


EDITORIAL



**Prof.dr. Martin Kersten (CWI):
Gold nuggets in Database Technology are
widespread over Europe.**

SPECIAL :

Activities in Databases.....9

CONTENTS

Editorial	1
Joint ERCIM Actions	2
The European Scene	8
Research Activities:	
Databases	9
Other activities	31
Technology Transfer	34
International Relations	35
Events	36
In Brief	39

Next Issue:

Human-Computer Interaction

Database technology for sequential processor architectures and for administrative applications is nowadays well-understood. The technology has moved from file processing by large numbers of individual programs to data models and data manipulation schemes analysed with mathematical rigor and encapsulated in portable database management systems ready for use.

Moreover, the trade-press covers relational systems for parallel processing platforms and object-oriented systems to deal with engineering applications. Deductive systems for knowledge-based processing are shipped (or flown in) from abroad. One may thus be tempted to conclude that the field has reached its technical summit and that what's left is merely an engineering art or a matter of salesmanship.

Despite its maturity as a technical field with a high economic impact for administrative systems, several core problems are not yet resolved. The kaleidoscopic picture presented in this issue of ERCIM News is proof of the major strands of open research issues, which are addressed both by theory and experimentation. For example, better understanding of the semantic intricacies of information systems in real-life situations is tackled by re-use of results from computational linguistics. Likewise the hard performance requirements posed by telecommunication systems require novel techniques in DBMS algorithms and architecture to evolve and to adapt to the quickly changing hardware scene. Generic real-time and multi-media DBMS packages are yet to be produced.

Users of the few relational product families - recognized as major market players - are not immune to these trends. They require increased functionality at reduced cost, be it support for images, efficient query processing for decision support systems or data mining applications. And the trade-press resonates with questions like: Is object-orientation the answer? Are the new products stable and performant? Or, for a more evolutionary track: When will the new version be available on the new hardware box? When will distributed multi-level transaction processing be realized?

Indeed, the challenge is to assess early and thoroughly the potential of new product offerings and to separate the marketing hype from the gold nuggets. The latter are often found in small streams in the back country, rather than in waterfalls surrounded by PR-driven tourism.

This issue underscores the wealth of expertise in database technology in Europe, readily available to tackle real-life problems and to work on new DBMS products with high economic potential. Don't hesitate to call upon the imagination and dedication of its researchers.

Martin Kersten

AEDIMA joins ERCIM

by Pere Brunet

AEDIMA (the Spanish Association for Informatics and Applied Mathematics) is a consortium of eleven Spanish universities and four research institutes of the CSIC (Consejo Superior de Investigaciones Científicas).

AEDIMA was founded in 1992, ending a process that had started in 1990. Indeed, at that moment ERCIM had several contacts with different institutes and universities from Spain, in order to incorporate a Spanish organization into the European consortium. It was concluded that no existing research organisation in the field was appropriate to be the sole Spanish representative. Consequently, a meeting of delegates from several Spanish universities and the CSIC was held at the Polytechnical University of Catalonia on 11 December 1991; Alain Michard was invited to the meeting. The idea of founding AEDIMA and joining ERCIM was agreed.

AEDIMA's goals include the promotion of research in informatics and applied mathematics, and the cooperation between Spanish universities and public research institutes. AEDIMA has a Board of Directors and a Scientific Board. The President of the Board of Directors represents AEDIMA at the ERCIM Directors Meeting. The President of the Scientific Board is the representative at the ERCIM Executive Committee. AEDIMA members signed the founding agreement between June and October 1992; during a meeting on November 24, 1992, the agreement for the AEDIMA constitution was signed and the Spanish consortium was formally established. Cor Baayen, Bob Hopgood and Alain Michard attended the meeting. AEDIMA became an official ERCIM partner at the Røros Directors Meeting (28 May 93). Presently, AEDIMA is in the process of becoming an AIE (Asociacion de interes economico) in order to join ERCIM-EEIG.

AEDIMA's universities and research Institutes are spread geographically across the whole of Spain, and they cover almost all of the ERCIM areas of research.

AEDIMA universities include the University of Barcelona (UB), the Autonomous University of Barcelona (UAB), the Polytechnical University of Catalonia (UPC), the University of Granada (UGR), the University of Malaga (UM), the University of Madrid -Complutense- (UCM), the Polytechnical University of Madrid (UPM), the University of Santiago de Compostela (USC), the University of Sevilla (US), the University of Valladolid (UVA) and the University of Zaragoza (UZ).

The AEDIMA research institutes include the Instituto de Automatica Industrial CSIC-IAI, the Instituto de Cibernetica CSIC-IC, the Centro de Estudios Avanzados de Blanes CSIC-CEAB and the Centro Nacional de Microelectronica CSIC-CNM.

The next paragraphs contain a short introduction to the structure and main research topics of AEDIMA's members. Presented figures correspond to 1991 information.

The University of Barcelona (UB) has a total of 81223 students and 3285 lecturers (375 full professors). Its annual budget is 26641 Mpesetas. Their main research areas related to ERCIM topics include Computational Linguistics, System and Control Theory, Numerical Analysis and Numerical Linear Algebra.

The Autonomous University of Barcelona (UAB) has a total of 34509 students and 2021 lecturers (218 full professors). Its annual budget is 14183 Mpesetas. Their main research areas related to ERCIM topics include Complexity and Algorithms: neural networks; Computer Architecture: support for operating systems and platforms, performance evaluation; Computer Graphics and relationship with Imaging; Distributed System Techniques: distributed operating systems, performance evaluation, satisfaction,

fault-tolerant systems; Image processing and Computer Vision: geometrical analysis and modelling, medical imaging, mathematical morphology; Numerical Linear Algebra; Parallel Systems: generic architectures, scheduling techniques, environments; Scientific Computation, Optimization; Signal Processing: acquisition and processing in biomedical fields; Symbolic Computation: group theory and computer algebra; System and Control Theory: advanced process control, predictive control, modelling and identification, fuzzy modelling; VLSI Design: design for testability and quality, CMOS technologies, simulation of semiconductor devices, silicon computers, novel architectures; Knowledge Engineering: object based representations, user modelling.

The Polytechnical University of Catalonia (UPC) has a total of 35714 students and 2282 lecturers (190 full professors). Its annual budget is 15000 Mpesetas. Their main research areas related to ERCIM topics include Complexity and Algorithms: average analysis and structural aspects; Computer Architecture; Programming Tools and Operating Systems for Multiprocessors; Distributed Systems Architectures; Computer Graphics: geometric modelling, rendering, visualization, volume modelling, constraint-based CAD systems, distributed systems; Computational Geometry; Concurrency; Databases: data models, database design, interoperable databases; High Speed Networking; Image Processing and Computer Vision; Information Systems: conceptual modelling of information systems; Knowledge Engineering; Fuzzy Mathematics and Artificial Intelligence; Natural Language Processing; Numerical Linear Algebra; Dynamic systems; Parallel Systems; Algorithms and Parallel Architectures; Scientific Computation: finite element methods, structural stability; Signal Processing; Formal Methods and Tools for Software Engineering; System and Control Theory: advanced systems for process control, modelling and simulation of system dynamics, robotics, bio-

engineering; VLSI Design and VLSI Architectures.

The University of Granada (UGR) has a total of 55531 students and 2986 lecturers (205 full professors). Its annual budget is 19262 Mpesetas. Their main research areas related to ERCIM topics include Complexity and Algorithms, Computational Linguistics, Decision Support, Image Processing and Computer Vision, Knowledge Engineering, Numerical Linear Algebra, Signal Processing, Software Engineering, VLSI Design.

The University of Malaga (UM) has a total of 24058 students and 1160 lecturers (105 full professors). Its annual budget is 7337 Mpesetas. Their main research areas related to ERCIM topics include Scientific Computation: differential equations and dynamic systems, finite element methods; Numerical Linear Algebra; Parallel and Distributed Systems, Systems Engineering; Knowledge Engineering; Performance Analysis; Image Processing and Computer vision; Databases; Computational linguistics; Interactive systems; Computer graphics.

The University of Madrid - Complutense- (UCM) has a total of 114735 students and 5312 lecturers (550 full professors). Its annual budget is 33388 Mpesetas. Their main research areas related to ERCIM topics include Concurrency: formal models, semantics, time and probabilities; Computer Architecture: automatic high-level synthesis of digital design, distributed parallel simulation; Logic Programming: definition and implementation of declarative languages; Software Engineering: knowledge based software engineering, reusability; System and Control Theory: mathematical models and process modelling.

The Polytechnical University of Madrid (UPM) has a total of 45509 students and 3020 lecturers (330 full professors). Its annual budget is 20570 Mpesetas. Their main research areas related to ERCIM topics include Systems and Network Security; Telematic Networks; Computational Geometry; Concurrency; Databases: mathematical modelling, deductive databases; Information Systems; Knowledge Engineering, Logic Programming, Neural Networks, Fuzzy Logic, Models for Knowledge Bases Verification;

Symbolic Computing; Numerical Linear Algebra, Orthogonal Polynomials; Parallel Systems; Finite Element Methods; Fault Tolerant Programming; Distributed Systems Programming; User Interfaces; Software Engineering; System and Control Theory; Design of Adaptive Control Systems.

The University of Santiago de Compostela (USC) has a total of 33339 students and 1675 lecturers (184 full professors). Its annual budget is 12402 Mpesetas. Their main research areas related to ERCIM topics include Scientific Computing: asymptotic methods in solid mechanics, numeric methods for PDEs, optimal control and design of distributed systems, free boundary problems; Knowledge Engineering and Applications; Man-machine Interaction Tools; Neural Networks and Diffuse Logic; Bio-engineering; Advanced Architectures and Hypercube Computers; Microelectronics and CMOS Technology; Multidimensional Signal Processing; Computational Linguistics.

The University of Sevilla (US) has a total of 55949 students and 3131 lecturers (241 full professors). Its annual budget is 18687 Mpesetas. Their main research areas related to ERCIM topics include Differential Equations, Numeric Simulation and Software Development; Biomedical Engineering; System and Control Theory; Dynamic Engineering Systems.

The University of Valladolid (UVA) has a total of 35858 students and 2200 lecturers (144 full professors). Its annual budget is 11439 Mpesetas. Their main research areas related to ERCIM topics include Complexity and Algorithms, Image Processing and Computer Vision, Scientific Computation, Signal Processing, and System and Control Theory.

The University of Zaragoza (UZ) has a total of 40841 students and 2275 lecturers (167 full professors). Its annual budget is 16091 Mpesetas. Their main research areas related to ERCIM topics include Complexity and Algorithms: Design and analysis of algorithms in distributed and parallel computing; Computer Architecture: architectural support for operating systems; Memory's Hierarchies; Multiprocessor Systems; Computer Graphics and Visualization: algorithms for computer graphics,

visualization techniques and their applicability; Concurrency: specification and verification of distributed and concurrent systems, specification and verification of systems with time and probability, correctness of real-time and distributed systems; Databases: theoretical work on logic, language and models; Decision Support: engineering applications; Distributed Systems Techniques: techniques for the design, analysis and construction of fault-tolerant, real-time distributed systems; Image Processing and Computer Vision: sensing, processing, presentation, analysis, synthesis, and interpretation of digital images and image-related digital information; Information Systems: specification management. CASE tool environment; Workflow Analysis and Design; Interactive Systems: user interface management systems; Formal Notations; Models of Interaction; Knowledge Engineering: object-based representations; Numerical Linear Algebra; Parallel Systems: automatic parallelization; Operating Systems for Parallel Architectures.; VLSI Implementation of Parallel Architectures; Performance Analysis: queuing networks; Signal Processing: adaptative and biomedical signal processing; Speech Recognition; Software Engineering: specification languages and their semantics; Programming Environments; Symbolic Computation: symbolic computation and functional programming integration; System and Control Theory: integration of model-based mathematical models with knowledge-based systems; VLSI Design: integrated design, simulation and test facilities for advanced VLSI.

The Consejo Superior de Investigaciones Cientificas (CSIC) is a large research Institute in Spain, with research centers distributed at many locations in Spain. Four of them are involved in ERCIM-related topics: the CSIC-IAI in Madrid, the CSIC-IC in Barcelona, the CSIC-CEAB in Blanes, Girona, and the CSIC-CNM in Bellaterra, Barcelona. The research areas of the Instituto de automatica industrial, CSIC-IAI are Computer Architecture, Computer Graphics and Visualization, Concurrency, Databases, Distributed Systems Techniques, Image Processing and Computer Vision, Knowledge Engineering, Parallel Systems, Signal Processing and System and Control

Theory. On the other hand, the Instituto de Cibernetica, CSIC-IC is involved in Complexity and Algorithms, Image Processing and Computer Vision, Signal Processing, System and Control Theory and VLSI Design. The Centro de estudios avanzados de Blanes, CSIC-CEAB is working in Knowledge Engineering and Software Engineering, whereas the Centro Nacional de Microelectronica, CSIC-CNM, is involved in VLSI Design, Analogic-Digital Design and Compound Semiconductors.

Please contact: Pere Brunet - AEDIMA
Tel: +34 3 401 6595
E-mail: brunet@lsi.upc.es
or: Anna Franco - AEDIMA
Tel: +34 3 401 7289
E-mail: aedima@fme.upc.es

HCM Funding for ERCIM Database Research Group

by Keith Jeffery

The Commission of the European Union has awarded 200,000 ECUs of funding to ERCIM's Database Research Group for a network of excellence. This complements rather nicely the ERCIM fellowships funded by the Commission's Human Capital & Mobility (HCM) Programme which include researchers in advanced database topics who are currently working in some of the ERCIM institutes alongside node people in the HCM network. The network funds will be used to facilitate more contacts between the database researchers in ERCIM institutes leading to joint actions. The Group already has an active series of workshops to ensure meeting of ERCIM database researchers with similar interests.

Earlier, a similar grant of 300,000 ECUs was awarded to the ERCIM Graphics Research Group.

Please contact: Keith G Jeffery -- RAL
Tel: + 44 235 44 6103
E-mail: kgj@inf.rl.ac.uk

ERCIM Symposium

Affordable Parallel Processing: ERCIM's first Knowledge Transfer Symposium

by Mike Jane

The first ERCIM Annual Symposium was held at the Rutherford Appleton Laboratory (RAL) on 10-12 November 1993. Affordable Parallel Processing for Industry and Commerce was the chosen "hot" topic.

ERCIM intends to run one Symposium per year on a current key topic which is particularly relevant to Small and Medium Size Enterprises (SMEs) in Europe. The main purpose of these events is to establish close links and working relationships between ERCIM and a wide range of SMEs, which will ideally lead to a higher level of competitiveness and prosperity for these companies in the world market place.

A total of 65 delegates attended the Symposium, half of them were representatives of SMEs from 9 countries, providing 10 of the 18 speakers. The other half of the delegates were members from ERCIM institutes and academia.

The programme offered a range of invited presentations covering the following areas:

- An Overview of Parallel Processing
- Case Studies of successful implementations of parallel processing, covering such areas as:
 - Car Crash Simulation

- Medical Image Management
- Porting Commercial Applications
- On-line Quality Inspection in Manufacturing
- A Parallel Oracle Database Server
- A Multi-Media Telecomm Server
- A PET Scanner Image Reconstruction System
- A Real Time Synthetic Aperture Radar Processor
- Hardware and Software Issues
- A Market Survey
- Specific Applications
 - CPU Intensive
 - I/O Intensive
 - Real Time
 - Fault Tolerant/Safety Critical
- Future Trends World-wide
- Thinking Parallel is Difficult
- Funding Mechanisms for SMEs in Europe

Three panel sessions on the final day addressed the needs of SMEs. These covered needs in Awareness, Suppliers' needs from End Users, and End Users' needs from Suppliers. Most of the discussion and debate, not surprisingly, occurred in these sessions and the very final session where the conclusions, recommendations and possible follow-up actions were drawn together.

Needless to say there were many diverse views expressed ranging from "parallel

processing has arrived and can be used without difficulty by all”, to the exact opposite that parallel processing was non-viable at present and existing systems were difficult to use!! However the Case Studies did give excellent examples of many successful industrial/commercial uses of truly affordable parallel systems, some of which cost as little as £ 2000.

The main conclusion was that there were opportunities for ERCIM in:

- heightening *awareness* so that SMEs can take up and exploit affordable parallel processing
- improving the *efficiency* of parallel systems, particularly in systems that are already used by SMEs
- support suppliers by *evaluating*, and if possible, buying their *products*, as well as offering SMEs *cheap access*

to new/special facilities at an ERCIM site

- setting up a networked *information service*, as well as gathering opinions and trends and publicising these
- providing both national and international *events and activities*
- *Joining with SMEs* to provide an interface to Brussels.

The proceedings of the Symposium including the main conclusions, recommendations and follow-up actions will be published in an ERCIM Report. The availability of this Report will be announced in the next ERCIM News.

Please contact Mike Jane - RAL
Tel: +44 235 44 5408
E-mail: mrj@ib.rl.ac.uk

1. ERCIM president Cor Baayen outlining the role of ERCIM.
2. Delegates from Germany, Spain and France at the registration desk.
3. Dr. Erik Hagersten, formerly SICS and now with Sun Microsystems.
4. The Symposium attracted participants from 9 countries.
5. RAL Director Dr Paul Williams welcomes the delegates.



ERCIM Network

ERCIM Computer Graphics Network

by David Duce

ERCIM is receiving support under the CEC Human Capital and Mobility programme for a network on computer graphics. The central aim of this project is to provide a focus for research activities in computer graphics in Europe which are particularly relevant to international standardization activities in ISO/IEC. The CEC support provides travel support and organization costs for a series of technical workshops over a period of three years.

A steering committee has been set up, consisting of one representative per ERCIM site involved in the network. The first meeting was held at CWI at the beginning of October and agreed a programme of workshops up to October 1994. During the first year of the network the work programme is organized into three tasks:

1. Graphics and Other Media

This task will look at object oriented frameworks for graphics and multimedia systems, combined functionality of graphics and vision systems, multimedia and multi-sensor systems, media and time, and constraints and structures. The main outputs from the task will be contributions to the definition of the emerging ISO/IEC standard, Presentation Environment for Multimedia Objects (PREMO). The task leader is Paul ten Hagen at CWI.

2. Graphics and Knowledge Engineering

The aim here is to look at the properties of graphics systems which are necessary in order to link to a reasoning system. Intelligent multi-modal interfaces also come under this task. The task leader is Klaus Kansy at GMD.

3. Model Visualization

This task encompasses parallel visualization algorithms, volume and surface modelling and model extraction. The main output is expected to be a deeper understanding of the relations between models, modelling and visualization and the role of parallelism in each. The task leader is Sabine Coquillart at INRIA.

A plenary meeting of the network is being planned for October 1994 at which progress over the first year will be presented and the detailed work programme for the second year will be formulated.

Please contact David Duce - RAL
Tel: + 44 235 44 5511
E-mail: dad@inf.rl.ac.uk


Call for Contributions

Welcome to 12th ERCIM Workshop in Sweden

**Electrum, Stockholm-Kista,
1-3 June 1994**

The 12th set of ERCIM workshops to be held in Stockholm is dedicated to the following areas:

- Distributed Virtual Environments
- Personal Computing & Communication Environments
- High-Performance Communication



The
Electrum
building
where SICS,
SISU and
the
workshop
are located,
10km from
Stockholm
City.
(Photo:
Sandström,
Graffiti)

Audio and video from the ERCIM workshops will be broadcast over the Internet Mbone worldwide. These sessions will be announced via 'sd' over MBONE. For further information about how to prepare for this, pick up the file:

isi.edu:pub/mbone/faq.txt.

The objective and scope of each of the workshops is described below.

High Performance Communication

This workshop is the first to be organized by ERCIM's High Performance Communication group. The objective of this group is to promote cooperative research on high-speed networks in European programmes and international exchange of experiences and results. The workshop aims to provide a hearing on different switching techniques for Gigabit networks. Participants are invited to discuss advantages and disadvantages with circuits, cells and packets as the basis for switching in such networks. Examples of interesting questions are:

- do multimedia traffic characteristics in general allow bandwidth gain by statistical multiplexing, or will peak rate allocation be predominant?
- can ATM be switched fast enough in a photonic network?
- is fast circuit switching an alternative to cell/packet switching?
- because of the opto/electronic barrier that must be crossed to parse the cell header?
- are ATM cell sizes too small for high speed data networking, i.e., too much per-cell overhead?
- how to provide variable length packet switching in high speed networks.

Programme co-chairmen:
Bjorn Pehrson - KTH/Teleinformatics
E-mail: bjorn@it.kth.se
and: Gunnar Karlsson - SICS
E-mail: kar@sics.se
and: Stephen Pink - SICS
E-mail: steve@sics.se
and: Peter Sjodin - SICS
E-mail: peter@sics.se

Personal Computing and Communication Environments

This workshop is dedicated to the discussion of key issues in distributed mobile computing. Examples of relevant issues are:

- mobile terminals (Walkstations)
- multimedia applications and services
- internetworking between mobile and fixed networks
- wireless infra-structures
- effects of multiple interfaces/multiple networks
- packet size trade-offs for integrating voice and data
- energy saving techniques
- data compression and forward error correction techniques in the face of high error rates for wireless networks
- upcoming European standards and future networks.

Programme co-chairmen:
Gerald Q Maguire - KTH
E-mail: maguire@it.kth.se
and: Frank Reichert - KTH
E-mail: reichert@it.kth.se
and: Jens Zander - KTH
E-mail: jenz@it.kth.se

Distributed Virtual Environments

The primary scope of this workshop is Virtual Environments as a cooperative medium in a distributed setting. The workshop aims to bring together researchers in the various disciplines needed to realize this. This spans a wide field, from behavioural/cognitive science to high speed protocol issues. The workshop includes, but is not limited to:

- Latency issues
- Models of interaction and collaboration
- Spatially coded sound
- Case-studies
- World/environment builders
- Learning and tutorial systems
- Global environments
- Real world applications
- Art and aesthetics
- Application

- Network, intermediate layers
- Requirements on enabling technologies and communication infra-structure.

Programme co-chairmen:
Yngve Sundblad - KTH
E-mail: yngve@nada.kth.se
and: Bjorn Pehrson - KTH
E-mail: bjorn@it.kth.se

Contributions may consist of Papers, Posters, Interactive Demos and Video Tape Demos. All contributions should be submitted to the chairmen of the relevant workshop.

Deadlines

Papers:

- 1 February 1994: Draft full paper due for simple review
- 1 April 1994: Notification of acceptance or rejection
- 1 May 1994: Final camera ready versions due for publishing

Posters:

- 1 May 1994: Draft poster due
- 10 May 1994: Notification of acceptance or rejection
- 19 May 1994: Final posters due

Interactive Demos:

Interactive demonstrations of research prototypes and tools are encouraged. Within limits, reasonable technical requirements will be met. Internet and ISDN connections will be available at the conference site.

- 1 April 1994: Submission demo description, including technical requirements.
- 15 April 1994: Notification of acceptance or rejection

Videos:

If enough material of sufficient quality is received, a special video session will be organised.

- 1 May 1994: Final video tapes due (VHS, either PAL or NTSC)
- 15 May 1994: Notification of acceptance or rejection

Please contact: Bjorn Pehrson - Swedish member ERCIM Executive Committee
Tel: +46 8 752 1481
E-mail: bjorn@it.kth.se
Fax: +46 8 751 1793

CAFE – Towards an Open and Secure Electronic Payment System

by Henk Nieland

The project CAFE (Conditional Access for Europe) – part of the European Commission's ESPRIT programme – is intended to design and actually demonstrate an 'electronic wallet', to be used as a pan-European device for consumer payments, access to information services and – if required – identification. The system is designed to be open, that is, easily usable for all types of users, and secure – a matter of special attention in a network environment. The CAFE project will conclude with a trial of the payment device in 1995. Possible future extensions include other systems requiring conditional access, for example, personal credentials (like passports, driver's licenses and house-keys) and medical information.

New results in public-key techniques have made it possible to construct practical systems that can maintain both the security and privacy of individuals. With public-key techniques, each participant possesses one secret and one public key. Only the participant knows the secret key, but the public key can be freely disseminated. This provides an added measure of security because participants need entrust no secret information to any other participant. Public-key methods are particularly suited for network communication with many participants.

In a preceding ESPRIT project, CRYPTO IC-CARD, a secure chip was designed that enables the use of these techniques. Utilizing this specialized hardware, CAFE provides the basis architecture for implementing these results in concrete, wallet-like devices and will produce prototypes of the devices for trial. In its basic form, the wallet has one or two buttons and a simple LCD screen and is able to communicate to payment terminals and other devices through an infra-red communication link. More elaborate

versions might include electronic debit and credit facilities and could be inserted into Personal Digital Assistants or personal computers, in order to access the financial data more flexibly.

The electronic wallet is a hand-held pocket-sized personal device. It can be used to pay at shops, vending machines, pay-phones, parking meters, toll roads, public transportation and even directly to other people. Because of the offline nature of the CAFE protocols, it is not necessary for the recipient of the payment to be connected to any sort of central authority or database. Payments can be made by telephone, and a special 'home ATM' device will allow users to load their wallets with money from their bank at home by telephone. Because of the public-key nature of CAFE protocols, the system is open to participation by not only multiple providers of goods and services, but also multiple issuers of electronic money.

One of the primary goals of the project is to provide the capability for payments in any country and in any currency. Users who remain in their home country will make payments in their own currency. But when they travel abroad, they will be able to use their wallets to pay in the local currency. The trial will be held in at least two European countries.

The wallet itself can be purchased just like any piece of consumer electronics, so different manufacturers can provide different styles and functionalities. A special guardian smart card chip inserted into the wallet provides ultra-high security against fraud. In its last phase, the project will develop at least two alternative devices, to demonstrate that it has defined a general architecture and not a specific hardware platform. The exact physical form of these devices has yet to be decided.

Possible migrations paths from smart cards will be considered. When used with

a payment terminal that has infra-red capabilities, a wallet need never leave its owner's hands. But initially, such terminals may not be widely available. One possibility under consideration is to design the wallets to hold a smart card that can be used by itself with reduced functionality in a smart card reader, or inserted into the wallet for full functionality at an infra-red terminal.

The CAFE Consortium brings together academic centres of research in public-key techniques with an array of companies involved in current payment devices. Included are manufacturers of smart cards, payment terminals, and computer chips. Other research organizations, from telecommunications providers to those specialized in consumer research and the financial industry, round out the Consortium. The CAFE project is carried out by: CWI (Centre for Mathematics and Computer Science, The Netherlands, co-ordinator), CardWare (UK), IFS (Institute for Social Research, Germany), Gemplus (France), DigiCash (The Netherlands), Ingenico (France), Siemens (Germany), SEPT (France Telecom & Post Research, France), Catholic University of Leuven (Belgium), Royal PTT Nederland (The Netherlands), SINTEF-DELAB (Norway), Aarhus University (Denmark), and the University of Hildesheim (Germany). ■

**Please contact: Ray Hirschfeld - CWI (technical manager)
Tel: +31 20 592 4049
E-mail: cafe@cw.nl**

Activities in Databases

Twenty-nine articles for the Central Theme in ERCIM News is something rather unusual. In general, our aim is to maintain a balance between contributions in and outside this section.

In this case, with a broad field as Databases being the Central Theme, we received on our call so many contributions that this posed an editorial problem. Of the few alternatives, we chose for inclusion of all articles in one issue, even if this would lead to a section of record length.

We have tried to minimize the undesirable effects of this by presenting the separate contributions as concise as possible, which sometimes required severe cuts. As a consequence, we have only included some other Research Activities in this issue.

The submitted articles will be published in the next issue. We hope that our efforts have led to an interesting and readable section.

COMFORT – Comfortable Performance Tuning

by Gerhard Weikum

The long-term goal of the project COMFORT is to automate, to the largest possible extent, the performance tuning of database systems.

Tuning of database systems critically depends on the expertise and experience of system administrators and other experts who are responsible for adapting a system to the specific characteristics of a given workload by the setting of system parameters. There is a strong need for simplifying the tricky job of human

administrators and ideally automating at least some critical tuning decisions.

Our approach is to derive appropriate tuning heuristics from quantitative performance models for individual tuning problems, and to incorporate such heuristics in an adaptive or "self-tuning" system architecture. We are investigating specific tuning problems, in particular those posed by multi-user parallel database systems, as well as architectural principles ensuring a system's automatic adaptation to the workload.

We have been addressing the following issues.

- **Data Placement in Multi-disk Systems:** Development of algorithms and system software that can effectively exploit the I/O parallelism of multi-disk architectures. Data placement tuning is a crucial factor in the performance of multi-disk systems. The research entails issues of data partitioning, data allocation, data migration and dynamic load balancing, and on-line reorganization.
- **Adaptive Load Control:** Load control is necessary to prevent data-contention or memory-contention thrashing, due to temporary load peaks. Virtually all commercial database systems limit the maximum number of transactions that are allowed to execute concurrently. However, this method cannot react to evolving workloads. Therefore we have been investigating methods that adapt the degree of multi-programming to the evolving transaction mix dynamically and automatically.
- **Benefit/Cost-oriented Parallelization of Complex Queries:** By incorporating the benefit/cost relationship of parallelized execution plans into the optimization and processing of complex queries, we aim at efficient, heuristic algorithms for execution plan selection and resource assignment based on benefit/cost ratio rather than response time speedup alone.
- **Processor Allocation for Inter- and Intra-Transaction Parallelism:** In a multi-user parallel database system, we can either service only few transactions in parallel, each of which is assigned many resources, or service

many transactions in parallel, each of which is assigned only few resources. We are aiming at intelligent heuristics for tuning the degrees of parallelism in transactions to the current load dynamically and automatically.

So far we have built a self-tuning storage system, consisting of a parallel file system, a buffer manager, a transaction manager, and a load control component. It runs on a Sequent shared-memory multiprocessor and on Sun computers. The system can either manage real data on a real hardware platform, or it can simulate the performance impact of unavailable resources (processors, disks).

COMFORT is supported in part by the UBILAB of the Union Bank of Switzerland. Other project members are: Christof Hasse, Axel Moenkeberg and Peter Zabback.

Please contact: Gerhard Weikum - ETH Zürich, Department of Computer Science
Tel: + 41 1 632 7242
E-mail: weikum@inf.ethz.ch

SAMOS – an Active Object-Oriented Database System

by Stella Gatzju and Klaus Dittrich

Active database systems aim at representing more real-world semantics in the database which would otherwise be hidden in applications. They are able to recognize specific situations and to react without direct explicit user or application requests. Such a system registers event-condition-action (ECA) rules which can be interpreted as: 'When the specified event occurs and the condition holds, then execute the action'.

We study the definition, detection and representation of events in the context of

the active object-oriented database system SAMOS.

Specification

An event can be regarded as a point in time, when some sort of reaction is required from the DBMS. In SAMOS we distinguish 'primitive' and 'composite' events. In the first case the event of interest is immediately the specified one, in the second case it results from a combination of other events. For primitive events, the point in time can be specified by an occurrence in the database itself or in its environment. For composite events, the point in time is specified by an event algebra expression composed of other events by means of constructors (conjunction, sequence, etc.).

SAMOS provides two further concepts for event specifications. First, an extension of the definition of events with a time interval, during which the event has to occur in order to be considered as relevant. Second, an event parameterization to pass information about the database state at event occurrence time to the condition or action part.

Management

The representation of defined events should enable efficient information retrieval and access. The event management may benefit from the underlying data model in that data and events are represented using the same constructs. Since SAMOS is an object-oriented database system, events are represented as objects.

Detection

In SAMOS, an event detector detects the occurrence of all defined events. Primitive events may be detected by, for example, monitoring the system clock. For composite events, the detector needs to know the sequence of all primitive events leading to the occurrence of the composite event. This sequence is represented as a Coloured Petri net. When a primitive event occurs, a place of the Petri net is marked. Then, possibly one or more transitions are fired and one or more places are marked. In this case, the occurrence of the appropriate composite event(s) is signalled.

For the prototype implementation of SAMOS we used the object-oriented DBMS ObjectStore. The components that implement the active behavior of SAMOS – currently the analyzer of event definitions, the event manager and the event detector – are located on top of it.

The project is supported by the Informatics Laboratory of the Union Bank of Switzerland.

Please contact: Stella Gatzju or Klaus Dittrich - University of Zürich, Database Technology Research Group
Tel: +41 1 257 4578
E-mail: {gatzjudittrich}@ifi.unizh.ch

Class-less Object Models for Engineering Design Applications

by Margret Gross-Hardt and Gottfried Vossen

In engineering, a designer working with traditional object models has to deal with the notion of class - not a natural notion in such an environment. At the University of Münster a class-less object model has been developed, using concepts from prototype-based languages.

Traditional models for object-oriented databases center around the notion of a class. Thus, they resemble the common approach of defining a schema first and then introducing instances. This receives justification by the fact that a large database is commonly designed before it is used. However, engineering applications can hardly be well-supported following this approach. Indeed, a designer in a CAD environment thinks in terms of concrete design objects he or she is working on, and not about the structure or behavior of a class of which these objects may be instances in the first place. It therefore makes sense to adopt con-

cepts and techniques from prototype-based programming languages to this context.

In the CLOOD project a dedicated object model has been developed, which combines recently proposed concepts from prototype-based programming languages with versioning. The model does not know classes, but objects, which may associatively be named by a user, can freely exist in a database; in addition, they can reference each other, reuse structure and/or behavior of other objects, and may exist in multiple versions.

The basic idea underlying the model is to have objects freely exist in a database. Objects still have a structure and a behavior, but can additionally be versioned. Thus, each object in our model has an identifier, an associated set of versions, and a set of messages it understands. Versions in turn have a number and a value, where values are elements of domains that are described by types. Furthermore, objects can be composed of other objects, and can interact. Finally, objects can inherit structure and or behavior from other objects.

Clearly, a declarative way of querying and updating a database is still desirable. To this end, we introduce the notion of a collection, which generally is an arbitrary set of objects or versions; however, it seems useful to identify a number of collections which have specific properties. For example, configurations are consistent sets of versions, or constellations are collections containing an object together with all its subobjects. Thus, a database can be grouped into a variety of application-dependent collections, which can then be subject to user operations.

Currently a language for defining and querying CLOOD-objects is under development. This language, CQL, allows the user in an easy and elegant way not only to query attribute-values of objects but also to select objects according to their common descriptions. Future work is directed towards the formalization of CQL and the integration of integrity constraints between versions of different objects.

Please contact: Margret Gross-Hardt or Gottfried Vossen - University of Münster
Tel: +49 251 9275 108
E-mail: {hardt,vossen}@uni-muenster.de

Translation, Optimization and Modification of Designs

by **Gottfried Vossen**
and **Christian Fahrner**

Current database design systems have several limitations. A novel design toolbox, under development at the Universities of Münster and Cottbus, will help to overcome these limitations.

Database design can be seen as a four-step process:

1. in a requirements analysis, all the information needed to model the desired application is collected
2. a conceptual schema is created describing the structure of the database independent of the underlying DBMS
3. the conceptual schema is translated into a logical schema describing the database structure that can be processed by a DBMS
4. a physical schema is derived from the logical schema, which describes the implementation of the database in secondary storage.

Parallel to this data-driven design a functionally driven design is made to model transactions that should be processed on the database. Data and functional design influence each other.

Over the last years several database design systems have been developed to support this process. Most of them use the Entity-Relationship model or one of its extensions for conceptual design, and translate conceptual schemas into relational, network, or hierarchical ones. Current systems have several limitations:

- the systems very often are just graphical interfaces (drawing tools) and do not really support the design of databases
- the number of models, into which the model describing the conceptual schema can be translated, is limited
- reverse engineering is not supported

- optimization is limited to a normalization of entire schemas
- operations on the data cannot be modeled.

In a project supported by the Deutsche Forschungsgemeinschaft and carried out at the University of Münster in collaboration with the University of Cottbus, techniques for translating, optimizing and modifying designs are under development and the results will be integrated into a novel toolbox for database design. So far, our emphasis has been on the development of new translation facilities which support:

- a translation of a rich semantic model into a number of target models
- direct transformations between pairs of target models
- a translation of schemas in a semantic model into user-defined new models
- reverse engineering.

An implementation of these facilities on SUN workstations is underway. ■

Please contact: Gottfried Vossen or Christian Fahrner - University of Münster
Tel: +49 251 927 5103 (5108)
E-mail: {vossen, fahrner}@uni-muenster.de

Database Research at GMD-IPSI

by **Wolfgang Klas**

At GMD's Integrated Publication and Information Systems Institute (IPSI), several research and development projects are running in the fields of object-oriented (O-O) database technology, multimedia extensions of O-O database systems, and interoperable heterogeneous database systems. Many of the projects are carried out in cooperation with other research organizations as well as with industry.

The research covers:

- **interoperability of relational and O-O database systems**
- **O-O database management support for computer-aided drug design**

- **transaction management support for cooperative applications**
- **the next generation of O-O database technology**
- **database management support for a hypermedia-based cooperative multimedia systems engineering workbench.**

IRO-DB – Interoperable Relational and Object Databases

This ESPRIT project intends to achieve interoperability of pre-existing databases and new O-O databases. It will concentrate on providing a C++ library to access in an integrated way heterogeneous databases supported by communication protocols to exchange O-O SQL commands and objects. Furthermore, tools to design and dynamically maintain integrated applications on large federations of heterogeneous databases will also be developed. They will provide uniform C++ interfaces on top of dynamically adaptable integrated schemas for data exchange and sharing, global transaction management and concurrency control, without violating the autonomy of the participating databases.

Initially, a system will be developed for only querying distributed relations and classes with mono-valued attributes and references on the basis of export schemas. The object interchange protocol OORDA – an O-O specialization of the generic Remote Database Access (RDA) standard – will support classes of complex objects possibly interrelated using object identifiers, and generalization hierarchies. Supporting communication software and procedures for interfacing local databases will be implemented on top of INGRES and MATISSE. A C++ library, and a first version of a schema integrator with a query translator will be provided as attractive tools to develop O-O applications using existing DBMSs.

Planned extensions concern the OORDA protocol (a rich set of type constructors, behaviour encapsulation, update and facilities to maintain atomic and concurrent transactions, the schema integrator (integrated schemas for multiple groups of users), more powerful mapping facilities between local database schemas and export schemas, and a more

sophisticated interoperable object manager. Also other database servers (O2, ONTOS) will be interconnected.

The project will provide an easy path for integrating the relational database technology to the new O-O one. A demonstrator will be built to evaluate the system using an engineering application from INTRACOM, the major telecommunication company in Greece. It will integrate an existing CIM database running on INGRES with an SQL production database and an abstraction of a CAD database developed on top of an O-O system. Further investigations of applications, notably in the banking domain, will be done by project partners (Ibermatica in particular).

Partners in the project are: Ibermatica (Spain), Intrasoft (Greece), Intellitic, Infosys, O2 (France), and GOPAS (Germany).

Receptor-Ligand Docking Database

Ligands are substances having biological effects by docking to proteins (receptors), due to high affinity and specificity. In the search for new ligands – a major pharmaceutical research area –, bioinformatic tools play a central role, in particular for rational drug design. In view of the recent explosive growth of data about proteins, there is a need for tools to represent the knowledge about receptor structure. The challenge is to provide these data with algorithmic and database methods as well as appropriate software tools for a computer supported analysis of receptor-ligand docking. This will allow a goal-oriented design of

diagnostics and pharmaceuticals – a perspective of enormous economic importance.

In RELIWE, a national project funded by the Bundesministerium für Forschung und Technologie and started in April 1993, the receptor-ligand interaction is studied. The data used is extremely heterogeneous and the database system must support not only the integration of this data, but also the integration of new datatypes and associative data retrieval. Receptor-ligand docking algorithms, model building tools for receptors and visualization tools, which are developed or provided by the project partners, must be connected to the DBMS. In the study of heterogeneous database support GMD-IPSI takes the leading role.

The other project partners are: the GMD Institute for Methodological Foundations (Sankt Augustin) the European Molecular Biology Laboratory (Heidelberg) and two industrial partners, BASF (Ludwigshafen) and Merck (Darmstadt).

TRANSCOOP

There is a need for software tools supporting multiple users who cooperate in some task, for example, several authors working concurrently on the same document. The support of cooperative tasks provided by available tools is mainly very superficial, e.g., by coordinating multiple users at the level of a window system, without having any understanding of the data. Also, existing transaction management concepts are unfortunately designed to fully isolate users from each other.

A major conceptual problem is to ensure consistency criteria for the data concurrently processed by multiple users. Another serious problem is the design of cooperative systems itself, a task involving several formalisms and techniques that have up to now hardly been studied in combination. Formal techniques are needed to capture the notion of cooperation at the design level and to allow for the specification of cooperative tasks. In general, a theory for cooperative transactions is needed, including aspects such as:

- distinct notions of correctness
- simulation models
- test generation
- compiler building and report generation
- correctness-preserving transformations.

The whole process needs to be supported by a coherent set of tools.

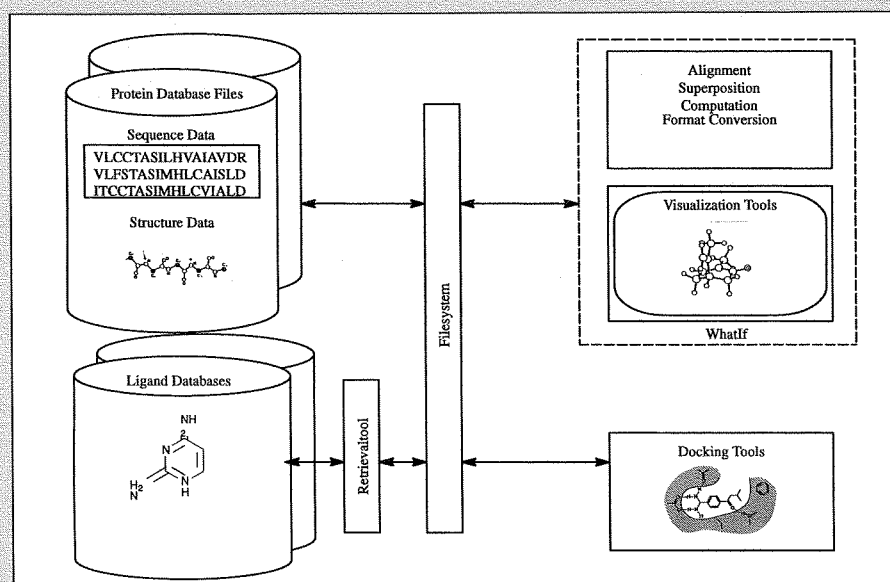
The ESPRIT-BRA project TRANSCOOP aims at the development of transaction management support for cooperative applications. Its objectives are:

- to identify the requirements for a cooperative transaction model in selected types of cooperative environments
- to develop a specification language for describing cooperative tasks
- to adapt and extend an existing transaction model towards the support of cooperative environments such that it meets the requirements given by selected application types
- to deduce the formal properties of the chosen model and formally describe it
- to integrate the cooperative transaction model into an existing prototype database system
- to implement a prototype which proves the applicability of the approach for a selected application scenario.

The other partners in TRANSCOOP are the University of Twente (The Netherlands) and VTT (Finland).

Please contact:
Wolfgang Klas - GMD
Tel: +49 6151 869 943
E-mail: klas@ darmstadt.gmd.de

Current data management in the receptor-ligand docking environment.



How to Handle Imprecise Medical Information

by Olga Pons

In a project of the University of Granada, in cooperation with the local Clinic Hospital, a fuzzy query and storage system for processing imprecise medical information is under development.

In the relational database approach, there are two ways to query a database: through relational algebra and through relational calculus. However, neither way can answer queries formulated in an imprecise manner, such as: 'Find all young subjects with a not too high salary'. The solution lies in using fuzzy database systems, connected to Artificial Intelligence in the sense that they allow the treatment of information in a more natural language.

We consider two levels of fuzziness:

- query of a classic database in imprecise terms
- representation of fuzzy data, in which case a new data representation is needed.

Our database system must be able to operate at both levels. The mathematical modeling of the imprecision is made through Zadeh's Fuzzy Sets Theory. In particular, linguistic labels are used to represent imprecise data, and queries are solved with the basic ideas of Possibility Theory.

The first goal of the project is the development of a fuzzy query system for non-fuzzy relational databases. In addition, a fuzzy storing system is required, so that the previous work for fuzzy queries can be adapted. The user software will be as general as possible, since we hope it will be portable to a wide range of medical problems. Therefore we will use SQL as data language, in the framework of ORACLE, operating on a SUN station under Unix.

Until now we have developed a theoretical model for fuzzy relational databases and implemented the fuzzy meta-knowledge base, where different kinds of fuzzy data types are considered. Furthermore we have almost finished a fuzzy SQL parser to query the system, which works on classic SQL.

The work is supported by the Spanish Ministry of Education and Science.

Project members are: M.A. Vila (project leader) and J.C. Cubero, J.M. Medina, O. Pons and M. Prados (all University of Granada), and J. Diaz (Clinic Hospital of Granada).

**Please contact : Olga Pons - University of Granada, department of Computer Science and Artificial Intelligence
Tel: +34 58 244 079**

Database Interoperability

by Felix Saltor

A project on the interoperability of a multi-database is carried out at the Universitat Politècnica de Catalunya (UPC), in collaboration with Austrian, German and French institutions. The research focuses on semantic issues.

When several pre-existing databases, designed independently and operating autonomously, are interconnected to form a federation of databases, so that a query to the federation produces a single, consolidated answer ('integrated access'), a number of problems arises. The federation, also called 'multi-database', is a distributed system of heterogeneous and autonomous databases, and the operation between the databases is a worldwide research topic. At UPC's department of 'Llenguatges i Sistemes Informàtics' research in the field focuses on semantic issues, particularly semantic heterogeneity, including semantic relativism, and evaluation and development of data models.

Until now, the power to derive external schemata from the database (or conceptual) schema was characterized, and

the semantic relativism of several models analyzed. The analysis of desirable structures for a Canonical Data Model was followed by the development of a framework of suitable characteristics. Several data models with respect to this framework were analyzed and, in order to satisfy these characteristics, an Object Metamodel and object-oriented data model BLOOM (BarceLona Object Oriented Model) was developed. Semantic enrichment of database schemas, including extraction of dependencies and conversion from relational to object-oriented, was achieved. Finally, techniques were developed for the integration of export schemas, including discriminated operations, schema transforming operations and structural discovery of resemblance of classes.

At present, work proceeds on:

- other techniques for export schema integration
- languages and implementation issues for BLOOM
- support of views in the BLOOM model
- interdependencies between databases
- architectures for database interoperability.

The project is funded by the Spanish PRONTIC programme and carried out in collaboration with the University of Vienna and the Technical University of Darmstadt, and in contact with GMD-IPSI and INRIA. The Spanish participants are: F. Saltor, J. Sistac, M. Garcia-Solaco and M. Castellanos (UPC), J. Samos and B. Campderrich (Univ. Rovira i Virgili, Tarragona).

**Please contact: Felix Saltor
UPC, Barcelona
Tel: +34 3401 7023**

Parallelism and Non-determinism in Active Database Application

by Stéphane Grumbach

The 5th ERCIM Database Research Group Workshop on parallelism and non-determinism for data intensive applications was held at INRIA, Rocquencourt on November 2-4, 1993. About twenty participants attended the workshop, coming from seven different countries. Among the non-ERCIM participants were researchers from Paris, as well as from abroad, including Greece and Germany.

Parallelism

Patrick Valduriez (INRIA) presented a tutorial on parallel database systems, which are becoming the primary application of multiprocessor computers. Valduriez compared the traditional shared-memory, shared-disk, and shared-nothing architectures of parallel database systems under the following dimensions: simplicity, cost, performance, availability and extensibility. Based on these comparisons, he made the case for the shared-something architecture which can provide a better trade-off between the various objectives.

Mohamed Zait (INRIA) presented his work on the effectiveness of optimization search strategies for parallel execution spaces. The cost of query optimization is affected by both the search space and the search strategy of the optimizer. In a parallel execution environment, the search space tends to be much larger than in the centralized case. The trade-off between optimization cost and parallel execution cost using the DBS3 parallel query optimizer was investigated. Zait proposed to reduce the optimization cost by controlling the search strategy. Randomized strategies, such as Toured Simulated Annealing, were extended to adapt them to parallel query optimization.

Carel van den Berg (CWI) presented the Goblin query processing architecture. Query optimization is a difficult and time consuming process. Goblin can efficiently take load balancing decisions and perform query optimization at run-time. Its primary characteristics are that it only exploits data parallelism, not pipeline parallelism, and uses a two-level query processing structure, in which a query is evaluated first on the summary data base and then on the stored fragments. A consequence of the declustering scheme used in the storage architecture is that a query is mapped into a union query over all possible fragment combinations, or tasks. These tasks are independent, run in parallel, and they are small enough to be executed in main-memory.

Non-determinism

Victor Vianu (University of California, San Diego) presented a tutorial on non-determinism in query languages. He showed that there are non-deterministic languages expressing low complexity classes of queries/updates, whereas no such deterministic languages exist. He reviewed extensions of Datalog with negations in bodies and/or heads of rules, with non-deterministic fixpoint semantics, and non-deterministic extensions of first-order logic and fixpoint logics, using the witness operator. Languages expressing exactly the queries/updates computable in polynomial time were exhibited. Considered problems of practical interest included: checking if a given program is deterministic, detecting coincidence of deterministic and non-deterministic semantics, and verifying termination for non-deterministic programs.

Fosca Gianotti (CNR-Pisa) presented a declarative approach to non-determinism in logic database languages. She addressed the issue of non-deterministic extensions and showed how a dynamic choice construct is capable of expressing interesting problems, such as several forms of negation, and ordering. Datalog augmented with the dynamic choice expresses exactly the non-deterministic polynomial-time queries. This leads to a complete characterization of the expressiveness of the dynamic choice, as well as a characterization of the class of queries NDB-PTIME.

Foto Afrati (NTU, Athens) investigated the expressive power of Datalog, and presented queries which are monotonic and Ptime, but not expressible in Datalog. She also studied the effect of various syntactic restrictions on the expressive power of Datalog and presented natural examples of programs which (i) require recursively defined predicates of arbitrarily large width, or (ii) require rules with arbitrarily many recursive calls.

Georg Lausen (Mannheim University) presented work on a language for database programming with states. Many approaches to the performance of update operations in deductive databases rely on a procedural semantics of rulebased languages. He claimed that updates can be specified in a purely declarative manner using model based semantics. The key idea is to incorporate states as first class citizens into the language. He outlined how to implement the computation of the perfect model semantics for update programs and presented some results relating the update language to temporal deductive database languages.

Zoë Lacroix (University of Paris-Orsay) discussed implicit definitions and unambiguous non-deterministic Turing machines. She presented a subclass of NP where one requires for each input at most one accepting computation, the class UP, of languages accepted by an unambiguous TM in polynomial time. On the class of finite ordered structures, $UP \cap co\sim UP$ has a logical characterization, namely IMP the class of implicit formulas.

Sergio Lifschitz (ENST-Paris) presented a joint work with Victor Vianu on a probabilistic view of datalog parallelization. He argued in favor of designing parallelization strategies with the aim of achieving good expected performance, in the context of given information about data. The pure parallelization algorithm is a good illustration, when no information is known a priori about the input: its performance is excellent on arbitrary, random inputs. For applications involving highly structured data, like complete binary trees, pure parallelization needs to be tailored further to the specific situation in order to be practical.

Please contact:
Stéphane Grumbach -- INRIA
E-mail: Stephane.Grumbach@inria.fr

The EDS Parallel Query Optimizer

by Patrick Valduriez, Eric Simon and Mohamed Zait

Design and implementation of efficient and effective advanced database management systems for parallel platforms is one of the grand challenges of database research. The INRIA Database Research group is active and worldwide successful in this area for many years.

EDS is an ESPRIT project started in 1989 by Bull, ICL, Siemens, ECRC and INRIA with the major goal of producing a parallel database server which exploits recent multiprocessor computer architectures to provide high-performance and high-availability database support at a much lower price than equivalent mainframe computers. The thrust of the project is that database management and parallel processing technologies are mature enough to be successfully combined and take a central position in mainstream commercial information systems of the 1990s.

The EDS database server targets business data processing applications with mixed workloads of low-complexity (OLTP-like) transactions and more complex decision-support queries because they contribute most to throughput demands. The industrial context of the project naturally led to the choice of standard SQL for providing database capabilities extended with deductive and object-oriented capabilities, called ESQL.

With ESQL, the programmer can write high-level queries that are more powerful and complex than relational queries. In particular, they may involve path expressions for navigating through objects and recursive predicates. Thus, to achieve the EDS database server's functionality and performance objectives, it is crucial to compile and optimize ESQL for efficient parallel execution.

We have built the optimizer following a decade of experience with commercial

optimizers. The main components of the optimizer are:

- which defines in an abstract way the
- alternative execution plans
- which predicts the cost of an execution plan
- which is used to select the best execution plan in the search space.

For each component, we have produced original solutions to address the new problems caused by the language, the parallel execution model and the need for extensibility. The EDS optimizer has been successfully integrated with the rest of the ESQL compiler.

Please contact:
Patrick Valduriez, Eric Simon and
Mohamed Zait -- INRIA Rocquencourt
+33 139 635 511
E-mail: {patrickv@sylvia.inria.fr,
simon@laure.inria.fr,
mohamed.zait@inria.fr}

O2 Technology contributes to ODMG Standard

by George Nissen

The Object Database Management Group (ODMG) has taken over large parts of the O2-SQL language in its standards recommendation.

O2 Technology, an INRIA spin-off marketing an object-oriented Database Management System conceived at INRIA, is the only European company to take part with voting power in ODMG, affiliated to OMG, all the others being from the US. ODMG has recently released a recommendation to standardize the access to objects. This recommendation, called ODMG-93, should play a role similar to SQL for relational DBMS. It is interesting to note that the OQL access language contained in ODMG-93 is very similar to the O2-SQL language promoted by O2Tech. This can be considered as an international recognition of the soundness of the approach of the problem by O2Tech, one of the very few

European companies to address this promising field.

Please contact: Guy Ferran -- O2 Technology
Tel: +33 130 847 777
E-mail: ferran@o2tech.o2tech.fr

WEA – a Persistent Object Manager for Distributed Applications

by Didier Donsez, Philippe Homond and Pascal Faudemay

The WEA object manager developed at MASI provides a seamless interface between C++ and a persistent store.

WEA is a C++ extension to develop object-oriented database applications. The interface offers a class library for collections, such as list, set and map, and can be easily extended by the application programmer. The current implementation supports a powerful transaction model based on the concept of a shared workspace.

Several workspaces can exchange services, such as to provide remote access to database objects over the network. As such, it supports both client-server architectures and cooperative applications (groupware). Cooperative transactions update objects versions, which are visible to other transactions within the group, and agree on the validation of some of them to become visible to external transactions.

WEA is an implementation of the Workspace architecture. The design of Workspaces is widely based on functionalities of modern operating systems, such as multi-threading or memory mapping. The concurrency control method is CallBack Locking.

This two-phase locking technique is well adapted to distributed architectures. By

caching the client read-lock, it favours successive reads of the same object by several transactions placed on the same client

An alpha version of WEA is available from the server ftp.ibp.fr for SUN4 machines running Solaris 2. It supports construction of applications distributed over a network of SUN workstations. ■

Please contact: Didier Donsez, Philippe Homond and Pascal Faudemay -- Laboratoire MASI - EquipeRapid, Institut Blaise Pascal, Universit Pierre et Marie Curie, Paris
Tel: +33 144 276 018
E-mail: {donsez,homond,faudemay}@masi.ibp.fr

The RAPID Data Store – a Real-time Front-end Database

by Antoni Wolski

Real-time processing of vast amounts of data – frequently occurring in an industrial environment – requires a new approach to database systems. This issue is addressed in the RAPID project at VTT (Helsinki).

In various industrial systems, vast amounts of sensor data need to be collected promptly and made available for processing by time-critical applications. Examples include control room applications of complex industrial processes (e.g., a paper mill operation), applications for remote operation control (e.g., of energy utilities), and various quality control and integrated manufacturing systems.

Contemporary disk-based database systems are not suitable for this purpose because they can not efficiently accommodate the rate of data-inserting transactions (hundreds of transactions per second). On the other hand, the builders of industrial systems are willing to use relational databases for processing

the collected data because standard interfaces (e.g. the SQL language) and various application development tools are available.

In the circumstances, a front-end database system is needed, whereby the data may be efficiently collected and used by real-time applications and then transparently migrated into a disk-based relational database.

The RAPID project aims at producing such a system for a broad range of applications. The RAPID Data Store System contains a main-memory database which can be accessed using an SQL-like query language and an API interface based on the X/Open SQL Command Language Interface.

The major challenge of a high data assimilation rate is overcome by appropriate main-memory-based data structures and optimized data insertion processing. Because the data usually comes in the form of a measurement time series, temporal operations on data are of great value. Analysis of the application domain resulted in discovering new types of temporal predicates, for example those having 'temporal memory', i.e. yielding results depending on some states in the past.

The RAPID Data Store is an active database beyond what is offered by

typical implementations. Database triggers may invoke actions in various components of an application system. The action invocations are distributed to both permanent and temporary action handlers depending on a run-time configuration of a system. For example, an application process may 'subscribe' to an action invocation, and cancel the subscription if it is not needed any more.

RAPID is a three-year, 8 person-years project, started mid-1992 and carried out at VTT Information Technology. It is jointly financed by the state research fund TEKES, VTT, and companies participating in the project management. ■

Please contact: Antoni Wolski - VTT Information Technology
Tel: +358 0456 6012
E-mail: Antoni.Wolski@vtt.fi



ROAr – a Distributed Office Automation System

by Simon Dobson, Liz
Hutchinson and Victoria Burrill

ROAr is an application-oriented development which supports the workflow of RAL's recruitment office. It uses a relational database which can be accessed from different hardware and software platforms, and which has multimedia capabilities.

We were approached to provide a system to semi-automate the Laboratory's recruitment process. A network of PCs offered the most cost-effective, timely and user-friendly solution, and provided us with the opportunity to explore the ease with which a highly distributed application for such machines could be created using readily-available software tools. A state-of-the-art graphical user interface was an integral factor in the system design.

Recruitment clerks have to deal with up to one hundred vacancy notices (VNs) simultaneously. Completing a VN requires around four hundred steps to be performed, with about three hundred documents involved (many of which are needed for each candidate for the post, each interviewee, each board member,...). In addition there are application forms, certificates and references.

The ROAr (Recruitment Office Assistant resource) system architecture is based on a file server and a number of workstations for the staff (which might be PCs, Macs or Unix workstations). The heart of ROAr is the Microsoft FoxPro relational database engine. The FoxPro development tools have been used to build basic data entry and browsing facilities. More complex graphical overview screens and the overarching application has been written using Visual Basic, and the two applications interact seamlessly via Dynamic Data Exchange Facility,

DDE, provided by the Windows environment.

The established recruitment case workflow has been captured in a semi-formal manner. This forms the basis of the ROAr interface, modelled using a cue-card metaphor: a series of cue cards represents the sequence of steps which must be carried out to fill a vacancy. On system start-up, the clerk may select any currently active VN and will be taken to the appropriate cue card in the sequence. Some steps can be performed out of order. This allows a great deal of flexibility whilst ensuring data integrity for each stage.

ROAr also provides overview facilities to allow the progress of work in the group to be observed. In addition, the system can gather statistics needed to comply with legal obligations or Laboratory policies. Summaries of each vacancy's progress are generated on a regular basis and mailed to the initiator.

ROAr has been designed with expansion in mind. There is considerable local caching of data to reduce network and server loads, and vital data is backed-up automatically. For the first prototype, most processing is performed on the workstations, but the architecture allows processing to be shipped-off to a database server with minimal changes. As the


system deals with personal data, security is a major concern: we are exploring the use of Microsoft's SQL Server running on Windows NT as a suitably secure database engine.

Planned extensions include image capture of candidates arriving for interview (which would then be placed directly onto the document given to the interview board) and of certificates and other application related documents. OCR capabilities to support the initial processing of application forms may also be needed. This requires the database and system architecture to support acquisition, storage, retrieval and transmission of multimedia data.

The semi-formal analysis has already highlighted optimisations which ROAr can introduce into the recruitment process and these have been incorporated in the database design: it is anticipated that a subsequent full formal analysis will expose still further improvements, along the lines of an 'intelligent assistant'.

Please contact: Liz Hutchinson - RAL
Tel: +44 235 44 5235
E-mail: {sd, liz, vab}@inf.rl.ac.uk

Timelines showing the status of job applications.



Activities on Multimedia Databases in MUSIC

by Stavros Christodoulakis

The Multimedia Systems Institute of Crete (MUSIC) is a centre of research, development, training and technology transfer in the area of multimedia systems, located in Chania. MUSIC cooperates with the Technical University of Crete as well as with FORTH, maintains strong links with local companies for technology transfer and cooperates with many Greek and international computer and communication companies in several projects.

MUSIC has extensive experience and expertise in the area of multimedia information systems. It currently leads (Technical Leader and Prime Contractor) a large consortium of the most important Greek forces (companies, institutes, universities) in the area of multimedia information systems, in a Greek Action for Multimedia Information Systems. The objectives of the consortium are to develop the infrastructure, experience, basic software platforms and application platforms in the areas of training, culture and tourism, to internationally compete

in the area of multimedia markets. MUSIC is also part of the ESPRIT Basic Research working group MIRO for multimedia information retrieval, part of the ESPRIT network of excellence IDOMENEUS for distributed multimedia information systems, and part of the ESPRIT- NFS cooperation project DAIDALOS.

MUSIC also has participated and currently participates in several R&D projects with emphasis on multimedia database management systems and multimedia information systems. The areas of these projects are object servers (ESPRIT- KIWIS), medical multimedia data servers (AIM-MILORD), multimedia information systems and applications (MULTIMEDIA-STRIDE), hypermedia systems (ESPRIT-HIFI) and multimedia cultural publishing (ESPRIT-MINERS and ESPRIT-IMAGINA). MUSIC is also developing a multimedia information system for the promotion of tourism in Crete, which is financed by the EEC SPA programme, and participates in the ESPRIT project MINOS (multimedia information systems for tourism).

The research interests of the staff include Multimedia Information Systems, Multimedia Publishing, Very Large Databases, Information Retrieval and Office Automation Systems. The staff is also currently working on the KYDONIA project, a long-term internal project of MUSIC which aims at the development of a high performance, distributed

multimedia information system. The architecture of this system is shown in the figure. The system is based on state-of-the-art multilevel storage hierarchies (including optical disks, jukeboxes of optical disks and scheduling support for tertiary storage), multiprocessor technology and a client-server environment. A storage manager layer for multilevel storage (jukeboxes of optical disks), an object-oriented (O-O) multimedia kernel, and an O-O application language have been implemented, and multimedia user interface and authoring tools (windows, toolbok, etc.) have been integrated on top of it. The existing system is used as a basic platform for multimedia applications as well as a testbed for the development of algorithms and software for the large scale multimedia information systems of the future.

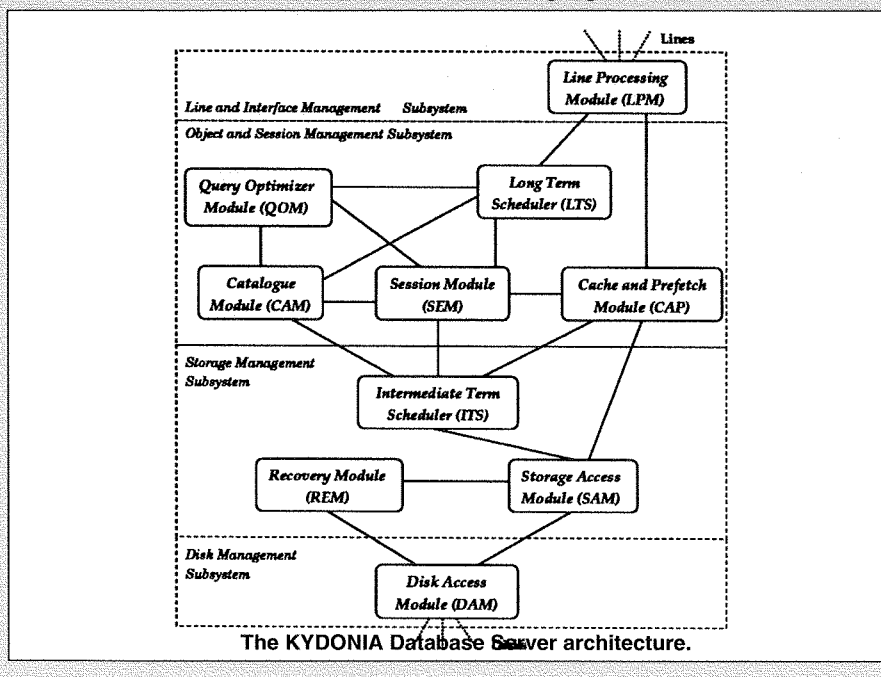
Please contact: Stavros Christodoulakis - MUSIC
Tel: +30 821 64803
E-mail: stavros@ced.tuc.gr

FIBONACCI – a Database Programming Language

by Antonio Albano

Fibonacci is a prototype of a statically and strongly typed object-oriented database language incorporating new mechanisms to model databases in terms of objects with roles, classes and associations.

In the last decade many database programming languages have been based on the object paradigm, and several commercial products have appeared that support object databases and languages. The approach adopted to move from relational DBMSs to object DBMSs is by evolution or by revolution. The evolutionary approach attempts to extend the relational data model with object-oriented (O-O) features. For example, this approach has been adopted by the designers of the Miro system and of the new data model supported by the SQL3



database programming language, expected to be completed by 1995. The revolutionary approach attempts to design ex novo an object data model supported by a new database programming language. Examples are Gemstone, ObjectStore, Ontos, O2, and Orion. Both approaches adopt the most consolidated mechanisms proposed in the last decade for object databases which however have been considered not completely satisfactory for modelling objects and associations between sets of objects.

A limitation of currently available object mechanisms is their inability to model the dynamic nature of real-world entities, which can change their role during their life and exhibit a role-dependent behavior. In addition, objects model entities of the domain of discourse, and an association between sets of entities is modeled by an object method which returns other objects.

Some of the problems inherent in this approach are:

- the enforcement of the constraint that the two methods of modelling an association are the inverse of each other is left to the programmer
- associations are not necessarily binary, and can have their own attributes; these aspects can only be modelled indirectly by means of methods
- associations cannot be defined incrementally without redefining the structure of existing objects
- operations on relationships as a whole are not possible in a straightforward way.

In order to overcome these limitations we are developing an object data model supported by the Fibonacci programming language, designed to integrate features, which can model naturally O-O database applications, into a general purpose programming language which is statically and strongly typed. The design of the language has drawn heavily on experience with the Galileo language, also developed at Pisa University.

The main contributions of Fibonacci are:

1. The definition of a novel mechanism to model objects. Objects are entities with an immutable identity and a mutable state, organized as an acyclic

graph of roles. A role is an entry to access the object it belongs to. Role types are organized into an inclusion hierarchy with multiple inheritance and can be implemented separately. Operators are defined to test for the existence of a certain role and the extension of the object with new roles.

2. The integration into a statically and strongly typed, polymorphic persistent database programming language of new mechanisms to model databases in terms of classes of objects and associations between classes. Classes are modifiable collections of values, on which constraints, such as inclusion or mutual disjointness, can be defined. Associations represent modifiable n-ary symmetric relations between classes, on which other constraints, such as cardinality, surjectivity, dependency and non-mutability, can be defined.
3. The design of a query algebra for bulk data (sequences, classes and associations) based on comprehension mechanisms which integrate aspects of relational and nested relational query languages with aspects typical of functional languages having the sequence as a data type.

The current implementation of Fibonacci (by Engineering Informatica S.p.A.) runs on Sun workstations and is written in Modula-3. It does not support all the features of the Fibonacci type system. The project is supported by ESPRIT BRA project FIDE2 (Fully Integrated Data Environment), CNR and the Ministry of Research. In addition to the author, other current project members are: G. Ghelli, R. Bergamini (University of Pisa), R. Orsini (University of Venice), M. Diotallevi and C. Brasini (Engineering Informatica S.p.A.). D. Blakeman, A. Marrara and L. Vinciotti gave an essential contribution to the success of the implementation.

**Please contact : Antonio Albano -
University of Pisa
Tel: +39 50 510269
E-mail: albano@di.unipi.it**

Advanced Database Systems at CNR

by Domenico Sacca

The project Advanced Database Systems, started in 1989 as a part of the CNR programme for Information Systems and Parallel Computation, focuses on intelligent databases, logic query languages, interaction with multimedia and with heterogeneous databases, and end-user methodologies and tools.

More specifically, the research concerns:

- innovative database systems supporting wider functionalities to meet the requirements of new applications, i.e. knowledge-based management systems
- expressive languages for knowledge-based applications
- efficient techniques for the compilation and execution of logic programs for data and knowledge bases
- the integration of multimedia databases (text, image, voice) with traditional ones and of heterogeneous databases on different sites
- general tools enabling the end-user to design and construct simple applications operating on his or her databases, and to query information systems.

The research is carried out by several academic and industrial partners. The results have been experimented through prototypes aimed both at validating these results and at providing bridges towards industrial exploitation. The main activities concern the design and implementation of a number of systems listed below, together with some concrete resulting products.

1. Significant modules for an O-O database system, including query and programming languages for the development of applications. A programming environment was created for the new, strongly-typed, O-O

database language Fibonacci; an object server will be forthcoming.

2. An advanced database system that significantly extends the functionality of current technologies by adding the inferential mechanism of logic programming and some typical abstractions of the O-O paradigm. Results include a prototype of an extension to a relational database system, based on interfaces designed to handle complex objects and logic constraints.
3. A system which extends relational database systems to cope with texts, images and cartography. A prototype was developed which adds a module for managing geometric data with ad-hoc advanced data structures to a relational database system; a database design environment for geographic information systems will be forthcoming.
4. An environment for the integration of multimedia data (images, sound and voices) with traditional databases on distributed architectures. The research has led to an initial prototype for querying heterogeneous database and an environment for the management of distributed transactions.
5. A cooperative, intelligent interface for casual users to bibliographic databases by modelling both domain knowledge and individual users. Results include an intelligent interface to bibliographic data, and a tool for modelling the characteristics and skills of the users of any interactive information source system, e.g., information retrieval systems.

■
Please contact: Domenico Sacca
DEIS DEPT.
University of Calabria
Tel: +39 984 494750

Static Analysis for Object-Oriented Databases

by Fosca Giannotti

Can Object-Oriented Database Management Systems be implemented on the basis of formal techniques? How far is a mature technology, based on formal techniques, capable of supporting the conceptual, logical and physical layers of such systems? Research at CNR tries to answer such questions by experimenting with a formal technique called 'abstract interpretation'.

The evolution of database programming languages, which provide a unifying framework for data definition and manipulation, opens up opportunities for technology transfer between the database and programming language areas. Persistent programming languages are a subclass of database programming languages, in which the persistent value data model is completely integrated in a programming language. In particular, if a persistent programming language is designed on the basis of formal semantics, formal methods can be exploited to design, transform and optimize database programs. A challenging line of research is the investigation of to what extent these formal methods can be successfully applied in the database context.

The emerging database technologies strongly demand new functionalities. The O-O Database area, for example, is now very active in defining the new features that a DBMS should provide in order to support an O-O data model. Proposals have been made for new protocols for concurrency control, long transaction management, new indexing mechanisms to support navigation, new algorithms for efficient access to persistent data and new organizations of secondary storage (prefetching, clustering). In our opinion, the combination of these events should lead to a new generation of DBMS where all the components, from the data model

to the physical design, are based on formal techniques that guarantee the safety of optimizations and transformations, too often based only on empirical experience.

In the programming language area, there is a consolidated class of formal techniques for static program analysis, known as 'abstract interpretation'. These techniques aim at gathering information about the dynamic semantics of a program to be used by compilers, partial evaluators and debuggers, or merely as documentation. For example, logic programming has successfully used abstract interpretation for detecting properties like groundness and goal independence, to provide efficient parallel execution of programs. Functional programming uses abstract interpretation to detect strictness and reference independence for efficient graph reduction and optimised garbage collection.

We are now experimenting this technology in the context of database programming languages. Our first experiment consists of defining an abstract interpreter for a persistent O-O programming language, namely a subset of the database language Galileo. The aim is to detect opportunities for safe parallelism of transaction operations scheduling. Transaction operations can be safely interleaved when there is no overlapping among their readsets and writesets. Given the text of the transaction and a representation of the data stored in the database (the conceptual schema), the abstract interpreter automatically derives, in finite time, an approximation of the readset and writeset of the analysed transaction. It is worth noting that the goal of the analysis is to count only accesses to shareable and modifiable persistent data structures. The result is passed to the scheduler before execution of transaction operations. In this way, we obtain a scheduler that implements a conservative two-phase locking protocol for a persistent programming language without the transaction programmer needing any further information.

A prototype of this tool has been implemented using the functional language Miranda and runs on a Sun workstation. It has been used successfully

in a few prototypical applications. Further experimental work is necessary to learn how to relate the granularity of the analysis with the peculiarities which characterize a class of applications. The aim of our project is more ambitious. We intend to investigate in a general setting how information on the dynamic behaviour of a transaction or a set of transactions can be used to improve the efficiency of an O-O Database Management System. In addition to the author, other current project members are G. Mainetto and G. Amato (CNUCE-CNR), P. Asirelli and P. Inverardi (IEI-CNR), and A. Wolf (University of Colorado).

Please contact: Fosca Giannotti - CNUCE-CNR
Tel: +39 50 593 285
E-mail: fosca@cnuce.cnr.it

Network Databases for Telecom

by Oddvar Risnes

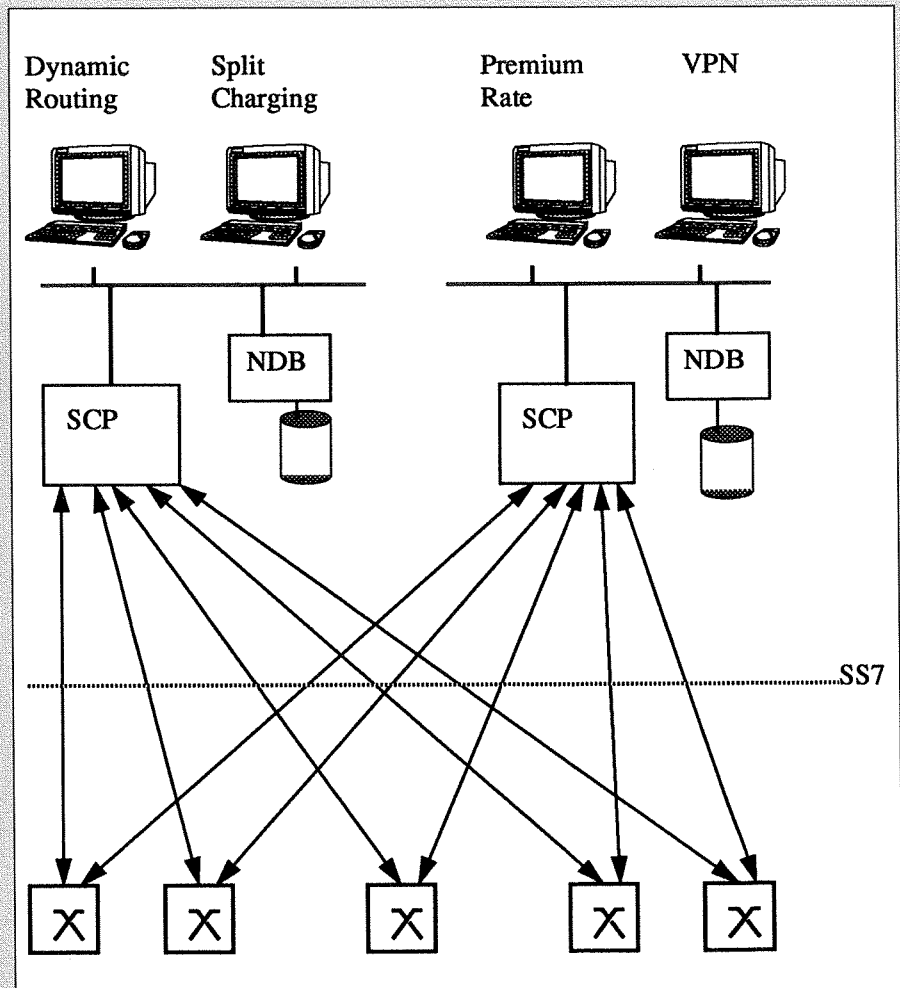
Existing database technology does not satisfy the demands by telecom industry. At SINTEF DELAB, research in Network Databases focuses on continuous availability and high performance transaction services.

Since the mid-eighties, the telecom industry has been looking into usage of database technology in network management systems, charging systems, and in future applications like intelligent networks. Network Databases has become a common term used for databases that are applied as a component in the telecommunication arena. Network Databases will in this context typically be used as an integrated part of a mission critical system.

Database technology has long proven its applicability as a cost-effective platform for building and maintaining information systems for administrative purposes. The existing technology, however, has so far not satisfied the operational requirements of telecom switching systems. For example, in the access to large real-time databases, service requirements such as response time, availability, transaction rates and fault-tolerance are not satisfied.

There is a growing understanding that services in the context of Telecommunication Management Network (TMN) and Intelligent Networks (IN) will heavily rely on large high-performance databases. Mobile telecommunication services (GSM, UPT, UMTS) have been stipulated to grow rapidly. There will also be a strong demand for highly available, high-performance Network Databases to support these mobile services.

Research at SINTEF DELAB in Network Databases for Telecom was initiated, and primarily financed, by Norwegian Telecom about five years ago. The requirements to network databases as



Network Database in Telecommunication System Architecture.

The Network Database Server (NDB) is depicted as a Service Data Point connected to a LAN-segment that holds a Service Control Point (SCP) and some specific application services.

specified by Norwegian Telecom Research were:

- transaction response time of 15 milliseconds for more than 95% of the TPC-B equivalent transactions
- transaction capacity of at least 1000 TPC-B equivalent transactions per second
- an accumulated unavailability at transaction service level of less than one hour over 30 years mission time.

The activities of SINTEF's Database Technology Group have focused on developing DBMS technology that meets these requirements. The focus has been on software based fault-tolerance, high availability, and continuous availability as opposed to the traditional hardware fault-tolerance focus.

Our research on Network Databases is based on a shared-nothing multiprocessor architecture. All nodes are equal and equipped with CPU, main memory, and

disks. The nodes are interconnected by a fault-tolerant, high capacity, and low latency communication network. We are now starting a new project based on ATM interconnection technology which gives the ability of wide area geographical distribution and still fulfilling the transaction response time requirements. The use of ATM technology also simplifies inter-operation with other telecom components.

The Network Database is distributed over the nodes in the multi-processor. This gives the required transaction processing capacity and the ability to scale the capacity incrementally. Data is also replicated so that if one node or site fails, another data replica of the lost data is still available. We have developed replication methods that can handle two or more replicas of the data at table granularity with one replica being the primary and the others being hot stand-by replicas. This facilitates fault-tolerance at a transaction service level.

To meet the availability requirements to Network Databases, we have focused on developing algorithms for seamless take-over in case of a failure and automatic re-establishment of the original fault-tolerance level after a failure without compromising transaction availability. The replication method we have developed enables the system to support database maintenance, reconfiguration, and self-repair operations to run in parallel with normal transaction processing.

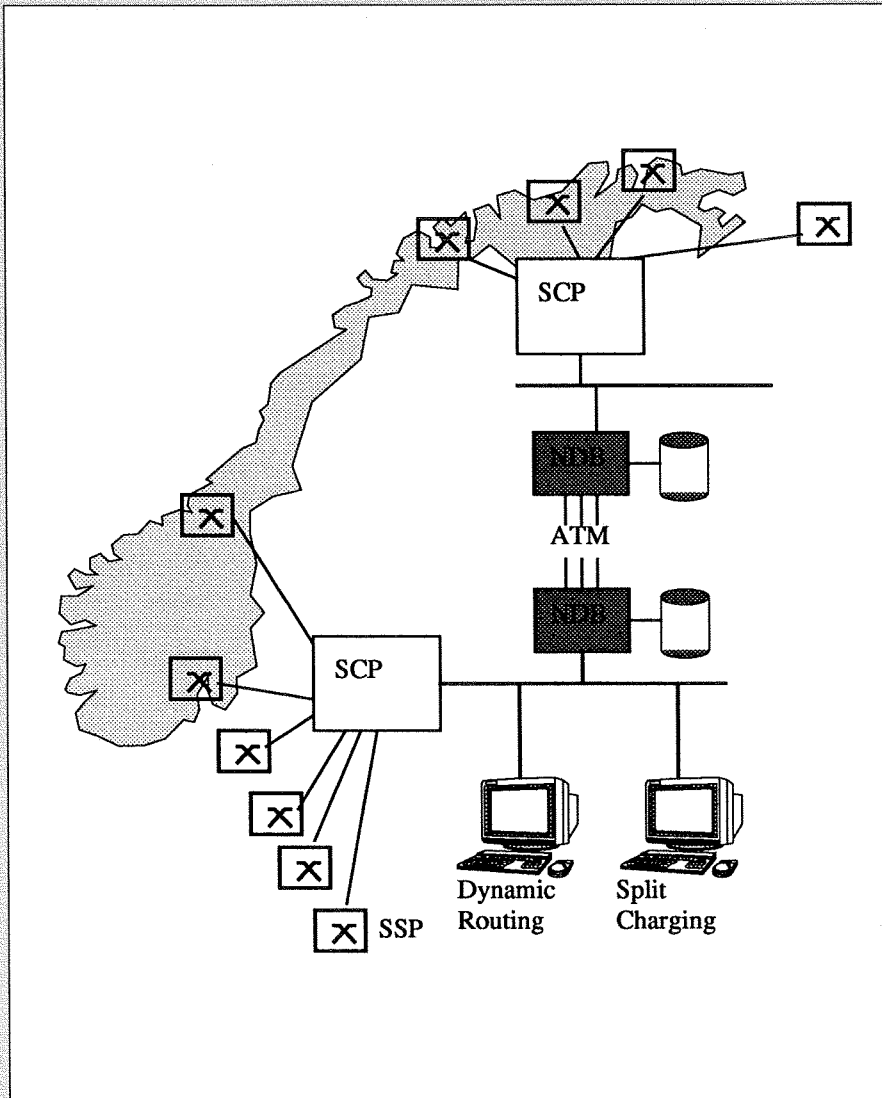
SINTEF DELAB will participate in a three-year programme of Norwegian Telecom to develop an experimental prototype Network Database with continuous availability and high performance capabilities. The goal is to demonstrate its applicability in handling large real-time databases in mission-critical applications. The 'hottest' application is switching telecommunication networks to support IN services like Universal Personal Telecom (UPT). Other goals are to demonstrate:

- high performance characteristics of high transaction loads (20.000 TPC-B transactions per second)
- real-time response (5 ms)
- extremely large databases (Terabytes)
- handling complex queries without degrading the transaction capacity and the real-time response
- characteristics of systems running around the clock all days of the year, which are less than 30 seconds unavailable per year.

In order to achieve continuously available services throughout the mission time of a system, the Network Database must support database maintenance, reconfiguration, and self-repair operations to run in parallel with normal transaction processing.

Please contact : Oddvar Risnes - SINTEF DELAB, Database Technology Group
 Tel: +47 73 59 27 61
 E-mail: Oddvar.Risnes@delab.sintef.no

Future Network Database Server Architecture.



Distributed Database R&D at SINTEF DELAB

by Tor Didriksen

Research on Distributed Databases has been a major area for SINTEF DELAB since 1985. After the successful development of the MIMER/STAR project, a number of applications for various clients has been completed. The total R&D effort in this area has been about 30 person-years.

The aim of the MIMER/STAR project was to develop a distributed DBMS prototype for MIMER Software AB in Sweden. The main focus was on transparency issues in distributed databases and how the system can provide a set of mechanisms allowing the user to trade between data independence, data availability and efficiency.

The REPLIC Demo Facility for Command Control Information was developed for Shape Technical Centre in Den Haag. The goal of the project was to demonstrate the use of common data definition, manipulation and communication standards to achieve interoperability between dissimilar equipment, running different operating systems and different database management systems. A further goal was to investigate a trade-off between the amount of availability and the degree of serializability which exists in replicated distributed databases. The REPLIC Demo Facility was based on standard off-the-shelf components and demonstrated interconnection of heterogeneous subsystems. The system used Ingres and Oracle DBMS, ran on MicroVAX and Sun hardware platforms using VMS and Unix Operating system. The communication protocol used was DECNET. The system was coded in ADA.

The TelSQL Application Platform was intended as a prototype DBMS platform for Telecommunication Management Network applications and developed for the Norwegian Telecom Research

Department. The project aimed for a new system architecture in which operation and management could be implemented independent of the telecommunication system. TelSQL was based on ANSI/ISO SQL and provided two extensions to the standard SQL language. First, a data distribution and replication mechanism, enabling transparent access to stored objects, coordination of a large number of sites, and transparent distribution so that local systems may run autonomous operations. Second, an enhanced trigger mechanism, enabling event and time triggering across databases participating in the distributed database. ISODE (ISO Development Environment) was used to implement communication protocols. The compiler and runtime system were implemented in Prolog, and ran on top of Oracle DBMS.

STAROS was a system developed for the Norwegian Defense Research Establishment, and integrated ideas from the Replic Demo Facility into the TelSQL system. The design focused on the trade-off between functionality and performance. Relaxing some of the replication and trigger functionality improved the performance by an order of magnitude.

Dibas is an ongoing project sponsored by the Norwegian Telecom Research Department. The basic idea is to partition a centralized database into fragments, and assign an owner to each fragment. Data is replicated according to user and application needs, so that all data access can be done locally at each site. Data imported from (owned by) other sites are read-only. This is implemented as a general extension to Sybase DBMS. There is no modification of the DBMS, and no modification of any user applications. A rule compiler produces Sybase SQL triggers and transact SQL statements. The triggers will abort transactions that try to modify read-only data. SQL statements select the desired subset of the data to be replicated at other sites. A separate distribution application performs asynchronous propagation of replicated data. A demo version of Dibas is running, and we expect system deployment in the first half of 1994.

Please contact:
Tor Didriksen - SINTEF DELAB
Tel: +47 7359 7079
E-mail: tor.didriksen@delab.sintef.no

PROSPECT Becomes Industrialized

by Per Holager

PROSPECT is a prototype database application generator, with the novel aspect that working application forms are generated from a tree, set up by 'clicking' in an Entity-Relationship (E-R) diagram. Thus, the specification of an application is a visible, comprehensible structure instead of being the aggregate of whatever was written into a stream of 'pop-ups' that are long gone and forgotten.

These ideas are now being industrialized by Sysdeco Innovation AS (Oslo) and SINTEF DELAB. The initial version of the product is a prototyping subsystem of Sysdeco's 'Systemator' suite of application development tools, consisting of an E-R diagram editor (Data Selector), a Form Structure generator and editor, and the Application Interpreter.

The E-R editor is a rather straightforward graphical editor specialized in handling E-R diagrams. The Data Selector editor makes it possible to set up and edit trees of Entities connected by Relationships selected from the E-R editor's window. A Data Selector diagram specifies the data from the data base that should be presented in an application form on the screen.

A Form Structure is a tree (with some cross-links) specifying the window, 'boards', variable fields, scroll-bars, menus and what not, that make up the form one sees on the screen. The Form Structure generator takes a Data Selector tree and produces a Form Structure suitable for displaying that data selection in the database. With the Form Structure editor one can move subtrees, delete and extend the specification, change the presentation of variable value of enumerated type from textual to 'radio button' style and change a scrolled table to a 'book' for presenting a set of values.

Finally, the Application Interpreter will instantiate a Form Structure to produce

the specified form image on the screen, keep a buffer of database values and presents these in the variable fields of the form, and also realizes a set of commands for navigating in the data, modifying individual fields, inserting and deleting records, instanciating more forms, etc.

For now, this product is for prototyping only because the data buffers cannot communicate with any 'real' database, only with a limited set of example data stored in the development dictionary. In a full-blown development system, one must also have a facility for programming one's own commands and for doing detailed editing of the two-dimensional appearance of the forms.

Please contact:

Per Holager - SINTEF DELAB

Tel: +47 73 59 42 09

E-mail: per.holager@delab.sintef.no

Data Mining – Treasure Digging in Databases

by Henk Nieland

More and more data is being stored in databases. These contain treasure troves of strategic information, which are difficult to unearth when using common database management systems. However, this is possible with the help of 'data mining', a recently developed technique, which unites elements taken from database theory, statistics and the theory of machine learning. Having expertise in all three fields, CWI recently launched a Data Mining project.

Databases grow very rapidly, both in size and in number. Estimates of the average growth in size range from a redoubling every three years to a tenfold increase per decade. Apart from the factual data stored in a database, the information derived from it is of particular interest. Take, for example, an insurance company which uses a database with date information on its car insurance. It will

be very simple for such a company to check how many of its clients have had accidents. In addition, the database contains other, implicit information of interest to the company, such as risk profiles, yet this information is much more difficult to trace.


Data mining is the technique of searching for such distinctive patterns in a database. First, we gather the information stored in the database into clusters, for example separating accident prone persons from safe drivers. Then we look for a distinctive description for each group in such a cluster. The patterns thus obtained provide a strategic understanding of the contents and coherence of the database. For example, the risk profiles found may be used to raise insurance premiums for high risk groups, while premiums for others may be lowered. Data mining helps us find a high level description of the database in the form of a classification into groups, each with their own distinctive features.

The current generation of database management systems is not suitable for data mining. The systems can be applied to test a presumed risk profile, but the profile has to be drawn up by the user himself. And if the user is not familiar with the profile, he will face too many possibilities to test them all. A data mining tool limits the number of possibilities by combining human insight (not all clusters and far from all patterns are significant) and smart algorithms. This also enables us to find unknown links within the database.

Data mining is already being applied in practice. It has enabled IBM, for instance, to make considerable cuts in the number of final (and expensive) tests of disk

drives by drawing up a profile which allows early identification of disk drives that are likely to fail in those final tests. Another example is the British banking house TSB, which, on the basis of data on old loans, defined general patterns facilitating the assessment of the credit worthiness of new customers.

Data mining has four key aspects. The first is gathering database information into clusters. The quality of the patterns found largely depends on the clusters being used. Once the clusters are made, a distinctive pattern will have to be found for each individual group. The patterns will then be put into a special representation language. One of the main aims of our research is to design a representation language which will speed up the process of finding easily understandable and adequate descriptions. In order to find the best description, a quality function is used, which indicates for each pattern how well it describes the cluster in question. The search operation can now be visualized as a hilly landscape, where altitude indicates quality levels and where the purpose is to spot the highest peak as soon as possible. An abundance of useful algorithms is available for that purpose. The final aspect is that databases always contain false information (noise), as a result of typing errors for instance. And some times information is missing because it was deemed unimportant at the time of compilation. To a large extent, noise can be removed with the help of statistical techniques like deletion of strongly deviating examples. Moreover, in practice, data mining has proved to work fairly well on databases containing noise. The absence of information, however, is basically an insurmountable problem. The only solution here is to work with



rules of probability, which indicate the chances of a certain object belonging to a certain category, if it meets certain conditions.

In mid 1993, CWI's database research group launched its data mining project with a view to creating a data mining environment for modern database management systems. Its design activities focus on the integrating results taken from statistics, data banks and knowledge banks, as well as machine learning, in order to guarantee both the efficiency of data mining and the reliability of its results. Since data mining is an interactive process, much attention is focused on designing a user interface.

Please contact: Arno Siebes or Marcel Holsheimer - CWI
Tel: +31 20 592 4139 (4134)
E-mail: arno@cwi.nl

Database Research at the University of Twente

by Peter Apers

The University of Twente's department of computer science includes a database group of about fifteen researchers. Its research interests are in the field of object-oriented, logical, and extended relational systems, in particular:

- development of well-founded (object-oriented) data models and tools
- design of and optimization in database systems, with a strong emphasis on distribution aspects
- database support for new (technical) applications.

Below some current research topics are highlighted.

TM

TM is an object-oriented model that includes standard object-oriented

features, such as arbitrarily nested attributes, methods, and inheritance. Specification of methods is done in a functional language, and set comprehension may be used for the description of, e.g., attributes. TM also allows for constraint specification in a first order logic language, and constraints and methods are subject to inheritance, just like attributes. The theory is described in a formal language based on the Cardelli type system, set theory, and lambda calculus. A logical query language allows a declarative formulation of ad-hoc queries on a TM database specification.

The integration of TM and LOTOS, a protocol specification language, will be undertaken in the ESPRIT-BRA project TRANSCOOP, which aims at providing computer support for cooperative applications. Some specific applications like Design for Manufacturing (DFM) will be used to check whether the integrated language suits its purposes.

Tools

A design tool for TM has been built that provides a graphical representation of TM specifications and is methodology independent. The tool is currently extended to include algorithms to analyze the specifications concerning internal consistency, safeness, etc. We are also working on translating TM to existing platforms (OO databases, DBPLs).

Another notable tool, the prototype generator, generates a quick-and-dirty prototype of a TM specification. Its architecture allows for prototyping in different target languages.

These tools are developed in the framework of ESPRIT project IMPRESS which aims at obtaining a reliable, distributed object server for technical and multimedia applications.

Complex Objects

The representation, manipulation and storage of complex objects have been studied in applications including CAD/CAM, multimedia, and cartography. A storage server for complex objects was implemented under the Amoeba distributed operating system. A SQL-like query language for TM is being translated into an algebra for complex objects based on the NF² algebra, including logical optimization of algebraic expressions. The emphasis is

on translation and optimization of nested queries.

Physical database design

DBMSs offer structures, like hashing and indexes, to allow storage of data and fast processing of queries and updates. Given, e.g., a query and a physical design, the query optimizer will select a fast way of handling the query. During physical design a database administrator has to decide which structures to use for a specific database and a specific load. Usually heuristics are used to obtain a good physical design. Current research focuses on combining heuristic rules and inter-action with the optimizer.

Applications

Databases are being used more and more beyond traditional administrative environments, demanding new capabilities, such as new data types and more powerful query languages. Moreover, many design decisions, including the traditional transaction concept and query optimization techniques, need to be reconsidered. Distributed processing (parallelism) may turn out to be essential for these new application areas.

Currently, we are focusing on Geographic Information Systems (GIS), in particular the optimization of queries. Possibilities offered by multimedia, interoperability, and data mining are taken into account.

Please contact: Peter Apers
Tel: +31 53 893 719
E-mail: sandra@cs.utwente.nl

LIKE – Linguistic Instruments in Knowledge Engineering

by Hans Weigand and Reind van de Riet

The design of information systems require - aside from traditional software engineering tools- linguistic instruments to describe and analyse the intended semantics of the application domain. LIKE is a national project to develop such linguistics tools.

Computer science has contributed much to the development of modern linguistics over the last three decades, leading to a branch of research called *computational linguistics*. The progress made enables researchers in information systems and communication technology to benefit from its results. This is the prime objective and working hypothesis of the LIKE project group in The Netherlands.

Information systems with (un)formatted data are becoming the fastest-growing resource in companies. However, finding relevant data items in either a text base or a formatted database is becoming harder - unless measures are taken to manage the information content properly and to assure accessibility from the start. In particular, document management systems and office automation tools do not address the problem of deviations in word use and word meaning.

In our view, the problem of controlling the meaning of words is becoming a key issue to improve the DBMS and EDI systems currently offered, since transmission of data is only useful when the sender and the receiver agree on their meaning. Data Dictionaries are of currently limited help, because they lack linguistic knowledge to sort out semantic and morphosyntactic aspects; they lack a formal basis to reason with the meaning descriptions; and they are usually 'closed systems'.

In the area of Information Retrieval, the sauri are used intensively to find information in a knowledge- and document base. However, it is still not well known how the methods used by linguists to attach meaning to words can be used for large collections in an efficient and flexible way.

Traditionally, in computer science many words are being used, which usually are only significant in the framework of the program/specification or design.

A central theme of the LIKE project is to profit from the application domain knowledge carried by its intended users and to provide tools to capture and reuse its formalized descriptions, i.e. *lexicons*. Concrete objectives are:

- item to apply formalisms(s) developed in linguistic theory to express meaning
- item to represent contextual information of the environments in which communication takes place
- item to develop a Lexical Management System for the management of large collections of words together with their meaning
- item to check for correct combinations of words (collocations).

Please contact:

**Hans Weigand -- Katholieke Universiteit
Brabant, Tilburg
+31 13 66 9111**

E-mail: weigand@kub.nl

or:

**Reind van de Riet -- Vrije Universiteit,
Amsterdam
+31 20 548 4177**

E-mail: vdriet@cs.vu.nl

Database R&D Activities at INESC

By Rogério Carapuça
and Amândio Vaz Velho

Exploring the object-oriented (O-O) paradigm in the database field, the INESC Information System Group has developed a semantic object model, including an application environment, investigated the building of a distributed object server, and participates in several other projects, e.g., Hospital 2000 (ESPRIT Special Action).

Object orientation has progressively eliminated the separation between databases and application programs and has blurred the distinction between analysis, design and implementation, thus bringing a seamless development process closer to realization. In our view, at least two prerequisites are necessary to achieve these goals:

- a uniform data model, which smoothly combines abstract modeling features (associations, roles, life cycles, etc.) with O-O programming features (classes, attributes, methods, inheritance, etc.)
- an integrated tool set, enabling uniform object manipulation in the domains of problem, solution and implementation.

Based on its previous extensive experience in relational DBMS, as well as in semantic data models and supporting CASE tools, our group aims to develop a full O-O seamless development method and environment that supplies those prerequisites. We have developed a semantic object model (SOM), which can be simultaneously used as an O-O analysis and design notation and as an O-O programming language. This is achieved by integrating traditional data modeling techniques within an O-O framework, while preserving principles as encapsulation and extendibility. Specifically, the concept of attribute, understood as instance variable in O-O programming languages, is semantically enriched in such a way that it becomes able to abstractly capture the semantics of

structural relationships, without any breakdown in encapsulation and extendibility. SOM supports the modeling of evolving real world entities, by allowing an object to change type during its lifetime. Such an entity, playing various roles, is a single object (with a single identity), which may have more than one type as well as change its set of types.

In creating a SOM application environment, we are developing an integrated tool set, SOM-CASE, for manipulating the SOM schema and generating some application components, as well as a software layer, SOM-DBMS, able to provide data manager services for objects defined with the SOM data model. SOM-CASE will include a schema editor, a browser and a user interface generator – all integrated graphical tools. They share a common repository and exhibit the same interaction paradigm. We classify SOM-CASE tools as semantic tools, meaning that they are able to explore the semantics of SOM concepts to improve interaction with users. The SOM-DBMS main modules are the parser, the repository and the translator SOM-C++. The translator converts one SOM object into various C++ objects, which, together, preserve the SOM semantics. In the present stage, SOM-DBMS relies on a commercial O-O DBMS to provide data manager services.

In order to replace in the long run this commercial DBMS, our group is also building a distributed object server, called SOAP (Servidor de Objectos e Aplicações). Two modules are presently being developed: a storage manager

(SOAP-SM) and a distributed transaction processing manager (SOAP-TM). SOAP-SM is based on versions, used to simplify the task of building recovery and increase the level of concurrence. Performance is improved by clustering facilities, and using flexible replacement policies for buffering pages and objects. SOAP-TM provides distributed transaction facilities. It will be integrated with an object manager, which allows information exchange between different sites, and an application programming interface that handles objects. The integrated environment obeys to the X/OPEN Distributed Transaction Processing model. Hence, at the end it will exhibit a modular architecture with open and standard interfaces.

The group also participates in several ongoing projects in neighboring fields. Examples include the ESPRIT Special Action Hospital 2000 which aims to build a complete software environment for Hospital Information Systems, and work on methodology engineering, meta-modeling and geographic information systems, as well as software engineering project management and software quality.

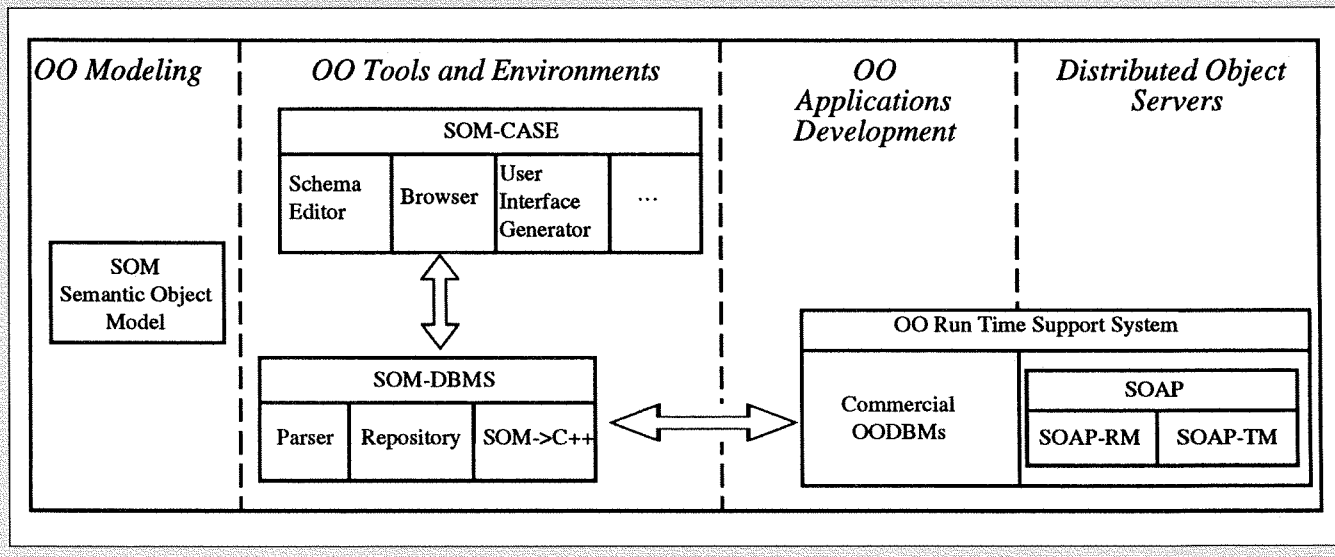
Please contact:
Rogério Carapuça - INESC
Tel: +351 1 310 0007
E-mail: carapuca@avila.inesc.pt

Applied Database Research at INESC

by João Barreira

Information systems based on powerful computational facilities have become an indispensable instrument for most institutions in the public and private sector. However, many existing information systems are slow and inefficient because of a lack of awareness of the internal strategy of the relational database management search system, and of the possibility of fine tuning the database parameters. INESC database management specialists substantially contribute to the nation-wide introduction of modern administrative systems, e.g., by training and consultancy.

The success of any institution operating in a competitive economic environment depends more and more in a critical way on well-operating information systems, based on powerful computational facilities. Sometimes these systems operate in a high risk situation, where efficient management is decisive in making or breaking a project. INESC has considerable experience with introducing sophisticated computational tools in society. As a result, our specialists have made decisive contributions to the definition of appropriate methods, taking





into account critical success factors, to implement modern information systems and analyze the various stages of their operating cycle from different technical, organisational, financial and human viewpoints.

In view of the constant requirement of information systems to make the maximum possible use of data (both image and voice), data distribution and information exchange with other information systems, client-server architectures and graphical interfaces, various database management systems have been studied and tested, in order to identify the most suitable system for each specific application situation. Special attention was given to the use of their utilities, 4GL programming languages, loading, handling and visualisation and data recovery, connections with other systems and tuning. These activities have led to training activities for

programmers and database managers, and to project support and consultancy, inside INESC as well as on behalf of companies.

The specific features of distributed information systems and the increasing complexity of distributed applications require new methodologies, models and abstractions in order to develop robust, easy to maintain and expandable distributed software. A well established concept in distributed processing is the client-server model. Within this model we considered the architecture for transaction processing – a concept enabling the maintenance of data integrity in a system subject to system node failures. In particular, we analyzed the architecture of Open/OLTP, which was developed for UNIX systems in accordance with the X/Open system model, with special reference to the following topics:

- facilities and flexibility
- supply of development interfaces with specific regard to supported languages, applications and database systems
- technical hardware and software requirements, integration with other systems, etc.
- performance figures
- robustness, specifically the maintenance of data integration caused by hardware and software faults.

INESC carries out automatization projects in the management sphere on behalf of institutions in the fields of tourism, personnel training and radio communication. The Directorate General for Tourism commissioned an information management system for national tourist resources, in which multimedia information on every national tourist resource can be kept up-to-date. On behalf of the Institute of Employment and Vocational Training (a body belonging to the Ministry of Employment and Social Security), a management information system was designed for the Institute's operational training programmes, enabling the automatic collection and processing of data related to the analysis of projects, as well as automatic management of decision making procedures. Finally, the time required for various procedures in an information system, used by the Companhia Portuguesa Rádio Marconi for 30.000 employees, was reduced by optimizing the formal database parameters and information access strategies. In particular, the work involved:

- analysis of data access information strategies
- definition of search indices
- restructuring of databases from single to multi-file
- adjustment to the formal parameters defining the database, for optimising the use of operating resources.

As a result, the program's slowest general selection time was reduced from 24 to only 15 hours and the calculation of the adjustment of subsidies from 3 hours to 22 minutes.

Please contact: João Barreira - INESC
Tel: +351 1 3100 075
E-mail: jmmb@avila.inesc.pt

AMOS – Active Mediators Object System

by **Tore Risch**

The classical relational database languages are not powerful enough for the manipulations needed, e.g., to build advanced models to filter and extract interesting information. AMOS (Active Mediators Object System), developed at Linköping University, Sweden, is an architecture to model, locate, search, combine, update, and monitor data in information systems with many workstations connected, using fast communication networks.

Future computer supported engineering, manufacturing, and telecom environments will have large numbers of workstations connected with fast communication networks. Workstations will have their own powerful computation capacities which store, maintain, and perform inferences over local engineering data- and knowledge-bases, or information bases. Each information base is maintained locally by some human operator and is autonomous from other information bases. Each information base will need a set of DBMS capabilities, e.g. data storage, a data model, a query and data modelling language, transactions, and external interfaces. The classical relational database languages are not powerful enough for the manipulations needed, e.g., to build advanced models to filter and extract interesting information. Facilities are also needed to support 'reactive' applications that sense changes in information, i.e. active database facilities.

AMOS (Active Mediators Object System) is an architecture to model, locate, search, combine, update and monitor data in information systems with many workstations connected, using fast communication networks. The approach is called 'active mediators', since it introduces an intermediate level of 'mediator' software between data sources and their use in applications and by users, and since it supports 'active' database facilities. A central part of AMOS is an

object-oriented (O-O) query language, AMOSQL, with O-O abstractions and declarative queries. The language is extensible, in order to allow for easy integration with other systems. This allows for knowledge, now hidden within application programs as local data structures, to be extracted and stored in AMOS modules.

A distributed AMOS architecture is being developed where several AMOS servers communicate, and where queries in a multi-database language are allowed to refer to several AMOS databases or other data sources.

We are working on the following classes of mediators:

- 'Integrators' that retrieve, translate, and combine data from data sources with different data representations
- 'Monitor models' that notify mediators or application programs when interesting data updates occur. A rule mechanism based on O-O queries has been added to AMOSQL. Work is in progress to use AMOS for industrial control applications
- 'Domain models' that represent application oriented models using application-oriented database operators
- 'Locators' that locate AMOS servers in a computer network.

Articles on AMOS can be ordered from the Department of Computer Science, Linköping University, 581 83 Linköping, Sweden.

**Please contact: Tore Risch
Tel: +46 13 282 541
E-mail: torri@ida.liu.se**

INTUITIVE – Interactive User Interface and Tools for Information in a Visual Environment

by **Peter Rosengren**

INTUITIVE handles access to a number of diverse and heterogeneous databases that contain various types of multimedia data. It includes advanced user interface facilities supporting the users accessing information. The toolset enables development of flexible and intelligent user interfaces with modes such as voice, pointing, mouse and keyboard.

INTUITIVE is part of an ESPRIT project in which an advanced user interface management toolset is being designed, supporting the task of information retrieval by enabling users to access information held in a number of different formats (for example, text, graphics, sound, video and pictures) and located in a number of source databases. It is aimed at companies and organizations with a need to access a number of diverse and heterogeneous databases that contain various types of multimedia data. Three applications are under development within the project. These applications are in the areas of medical information retrieval, emergency and damage control on board ship, and office documentation. Accompanying the advances in multimedia database storage technology there is an increasing need for advanced user interface facilities to support the user in accessing information.

The INTUITIVE toolset enables development of advanced information retrieval systems to develop efficiently flexible and intelligent user interfaces – exploiting a range of interaction modes (such as voice, pointing, mouse and keyboard). Support for multiple distributed standard databases, with meta-level

data dictionaries providing the mechanism for integrating customer-supplied data within an INTUITIVE environment. Internal client/server architectures and virtualisation of platform specific services allow the work environment to be configured for a range of current-day and future hardware configurations. Those are two database oriented features to be mentioned.

Please contact: Peter Rosengren - SISU
Tel: +46 8 752 1600
E-mail: peterros@sisu.se

EMOS – Enhanced Multimedia Object System

by Lars-Aake Johanson

The Enhanced Multimedia Object System (EMOS), to be delivered by the three-year ESPRIT project LYNX, supports long multimedia data storage needed for real-time video and audio digital data. EMOS can be incorporated in different distributed programming environments allowing clients to access different types of servers. An Object Request Broker provides a uniform view.

The LYNX project makes strong use of the partner's results from previous in-house, national and ESPRIT projects to deliver an Enhanced Multimedia Object System (EMOS) using advanced hardware and software providing powerful, and cost-effective, object-oriented multimedia support. This will be incorporated into any kind of distributed programming environment, allowing clients to access many types of servers.

The kernel of the LYNX EMOS is the Object Data Server (ODS), an ODBMS with enhanced support for long multimedia data, as required for the storage of real-time video and audio digital data. The ODS provides a Canonical Model and Declarative Query Language. It supports C++ and will be

fully portable to a wide range of UNIX platforms. Clients will run on UNIX or DOS machines. An Object Request Broker will provide a uniform view of the system to application developers, having EMOS servers as one type of component. Transparent access to objects stored in different systems will be provided. As performance and price are crucial selling factors in the target markets, faster and cheaper hardware is required. The LYNX Object Storage Processor will be developed for this purpose, based on existing multiprocessor APD machine. Multimedia I/O client devices will directly communicate with the Storage Processor, under the control of the ODBMS. The EMOS system will be validated by a demonstrator that has been selected among several application areas in the industrial and business fields. The project has just passed the first year of the three year project.

Please contact:
Lars-Aake Johanson - SISU
Tel: +46 8 752 1600
E-mail: johansso@sisu.se

MILORD

by Ulf Wingstedt

The MILORD project produces a unified framework integrating functions for database management and teleconferencing services in a Health Information System. This is an environment characterized by a large number of distributed information sources and different user categories with varying information needs. The MILORD solution is a federated system for information exchange between a number of autonomous but interconnected subsystems.

The work addressing cooperation and communication issues in the AIM project MILORD, has concentrated on providing specifications of requirements and architecture for a federated HIS (Health Information System) environment. The objective is to provide a unified framework integrating functions for database management and teleconferencing services in a HIS. This work

has also included the adaptation and reuse of results from ESPRIT-2 project KIWIS.

The work performed by SISU has been focused on the definition and description of a distributed information systems architecture for the HIS environment. To achieve this, the principles and techniques of federated databases and information systems have been chosen. The motivation for this choice is that the HIS environment is characterized by a large number of distributed information sources and different user categories with varying information needs. A federated system which can provide information exchange between a number of autonomous but interconnected sub-systems is thus appealing. At an early stage in the project, the most basic concepts and functions relevant to the MILORD environment were analyzed and modeled in close cooperation with the medical partners of the consortium, so as to provide a basis for the development of the MILORD architecture.

The first description (specification) of the federated architecture has been presented as a three-layered reference model, the intention being to separate the higher-level database and application functions, the client/server interactions, and the communications infrastructure from each other. The upper (application) layer of the model includes the schema architecture of the MILORD federation, consisting of a collection of schemata used to define interfaces between the nodes (e.g., different clinics or other organizational units) in the federation. The function of these schemata is thus to define the information objects that may be exchanged or shared among nodes in the federation. Another important aspect of the application layer is the description of user functions in terms of both teleconferencing and database access. The two lower layers of the model include description of the interfaces to the MILORD system components, such as the interactions between client/server processes supporting local and remote database access as well as teleconferencing.

Please contact: Ulf Wingstedt - SISU
Tel: +46 8 752 1600
E-mail: ulf@sisu.se

Computational Steering

by Henk Nieland

Ever since computers induced the massive production of scientific data, visualisation has been a powerful aid for the analysis of these data. CWI's recently started Scientific Visualisation project focuses on computational steering.

Graphical display of a series of measurements is already high school routine. Somewhat more complicated, but still quite common, is the display of the temperature distribution in an iron plate as a mountain landscape above the plate or using colours. In practice spatial phenomena frequently require visualization. For example, the flow around an aircraft wing can be visualized by drawing in each point the velocity vector and representing the air pressure in colour. However, several phenomena are described in higher-dimensional spaces and hence difficult to imagine. An example is a statistical cluster analysis. Insight in such phenomena is

increased by visual inspection of certain 2D or 3D cross-sections.

The dramatic increase in computing power and possibilities for graphical display has enabled the performance of realistic simulations, based on mathematical models and using form, colour and motion. Combined with man's visual abilities this yields a powerful research tool, particularly suited to tracing and monitoring structures which otherwise would remain hidden.


In common practice, on the basis of given input data a simulation generates output data which is visualized and analyzed, often leading to adjustment of the input data, a subsequent simulation, etc. Computational steering seeks to integrate the specification of input data, simulation and visualisation of output data into one system, thus providing the researcher with a direct feedback between changes in the input and their effect on the output. Because of the increasing complexity of simulated models the need for computational steering techniques will rapidly grow.

Navigation is a central concept in such an integrated system. By systematically

changing the input as a function of time and guided by the visualisation system the researcher tries to reach a certain goal, for example a function's local maximum. The computational steering system should also contain tools for 'logging', enabling backward searches by keeping track of the input data as a function of time, and 'heuristics', in which input variables are related to certain restrictions on the output. Another important function of the system is that of a 'watch-dog': usually the researcher is not interested in all details of the simulation process, he only wants to be warned if certain conditions are satisfied, such as convergence or instability.

At present CWI studies, in cooperation with the visualisation centre of the Dutch Energy Research Centre (ECN), the visual specification and monitoring of input data, in particular the representation of a scalar function of several variables as a matrix of orthogonal, 2D cross-sections (HyperSlice). These 2D-slices lend themselves very well to interaction via direct manipulation. Several navigation techniques have been defined, for example the location of maxima, user-definable edit paths and contouring. These simple techniques are valuable aids for traversal of the N-dimensional variable space.

Please contact: Robert van Liere - CWI
Tel: +31 20 592 4118
E-mail: robertl@cwi.nl



HyperSlice (CWI-ECN)
visualisation of a scalar function of four variables as a matrix of orthogonal 2D cross-sections, centred around a point in 4D-space which the user can vary with the mouse.

CWI Participates in Air Pollution Modelling

by Henk Nieland

As a part of its multi-disciplinary research programme in Environmental Mathematics, CWI develops in cooperation with several institutions in The Netherlands new efficient algorithms for air pollution models.

Early in 1992 CWI embarked on a large-scale research effort in environmental mathematics. The programme's contributing disciplines include numerical mathematics and computational fluid dynamics, population biology, statistics and system & control theory. An important part of the programme deals with the mathematical modelling of pollution effects in the atmosphere and beyond. Air pollution caused by industrial emissions of toxic chemical substances is a major environmental problem. A well-known class of models – Long-Range Transport of Air Pollution, or LRTAP-models –, containing one equation for each separate chemical substance, describes advection (transport by winds), diffusion and the mutual chemical reactions. Realistic LRTAP-models, being 3D systems of partial differential equations of the advection-diffusion-reaction type rank

among the computationally most expensive models in environmental research and large-scale scientific computing.

In the EUSMOG project CWI develops, in cooperation with the National Institute of Public Health and Environmental Protection (RIVM), efficient algorithms for LRTAP-models to be used in regional smog forecasts on the European scale. Numerical issues involve numerical advection schemes, local grid refinement, stiff solvers and numerical software development.

In a related project, CIRK, CWI cooperates with RIVM, the Institute for Marine and Atmospheric Research of Utrecht University (IMAU) and the Royal Netherlands Meteorological Institute (KNMI) in the mathematical modelling of global transport and chemistry of trace constituents in the troposphere. Development and use of such models are essential in the analysis and prediction of expected changes in the atmospheric chemical composition due to ever increasing emissions of polluting substances. The project's specific goal is to develop a 3D model for the global troposphere, including exchange with the stratosphere, for predictions over extensive periods of time. The model will consist of a large system of time-dependent partial differential equations of the advection-diffusion-reaction type, supplemented with subgrid scale physical parameterizations.

Please contact: Jan Verwer - CWI
Tel: +31 20 592 4096
E-mail: janv@cwi.nl

Aircraft Noise Prediction Model

by Truls Gjestland


Most Aircraft Noise Prediction Models assume that the airport is located in an infinitely wide, flat area covered with soft grass. In a mountainous country like Norway, where the ground may be frozen half the year, this simplification often leads to large discrepancies.

The American based Noise Prediction Model, INM, is widely used for calculation of noise in the vicinity of airports. A number of countries routinely base their noise zoning for land use purposes on the results from this prediction model. Like most other noise prediction models, INM, assumes the world is flat and soft! For modelling purposes the airport is considered surrounded by an infinite plane with absorption corresponding to that of a soft grass covered area. Even Columbus knew the world was not flat. A quick look at a map of Norway will show that most airports are surrounded by mountains and water, and not soft grass covered plains.

SINTEF DELAB is currently developing general routines to be used with noise prediction programs where terrain parameters are taken into account. GIS based parameters are used to correct for shielding, real surface-to-aircraft distance and modified absorption (i.e. total reflection from water). These routines will yield correction factors of the order of 5-10 dB in extreme cases, and will improve the noise prediction models considerably.

In addition to developing terrain correction procedures, SINTEF DELAB has also been invited to participate in the further improvement of the INM program. A new generation program, INM 5.12, will be established by the US Federal Aviation Administration. SINTEF DELAB is the only non-American institution involved in this project.

Please contact:
Kare H. Liasjo - SINTEF DELAB
Tel: +47 73 59 27 51
E-mail: kare.liasjo@delab.sintef.no



Hand-Gesture Interaction Model for 3D User Interfaces

by Monica Bordegoni

Typically, in complex 3D and Virtual Reality systems a scene is represented resembling reality as close as possible. Scene objects have a realistic appearance and their behaviour is defined by real world's laws. The user interface of such systems has to be transparent, so that users get the impression of directly interacting with application objects, rather than doing it via a computer. A Model-World metaphor, where the world of interest is explicitly represented, and there is no intermediary between user and world, seems appropriate.

The scene appearance (output) is required to be perceived as it was real, as well as users' interaction with the system (input) needs to closely resemble reality. Human's behaviour in the real world tends to be reproduced when users interact with the virtual world. These systems require highly interactive I/O techniques and appropriate input devices, offering naturalness and intuitiveness.

The purpose of the research has been to stud a suitable interaction model, combined with the use of novel input devices, for interacting with 3D user interfaces of complex 3D and Virtual Reality applications. The system adopts a Tool-Object syntax-based interaction. This interaction structure syntactically simulates an action sequence typical of human's every day life: one picks up a tool and then uses it on an object.

Moreover, the conventional mouse-menu pair has been replaced with hand gestures, performed by using a hand-input device, as a way of interaction with user interfaces. Hand gestures are used for selecting tools (functions) and manipulating objects.

Finally, a proper feedback to users' interaction is important in that it helps reducing the users' cognitive load, providing information on the ongoing dialogue.

A system has been implemented for teaching and recognizing gestures performed by means of a hand-input device (a VPL DataGlove). For increasing the capability of the interaction methods, a force input device (a Spaceball) has been used in parallel with the hand-input device.

Please contact: Matthias Hemmje
Tel: +49 6151 869 844
E-mail: hemmje@gmd.de

Specification of Complex Systems

by Axel Poigné

The research group "Specification of Complex Systems" has been newly formed within the GMD-Institute for System Design Technologies with the objective to study systems combining software and hardware as well as, for instance, mechatronic components. In particular there are studied systems which have analog and digital components, are reactive and time critical, and are embedded.

Such heterogeneous systems need a new design methodology. Components of such systems can no longer be considered in isolation due to trade-offs (for instance, software solutions may be more flexible but hardware provides the necessary performance). We need design methods where initial system analysis and modelling is independent of target technology, but which support partitioning and mapping to appropriate technology according to needs of functionality and performance. Such design methods presuppose an interaction of the various design disciplines involved to enable well-founded, flexible and innovative decisions. Hence our research emphasizes the integration of established specification and analysis methods.

We now pursue the following themes:

- We investigate the interaction of realtime languages such as ESTEREL and LUSTRE with object-oriented design techniques while reconstructing an industrially designed mass-flow meter. The more long-term goal is to define, based on the experience gained, a language which supports partitioning in software and hardware codesign.
- Complementary, simulation and formal verification techniques are experimentally applied to validate parts of the design of the mass-flow meter. In the long-term, we aim at the integration of established technology into an "analysis-workbench" as part of a design environment supporting co-design of systems.

The group consists of about 15 researchers with roughly 7 researchers focussing on the particular themes sketched above.

Please contact: Axel Poigné
Tel: +49 2241 14 2440
E-mail: poigne@gmd.de

Research Programme TELECOM 2005

by Truls Gjestland

The Norwegian Research Council has launched an ambitious research program called TELECOM 2005. The primary goal for this program is to increase productivity and export in Norwegian industries. The program consists of several parts: post-graduate education, a strategic program for generic technologies and direct R&D-contracts.

Over the next eight years 42 new candidates will receive their Ph.D. (dr.ing) in telecommunication. These candidates come in addition to the normal post-graduate programs. A new one-semester course will also be offered at the Norwegian Institute of Technology, aimed at "re-educating" people with a degree in telecommunication. The idea is that this course should be taken every five to ten years by engineers and scientists involved in telecommunication R&D in Norway. The plan calls for about 30 MNOK (million Norwegian crowns) per year for these activities.

The strategic technology program will focus on basic technologies, first of all for mobile communication and high capacity networks. According to the plans, an extra 20-30 MNOK per year will be spent mainly for fundamental work in research institutes.

Four areas have been identified for increased R&D support: high capacity networks, maritime satellite communications, telecommunication in the transport business and application of telecommunication in small and medium sized businesses. National research grants to these activities are supposed to be increased by about 30 MNOK per year. The money will be mainly granted to industry to help them finance projects in cooperation with Norwegian research institutes. The total program has just been started and will be in operation from 1994.

Please contact:

Truls Gjestland - SINTEF DELAB

Tel: +47 73 59 26 45

E-mail: truls.gjestland@delab.sintef.no

TECHNOLOGY TRANSFER

Highlights of the World Transputer Conference '93 in Aachen

by David Johnston

This year's World Transputer Conference (WTC) joined up with the annual German transputer event, called the Transputer Anwender-Treffen (TAT). The combined event took place at the Eurogress Centre in Aachen 20-22 September, 1993.

The conference was run as five concurrent streams. Keynote talks and a closing panel discussion formed the plenary sessions. Hardware and software demonstrations from commercial suppliers and academic poster sessions took place in an exhibition area. It is rumoured that early silicon of the fabled T9000 was being demonstrated! A two-day tutorial programme followed the main conference.

Overall, the conference was a great success. I attended as a delegate and picked up a lot of positive feedback particularly from the some of the 55 young European researchers who would normally have been unable to attend, had they not received funding from the CEC! There were around 500 attendees, with many gratifyingly from the less affluent parts of Europe.

Transputers On The Road, Dr. Uwe Franke, Daimler Benz

The use of transputers to autonomously drive a car was demonstrated. I was impressed - believing that autonomous driving was still some time away. The trick is to scale the size of the problem to the hardware, which has to be both commodity cheap and not add significantly to the price of a car. To reduce the computational power required, there is intelligent selection of which parts of the video image are to be analysed and also careful selection of the visual clues used to identify cars, road signs and road markings etc. The system has already clocked up around 2000 km of autonomous driving on the autobahns of Germany. There were lively discussions amongst the delegates on the legality of this!

Architectural Advances in The T9000 Roger Shepherd, Immos

Two aspects of the architecture of the T9000 were treated: the cache and shared channels. Such an admirably clear explanation was given of the design decisions in developing the cache of the T9000, that one could see no other way! However, predictability of performance for real time applications is lost with a cache. The T9000 solution is to use half the cache as fast on-chip memory and the other half as true cache. We gain from predictable performance when required as well as general cache acceleration of a statistical nature. Instead of having to place critical data in-on chip memory as for the T800, the cache tables for this "internal memory" are set up to refer to the critical data - notionally in off-chip memory. The result is that instead of moving the data to the on-chip memory, the on-chip memory is cleverly moved to where the data is.

Please contact: David Johnston - RAL

Tel: +44 235 44 6738

E-mail: david@inf.rl.ac.uk

CNR Delegation Visits the Foundation for Science and Technology of the State of New York

by Franco Denoth

In October 1992 an agreement for collaboration between the Centers for Advanced Technology (CATs) of the State of New York and the Institutes of CNR was signed by Governor Mario Cuomo and the then President of CNR, Professor Luigi Rossi Bernardi. The scope of the agreement was to facilitate cooperation between the academic and the production worlds at an international level.

The signing of the agreement has been followed by visits from CNR delegations to study the role played by the CATs in the transfer of results and know-how between the academic and industrial communities. These Centers are sponsored by the State of New York with the objective of establishing contacts and setting up collaborations between university and business, and of developing commercially viable

technologies. The aim of a visit by a delegation in July '93 was to evaluate to what extent it would be possible to set up similar structures in Italy, in order to render the activity of CNR in the transfer of know-how to industry and to the business world in general more flexible and efficient.

One of the CATs visited by this delegation of particular interest for IT activities was the CASE Institute (Computer Applications & Software Engineering) at Syracuse University. The following research topics were discussed and identified as of interest for the agreement for cooperation:

- Software Engineering: Formal Methods; Specification Languages; Re-engineering and Reuse
- Multimedia: Applications; Instructional technology; Network Services; Databases

- Others: Multilingual Translation; Image Processing; VLSI; Teleteaching.

The initial forms of collaboration will include the exchange of researchers and students, the setting up of joint research projects and the organization of joint conferences/workshops. In order to promote the exchange of information, it has also been decided to encourage the setting up of specific information services on both sides which can be accessed through the network (gopher servers) and to activate distribution lists which facilitate the exchange of information between groups with common interests.

Please contact: Franco Denoth - IEI-CNR
Tel: +39 50 593430
E-mail: denoth@iei.pi.cnr.it

Franco-Russian Research Institute Opened in Moscow

by Pierre Népomiatshy

On November 19, 1993, the Franco-Russian A.M. Lyapunov Institute was inaugurated in Informatics and Applied Mathematics. This institute, created by a scientific agreement between INRIA and Moscow State University, is installed inside Moscow University, but is open to all Russian and French scientists working in the framework of precise collaborative actions and in the specified domain. It has been equipped with French workstations (DPX 20).


The Lyapunov Institute is starting with five joint scientific projects :

- Numerical Solution of Partial Differential Equations and applications in aeronautics and space
- Numerical Solution of Fluid Dynamic Problems

- Modelling and Evaluation of Computer Systems
- Probabilistic Problems in Computer Networks
- Parallelization and Parallel Compilers
- Object-oriented Operating Systems.

Please contact:
Pierre Népomiatshy - INRIA
Tel: +33 139 635 646
E-mail: pnepo@nuri.inria.fr

The inauguration of the Franco-Russian A.M. Lyapunov Institute, from left to right: V. Sadovnichy, Rector of Moscow University, A. Bensoussan, Director of INRIA, P. Morel, Ambassador of France in Moscow and the Vice-Rector of Moscow University. (Photo: INRIA)



RIACA – a New Institute for Computer Algebra Research

by Henk Nieland

The Research Institute for Applications of Computer Algebra (RIACA) was officially founded on October 1, 1993, following an initiative of the Foundation for Computer Algebra in The Netherlands (CAN), as a cooperative effort between CAN, the Kurt Gödel School (formerly RISC-Linz, Austria) and CWI.

The institute is supported by the Ministry of Education and Research, and by the Foundation SMC (to which CWI belongs) and CAN. RIACA is located at the Science Park Watergraafsmeer in Amsterdam – as are CAN and CWI. The expertise Centre CAN was founded in 1989. Its goal is to promote the use of computer algebra packages in The Netherlands, including their support. CAN also coordinates seminars and performs case studies on commission. The Kurt Gödel School, recently emanated from the University of Linz, carries out research in symbolic computation.

RIACA intends to play a pivotal role in research and development of applications of Computer Algebra in mathematics, science and engineering. It aims to be internationally oriented. It can provide financial support for about eight (temporary) scientific positions and host a few additional visitors who come with their own financial support. All in all, RIACA can host about twelve researchers. It is a good location for workshops and conferences.

If you want to have further information about RIACA, if you have ideas of projects that may be pursued in cooperation with the institute, or if you are interested in visiting RIACA for a shorter or longer period within the next two years, please contact the Scientific Director of RIACA.

Please contact: Arjeh Cohen - RIACA
Tel: +31 20 560 8470
E-mail: amc@win.tue.nl

AQuIS '93 – International Conference on Achieving Quality In Software


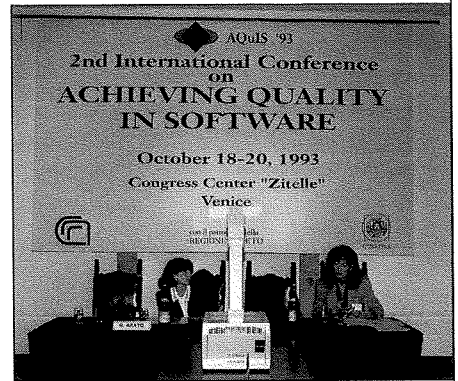
by Antonia Bertolino

From 18 - 20 October 1993, 80 delegates from 12 countries gathered on the Giudecca Isle, just off the fascinating Piazza S. Marco, in Venice, to attend AQuIS '93: the second in a biennial series of International Conferences on "Achieving Quality in Software", established in 1991 by IEI-CNR in collaboration with the Qualital Consortium, Pisa. The aim of AQuIS is to bridge the gap between theory and practice in the still immature field of software quality in order to speed up and render more effective the transfer of quality techniques into industrial practice.

The conference programme included sessions on: management issues, which included surveys on factors affecting the roles of software managers and professionals; effective and up-to-date methods, tools and techniques for achieving quality; metrics, in which various studies on measuring relevant aspects of software quality were reported; process improvement, reflecting the actual trend that focusses attention on the software development process; experience and technology transfer, in which important data and experiences from the industrial world were reported. Copies of the proceedings can be requested from: Piero De Risi, Consorzio Qualital, Piazza del Pozzetto 9, 56127 Pisa. Fax: +39 50 541753.

The objective of the conference was to assess currently viable techniques for quality achievement. Both important research results but also lessons from real-world experiences were shared among the participants, stimulating a common awareness of current issues and a willingness to try the directions indicated. The general impression, gained

informally during the discussions and also formally from questionnaires



compiled by the participants, is that a nonhomogeneous group of people, working in different aspects of the quality field, with different scopes and objectives, is beginning to form a coherent group that can work together in the pursuit of software quality. The natural conclusion of AQuIS '93 was thus a general commitment to meet again in the future to compare and discuss further progress in the field.

Please contact:
Antonia Bertolino - IEI-CNR
Tel: +39 50 593478
E-mail: bertolino@iei.pi.cnr.it

EPLP – Workshop on Portuguese Language Processing

by Isabel Trancoso

The first Workshop on Portuguese Language Processing (EPLP - Encontro de Processamento da Língua Portuguesa Escrita e Falada) took place in Lisbon on 25-26 February 1993. This workshop was intended as a forum for exchanging experiences between two communities: the natural language processing community and the speech processing community.

The integration between the two scientific areas aims at the construction of a comprehensive model of the "cognitive chain" which links speech to reasoning via NL. Although the need for integrating these two areas has long before been recognized, more visible efforts towards this goal have only occurred quite recently, one of them being the creation of the ESPRIT Network in Language and Speech (ELSNET).

The workshop included 4 tutorial sessions, 7 technical sessions and a final panel session. The Scientific Committee, formed by 15 experts from several Universities and research centers in Portugal selected around 30 out of 40 submitted papers for presentation on the basis of an extended abstract.

The workshop was very successful, having fulfilled its objectives both from the scientific point of view and from the point of view of establishing fruitful contacts between the research communities engaged in the automatic processing of the Portuguese language. ■

Please contact: Isabel Trancoso - INESC
Tel: +351 1 3100 268
Fax: +351 1 525843
E-mail: imt@inesc.pt

CALL FOR PAPERS

ADB '94 – International Conference on Applications of Databases

Vadstena, Sweden,
21-23 June 1994

Successful database research and development has made many general database systems available. The industrial offerings range from traditional large mainframe systems, through workstation & PC servers, to mass-produced \$100 personal systems. Users, application designers, and the software industry strive to get the most from this technology. Database researchers and developers try hard to further satisfy application needs. The conference aims at developing synergy between these communities. It is intended as the forum to explore innovative applications of databases and innovative database services for specific applications.

Topics – Applications (not limited to):

- General Enterprise Management
- Financial and Business Decision Support and Reporting
- Spreadsheets
- Interpersonal Communication
- Application Development
- Telecommunications
- Manufacturing
- Consumer Services
- Mathematics
- Applied Sciences & Engineering
- Healthcare
- Law & Humanities
- Training and Education
- Leisure & Entertainment.

Topics – Services (not limited to):

- Interoperability of databases, legacy applications
- Visual Database Languages
- Databases on multicomputers, parallel computers, and data highways
- Database Mobility Agents and Mediators
- New storage media and technologies
- Text, Image, Sound, and Video in databases
- Object Management and Libraries
- Data migration to databases
- Database performance, benchmarking, and tuning
- Large Scale Information Retrieval and Mining
- Time, and Event Management and Monitoring
- Customization of database manipulations
- Use of standards
- DBMSs on new OSs, FMSs, and Transaction Managers
- Security.

We solicit research, industrial and user's papers, panel proposals, and tutorial proposals.

Deadlines:

14 February 1994: Submission of papers (address below)
1 April 1994: Notifications of acceptance ■

**Please contact: ADB-94 Secretariat,
Anne Eskilsson, Dept. of Computer
Science, Linköping University, 581 83
Linköping, Sweden**
E-mail: adb94@ida.liu.se

EUREKA Conferences

Lillehammer, Norway,
14-15 June 1994

During the arrangement "Vision Eureka" at Lillehammer in June 1994, one of the conferences will focus on Information Technology for Environmental Management. This particular conference is scheduled for 14 and 15 June. The programme includes presentations of four EUREKA projects:

- ENSIS – An Environmental Surveillance System for the Olympic Winter Games in 1994
- EUROMAR – An umbrella project for development, application and exploitation of marine environmental technology
- EUROENVIRON – A project for innovative R&D to tackle Europe's major environmental problems
- MEMbrain – A decision support system for emergency management.

In addition there will be a number of parallel workshops where the participants will get a chance to present their work, and to discuss particular details of a complete environmental management system.

Another important aspect about the conference is the PARTNERING process. This is a formalized method to try to bring people with common interests together. Participants will be encouraged to complete a questionnaire describing their scientific profile, business plans, liaisons they are looking for, etc. The information will be distributed to all the participants ahead of the conference. This will give each one a chance to locate potential partners, and the conference secretariate will make the practical arrangements necessary for these people to meet. An invitation leaflet with practical information and a registration form has been distributed to all the ERCIM member institutes.

Please contact:
E-mail: eureka@delab.sintef.no
or: Truls Gjestland - SINTEF DELAB
Tel: +47 73 59 26 45
E-mail: truls.gjestland@delab.sintef.no

WTC'94 – World Transputer Congress 1994

Como, Italy, 5 -7 September 1994

The Transputer Consortium (TTC) is pleased to announce that the World Transputer Congress 1994 (WTC '94) will be held on 5 -7 September 1994 at the Villa Erba, Como, Italy, in conjunction with the Italian Transputer User Group (ItTUG), which is due to be formed in early 1994. WTC '94 will act as the Inaugural Meeting for ItTUG. WTC '94 is the leading international Transputer Conference and Exhibition.

WTC '94 will be the first major conference where significant applications of the new T9000 transputer and its associated technologies (e.g. packet routers) will be extensively reported. An associated exhibition attracting the world's leading suppliers of transputer-based and other relevant hardware, software and application products will be held at the Villa Erba Exhibition Centre.

Tutorials and Workshops:

Before the World Transputer Congress 1994, we shall be holding tutorials on the fundamental principles underlying transputer technologies, the design paradigms for exploiting them, and workshops that will focus directly on a range of specialist themes (e.g. real-time issues, formal methods, AI, image processing ..). We welcome suggestions from the community of particular themes that should be chosen for these tutorials and workshops. In particular, we welcome proposals from any group that wishes to run such a tutorial or workshop.

Deadlines:

1 March 1994: Submission of papers, tutorial proposals and Workshop themes.
14 May 1994: Notification of acceptance

Please contact: Susan C Hilton - RAL
Address: Building R1 Rutherford
Appleton Laboratory
CHILTON, DIDCOT, OXON. OX11 0QX
Tel: +44 235 44 6154
E-mail: sch@inf.rl.ac.uk

CALL FOR PARTICIPATION

EUROGRAPHICS 1994

Oslo, Norway,
12-16 September 1994

Eurographics '94 will be held on September 12-16, 1994 in the Conference Centre associated with the SAS Hotel in the centre of Oslo. The conference is concentrating on the three major R&D areas:


- Visualisation
- Graphical Interaction
- Algorithms,
- the focus on applications in:
- Computer Aided Engineering
- Geographical Information Systems
- Multimedia.

As always, although this will be the main thrust of the conference, participation will be accepted in any area of computer graphics and related application areas.

Tutorials on Computer Animation, Broadcast Commercials, PEX, Open GL, Fractals, Object-Oriented Computer Graphics, Graphics Hardware, Tools for the Artist, Photorealistic Image Generation, User Interface Design, Virtual Reality, Visual Communication and Multimedia, Visualisation in Scientific Computing are planned.

As always, a Slide and Video Competition is taking place. Up to 5 slides and 5 videos can be entered to the two separate competitions. In this case, the deadline for submission is 31 May, 1994. PAL is the preferred encoding on either VHS or U-matic.

Please contact: Conference Secretary,
EG'94 Box 124 Blindern N-0314 Oslo,
Norway
Tel: +47 22 06 73 00
E-mail: eg94@si.sintef.no



Steven Pemberton
(Photo: CWI)

CWI – **Steven Pemberton**, researcher in CWI's Algorithmics & Architecture department, has been appointed as editor-in-chief of ACM-SIGCHI Bulletin, the organ of ACM's Special Interest Group on Computer-Human Interfaces. ACM was until recently strongly US-oriented. In trying to acquire a more international profile, a European office was recently opened in Brussels and for the first time the ACM-conference on Computer-Human Interfaces, INTERCHI'93, took place outside the USA (in Amsterdam). Pemberton's appointment is in line with this policy. He joined CWI in 1982 after having held a research position at the University of Brighton (UK). His main interest lies in software ergonomics (user interfaces).

INESC – **Professor José Tribolet** recently visited Macao. The aim of this visit was to initiate a series of contacts with the objective of stimulating co-operation between Portugal, Macao and China in the areas of science and technology, in addition to the business application of information technology and telecommunication. His visit included contacts with the Governor of Macao, the United Nations Software Institute and the "Oriente Foundation" as well as discussing the possibility of establishing joint ventures between China, Macao and Portugal. All signs show that the conditions to enable Portuguese co-operation with their counterparts in Macao and China in the field of information technology and telecommunications through staff exchange, establishing business ventures and technological co-operation, are beginning to make themselves felt.

FORTH – **Dr. Jan Tretmans** has recently started his second period as an ERCIM Research Fellow at FORTH-ICS in Heraklion, Crete, Greece, having spent the first period at SINTEF, Norway. Dr. Tretmans has graduated from the

University of Twente (The Netherlands) having worked in the application of formal techniques (process algebra, transition systems) to the testing of distributed systems. At FORTH he is currently working with Prof. Courcoubetis. His aim is to combine FORTH's knowledge about the verification of real-time systems with his experience in formal testing and SDL gained at SINTEF. The research plan is to apply the real-time verification techniques to the validation of SDL systems with timers. Currently, timers, although important, are usually not considered in formal SDL verification. A second item is the investigation of testing and testing equivalence of real-time systems. The current verification algorithms are mainly based on bisimulation equivalence, which is convenient in the area of verification, but for testing, testing equivalence is usually considered more appropriate.



Rainer Berling,
Managing
Director
of SICS
(Photo: SICS)


SICS – **Rainer Berling**, M. Sc. and M.B.A., is the Managing Director of SICS since the first of December 1993. His industrial experience ranges from systems engineer and project leader to department manager in fields such as airplane systems, computing systems and telecommunication systems. His ambition is to keep the position of SICS in combining high scientific competence with the ability of deploying this competence rapidly in cooperation with industry and other research centers. Another challenge for SICS at this moment is to contribute effectively to the increase of the competitive capability of Swedish SME's.

CNR – **An important meeting**, attended by many of the most prominent representatives of the Italian academic community working in applied mathematics, engineering and computing sciences, was held on 22 November

1993, in Rome, to discuss the future prospects for research at CNR in IT. The debate is part of a wider discussion that aims at the reorganization of CNR with the creation of new structures to guarantee a more efficient and effective coordination of research work throughout the country. Attention is focussed on the need to establish national institutes for specific fields of research. In particular, in the IT field, CNR researchers highlighted the problems that derive from the current division of the various Institutes and Groups working in this sector under four different scientific supervisory committees. They feel that the establishment of a National Institute for Information Technology with the mandate to coordinate all the CNR research activity in this area would greatly assist in overcoming the problems that derive from the current dispersion of efforts and initiatives.

European Research Consortium for Informatics and Mathematics

ERCIM



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

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

Central Office: ERCIM Domaine de Voluceau Rocquencourt B.P. 105 F-78153 Le Chesnay Cedex FRANCE	Central Editor: Henk Nieland	E-mail: henkn@cwi.nl	Telephone: +31 20 592 4092
Lay-out and Design: Helena du Toit	Local Editors: Wernher Behrendt (RAL) Lars Bergman (SISU) João Bilhim (INESC) Truls Gjestland (SINTEF) Päivi Johansson (VTT) Siegfried Münch (GMD) Pierre Népomiatshy (INRIA) Henk Nieland (CWI) Carol Peters (CNR) Constantine Stephanidis (FORTH)	wb@inf.rl.ac.uk lars@sisu.se jafb@eniac.inesc.pt truls.gjestland@delab.sintef.no Paivi.Johansson@vtt.fi siegfried.muench@gmd.de pierre.nepomiatshy@inria.fr henkn@cwi.nl carol@vm.iei.pi.cnr.it cs@csi.forth.gr	+44 235 44 5281 +46 8 752 1613 +351 1 572 895 +47 73 59 26 45 +358 0 456 5936 +49 2241 14 2303 +33 1 39 63 5646 +31 20 592 4092 +39 50 593 429 +30 81 22 64 87

AEDIMA				
Asociacion Española de Informatica y Matematica Aplicada	Centrum voor Wiskunde en Informatica	Consiglio Nazionale delle Ricerche	Foundation of Research and Technology – Hellas	Gesellschaft für Mathematik und Datenverarbeitung mbH
UPC Av. Gregorio Marañon, s/n E-08028 Barcelona	Kruislaan 413 NL-1098 SJ Amsterdam	IEI-CNR Via S. Maria, 46 I-56126 Pisa	Institute of Computer Science P.O. Box 1385 GR-71110 Heraklion, Crete	Schloß Birlinghoven P.O. Box 1316 D-53731 Sankt Augustin
Tel: +34 3 401 72 89 Fax: +34 3 401 62 10	Tel: +31 20 592 9333	Tel: +39 50 593 400 Fax: +39 50 554 342	Tel: +30 81 22 93 02 Fax: +30 81 22 93 42	Tel: +49 2241 14 -0 Fax: +49 2241 14 2889
		RAL	SICS	 
Institut National de Recherche en Informatique et en Automatique	Instituto de Engenharia de Sistemas e Computadores	Rutherford Appleton Laboratory	Swedish Institute of Computer Science	Stiftelsen for Industriell og Teknisk Forskning ved Norges Tekniske Høgskole
B.P. 105 F-78153 Le Chesnay	Rua Alves Redol 9 Apartado 10105 P-1017 Lisboa Codex	Chilton DIDCOT GB-Oxon OX11 0QX	Box 1263 S-164 28 Kista	SINTEF DELAB N-7034 Trondheim
Tel: +33 1 39 63 5511 Fax: +33 1 39 63 5330	Tel: +351 1 310 0217	Tel: +44 235 82 1900 Fax: +44 235 44 5385	Tel: +46 8 752 1500 Fax: +46 8 751 7230	Tel: +47 73 59 26 00 Fax: +47 73 59 43 02
				Technical Research Centre of Finland VTT Information Technology P.O. Box 1200 FIN-02044 VTT Tel: +4358 0 456 6041 Fax: +358 0 456 6027