond.ian.pv.cnr.it
or: Franco Brezzi - IAN-CNR
Tel: +44 39 382 21185
E-mail: brezzi@ipvian.bitnet

Research in Informatics at IASI-CNR

by Michele Missikoff and Gianna Cioni

IASI – Istituto di Analisi dei Sistemi ed Informatica – was founded more than a decade ago as the result of the transformation of a research centre, active in the areas of Informatics and Control Theory, which had been constituted in 1969 within the Faculty of Electronic Engineering at the University of Rome "La Sapienza". In this transformation, the scientific traditions of the centre were maintained. For this reason, at IASI, research in Informatics represents half of the scientific activity in course.

An attempt to describe the wealth of IT activities at the Institute in a few words risks resulting in a mere check-list of topics. We thus prefer to focus on the rationale behind the organisation of the research. Although there is general consensus that basic research represents the core of our activity, applied research is seen as being extremely important for two complementary motives: it permits the quality of theoretical results to be verified; it provides the opportunity to maintain close contacts with industry and thus assist the development of the national productive system. In this perspective, research in Informatics at IASI addresses three domains of problems: theory, systems and tools, and applications.

In the domain of theoretical research, the Institute has a strong on-going activity in the traditional fields of languages, algorithms, and data structures. In these fields, particular attention is given to parallelism and management of complex data structures. Another domain with a strong tradition of theoretical research is that of data and knowledge bases. Our research here addresses problems of design and conceptual modelling, query processing, and methods for performance evaluation. In both domains, semantics has been an important topic of investigation.

In the domain of systems and tools, a lot of work has been done in the data and knowledge base area. In the last decade, there has been much experimentation and prototypes of several systems have been implemented, i.e. relational, statistical, geographic, deductive, and object-oriented databases, as well as expert system environments. The development of tools has particularly interested the layout of graphs and visual interfaces for knowledge bases. With respect to programming languages, prototypes of object-oriented programming environments, automated deduction, and rewriting systems have been developed. The Institute is also very active in the area of computer networks, where methods and tools for the assessment of different topologies and network organisations are being studied.

In the domain of applications, our activities have traditionally addressed problems related to transportation systems; our interest in Air Traffic Control goes back to the very beginnings of the Institute. Railway Traffic Control and Urban Traffic Systems are now also being studied and prototyped. Medical Expert Systems represent another important application area.

The above activities are carried out within several international projects, such as ESPRIT, and in cooperation with other national and international research organisations. ■

Please contact:
Michele Missikoff - IASI-CNR
Tel. +39 6 7716 422
E-mail: missikoff@iasi.rm.cnr.it
or: Gianna Cioni - IASI-CNR
Tel. +39 6 7716 436
E-mail: cioni@iasi.rm.cnr.it

Human-Computer Interaction

Accessibility of B-ISDN Terminals and Services by People with Special Needs

by Constantine Stephanidis

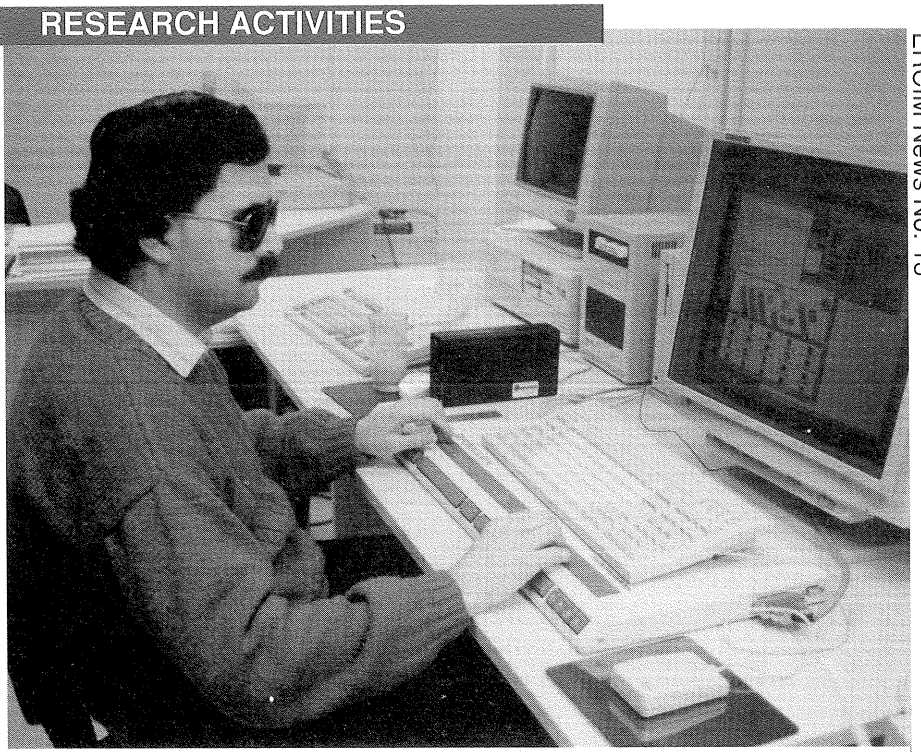

A blind user interacting with a multi-media workstation at ICS-FORTH. (Photo: FORTH)

IPSNI (R1066 - Integration of People with Special Needs in the IBC) has been one of the RACE projects that addressed the problem of accessibility of future B-ISDN terminals, services and applications by people with special needs (disabled and elderly people).

Emphasis has been given to the goal of facilitating, where possible, direct access by people with special needs to general purpose terminals and services through appropriate and modular adaptations. The adopted approach has been:

- To analyse the characteristics of the B-ISDN environment and to identify the requirements and difficulties of different categories of people with special needs in accessing terminals and services.
- To investigate technical solutions in favour of people with special needs at the level of the B-ISDN terminal, following functional and technical descriptions of a proposed generic multimedia terminal set-up.
- To examine potential alternative solutions for the accessibility of B-ISDN terminals and services by people with special needs taking into account aspects of user interaction with the B-ISDN environment and focusing on issues of redundancy and transduction of the various information types.
- To define the functionality of an adapted B-ISDN terminal in the form of functional specifications and to evaluate and verify the proposed terminal design specifications.
- To provide preliminary guidelines for the development of special services for those cases where terminal and/or service adaptations are inadequate or not feasible.

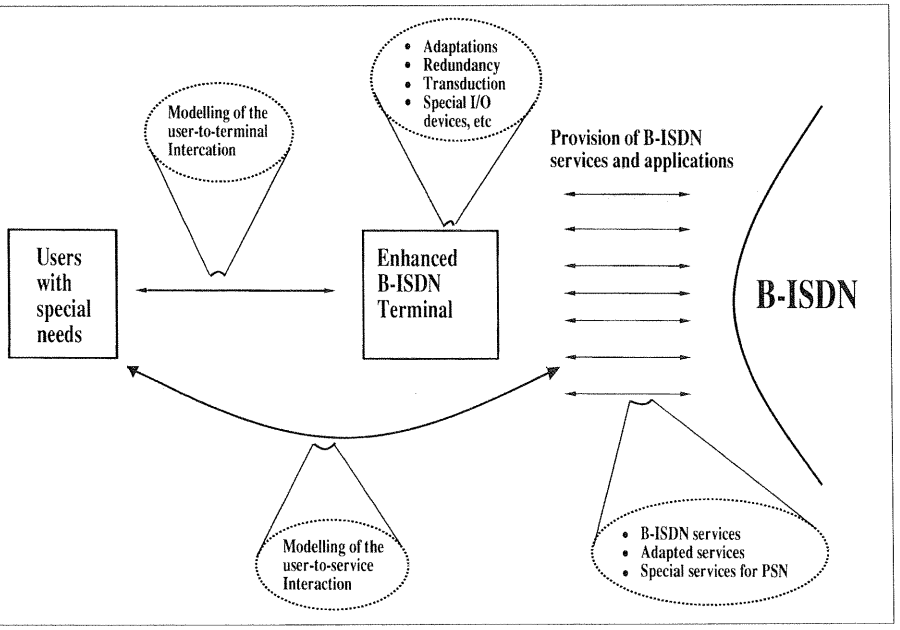

Intervention in the design of the B-ISDN terminals and services facilitating their accessibility by people with special needs.

In order to satisfy these objectives the IPSNI consortium has:

- identified relevant B-ISDN environment issues regarding the provision of accessibility through alternative technological solutions, including redundancy and transduction, and assessed the corresponding user performance characteristics in mono- and multi-modal interaction of people with special needs;
- defined the functional characteristics of the various categories of people with special needs in relation to the B-ISDN environment (User Profile), and provided guidelines for design engineers regarding issues of accessibility to telematic terminals and services by people with special needs;
- constructed a Functional Block Model (FBM) of the B-ISDN terminal which

provides a comprehensive view of the processes involved in the user-to-terminal interaction, a set of additional models of user-to-terminal and user-to-service interaction, and a complete set of Functional Specifications of a B-ISDN terminal aimed at design engineers who would need to introduce a minimum set of functions at the terminal level, in order to render it accessible by people with special needs;

- analysed issues concerning the user tasks, the interaction techniques and the physical implementation of them in a multimedia environment;
- evaluated the accessibility of some foreseen B-ISDN services and applications through the proposed terminal architecture and the pre-defined user tasks;
- examined the possibility of introducing special services and applications targeted specifically at certain categories of people with special needs.

FORTH-ICS has been responsible, amongst other things, for the development of a set of descriptive, predictive and characterisation models which aid the understanding of the functionally important processes associated with the user-to-terminal and user-to-service interactions and which establish, in a compatible way, the types of intervention that may be achieved either at the terminal or at the service level in order to render the B-ISDN services accessible by a particular user group.

An important added value of this modelling exercise is the generality of the models, which allow their use by B-ISDN designers in the wider context of the B-ISDN environment and the RACE programme. This means that guidelines and design criteria can be deduced, not only in the specific context of B-ISDN users with special needs, but for the design of a better, more functional and more versatile B-ISDN terminals for the benefit of the user population at large.

Additionally, FORTH-ICS has conducted a methodological analysis of user tasks in a multimedia environment, based on service definitions proposed by other RACE projects and taking into account current R&D efforts, focusing on issues related to the dialogue control in a multimedia environment, and proposed a mechanism for selecting interaction techniques and devices for the construction of user interfaces accessible by people with special needs.

The publication and dissemination of the deliverables of the IPSNI consortium is hoped to influence B-ISDN designers in designing terminals and services which are modularly adaptable, or directly accessible by most users with special needs. The IPSNI consortium partners are CNR-IROE (Italy), KUL (Belgium), Alcatel-FACE (Italy), FORTH-ICS (Greece), IRV (The Netherlands), University of Dundee (UK), VTT (Finland). ■

Please contact:
Constantine Stephanidis - FORTH-ICS
Tel: +30 81 226487
E-mail: cs@csi.forth.gr

VLSI Design

Testing Integrated Circuits

by Heinrich Theodor Vierhaus

Testing of complex integrated circuits has been a problem for a long time due to ever-increasing circuit complexity and the reduced controllability and observability of interior circuit nodes. Since the arrival of high-speed CMOS circuits, the general problem has got a new dimension due to the complex fault behaviour of this technology. The partners of the ESPRIT III Basic Research Project ATSEC (Advanced Test Generation and Testable Design Methodology for Sequential Circuits) have started a joint venture to find methods for automatic test pattern generation that provide a reasonable fault coverage for all realistic kinds of defects.

Despite all advances in IC technology, defect-free production runs are not yet possible. Inevitably a certain percentage of ICs, will have physical defects. Hence testing an IC before inserting it in a complex system is still a necessity. IC testing is an extremely expensive process. Therefore complex ICs can be of commercial value only if the test is short and cost-efficient, but still safe and reliable. To achieve this objective, methods of testable design have to be followed. Still then the generation of an efficient set of test patterns covering all relevant defects is a formidable task. Due to the complexity of possible fault behaviours, all CAD-methods supporting the generation and the validation of test patterns have to be based on fault models that simplify the actual fault behaviour.

In traditional approaches it is assumed that under fault conditions just one internal node is either solid "stuck-at-1" or "stuck-at-0". Almost every commercially available ATPG tools is based on this simple model. Thereby ATPG for combinational logic is well understood, but many problems remain to be solved for sequential circuits. In CMOS circuits, defects will often not cause a stuck-at-fault, but will result only in delayed or impossible transitions between specific logic states. Even worse, such defects may turn out as "intermittent faults". Hence test patterns have to be expanded into test sequences for specific transitions.

The combination of ATPG for sequential circuits and advanced fault models is now tackled by a consortium guided by GMD's CAT (Computer-Aided Test) group in the System Design Institute. The other partners are University of Duisburg, University of Twente (Enschede, Netherlands), Politecnico di Torino, Université de Monpellier II (Institute LIRMM), and Oxford University. Katholieke Industriele Hoogeschoool West-Vlaanderen, Oostende (Belgium) and the Slovak Technical University in Bratislava will contribute as associate partner and subcontractor, respectively with focus on prototype tool evaluation and testable design methods.

As a result, a modular ATPG system for transition faults and delay faults in sequential circuits is targeted, which will help to create reliable test patterns in a cost-efficient way and reduce the necessary overhead for testable design. ■

Please contact:
Heinrich Theodor Vierhaus - GMD
Tel: +49 2241 14 2633
E-mail: vierhaus@gmd.de

Computational Morphology

Closed Object Boundaries from Scattered Points

by Henk Nieland

New techniques have been developed for the construction of closed boundaries from scattered points lying on an object boundary. The work was completed at CWI as a part of the National Informatics Facility's IIICAD project.

The research, initiated at the University of Leiden and elaborated in a Ph.D. thesis by Remco Veltkamp (CWI and University of Eindhoven), concerns a basic problem in computational morphology (the analysis of form by computational means). In many applications in geometric modelling, computer graphics, object recognition, distance map image processing and computer vision, input data is available in the form of a set of 2D or 3D coordinates that are points on the boundary of an object. For many applications in fields such as the initial design of an artefact, numerical analysis or graphical display, it is essential to have an unambiguous representation of the whole boundary available, preferably at various levels of detail and smoothness. If we can utilise a priori knowledge about structural relationships between the points, for example if we know that the object is a car, considerable time may be saved in constructing such a representation. A serious drawback is, however, that we have to adapt the construction method for each case separately. Moreover, there are several instances where no such relationships are known (the points are "scattered").

Veltkamp has developed new general techniques to construct from a set of scattered points the closed boundary of an object and to manipulate with it. The set of points is structured by defining a suitable neighbourhood of each pair of points and interconnecting points if no other points lie in this neighbourhood. Based on this graph, an algorithm for the construction of a polygonal (2D) or polyhedral (3D) boundary is given. The algorithm gradually removes the "loose" connections in the graph – a method similar to a sculptor's job – ending up with a uniquely defined boundary. From this boundary a hierarchy of approximations is derived, which may have considerable practical importance. In real-time applications, such as a robot avoiding obstacles, a global impression of the object may suffice. In the design of cars, aircraft or ships, on the contrary, a realistic image is important. The latter applications also require smoothing techniques. Veltkamp has also developed new schemes for a smooth piecewise cubic boundary, by splitting Bézier triangles into sub-triangles. ■

Please contact: Remco Veltkamp - CWI
Tel: +31 20 592 4118 or +31 40 47 4579
E-mail: remco@cwi.nl

1468 scattered vertices

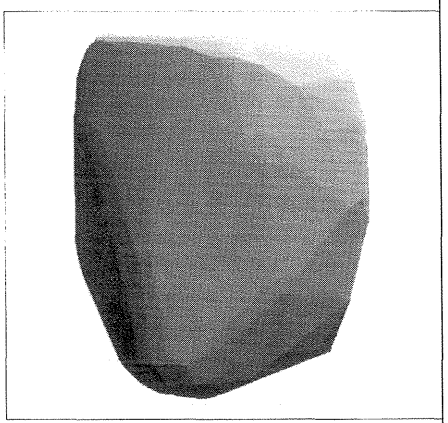

The convex hull, consisting of 255 vertices and 504 triangles

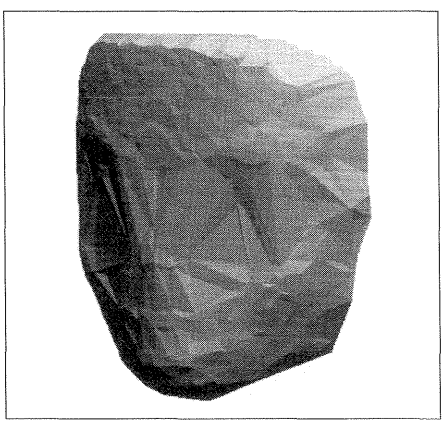

1019 vertices and 2034 triangles

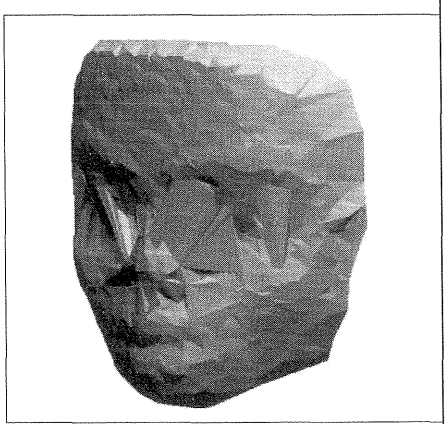

1337 vertices and 2670 triangles

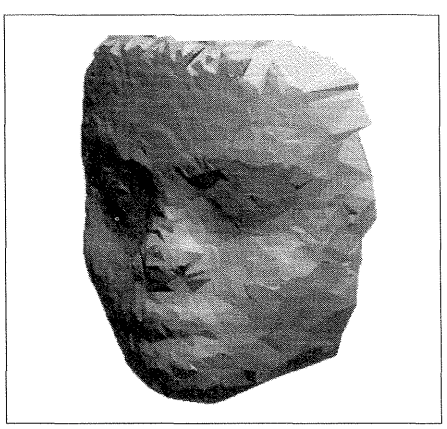

Final boundary with 1468 vertices and 2930 triangles

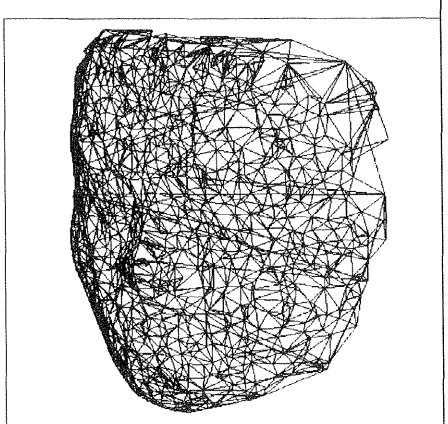

The corresponding underlying "wireframe"

Reconstruction of a mask from a collection of laser-range points (3-D).

Distributed Systems Techniques

DIVA Seeks to Create the Virtual Office

by Thomas Berlage

The DIVA (Dynamic Interfaces for Cooperative Activities) project has been started in the GMD - Institute for Applied Information Technology. The goal of the project is to create user interfaces for cooperative applications in distributed computer networks.

These applications are envisioned to reduce environment pollution caused by commuter traffic, as well as improve international cooperation, but a high-quality user interface is crucial for achieving these benefits. To transparently access cooperative facilities, the virtual office metaphor is used, which mimics social behaviour in the local environment. The GINA framework developed in an earlier project will be used to quickly prototype such an environment.

Many of the basic technologies for Computer Supported Cooperative Work, like broadband networks and multi-media workstations, have become available in the last few years. A lot of research work, especially in the user interface area, is still required to make them truly usable. The problem is to find a user interface that communicates a maximum of information through a rather small computer screen, but does not get in the way itself.

As a very natural metaphor, a virtual office represents the relevant structure of cooperation for each individual user. Coworkers can be located throughout the world, but are in close proximity in the virtual office. Communication is established by peeking into the other's room or by accident if two people work in the same project room on related documents. The goal is to mimic social behaviour in a real office to avoid the distracting effects of the telephone, for example. Communication itself is performed through automatically established audio/video links and shared multimedia applications.

The challenge for the user interface is to represent actions of co-workers so that they can be noticed subconsciously ("shared awareness"). There must be a transition from the limited personal desktop metaphor, where only one user induces changes, to a metaphor including the surrounding organisation composed of humans. History-based mechanisms are important in this context, as they allow notification to be deferred. Similarly, new animation mechanisms have to be developed to effectively visualise state transitions caused by other partners. The success of the metaphor depends on the integration of many detail solutions.

The project builds on previous experiences at GMD. The object-oriented application framework GINA, which was built by the DIVA group, will be used as the tool to rapidly prototype demanding graphical interfaces. Ongoing research in the Human Computer Interaction Group on ergonomic multimedia interfaces will also influence the project. DIVA will be part of the POLIKOM program launched at GMD.

■

Please contact: Michael Spenke - GMD
Tel: +49 2241 14 2642
E-mail: spenke@gmd.de

Networks & Security

CAFE: Conditional Access For Europe

by Henk Nieland

The future "open market" for supply and use of information services through a computer network requires special security and privacy measures. These aspects are dealt with in the ESPRIT project CAFE (1993-1995), coordinated by CWI.

By now it is technically feasible to construct information service systems, in which the demands of suppliers and users with respect to security and privacy can be satisfied. For realisation in practice some ingredients are needed. In the ESPRIT project CRYPTO IC-CARD, a secure micro-controller chip containing special hardware support for advanced "public key" coding techniques, was developed. Such a chip needs to be embedded in a miniature "pocket workstation" if it is to have display and data entry means to allow its owner to know and control what it does. Furthermore, sophisticated protocols need to be designed and implemented to realise the potential of such coding.

CAFE comprises of the following activities:

- research and design of public-key secure protocols;
- development of hardware and software for a trial by major financial institutions;
- modelling and specification;
- establishing an architecture for consensus forming and normalisation;
- assessment of needs and desires of end-users, service providers, financial institutions, public sector organisations and consumer interest organisations.

CAFE is a major step forward compared to, e.g., approaches to payment using prepaid IC cards, which are only suited to closed systems with low security. For conditional access more generally, very primitive calculator-like devices employing simple coding to secure the password process is all the state-of-the-art has to offer. CAFE aims to select, specify, construct, integrate, trial and evaluate a basic payment system, and to demonstrate feasibility and establish consensus within the European IT-industry. Cryptographic techniques achieving electronic privacy are described by David Chaum in the Scientific American, August 1992, p.96-101.

CAFE runs for three years (1993-1995), with an effort of 56 man-years. Participants are: CWI (coordinator), Digicash and PTT-NL (The Netherlands); Cardware (UK); Gemplus, SEPT and Ingenico (France); SINTEF-Delab (Norway); Institut für Sozialforschung Frankfurt, Institut für Informatik Hildesheim, and Siemens (Germany); and the Universities of Aarhus (Denmark) and Leuven (Belgium).

■

Please contact: David Chaum - CWI
Tel: +31 20 592 4169
E-mail: chaum@cwi.nl

Dynamical Systems

Aspects of Resonance in Dynamical Systems

by Igor Hoveijn

Various phenomena in nature can be modelled by Hamiltonian dynamical systems and their time-dependent perturbations. In a research project of NWO's National Activities in Mathematics programme, the global behaviour of such systems is studied by concentrating on the system's resonances.

Dynamical systems can be used to model various phenomena in nature, for example kinetics of chemical reactions, fluctuations in the number of a biological species, the motion of planets, atmospheric flow, electronic oscillators and vibrations of mechanical devices. The state of such a system is represented as a point in the phase space of all possible states. A differential equation (continuous time) or a difference equation (discrete time) on the phase space defines the evolution of the system.

In general it is much too difficult to give a global description of this evolution, not only because the phase space can be high-dimensional, but also due to the possible existence of chaotic phenomena. Therefore we first turn to a local description, in the neighbourhood of a stationary point. We simplify the system by a so-called normalisation transformation of the coordinates.Doing so we introduce an error and we consider the normalised system as a perturbation of the original system. The form of the normalised system is determined by the system's resonances.

In particular the study of Hamiltonian systems, consisting of a set of ordinary differential equations in the form of Hamilton's classical equations of motion, is of great practical and theoretical interest. As is well-known, for a system with one degree of freedom the classical Hamilton equations of motion can be solved by a canonical transformation to action-angle coordinates, the action being the sought-for "integral" of the system. Normal forms of Hamiltonian systems have at least two independent integrals. If a Hamiltonian system has as many integrals as degrees of freedom the system is called "integrable": all possible evolutions can be characterised. This is exceptional however. The question is: do we, by normalisation, simplify the system such that it becomes integrable, but at the same time throw away the interesting dynamics? The answer is "yes" for two degree of freedom systems. Therefore we have considered three degree of freedom systems. It is not a priori clear that a third integral exists. Hamiltonian systems can locally be labelled by the type of resonance. We have proved the non-integrability of the normal form of a particular resonance by considering the flow of the system on a subset in the phase space, where the system behaves as if it were integrable (see Figure 1).

A large class of problems can be characterised by periodically perturbed Hamiltonian systems. In particular we have considered a system of two coupled oscillators with damping and periodic forcing. There are many possible resonances, but we focused on the so-called sumresonance case, where the sum of the frequencies of the oscillators is equal to the frequency of the forcing. For a certain range of frequencies around the sumresonance the origin remains a stable point. By adding a linear damping term this range will change. However, by reducing the damping to zero we do not find back the original stability range. This non-uniform behaviour is caused by singularities on a certain "critical" surface in parameter space. Although the normal form of this system is somewhat simplified by a discrete symmetry group, we are left with a complicated bifurcation problem. We have determined the bifurcation diagram of the origin. The complicated dynamics of such systems is shown in Figure 2. ■

Please contact:
Igor Hoveijn - Utrecht University
Tel: +31 30 53 1527
E-mail: hoveijn@math.ruu.nl

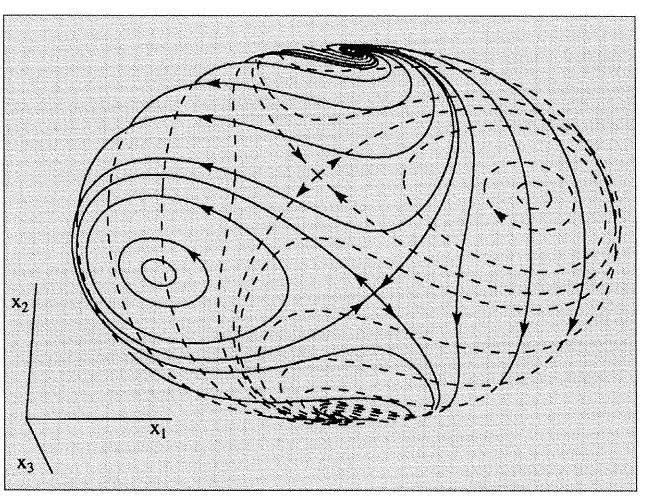

Figure 1: Orbits on a special set in phase space where the system behaves as if it were integrable.

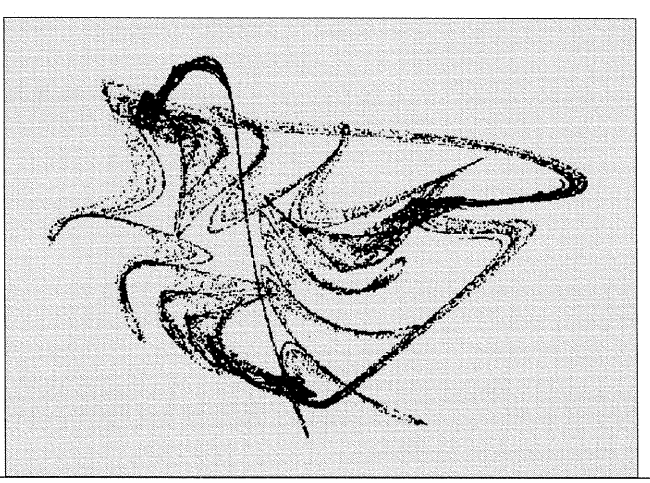

Figure 2: Stroboscopic picture in four dimensions of a periodically forced system of two coupled oscillators. The position in phase space is plotted every $2/\pi$ seconds.

Multiprocessor Modular Architectures for Diagnostic Imaging

by Luciano Azzarelli, Ovidio Salvetti, Enzo Gontero and Dario Nari

A multiprocessor multibus system, with parallel and real-time processing capabilities that can be used for diagnostic imaging, has been implemented in a collaboration between IEI-CNR and Scriba srl, Turin. The system architecture has been designed to optimize applications which need real-time image analysis.

Feature classification is an important and common problem in image processing. Depending on the complexity of the image to be analysed, classification requires fast memory for data storage and considerable computing time for algorithm execution. If the application requires a real-time feature extraction and classification, then high-performance computing systems must be used. Two levels of parallelism are required: an algorithmic level, to execute the image analysis algorithms; a processing level, to control the different operational stages of the process, e.g. multidimensional signal acquisition, data management and interaction with all other components of the production chain.

Real-time feature classification offers undoubted advantages to many industrial applications because it permits full quality control during the production process. A multiprocessor system can perform pattern segmentation and analyse the digital images acquired using different or parallel transducers.

Within the CNR Finalized Project "Information Systems and Parallel Computing", we have developed a system architecture that can manage both system-dependent and application-dependent aspects. In particular, the relations between the system and the external environment depend on: the data category, which determines interconnection and integration schemes between different kinds of information; the processing module and procedure category, which determines operating system capabilities for information and process management; the intelligent process category, which defines the system knowledge for data interpretation and understanding; the communication category, which defines the interface reference models to standardize data, information and process exchange.

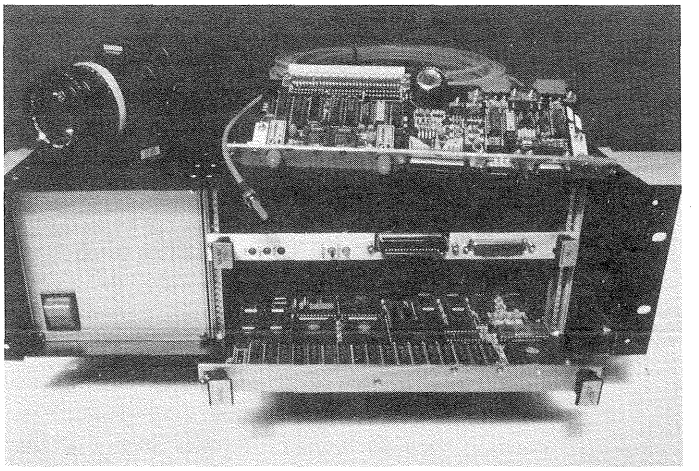

Figure 1: Some hardware modules of the multiprocessor multibus system for diagnostic imaging. (Photo: IEI-CNR)

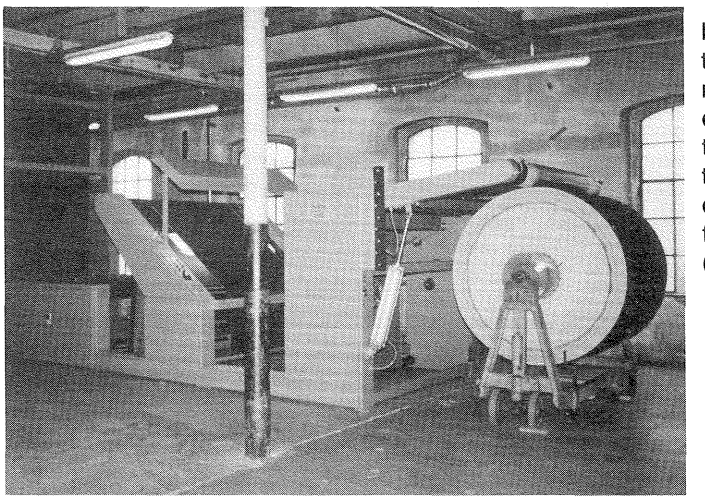

Figure 2: Part of the prototype machine developed to perform real-time quality control of denim fabrics. (Photo: IEI-CNR)

The main components of the software system are: a high-level man-machine interface; an image processing machine; a system-to-system and system-to-host communication kernel. Figure 1 shows some hardware modules of the system.

A specialization of the architecture has been experimented in the EC-BRITE Project: "An automatic on-line system for the detection, evaluation and mapping of defects and the monitoring of variations in finished fabrics", where a system for fabric quality checking has been developed.

The main functions executed by the system for defect recognition and classification are data acquisition, photometric correction, spatial and photometric segmentation, compacted data analysis and defect classification. Defect recognition, which is performed on a statistical basis, is preceded by a self-learning phase in which the specific properties of the material under investigation are defined. During acquisition, images of the fabric are transformed into data matrices by analog to digital conversion. The input data are processed to eliminate any artifacts produced during this stage. The data are then re-processed to evidence the presence of defects. Figure 2 shows part of the machine which has been developed for real-time quality control of denim fabrics. ■

Please contact:
Luciano Azzarelli - IEI-CNR
Tel: +39 50 593421
E-mail: azzarelli@vm.iei.pi.cnr.it
or: Ovidio Salvetti - IEI-CNR
Tel: +39 50 593455
E-mail: salvetti@vm.iei.pi.cnr.it
or: E. Gontero, D. Nari - Scriba srl, Turin
Tel: +39 11 728069
Fax: +39 11 793490

The TeleServe RDBMS Server

by Petter Moe and Per Holager

TeleServe Transaction Technology is a start-up company undertaking to produce continuously available SQL servers of record-breaking speed and capacity.

The business idea is to market these for use in the telecommunications network, particularly in "Intelligent Network" applications. The servers may hold the billing repository, the location register for mobile telephony and Universal Personal Telephone services, etc. The main efforts of TeleServe are directed towards developing the SQL DBMS software, the necessary hardware is to be produced by a partner company. TeleServe is a spin-off from joint projects at the Norwegian Institute of Technology (NTH) and its associated contract research organisation (SINTEF), in close cooperation with the Norwegian Telecom Research.

One TeleServe machine will consist of tens or hundreds of nodes. Each node will consist of a computer board with a moderate number of hard-disc units. Each node will have a handful of communication lines to other nodes. The components chosen are high-performance, mass produced units. The nodes are divided in two groups to be placed in separate locations, to prevent one disaster from stopping the system. Each node in one group should have a direct line to a "buddy" in the other, implying a formidable total communications capacity between the two sites. All nodes will execute the same software, but in a highly asynchronous fashion, exchanging messages containing commands, data and status information.

Applications will see this system as one integrated ISO SQL server with a number of access ports. The server is specifically designed to meet the fault-tolerance requirements of critical telecommunications application, i.e. an accumulated "down-time" of thirty minutes over an operative period of thirty years, at the user service level. One system will typically have a throughput of thousands of transactions per second. At the same time, it must be able to respond to real-time requests with ten milliseconds. Such extreme requirements can not be met by ameliorating existing DBMS designs, they have to be considered in every step of the design and implementation process. Modern data base algorithms have been enhanced to utilise the TeleServe hardware structure.

For failure tolerance, a primary and at least one backup copy is kept of every row of data. A backup copy will normally be kept on the "buddy" of the node holding the primary copy. A hashing scheme is used for distributing the primary copies of the rows of one (large) table over all nodes in the system. This distribution permits parallel execution of large jobs and gives load balancing over the nodes: Whenever possible, operations are executed on the node holding the data.

Basically, all correction of failures is done by "rolling back" the affected transactions and letting the software try again. Switching to backup data and backup processes may precede this. When a primary copy is permanently lost, a new backup will be produced by a background process, while the one existing copy is available for service. All hardware and software components in the system can be replaced with only a limited reduction of the system capacity. Failure detection is done along several "lines of defence". First, Ada is used to increase the system's ability to detect programmer errors. Then a process model with protection is used to let the operating system track illegal accesses. Lastly, messages undergo extensive consistency checking, both by the sending, transmitting and receiving processes. ■

Please contact:
Per Holager - SINTEF DELAB
Tel: +477 59 42 09
E-mail: Per.Holager@delab.sintef.no
or: Petter Moe - TeleServe AS, Norway
Tel: +477 59 65 40
E-mail: petter.moe@teleserve.no

INESC Produces 5000 Boards for IBM

by João Bilhim

Following the signing of a technology transfer contract, INESC will be supplying 5 thousand boards to IBM. The first 250 "backplane" adaptor boards were shipped last week from INESC's Electronics Production Centre to IBM's production facilities in Greenock Scotland.

For INESC this is a "small-giant step" as this is the first time it has supplied IBM with series production as well as being the first time IBM has received "non-clean" technology from a European supplier, i.e. technology not requiring washing and therefore environment friendly which is a sign of the growing ecological concerns of manufacturers of electronics products.

The main objective of the IBM contract is the transfer of technology with the objective of providing support to small and medium sized Portuguese companies having or acquiring the capacity to be IBM suppliers.

INESC's Technology Centre for Electronic Production will initially transfer less complex technology and help Portuguese industry to improve the quality of its production. When the small and medium sized companies attain a specific level of quality and have been properly certified they can become IBM suppliers or supply other multinationals on their own account.

INESC has manufactured the boards to IBM specifications and has also produced the customised software and hardware for testing requirements for the multinational company in question. ■

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

COST 229 WG4: Workshop on Massively Parallel Computing

Funchal, Madeira, 15-17 April 1993

The aim of this workshop is to promote the exchange of experiences and the interaction between groups involved in the areas of Massively Parallel Computing.

The Workshop will be organised as follows:

- Invited tutorials providing the state of the art in selected areas;
- Presentations containing new results of research work;
- Round tables to discuss and present open problems.

■

Please contact: Antonia Casimiro - Universidade de Madeira
Tel: 351 91 22 24 18/8
Fax: 351 91 23 02 43

4th Eurographics Workshop on Visualization in Scientific Computing

Abingdon, UK, 21-23 April 1993

Over the last three years, the Eurographics Working Group on Visualization has organised successful annual workshops at Clamart (France), Delft (Netherlands) and Viareggio (Italy). The next meeting will be held at Cosener's House, a conference centre owned by RAL (UK). The Workshop is an activity of the Eurographics Working Group on Visualization in Scientific Computing, and is organised jointly by RAL and the Numerical Algorithms Group Ltd.

The aims of this workshop, the 4th in the series, are:

- to present up to date results
- to promote exchange of views between researchers and users and
- to stimulate further advances in the field of visualization.

The format will consist of a number of selected presentations interspersed with focussed discussion/work sessions.

Topics:

- Visualization techniques
- Data interpretation
- The human interface to visualization systems
- Reference models, system architectures and standards for visualization
- Applications and case studies of visualization

■

Please contact: Julian Gallop - RAL
Tel: +44 235 821900
Fax: +44 235 445831
E-mail jrg@inf.rl.ac.uk

REX – a Decade of Concurrency

Noordwijkerhout, The Netherlands, 1-4 June 1993

This event marks the end of a ten year period of cooperation on the foundations of concurrency between CWI and the Universities of Leiden and Eindhoven. The REX-project (Research and Education in Concurrent Systems) was initially funded by the Netherlands Research Foundation for Computer Science (SION), and from 1988 on by the National Informatics Facility programme. Its goal was to contribute to the cross-fertilisation between formal methods from the fields of syntax, semantics and proof theory.

The preceding six schools/workshops covered the following topics (in brackets the proceedings' serial numbers, all published as Springer Lecture Notes in Computer Science):

- 1985 Current Trends in Concurrency (LNCS 224)
- 1988 Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency (LNCS 354)
- 1989 Stepwise Refinement of Distributed Systems — Models, Formalisms, Correctness (LNCS 430)
- 1990 Foundations of Object-Oriented Languages (LNCS 489)
- 1991 Real-Time: Theory in Practice (LNCS 600)
- 1992 Semantics: Foundations and Applications (to appear)

At the closing event, taking place under the auspices of the European Association for Theoretical Computer Science, 16 leading researchers from the USA and Europe (including ERCIM partners INRIA and SICS) will present reflections and perspectives concerning the field of concurrency. Directors of the school are: J.W. de Bakker (CWI and Free University, Amsterdam), W.P. de Roever (Christian Albrechts Universität, Kiel) and G. Rozenberg (Leiden University).

■

Please contact: Mieke Bruné —CWI
Tel: +31 20 592 4249
E-mail: mieke@cwi.nl

Advanced Workshop on Programming Tools for Parallel Machines

Lecce, Italy, 23-25 June 1993

The purpose of the workshop is to present and discuss the state-of-the-art and trends of programming tools – languages, operating systems, run-time and measurement systems, etc. – for parallel machines. The workshop is aimed at all those interested in parallel processing. The programme of the workshop includes eleven invited speakers. A concluding panel session in which all the speakers will participate is to be broadcast by satellite in Europe and USA/Canada. The Workshop is jointly sponsored by CNR and the University of Lecce.

The following papers will be presented:

- Introduction and Overview of Models and Languages: D. Skillicorn, Canada
- Tools and Compilers for sparse unstructured and adaptive problems: J. Saltz, USA
- Collection-oriented Models: J. Sipelstein, USA
- Functional Languages - A Concrete Example: SISAL: J. McGraw, USA
- Message Passing Environments: D. Walker, USA
- Shared Memory Environments: B. Bershad, USA
- Portable Parallel Programming Environments: A. Hey, UK
- Are Models and Hardware Converging?: D. May, UK
- Applications on High Performance Fortran: G. Fox, USA
- Heterogeneous Network Computing: J. Dongarra, USA
- Advanced Computing Environments: Future Trends: P. Messina, USA

■

Please contact:
Shona Dryburgh - TRE EMME Congressi
Tel: +39 50 44154
Fax: +39 50 500725
or: Domenico Laforenza - CNUCE-CNR
Tel: +39 50 593270
E-mail: free@icnucevm.cnuce.cnr.it

COSIT'93: European Conference on Spatial Information Theory

Elba Island, Italy,
19-22 September 1993

The COSIT Symposium will concentrate on theoretical aspects of spatial information and deal with all aspects of "large scale" or "geographic" space and description of objects, processes or events in large scale space. Spatial information theory is the basis for the construction of GIS but necessary for other uses of geographic information. The use of the technology reveals a large number of interesting research questions, which require an interdisciplinary approach to their solution. COSIT intends to bring together researchers from many different disciplines.

Topics:

- Structure of geographic information
- Languages of spatial relation
- Cognitive structure of space
- Time in geography and geographic information
- Spatial and temporal reasoning in geographic space

- Data quality aspects of geographic information
- Treatment of incomplete or imprecise spatial data

- Computational geography
- Spatial analysis and modeling
- Simulation of processes in geographic space

- User Interface design
- Design of generally useful elements for GIS interfaces
- Metaphors for GIS
- Virtual Reality

- Presentation of spatial information
- Cartographic generalization

- Spatial data integration
- Cooperative work with spatial information

The conference will be the place to present results from recent work and discuss advanced "work in progress". Three invited lectures by eminent scientists in the field will be held. One day will be devoted to papers discussing aspects of cultural differences in spatial information treatment. Panels will discuss relevance of research questions and compare possible approaches.Before the conference, tutorials introducing the topics of the conference will be organised.

■

Please contact:
Irene Campari - CNUCE-CNR
Tel: +39 50 593320
E-mail:irene@Icnucevm.cnuce.cnr.it
or: Andrew Frank - Technical University Vienna
Tel: +43 1 588 01 3786
E-mail: afrank@email.tuwien.ac.at

8th International Conference on Finite Elements in Fluids

Barcelona, Spain,
20-24 September 1993

The original objective of this conference series was to bring together researchers engaged in the application of finite element methods to the solution of fluid flow problems in all areas of science and engineering. The expansion of the subject since that time has drawn together many related methodologies. The main goal of the 8th Conference on Numerical Methods in Fluids is to address the outstanding challenges in flow simulation.

From the strong belief that the subject progresses through interaction between scientists and engineers, the organisers are convinced that an efficient way to address the above challenges is to bring together all researchers engaged in the applications of modelling and computational methods to the solution of flow problems.

Topics:
Papers related to the conference objectives will fall into the following subject groups:

- Acoustics
- Biological flows
- Combustion
- Euler solvers
- Error estimation and adaptivity
- Finite volumes
- Flow control
- Fluid-structure interaction
- Free surface flows
- Geomechanic flows
- High performance computing in massive parallel machines
- Lagrangean methods
- Magnetohydrodynamics
- Mesh generation
- Metereology
- Multigrid methods
- Multiphase flows

- Navier Stokes solvers
- Non Newtonian flows
- Optimum design and inverse problems
- Reacting flows
- Shallow water
- Time integration schemes
- Turbulence models
- Waves

Please contact: FEMIF'93 Secretariat
Tel: +34 3 205 70 16
Fax: 34 3 401 65 17

Future Trends of Distributed Computing Systems

Lisbon, Portugal,
22-24 September 1993

The fourth workshop in this series will deal with all aspects of distributed computing systems, especially future trends of various aspects of distributed computing. Previous workshops were held in Hong Kong (1988), Cairo (1990) and Taipei (1992). The workshop is sponsored by the IEEE Computer Society Technical Committee on Distributed Systems.

Topics:

- High Speed Networking and Protocols
- Real Time Distributed Systems
- Multimedia Systems
- Systems Architecture
- Distributed Databases
- Software Engineering for Distributed Systems
- Formal Techniques and Modelling
- Performance Evaluation
- Distributed Software Infrastructures
- Intelligent Distributed Systems
- Cooperating Autonomous Systems

Please contact: Paulo Verissimo - INESC
Tel: +351 1 3100 000
Fax: +351 1 525 843
E-mail: trends@inesc.pt

CALL FOR PAPERS

EP94: Electronic Publishing, Document Manipulation and Typography

Darmstadt, Germany,
13-15 April 1994

This conference will be the fifth in a series of international conferences organised to promote the exchange of novel ideas in the area of computer manipulation of documents. It will take place during the same week and at the same location as the third Conference on Raster Imaging and Digital Typography (RIDT), the TEP94 workshop (Teaching Electronic Publishing and Digital Typography), and the PODP94 workshop (Principles Of Document Processing).

The first two conferences in the series, EP86 held in Nottingham, England, and EP88 in Nice, France, concentrated mainly, on the specific aspects of the production of documents by computer, from composition to printing. EP90, which was held in Washington D.C., adopted a broader definition of the term Computer Assisted Publication, and accordingly, included more materials on new topics such as the application of data-base techniques to document handling, hypertext and hypermedia systems, and document recognition and analysis. EP92, held in Lausanne, Switzerland, confirmed the trend for documents to impact more and more areas in computer science. EP94 is expected to follow this trend.

The Conference is sponsored by ERCIM.

Topics:
Authors are invited to submit papers describing original research results on the following topics:

- Modelling and representation of documents; Document structures; Integration of text, images, graphics, sound and video; Integration of document manipulation systems with other software tools; Standards: evaluation and implementation; Active documents, documents as user agents; Object oriented approaches.
- Document management; Document preparation systems; Hypermedia and multimedia: production, editing and visualisation; Large document collections; Distributed documents: parallel algorithms, multi-user documents; User-machine interfaces.
- Document recognition and interpretation; Structural recognition of documents; Filtering and image handling techniques; Multi-lingual documents; Semantic document structures; Indexing and retrieval techniques.
- Typography and graphics; Graphics and imaging; Pagination and layout.

Deadlines:
16 August 1993: Full papers
15 October 1993: Notification of acceptance or rejection
3 December 1993: Final versions of accepted papers

Please contact: EP94 - GMD
Tel: +49-224114-2473
Fax: +49 224114-2618
E-mail: ep94@gmd.de

Multigrid Course

by Wolfgang Joppich and Barbara Steckel

Multigrid methods have been established as being among the most efficient techniques for solving complex elliptic equations: in general, a multigrid solution is obtainable in a number of operations proportional to the number of unknowns. Multigrid methods are not only very fast on sequential computers, their inherent locality also allows an efficient parallelization. The importance of multigrid was also demonstrated by the fact that for a Multigrid Course, held at the GMD, 22 - 26 June 1992, more applications were received than the number of available places.

The main purpose of the course was to inform on multigrid approaches and to

provide an understanding of basic multigrid principles. On completing the course the participants should have been able to write simple multigrid programs and to understand advanced literature. In addition, an overview of multigrid applications and recent research activities was given. The course was designed for those interested in solution methods for partial differential equations, including numerical analysts, physicists, engineers and those who need to decide on concepts for large numerical software packages.

In the lectures the following topics were covered: basic elements, performance analysis, adaptive methods, parallel aspects, multigrid-software and non-PDE multilevel computations. Multigrid programs developed at the Institute II of GMD were presented, e.g. for process simulation and for the incompressible Navier-Stokes equations. Copies of the slides are available as GMD-Arbeitspapier 690.

Please contact: Wolfgang Joppich or Barbara Steckel - GMD
Tel: +49 2241 14 2748 or 14 2768
E-mail: Barbara.Steckel@gmd.de

IN BRIEF

GMD – **GMD took over the convenorship** of the Working Group of Security Technology (WG-SEC) of RARE, the association of the European academic research networks. The overall objective of this group is to promote the implementation of security services in European networking, especially through pilot projects. WG-SEC had its first meeting in early November 1992 in Pisa, Italy, together with the European Network Service Conference.

The objectives of this group are to act as a focus for user requirements for security technology and secure applications. Pilot services will be promoted which demonstrate the viability of security techniques, i.e., techniques for authentication, identification, confidentiality and mechanisms for establishing hierarchies of certification authorities.

RAL – **It has recently been confirmed that a consortium** led by Longman Cartermill, and including library and technical staff from CWI and RAL, has been awarded a contract under the Libraries Section of the CEC Telematics Programme. The RIDDLE project, expected to last 20 months, will investigate the feasibility of using automatic methods of capturing the contents pages of technical journals with a view to including the information on the individual items in each issue in an online library catalogue. The results will not be specific to any one library catalogue but will be widely applicable. Scanning and OCR techniques will effect the data entry, and the Standard Generalised Markup Language (SGML) used as the medium for tagging the contents page entries prior to mapping these into the catalogue. The tagging will be carried out by program, making use of presentation cues such as position, font size and type. It is anticipated that some pilot development will be undertaken, involving the library catalogues at CWI and RAL. Further information can be obtained from Frank Roos, CWI, or Eric Thomas, RAL.

INRIA – **A framework agreement** of scientific cooperation was signed by INRIA and CRIM (Centre de Recherches en Informatique de Montreal, Quebec). The main topics of this cooperation will be robotics and software engineering.

ILC-CNR – **The kick-off meeting of the Management Board of EAGLES** (Expert Advisory Group on Language Engineering Standards), a new CEC initiative in the framework of the Linguistic Research and Engineering (LRE) programme, was held 7-8 January 1993. The aim is to accelerate the provision of common functional specifications for the development of large-scale linguistic resources, such as electronic text collections (corpora), computational lexica, formalisms, speech, and also evaluation and assessment.

EAGLES is coordinated by Prof. Antonio Zampolli and counts on the active participation of more than 30 research and industrial organizations from most EC countries. The Group is to present its findings in the form of guidelines in Summer '95. For further information please contact Nicoletta Calzolari, ILC-CNR (E-mail: eagles@icnucevm.cnr.pisa.it).

CWI – **Proceedings SCAFI'91 appeared.** The first of a series of seminars dedicated to Studies in Computer Algebra for Industry, was held in December 1991 at CWI, followed in 1992 by three similar courses (Bath, Brussels, Amsterdam). Proceedings of the first seminar, edited by Arjeh M. Cohen (CWI), have now appeared ("Computer Algebra in Industry: Problem Solving in Practice", Wiley, ISBN 0-471-93829-7). The SCAFI seminars are organized in the framework of the ERCIM Advanced Training Programme.

INRIA – The rector of Moscow State University, **Prof. Victor Sadovnichy**, visited INRIA on January 29. He discussed the possibility of creating a "Franco-Russian Center in Applied Mathematics and Computer Science" with **Prof. Alain Bensoussan**, President of INRIA, which should be located inside the building of Moscow University.

INESC – The Swedish Minister of Science and Technology visited INESC, Lisbon, on 9th November last. **Mr Per Unckle**, accompanied by the Portuguese Minister of Science and Technology, visited the Technology Centre for Electronics Production and the Centre for Telecommunications Operations and Management, the MONICAP Project and the Centre for Integrated and Multimedia Electronics Systems.

INESC Directors, **Prof. Lourenco Fernandes**, **Prof. Alves Marques** and **Prof. Luis Vidigal** accompanied the visit in question. Professor Lourenco Fernandes suggested the inclusion of Swedish companies in the Lisbon Science and Technology Park as a manner of encouraging cooperation between the R&D institutes of both of countries. Mr Per Uncle, in turn, expressed the opinion that this visit provided a good opportunity for the start of closer and more regular cooperation between the two countries.

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 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:	Email:	Telephone:
Helena du Toit	helena.dutoit@inria.fr	+33 1 39 63 55 39
Local Editors:		
Lars Bergman (SICS)	lars@sisu.se	+46 8 752 1613
João Bilhim (INESC)	jafb@eniac.inesc.pt	+351 1 572 895
Wernher Behrendt (RAL)	wb@inf.rl.ac.uk	+44 235 44 5281
Truls Gjestland (SINTEF)	truls.gjestland@delab.sintef.no	+47 7 59 26 45
Siegfried Münch (GMD)	siegfried.muench@gmd.de	+49 2241 14 2303
Pierre Népomiastchy (INRIA)	pierre.nepomiastchy@inria.fr	+33 1 39 63 5646
Henk Nieland (CWI)	henkn@cwi.nl	+31 20 592 4092
Carol Peters (CNR)	carol@vm.iei.pi.cnr.it	+39 50 593 429
Constantine Stephanidis (FORTH)	cs@csi.forth.gr	+30 81 22 64 87

Centrum voor Wiskunde en Informatica

Kruislaan 413
NL-1098 SJ Amsterdam

Tel: +31 20 592 9333
Fax: +31 20 592 4199

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
Postfach 1316
D-5205 Sankt Augustin 1

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 7521520
Fax: +46 8 7527230

Stiftelsen for Industriell og Teknisk Forskning ved Norges Tekniske Høgskole

SINTEF DELAB
N-7034 Trondheim

Tel: +47 7 59 26 00
Fax: +47 7 59 43 02

Directeur de Publication: A. Michard
Publication périodique réalisée par GEIE-ERCIM
Dépôt légal 31/03/93