


ERCIM NEWS

European Research Consortium for Informatics and Mathematics

Number 18

July 1994

EDITORIAL



David Barstow, Head of Computer Science Research at Schlumberger, Montrouge, considers effective representation of domain knowledge as a key enabling technology for Information Systems in the next century.

SPECIAL :

The Use and Capture of Domain Knowledge for Information Systems 8

C O N T E N T S

Editorial	1
Joint ERCIM Actions	2
The European Scene	5
Research Activities:	
Domain Knowledge for Information Systems	6
Other Activities	19
Technology Transfer	27
International Relations	30
Events	31
In Brief	35

Next Issue:

High Performance Communication

The importance of domain knowledge in information systems was first recognized in the 1970's, with the observation that there is a trade-off between power and generality: the power of an information system can be greatly increased by sacrificing its generality, that is, by restricting it to a specific domain. This observation led to a large body of research and experimentation with knowledge-based techniques, as well as the development of a number of commercially successful systems for domain-specific tasks. In all of these cases, the emphasis was on encoding domain knowledge explicitly in a form that a problem solver could exploit.

The importance of domain knowledge in the practice of software engineering was recognized in the 1980's. An important study by Curtis, Krasner, and Iscoe identified three primary causes of software engineering problems in industrial projects, and the first of these was the lack of knowledge about the application domain. At about the same time, the field of requirements engineering came into being, relying in part on the use of knowledge representation techniques to encode information about the application domain for a given software project. Domain knowledge has also begun to play an important role in software reuse (see ERCIM NEWS, July 1993). For example, recent work has focused on the development of domain-specific libraries and architectures. Several of the articles in this issue address such uses of domain knowledge in software engineering.

In the 1990's, we have begun to see what may be the ultimate role for domain knowledge in information systems, namely to enable the end user to exploit sophisticated systems without requiring him or her to be computationally sophisticated. One example is the use of domain knowledge for information retrieval, as discussed in several articles in this issue.

Thus, we now know that techniques for the effective representation and use of domain knowledge are among the key enabling technologies for the information world of the next century. The projects described in this issue are helping to lead the way to that world.

David Barstow

A Tribute to our Friend and Colleague Professor Cor Baayen

It is with much regret that I announce that, after 6 years of superb chairmanship, ERCIM's first President, Professor Cor Baayen, has decided not to remain in this position for a new term of office. On behalf of all the directors of the ERCIM partner institutes, I should thus like to express our deep and lasting gratitude to Cor for his many efforts over the last years. Not only was he one of the founders of ERCIM, but as President he has fully devoted his en-

ergies and his talents to promoting its cause and to inspiring that spirit of solidarity, friendship and efficiency which has made ERCIM a model for many similar endeavours in Europe.

Nobody could have been better qualified to exercise this chairmanship. Firstly, Professor Baayen has had intensive experience in the field, going back to 1980, both as Director of CWI and as Director of the Mathematical Foundation in The Netherlands. Secondly, his was the task of steering the evolution of a mostly mathematical center (Mathematisch Centrum) towards a pluridisciplinary institute, CWI, whose role as national "Centre d'Excellence" in Computer Science had been decided by the Dutch Government. The successful synergy that has been achieved between mathematics and computer science has become a corner stone of CWI activities and has meant that it has been

able to secure its central position in the Dutch research world. This synergy is also the spirit of ERCIM, which faces even broader challenges, since it brings together many institutes with different statutes and organization. Through his natural authority, his sense of consensus combined with his many talents, Cor Baayen has succeeded in shaping ERCIM into the best possible instrument to guide the creation of an integrated European Scientific Community in informatics and mathematics.

Needless to say, all the ERCIM directors will keep in mind his message and will continue his work.

Thanks Cor, and as "President d'Honneur" of ERCIM stay on our side.

On behalf of ERCIM Directors

Alain BENSOUSSAN

6th ERCIM Workshops at SICS

From 2-4 June 1994 SICS organised three ERCIM workshops at its site in Kista, Sweden. These were the last in the current series of ERCIM Workshops and differed from the others by a large participation of researchers coming from organisations outside ERCIM. Some participants even came from US and Japan. Brief descriptions of the objectives and results of these workshops are given here below. The Proceedings will be published by ERCIM. ERCIM has now changed its policy with respect to the organization of workshops. Information on the new programme, which aims at assisting and encouraging specialist groups in the organization of scientific events in strategically important areas, is given on page 4.

Mobile Computing and Communication

by Frank Reichert

The current state of the Walkstation project was presented by G. Maguire, F. Reichert, Ch. Roobol, and D. Kerek from KTH, Sweden. The design of flexible and efficient future mobile communication systems is a major challenge. The Walkstation Project involves researchers from different areas in order to find a solution via a global system approach.

Areas being studied include mobile adaptive applications, mobile systems man-

agement and security, integration of mobile stations into existing fixed networks (Internet), future communication architectures, use of multiple wireless interfaces, dynamic resource allocation and media access, direct sequence CDMA and low power VLSI integration. Mobile users expect global network connectivity, mobility transparent applications, a quality of service comparable to that of fixed networks, and of course low costs. Satisfying these requirements leads first to the need for integration of existing and future communication networks for indoor local area networks and outdoor wide-area networks to provide global connectivity.

In the Walkstation project mobile users may equip their mobile terminals with multiple wireless interfaces, which are controlled and selected by a system management function to meet criteria such as

achievable throughput and delay, real-time requirements, usage costs, and impact of selected communication link on battery lifetime and currently available power.

Because power consumption is of major concern for mobile terminals, several investigators have proposed the use of flexible transceivers for adjustable bit-rates on the same wireless link in order to reduce power. In the Walkstation project, the variety of networks (TDMA/CMDA-based, GSM, Mobitex, DECT, etc.) and the difference of environments is harmonized at the network layer. The goal is to provide mobility transparent communication services to traditional, non-adaptive applications. These are electronic mail, file-transfer, Information browsers and other applications common to users with PC or UNIX workstations on the Internet.

Not all differences of mobile networks can be hidden from user applications. Parameters such as available bandwidth, latency and setup time are crucial for the responsiveness and performance of multimedia applications. Therefore new adaptive applications will react to the available quality of service by changing the presentation of user objects from full graphical display to simple character based interaction. In addition, the investigation of new caching and pre-fetching methods will help to provide better performance.

Users will not accept expensive and heavy mobile computing and communication systems with short battery life times. On the computational side, new portable "personal data assistants" (PDAs) show the advances made in portable computing. These demonstrate impressively what is achievable with today's state-of-the-art computer technology. These devices have enough computational power to perform complex tasks such as stylus and voice recognition. On the communication side, most PDAs are designed to support new credit-card sized interfaces which are plugged into PCMCIA slots.

Therefore, an important research area in the Walkstation project is the investigation of new digital, highly integrated radio interfaces with low cost, small size and low power consumption. Direct Sequence (DS) CDMA is the basic media access scheme used in the our project.

Distributed Virtual Environments

by Kai-Mikael Jää-Aro

The keynote speaker Myron Krueger, a well-known personality within the virtual worlds community, started off by demolishing a number of myths about Virtual Reality. He asserted that even though the field in principle exists since the early 1970's, very little progress has in fact been made since then.

Nevertheless he was still actively pursuing his ideas about video-based artificial realities and was also planning excursions

into immersive virtual environments. The program then mainly presented results from the COMIC project, with participants from the universities of Nottingham, Lancaster and Manchester, as well as SICS and DRAL, exploring spatial models of interaction in three-dimensional data spaces, often employing versions of the DIVE system developed at SICS. Many papers concerned embodiments of users – different visual representations of users – to indicate their presence in a virtual environment, their attention, computational capabilities and other attributes, as outlined in a paper presented by Steve Benford. He and Johan Mariani had made video clips showing different kinds of Populated Information Terrains, which employ three-space for document layout and data retrieval and embodiments of the users to display which people were browsing what.

Chris Greenhalgh presented MASSIVE, a system which permits people with ordinary text-based terminals to interact with virtual worlds and other users in them, using different types of embodiments to indicate what kind of interface the users were accessing the world through and, consequently, what types of communication could be initiated with them. Martin Prime from DRAL discussed requirements for incorporating sign language in virtual environments and what that would require of the user embodiments.

The presentations that stirred the greatest surprise and interest and while not in any way connected with the COMIC work, still showed the most advanced embodiments, were made by Yasuichi Kitamura and Nobuyoshi Terashima from ATR Communication Systems Research Laboratories in Kyoto. They presented a teleconferencing system developed at ATR which features tracking of facial expressions and body movements and mapping them to geometric models of the participants' bodies, so that the two participants could share a table top as common space and manipulate common virtual objects.

High-Performance Communications

by Peter Sjödin

The objective of this workshop session was to provide a hearing on different switching techniques for gigabit networks. The participants were invited to discuss advantages and disadvantages with circuits, cells and packets as basis for switching in such networks.

Questions of interest:

- Is fast circuit switching an alternative to cell/packet switching?
- Can ATM be switched fast enough in photonic networks?
- How may variable-length packet switching be provided in high speed networks?
- Do multimedia traffic characteristics in general allow bandwidth gain by statistical multiplexing or will peak rate allocation be predominant?

The workshop was attended by researchers from Japan, Europe and USA. The presentations described experiences from and experiments with existing high-speed networks, as well as new ideas for network protocols and architectures. Three architectures for high-performance networks have been presented. L. Ramfelt (KTH, Sweden) presented an all-optical network architecture which is intended for multiprocessor interconnections. H. Brynhi (NTR, Norway) presented a desk and local area network built around the scalable coherent interface. The dynamic synchronous transfer mode was presented by Bohm, Hidell and Lindgren (KTH, Sweden) as a networking architecture based on fast circuit-switching.

Four empirical studies relating to the realization of ATM networks were presented. B. Davie (Bellcore, USA) described the complexity encountered for ATM-hardware in the Aurora testbed. He discounted the per-cell processing costs and bandwidth inefficiency, and showed that the biggest limitation to reach higher performance is movement of data. D. Morris (Univ. Leeds, UK) described performance problems his group

has encountered with running TCP/IP protocols over ATM. The throughput could be as low as 10 percent of the throughput of Ethernet. The issue lead to a fruitful discussion on the sensitivity of the parameter settings for such a combination of protocols.

G.P. Balboni (CSELT, Italy) reported on the effects of statistical multiplexing for today's applications. Y. Miyake (KDD, Japan) presented the implementation of a rate control mechanism for high-performance networks and its performance. A. Lazar (Columbia U., USA) presented a unified framework for packet and circuit-switched network. The high-performance communications session was closed by two presentations on networking: D. Farber (U. Penn., USA) covered suggested experiments on the global scale and B. Stockman (KTH, Sweden) gave an overview of the European scene with ongoing projects like Beteus, PET and SGN. ■

Please contact:
ERCIM Office
Tel: +33 1 39 63 53 03
E-Mail: ercim@inria.fr

CALL FOR PARTICIPATION
ERCIM Advanced Course

Partial Differential Equations and Group Theory

New Perspectives in Computer
 Algebra and Applications
 to Engineering Sciences

Bonn, Germany, 24-28 October 1994

In line with previous courses held at INRIA in November 1990 and GMD in 1992, the ERCIM Advanced Course "Partial Differential Equations and Group Theory", held at CWI in March 1993, was quite successful with 40 participants coming from 13 countries of western Europe. Accordingly, it will be repeated this year at GMD, Bonn, from 24-28 October 1994, with a slightly different programme, including more applications. The course is supported by INRIA and GMD.

Sponsored ERCIM Workshops

ERCIM, the European Research Consortium for Informatics and Mathematics has decided to sponsor up to six high quality internationally recognised Workshops per year.

ERCIM does not wish to set up a Workshop in an area where there already exist organisations willing and able to run such Workshops. An additional Workshop might be the last straw that makes the other workshops non-viable. The net effect would then be to reduce the number of Workshops in the area, and this could prove counterproductive.

The ERCIM approach is, therefore, to collaborate with existing groups to produce such Workshops jointly. ERCIM is willing to offer to supplement the costs of a self-financing workshop, organised in collaboration with an external group. The additional funding provided by ERCIM should be used to enhance the workshop by, for example, increasing the set of external speakers supported. The funding on offer is in the order of 2000 ECU for each of six workshops per year. This is not seen as a revenue gaining exercise for ERCIM so that, in case such workshops generated profits, the profit

As a basic motivation for organising such a course, one must notice that the widespread importance of computer algebra today hides the fact that most of the algorithms concerning PDE only use mathematics that were in fashion 50 years ago. Indeed, the formal theory of PDE and groups of transformations has been created and developed during the years 1960 to 1975 by Donald Spencer in the USA. However, though this work supersedes the classical approaches (Maurice Janet, Elie Cartan), it is still largely unknown by mathematicians and has never been applied. Also, it is a fact that, during the last 10 years, people have been trying to study nonlinear phenomena by means of formal algebraic or geometric techniques instead of numerical or functional ones.

would remain with the external organisation. Publication of the proceedings is left to the workshop organising committee but ERCIM would be prepared to help if appropriate.

In addition to the funds from ERCIM to enhance the event, the benefits that a Workshop might see in such an arrangement would be:

- Use of an ERCIM site for hosting the event. This would normally include meeting rooms free of charge, some administrative support and help with photocopying. Details may vary depending on the ERCIM site.
- Commitment from ERCIM to help organise and run the event.
- Publicity from ERCIM within the ERCIM members to raise the level of participation.
- Production of the interim Proceedings made available to attendees (it is assumed each such Workshop will have its own procedure for final publication).

The intention is to have a Call for Proposals at least once a year in the ERCIM Newsletter and elsewhere. Bids will have to indicate how the Workshop meets criteria set by ERCIM for awarding the sponsorship. For further information, ■

Please contact:
ERCIM Office
Tel: +33 1 39 63 53 03
E-mail: ercim@inria.fr

Therefore, the purpose of the course is to give a self-contained introduction, at a graduate level, to these new tools, while illustrating them with specific examples coming from various branches of engineering sciences and applied mathematics. In particular, the course presents for the first time the long-awaited group theoretical unification of the mathematical models of elasticity, heat and electromagnetism, allowing for a unique finite elements formulation. ■

Please contact:
Jean-François Pommaret – ENPC
Tel: +33 1 49 14 35 85
Fax: +33 1 49 14 35 86
or: Fritz Schwarz – GMD
Tel: +49 2241 14 27 82
E-mail: fritz.schwarz@gmd.de

GMD and Smith System Engineering Bring Parallel Computing to Industry

by Siegfried Münch

A vicious circle is the reason why industry hardly uses any parallel computers though only these computers are able to provide the computing capacity required for extensive realistic simulations in industrial development. Industry doesn't purchase parallel computers since important parallel production software is still missing and since it is therefore very difficult to assess whether parallel computers will in fact meet industrial requirements. Software houses are afraid of the cost involved in the parallelization of real production software as long as industry does not order such software. For this reason, the Commission of the European Union funds the project EUROPORT by granting 16 million ECU (total project volume: 28 million ECU) with the intention to further the confidence in the efficiency and profitability of parallel computers and to introduce large parallel systems into industrial use. For this purpose, a rep-

resentative choice of industrial programs will be parallelized. About 90 organizations from Western Europe are involved in this project.

On 12 and 13 January, 1994, the management teams of GMD and Smith System Engineering organized a first meeting of the EUROPORT project partners at GMD-Birlinghoven. Dr. Anne de Baas, representative of the European Union and official project officer, opened the meeting by presenting structure and objectives of EUROPORT. About 120 people participated in the event, among them the coordinators of individual activities and representatives of all project partners. More than 50 % of the participants came from industry, primarily automotive industry (Mercedes, BMW, Fiat, Ford), aircraft industry (British Aerospace, Rolls Royce, Aerospatiale, Saab) and chemical and pharmaceutical industry (Bayer, Merck, ENEA, Unilever) together with software houses (McNeal Schwendler, Intes, Genias, Pallas, POLYFLOW, NA Software, SIMULOG).

The two years' project consists of two parts:

EUROPORT-1 focuses on numerical fluid and structure dynamics, for example, simulations of aircraft flow conditions and crash simulations. Seven widespread commercial programs and three industrial in-house developments will be parallelized.

EUROPORT-2 covers a great variety of fields by 15 large application programs,

in particular numerical chemistry (e.g. agent development, simulation of oil deposits and oil production, numerical electrodynamics, radiotherapy, very large databases, animation for movies and television).

GMD has been entrusted with the management of EUROPORT-1. Smith System Engineering in Great Britain manages EUROPORT-2. GMD is responsible for nine project subgroups, Smith for ten. Each group parallelizes at least one large program. The group consists of the owners of the programs, parallelization experts and industrial program users. EUROPORT combines experience and expert knowledge from industry, software houses, research institutes and universities. Operators of parallel computers complete this spectrum.

It is the aim of the project partners to stimulate the interest in the production and portation of software on parallel computers by demonstrating the profitability and the competitiveness of parallel computing. Circulating the project results beyond EUROPORT is to increase the confidence of industry in parallel supercomputing.

Considerable improvements of the cost-benefit ratio in the industrial use of parallel programs are to convince software and hardware providers and industrial users of the efficiency of this technology.

The division of EUROPORT in only two subprojects is to keep the overall cost as low as possible, to avoid duplicate work and to secure the compatibility of results. The main tasks to be accomplished by GMD and Smith are subject-specific monitoring of the project, review of results, performance measuring and evaluation of final products and circulation of results in science and industry.

GMD's Institute for Algorithms and Scientific Computing is responsible for the management of EUROPORT-1. Smith System Engineering provides unbiased consulting for effective use of scientific and technological results.

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



"EUROPORT kick-off meeting" – Dr. Anne F. Baas from the Commission of the European Communities presents the subject of the meeting (Photo: Münch)

The Role of Domain Knowledge

by Gordon Ringland

From the solutions to the coding disasters of the 60's and 70's to the solutions to the requirements calamities of the 80's and 90's.

By the early 80's it was widely recognised that the major problems in improving software quality and productivity were concentrated in the early part of the software development lifecycle. This, in large measure, was the result of work on structured programming languages and structured development methods which ameliorated many of the coding problems analysed by Brookes while still leaving the problems of the early life cycle.

A small but significant number of studies were conducted during the 80's [Curtis, Krasner and Iscoe (1988), Communications of the ACM 31:1268-128; Adelson and Solway, IEEE Trans. Softw. Eng. 11:1351-1360]. These studies were distinguished by being problem-driven rather than technology-driven and so empirical studies were undertaken to assess the factors reducing productivity and quality in significant samples of a wide variety of large industrial projects.

The most significant problems were:

- (1) the "thin spread" of domain application knowledge
- (2) changing and conflicting requirements
- (3) communication and coordination breakdowns

In the ESPRIT project F3, RAL has also conducted empirical studies on the analysis, refinement and validation of input requirements, with a very capable and experienced Requirements Engineer at the Spanish company Ceselsa. These studies showed how domain knowledge is used, and confirmed how valuable it is to the Requirements Engineer. For further details see the accompanying article by Simon Lambert.

The importance of domain knowledge

Curiously, in Curtis et al. problem (1) is attributed to cognitive issues while (2) is attributed solely to the business milieu, though the main body of the text convincingly argues that much of (2) "changing and inconsistent requirements" is caused by (1) "inadequate application domain knowledge". With some advantages in hindsight in F3 we have been able to analyse the issue more clearly. Major problems with requirements arise from a dearth of domain application knowledge and a lack of knowledge of the business and its environment. Both are domain knowledge - in F3 the know-

ledge of a business and its environment is mainly held in the Enterprise Model (EM) while the application domain knowledge together with the strategies and tactics of an idealised Requirement Engineer is held in the tool and model for Early Requirement Capture Analysis and Validation (ERCAV). Note that the strategies and tactics of the able and experienced Requirements Engineer are themselves another domain knowledge base, the domain being Requirements Engineering itself.

Knowledge acquisition from requirements

It is a fact that a requirements document itself contains domain knowledge. It may be inconsistent, incorrect and incomplete but it is better than nothing. The requirements can provide clues for extension of the knowledge base, and suggest

The Use and Capture of Domain Knowledge for Information Systems

Contents :

- *The Role Of Domain Knowledge*, page 6
- *"From Fuzzy to Formal" - A New Approach to Requirements Engineering*, page 7
- *Formal Methods for Financial Products*, page 8
- *MOBAL: A System for Building Domain Models*, page 9
- *A Domain Model for Assisting Information Retrieval*, page 10
- *Electronic Dictionaries and Text Corpora*, page 10
- *Discovering Knowledge for Integrating Database Schemas in Interoperable Databases*, page 11
- *Acquisition of Domain Knowledge for Building an Explanatory System*, page 12
- *Interactive Self-Explanation - Making Domain Knowledge Explicit*, page 12
- *Reusability Aspects in Transactional Workflow Management*, page 13
- *Automatic Page Layout Applications in Catalogues, Newspapers, and Electronic Publications*, page 13
- *Abductive Information Retrieval of Multimedia Data*, page 15
- *Knowledge-based Access to Multimedia Document Bases*, page 16
- *Using Domain Knowledge in a Multimedia Toolkit*, page 17
- *A Software Test Pilot for Performance Assessment Support*, page 18

that some sort of “semi-automated” knowledge acquisition may be possible. It is, for example, easy to detect a term unknown to the knowledge base. It is then possible to prompt the Requirements Engineer (or domain expert) for its position in the ontology which is part of the knowledge base, and to prompt for its attributes (which distinguish it from other entities). The situation is more complicated for deeper knowledge such as the purpose of an operation. One approach is to relate this to knowledge about the use of the target system. The question becomes, what function can the operation support?

There is an excellent example in the ATC case study of a requirement prompting further knowledge acquisition. One of the original requirements states that “The system shall enable the operator to draw lines on the screen using the cursor.” This requirement is clearly ambiguous as it stands. Following the Task Model of the Requirements Engineer, the first step is to give a definite reading so that its consequences can be assessed. One reading is that the Requirements Holder is asking for a general-purpose drawing tool. The requirements engineer uses knowledge of the domain of software applications development to realise that, with such a reading, to meet the requirement would be very expensive in terms of time and money. Therefore the requirements engineer is stimulated to go back to the requirements holder and investigate why this requirement is needed. It turns out that it is to enable the controller to mark military exclusion zones on the screen, and all that is needed the facility to draw polygons - a much simpler requirement.

The economics of domain knowledge acquisition and archiving

There is some empirical evidence [Walz, Elam and Curtis, Communications of the ACM 36:63-77] that for teams working on projects which are fairly new to the team members, about 50% of effort and elapsed time goes into acquiring domain knowledge. However, this effort is not so attributed, is not planned for, and no training is given. Very little of the know-

ledge is recorded, and even more effort is used to partly reconstruct it. Clearly a well structured approach to acquiring and archiving the knowledge would produce far more valuable results, and at a lower cost than current practice.

There is a very telling quote in Curtis et al.: “Someone had to spend a hundred million dollars to put that knowledge into my head. It didn't come free.” This hints at the benefits to be had from REUSABILITY. Clearly the domain knowledge bases developed for the Ceselsa ATC case can be almost completely reused in the requirements process for any ATC system, be it in Spain or the USA. Further some of the domain knowledge bases since they are appropriately structured can be used in other applications e.g. Requirements for radar systems.

Other important applications for domain knowledge

Another important future reuse for domain knowledge bases is in the emerging field of Knowledge Based Software Engineering (KBSE). One might crudely characterise KBSE as attempting a huge enrichment of the CASE approach in which far more expressive knowledge representation and much deeper reasoning give the framework for a deep level of semantic processing.

Although we may not have ‘industrial strength’ examples of KBSE for another five years or so, the background and goals of this approach are very promising. In addition to procedures and support for the early life cycle, KBSE seeks to perform domain-specific program synthesis which should result in considerable improvements in both quality and productivity. It is possible that software engineering will largely evolve to become the discipline of capturing and automating currently undocumented domain knowledge.

Please contact:
Gordon Ringland – DRAL
Tel: +44 0235 44 5557
E-mail: gar@inf.rl.ac.uk

‘From Fuzzy to Formal’ – A New Approach to Requirements Engineering

by Simon Lambert

The Rutherford Appleton Laboratory is working on an ESPRIT project F3 (‘From Fuzzy to Formal’) which is concerned with Requirements Engineering. The Knowledge Engineering Group at RAL has adopted a knowledge-based approach to the problem, and is investigating the application of domain knowledge in the process of early requirements capture, analysis and validation.

In many commercial companies producing systems for clients, the starting point is a statement of requirements received from the Requirements Holder which may be ambiguous, incomplete, inconsistent, or flawed in many other ways. One of the jobs of the Requirements Engineer is to detect these problems and negotiate their solution with the Requirements Holder. This task needs knowledge of the domain as well as experience of requirements engineering itself. In the F3 project, RAL has worked closely with one of the user partners, the Spanish company Ceselsa, and their Air Traffic Control group. The terminology they use is to call the initial set of requirements “Blue” and the final, validated set “Green”. Knowledge acquisition by RAL with one of their top Requirements Engineers has produced a task model of the Blue-to-Green process and highlighted the role of domain knowledge within it. The task model includes steps such as “Check benefit to user and project”, “Distinguish and limit alternatives”, “Establish relations with other requirements.” Domain knowledge is ubiquitous in the task: for example, the requirement that a piece of apparatus must

conform to a stringent military standard may be suppressed if it is known that the apparatus is to be used in an aircraft with human operators, who could not themselves survive the conditions that the military standard specifies. This justification can then be presented to the requirements holder for negotiation purposes.

The Knowledge Engineering Group at RAL is working on the development of a knowledge-based assistant to help the requirements engineer in the process of early requirements capture, analysis and validation by applying domain knowledge to a set of requirements, identifying problems, and where possible suggesting amendments to the requirements and providing justifications for these amendments. This assistant uses generic knowledge of the domain of requirements engineering itself to detect problems with requirements and to suggest

solutions. These problems can be of a wide range of types, including missing or underspecified timings or quantifications, poor structuring of the requirements, inappropriate output directed to users, etc. The system has been developed for the domain of air traffic control, and is being applied to other domains within the F3 project to validate the generality of the approach.

Please contact:
Simon Lambert – DRAL
Tel: +44 235 44 5716
E-mail: scl@inf.rl.ac.uk

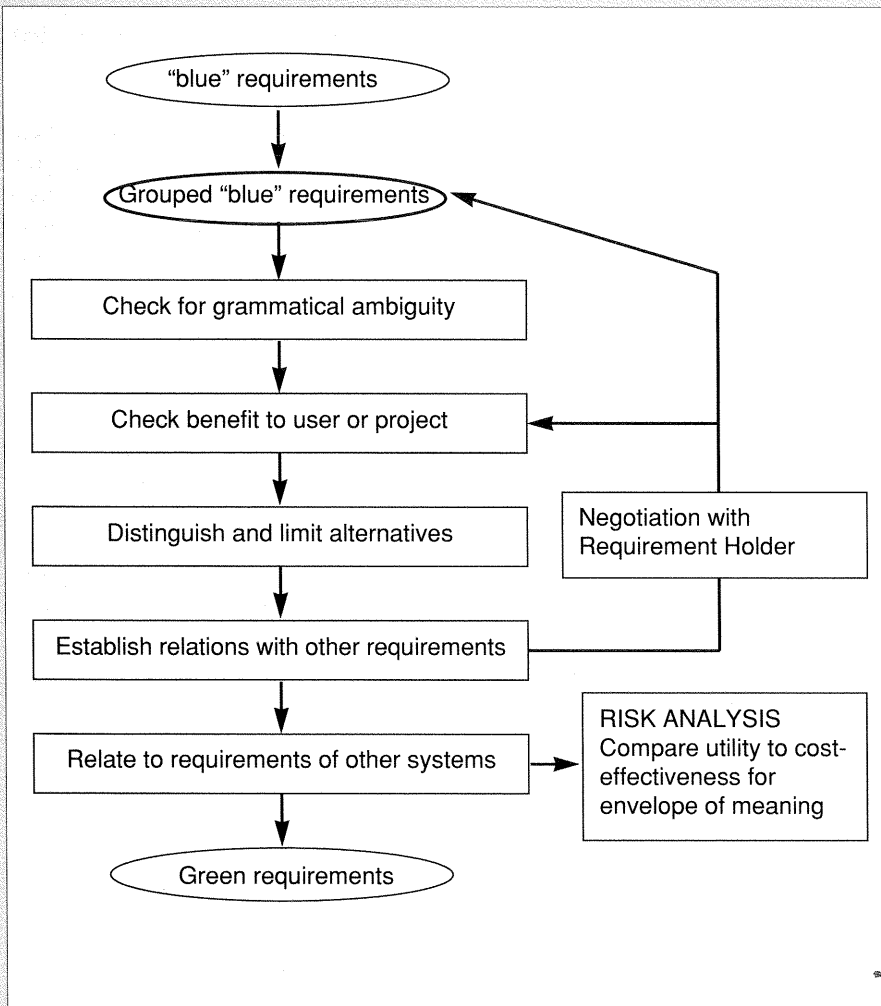
Formal Methods for Financial Products

by Arie van Deursen

Financial products of banks are subject to constant change. In order to deal with this, the Dutch Bank MeesPierson generates its software automatically from high-level product descriptions. CWI has applied its expertise in algebraic specification techniques to formally specify the Bank's product description language.

Bank MeesPierson, Rotterdam, is a merchant bank offering a wide variety of financial products. The complicating factor for the bank's automated systems -- which perform contract administration and provide management information concerning the bank's on- and off-balance position, interest and exchange rate risks, etc. -- is that the group of products is constantly changing: new types of products are introduced, and products that went out of fashion have to be withdrawn. How should the software construction process be organized such that the systems can deal immediately with such changes? The answer given by Bank MeesPierson is to describe the essence of their products in a high-level language -- called Risla --, and to generate the software automatically from these product descriptions. To that end, a small domain-specific language was designed, especially suited to describe so-called cash-flows following from products.

As the language was going to play a crucial role in the bank's automated systems, it was decided to define Risla formally. A joint project with CWI was started. In particular, CWI's experience with algebraic specifications of languages was exploited. Afterwards, the bank considered this a very fortunate decision. Specifying Risla formally re-



Subtasks of requirements engineering task

vealed several potential sources of vagueness (for instance concerning the fundamental data types of the language!) and had a major impact on the final Risla design.

The so-called ASF+SDF formalism was used to specify the syntax and semantic entities of the application language. This language definition formalism was developed at CWI in the framework of the ESPRIT project GIPE (Generation of Interactive Programming Environments, 1984-1993). The basic idea is to construct a development environment for formal language definitions.

The development of this specification was significantly helped by the ASF+SDF Meta-Environment, an interactive development environment for ASF+SDF itself. Its main merit is that the system's interactive qualities makes users unaware of the fact that each modification they make to their definition has major impacts on the generated environment. The system gives immediate feed-back while editing specifications, warning the user whenever an error is detected. Moreover, it can be used to execute and test specifications (using conditional term rewriting).

At the moment, the bank is implementing the software generators translating Risla to COBOL. By October 1994, all interest rate products must be dealt with exclusively by software generated from Risla descriptions.

Please contact:
Arie van Deursen – CWI
Tel: +31 20 592 4007
E-mail: arie@cw.nl

MOBAL – A System for Building Domain Models

by Werner Emde

MOBAL is a sophisticated system for developing, validating, and maintaining operational models of application domains. It integrates a manual knowledge acquisition and inspection environment, a powerful inference engine, machine learning methods for automated knowledge acquisition, and a knowledge revision tool.

By using MOBAL's knowledge acquisition environment, a user can incrementally develop a model of an application domain in terms of logical facts and rules, i.e., in a representation that is much more powerful than attribute-based formalisms. The entered knowledge can be viewed and inspected in text or graphics windows, augmented, or changed at any time. The built-in inference engine can immediately execute the rules that have been entered to show the consequences of existing inputs, or answer queries about the current knowledge. MOBAL also builds a dynamic sort taxonomy that provides information about the objects that have been used in a domain. The system can automatically construct a predicate topology reflecting the inference structure of a knowledge base. Machine learning methods can be used to automatically discover rules based on the facts that have been entered, or to form new concepts. If there are contradictions in the knowledge base due to incorrect rules or facts, there is a knowledge revision tool to help locate the problem and fix it by recording exceptions or modifying rules.

MOBAL is a system constructed according to the balanced cooperative modeling paradigm that regards knowledge acquisition as a continuing interactive

process of model construction. Consequently, in MOBAL, the distribution of work between system and user is flexible: the user can input all the knowledge needed for problem solving in a particular domain by hand, or can start by simply entering facts and using the various structuring and learning tools of MOBAL to extend this initial knowledge. Learning results are immediately incorporated into the knowledge base where they can be modified further. This contrasts with one-shot learning systems that run like a compiler and do not support the cyclic acquisition process. Consequently, use of MOBAL is highly interactive, and does not need to follow a fixed script.

MOBAL is applicable to all types of classification or advice-given problems, such as medical or technical diagnosis, assignment of objects or cases to a certain class, type or group, advice based on constraint violations, ... It is especially suitable if the domain requires the representation of time-dependent or other relational data, and in domains where frequent interactive changes are necessary. In the ESPRIT project MLT, MOBAL has been applied for example to the problem of diagnosing and treating Maldecensus testis diseases (with ICS Forth), and to the problem of deciding whether a given user has certain access rights in a telecommunication network (with Alcatel Alsthom Recherche).

Resulting from work in the ESPRIT basic research project ILP, MOBAL 3.0 now offers an "external tool" facility that allows other (ILP) learning algorithms to be interfaced to the system and used from within the same knowledge acquisition environment. The current release of MOBAL includes interfaces e.g. to GOLEM by S. Muggleton and C. Feng (Oxford University), GRDT by V. Klingspor (Univ. Dortmund) and FOIL6.1 by R. Quinlan and M. Cameron-Jones (Sydney Univ.).

GMD grants a cost-free license to use MOBAL for academic purposes.

Please contact:
Stefan Wrobel – GMD
Tel: +49-2241-14 26 70
E-mail: stefan.wrobel@gmd.de

A Domain Model for Assisting Information Retrieval

by Inge Nordbø

The aim of an information retrieval system is to retrieve references to documents that are relevant for the user. Such systems have access to large bibliographic databases containing document descriptions. Their retrieval system offers a language for expressing search queries consisting of keywords that are matched against the document description. The main problem for the user of a traditional information retrieval system is to select the set of terms which describe the information need in a way that matches the description in the database. We have provided the user with an intelligent assistant KNOWIT which assists in selecting a set of search terms within the selected domain. Based on the user query, the system extends it to also include terms that are possibly of interest to the user.

The KNOWIT prototype demonstrates that it is feasible to develop a knowledge based system that improves the information retrieval system ESA-QUEST,

based on a knowledge base that contains knowledge of satellite communication. The evaluation shows that the number of relevant references with KNOWIT is significantly higher than the number of relevant references retrieved without KNOWIT.

KNOWIT - overview architecture

The KNOWIT front end (see figure) uses a domain model structured as a semantic network for expanding the search query. The domain is modelled according to four perspectives that seem to be important in the domain of satellite communication - Component (e.g. antenna), Process (e.g. multiplexing), Function (e.g. facsimile communication), Technology (multichannel communication). The terms in the domain are classified according to these perspectives and related to each other with semantic relations (e.g. has-subclass, has-subprocess, influences, involved-in). In addition, these perspectives are described with properties. We have identified 17 different semantic relations which was useful in the modelling of the domain. In this way the application knowledge base can be viewed as an extended thesaurus, where there are more relations between terms, and where the relations are more specific than is usual in a thesaurus.

The capturing of the knowledge base was done in cooperation between knowledge engineers and a domain expert. The Knowledge Engineering Workbench - KEW, was used to support the acquisition process by providing facilities for

representation and graphical editing of the network.

Please contact:

Inge Nordbø - SINTEF-DELAB

Tel: + 47 73 59 70 79

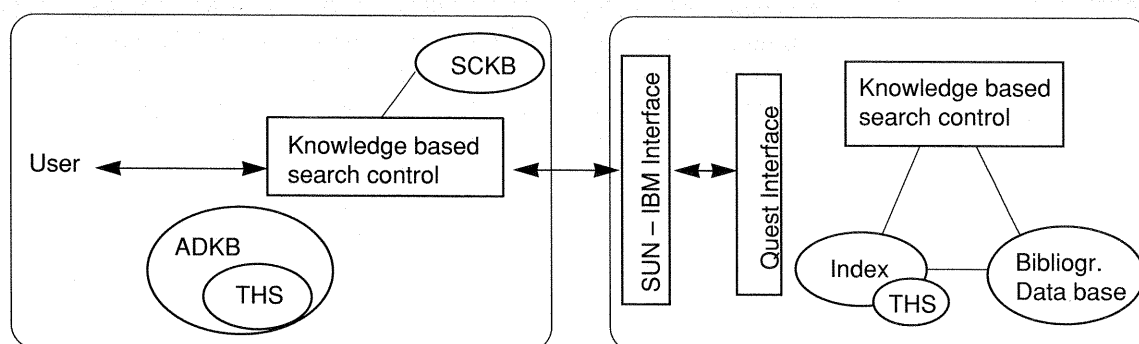
E-mail: inge.nordbo@delab.sintef.no

Lexical Knowledge Bases

Acquiring Information from Different Sources: Electronic Dictionaries and Text Corpora

by Nicoletta Calzolari and Carol Peters

One of the major current requirements in real world natural language processing applications is the construction of suitably structured and sufficiently exhaustive computational lexicons or lexical knowledge bases (LKBs). However, this is an extremely difficult, time consuming and expensive task, especially if it is necessary to start from scratch. For this reason, in recent years, researchers have begun to examine the potential of already existing resources, such as machine readable dictionaries and text archives, to see whether it is possible



KNOWIT - overview architecture

ADKB - Application Domain Knowledge Base
SCKB - Search Control Knowledge Base
THS - Thesaurus

to develop procedures that can capture different kinds of lexical information from them, representing it in a computationally tractable and reusable formal framework.

The Istituto di Linguistica Computazionale (ILC-CNR), Pisa, has played a major role in two CEU projects: Esprit BRA ACQUILEX I and II and ET-10/51, which have the extraction of morphological, syntactic, semantic and pragmatic information from machine readable dictionaries and text corpora and its representation in an LKB, using a typed feature structure formalism (TFS), as their main objective. The approaches adopted by these two projects are, to a certain extent, different but complementary. The first has concentrated on deriving a formal representation of meaning, while the second has had as its prime concern a representation of the combinatorial properties of lexical items. ACQUILEX-I, a large project involving a number of European languages and more than ten traditional dictionaries (for English, Dutch, Spanish and Italian), has focussed on classifying the headword in terms of the particular semantic features that can be derived from the genus items and differentiae in dictionary definitions, and on then establishing taxonomic chains throughout the dictionary, representing the lexicon as an inheritance network. A common type system has been developed at the top level of the hierarchy for the four languages.

ET10/51 on the other hand has concentrated on the analysis of one very particular dictionary taken from the Cobuild range. Cobuild dictionaries have two special features: (I) they are compiled on the basis of the evidence provided by a very large corpus of contemporary English; (II) the definitions appear as natural language sentences, i.e. with the definiendum inserted in its typical sentential context. The underlying assumption is that words only have sense in context and thus the headword is presented systematically in the definition in its typical syntactic and lexical/semantic environment. We have thus developed procedures which aim at classifying the lexical item mainly in terms of the syntactic, lexical and semantic characteristics of its

context, e.g. the selectional restrictions or preferences encoded on its arguments.

Clearly, both kinds of information, the purely semantic and the syntactic-semantic information, are needed by any NLP system. The contextual information which, for English, is partially derivable from the Cobuild dictionary, for other languages could be derived through an analysis of text corpora. This is the new approach taken in ACQUILEX-II, where we are now experimenting robust methodologies for corpus tagging and parsing as preconditions to the development of successful procedures for sense tagging, sense disambiguation, and the extraction of useful syntactic, lexical and semantic data from the corpus for the LKB.

DELIS, a third European project in which the ILC collaborates, is also working in the direction of deriving lexical information from large text corpora. The aim is to build tools geared towards automating the corpus exploration and data acquisition stages, and to design a model of the lexical entry for the LKB, again using the TFS formalism, which can represent adequately the linguistic generalisations that emerge from the corpus exploration.

On the basis of the work done so far, in particular in ACQUILEX, we are now experimenting the first implementation of what should become a complete semantic network for Italian (the noun taxonomy is almost complete). It is our aim to organize this network so that it is compatible with and mappable to Wordnet (developed in Princeton by G. Miller and publicly available), so that corpus analyses performed using such tools can be really comparable. ■

Please contact:
Nicoletta Calzolari – ILC-CNR
Tel:+39 50 560481
E-mail: glottolo@icnucevm.cnuce.cnr.it
or Carol Peters – IEI-CNR
Tel:+39 50 593429
E-mail: carol@vm.iei.pi.cnr.it

Discovering Knowledge for Integrating Database Schemas in Interoperable Databases

by M. Castellanos, F. Saltor and M. Garcia

The widespread use that databases have experienced during the last decade, has motivated the need to make them interoperate. This means, to make them cooperate sharing their data, while preserving their autonomy (as much as possible).

A main problem with interoperability is the schema integration process that has to be performed for building adequate interfaces to the user. Such interfaces make possible for the user to access autonomous and heterogeneous databases as if they were just a single one. Particularly, semantic heterogeneity, where equivalent or related concepts are represented in different ways in the various databases, add a high level of complexity to the integration task. If the schemas of the databases have to be integrated, semantic relationships between their objects have first to be identified, and this in turn requires to detect the resemblances that exist between them. All this, requires a deep knowledge of the meaning of the databases, which often does not exist. Moreover, the schemas being expressed in traditional data models do not help because of their semantic poorness.

One solution could be to upgrade the semantic level of the schemas by means of a semantic enrichment process. Such a process consists of two phases. In the first one, implicit or hidden semantics is discovered from different sources of knowledge: the human being (DB administrator and users), application programs, database extension, and so on.

Analyzing database extensions is a particularly attractive case of the machine

learning problem (general rules inferred from instances) where the simple and regular form of data leads to exact solutions (in contrast to inductive learning). From their huge amount of implicit semantics in the form of dependencies, the functional, inclusion, exclusion and complementariness kinds convey the semantics necessary to convert the schema to a richer one. Our algorithms to discover dependencies put emphasis on minimizing the number of accesses to disk required to analyze the extension of the database. Basically, some intelligent choices are made to minimize the number of candidate dependencies and to determine the order in which they are analyzed.

The idea that is behind is that data dependencies can be interpreted in terms of the rich abstractions of our O-O model called BLOOM. And this is precisely what is done in the second phase of the semantic enrichment, that is, to convert the original schemas together with the discovered knowledge to rich BLOOM schemas.

This enrichment is helpful for the integration, because the structures of the generalization/specialization (where several kinds of specialization are distinguished) and aggregation (where several kinds of aggregation are distinguished) semilatitudes of the BLOOM schemas make possible to set up an strategy which guides the comparison process, indicating at every moment which classes have to be compared and a criteria to determine how similar the classes are.

However, the applications of dependency discovery is not limited to integration. It applies also to other domains like database design (design by example), natural language interfaces, database migration, query optimization, reverse engineering and some more.

Further work on knowledge discovery is planned to analyze application programs with embedded SQL to identify operations and encapsulate them as methods of the classes identified in the first part of the semantic enrichment.

Please contact:
M. Castellanos – AEDIMA
Tel: +34-3 401-7328
Email: {castellanos, saltor, mgarcia}@lsi.upc.es

Acquisition of Domain Knowledge for Building an Explanatory System

by Rose Dieng

The Acacia Project is a multidisciplinary project that aims at helping the knowledge engineer and the expert during the knowledge acquisition phase. We work on a methodology and on a tool called Katemes (Knowledge Acquisition Tool for Explainable, Multi-Expert Systems), with the following objectives: (a) allow knowledge acquisition from multiple experts, (b) exploit the specific aspects of some classes of problems such as design problems, (c) help to improve the quality of the explanations that the future expert system will provide the end-user with, and (d) allow the validation of the acquired knowledge before implementation of the final expert system.

This article focuses on acquisition of knowledge useful for building an explanatory system and we will describe an application in telecommunication network design.

Explanatory knowledge acquisition

Cooperative explanations of a knowledge-based system (KBS) must be prepared as early as the knowledge acquisition phase: explanatory knowledge, not necessarily

useful for the problem solving but important for the future explanations, must be elicited from experts and end-users. Our approach was to determine what explanatory knowledge consists of, how to acquire it according to the used knowledge acquisition methodology, and how it influences the explanation module specification and implementation.

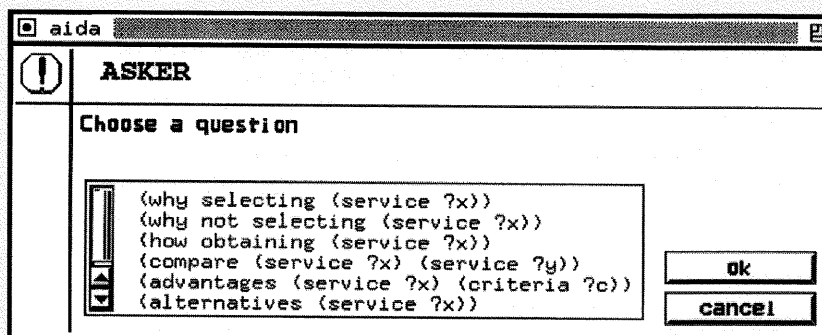
We first proposed a model of explanatory interaction: the user asks a request to the KBS that analyses it, selects an explanatory strategy, builds the explanation, communicates it to the user that evaluates the explanation and exploits it if accepted; in case of failure, the system can return to a previous step.

After determining the explanatory knowledge needed in each phase, we proposed several methods for acquiring it:

- extension of KOD for explanations
- extension of KADS-I through a library of "generic explanations"
- extension of KADS cooperation model.

We also proposed a method, independent of KADS or KOD, for determining:

- the characteristics of the different types of potential users
- the KBS problem-solving method influence on users' explanatory needs and system explanatory capabilities
- potential questions, explanation types and explanatory strategies useful for the different types of users
- domain knowledge enabling the construction of such explanations
- media for users' requests and explanation communication
- explanation evaluation criteria and strategies for palliating explanation failures.



Question menu of "OrientExplain"

Application: telecommunication network design

We built an explanation module for Orient, the CNET platform for assistance to telecommunication network design. Elicitation sessions with experts and end-users and a study of design task helped us to acquire the needed explanatory knowledge on: - the domain itself (i.e. networks, services, etc.), - the network design methodology: mandatory constraints to be satisfied by the network, criteria to be privileged if possible, multiple alternatives leading to different solution variants. - the potential end-users (designers evaluating the solutions proposed by the system vs consultants giving advices to clients): their tasks, objectives, competence level and explanatory needs, their possible questions to the system and preferred explanatory answers), - the different explanatory strategies for offering such answers, - Orient, the platform to be connected to the explanatory module, - the communication media: menus and explanatory text.

We specified and implemented OrientExplain that exploits such elicited explanatory knowledge: based on a planner of explanatory text, OrientExplain focuses on the choice of services (e.g. Transpac) and answers such questions as:

- Why have (resp. haven't) you chosen the service X ?
- What are the advantages (resp. drawbacks) of X (relatively to Y, for a criterion C) ?
- What are the advantages (resp. drawbacks) of X in the solution S1, relatively to Y in the solution S2 ?
- Why wasn't the criterion C taken into account ?
- How is realized the function F ?
- What is the component C useful for ?
- Compare the solutions S1 and S2.

Conclusion

This example shows how knowledge can be captured for the building of an explanatory system that constitutes a specific information system. ■

Please contact: Rose Dieng – INRIA
Tel: +33 93 65 78 10
E-mail: dieng@sophia.inria.fr

Interactive Self Explanation – Making Domain Knowledge Explicit

by Simon Lambert

The I-SEE project (“Interactive Self-Explanation Engine”) is an ESPRIT project with the Rutherford Appleton Laboratory as one of the partners. It aims to produce a customizable toolkit for enhancing information systems with explanation capabilities, and one of the approaches being used is the representation of knowledge about the application domain and the system itself.

The basis of this approach is that even a well engineered application will probably not contain explicit representations of all the knowledge which is not strictly required for its own reasoning, yet is essential for producing accurate and understandable explanations. For example, rule-based reasoning of the kind commonly found in expert systems may compress several inference steps into one rule, and does not include any way of explaining the meaning of the antecedents and consequents of the rules. I-SEE is founded on the idea of augmenting an application with explicit models which permit a rich variety of explanations. An important consequence of this is the possibility of “retrofitting” explanation facilities to existing applications, a prospect which clearly has great exploitation potential.

The domain knowledge held in these models is manipulated by a set of primitive “explanation techniques” which are built up into explanation plans designed to meet the user's information need. Each primitive technique accesses one type of knowledge in the domain model: for example, the system's reasoning chain which led from data to conclusions, or the deep causal knowledge underlying the reasoning (expressed as a network of additional rules), or relationships between

entities in an IS-A hierarchy with their similarities and distinguishing features.

The I-SEE approach is being applied to two applications, a knowledge-based system for management of sewage treatment plants, and a software package for modelling oil spills. The domain knowledge is organized into a number of sub-models, such as a functional model, a structural model, and a causal model. The genericity of the approach is important, and is achieved by having the primitive explanation techniques be domain-independent and available to the developer of a new explaining system as a set from which the most useful may be selected. The I-SEE project also aims to produce guidelines or methods for the developers of domain models. The project is an example of the way in which value can be added to an application by utilization of domain knowledge. ■

Please contact:
Simon Lambert – DRAL
Tel: +44 235 44 5716
E-mail: scl@inf.rl.ac.uk

Reusability Aspects in Transactional Workflow Management Architecture - the TransCoop Project

by Jari Veijalainen and Aija Palomäki

The ESPRIT III BRA project TransCoop aims at developing transaction management support for cooperative applications. The project starts with a study of transactional characteristics and requirements of different kind of cooperative application domains, among them the workflow systems. Based on this analysis, a high-

level specification language and cooperative transaction management architecture is to be developed and a prototype implemented. One important goal of the architectural development is the reusability of existing applications, cooperative transaction specifications and tools.

The TransCoop project aims at the development of methods and software tools supporting cooperative users. Within this field, each project partner (see below) has an application area of its own, workflow systems being that of VTT.

Workflows can be seen as abstractions of business processes. The processes consist of individual tasks interrelated by data, temporal, or order dependencies. Data dependent tasks share data. Temporally dependent tasks have an ordering with a time slice (e.g. "Task T2 must start 3 seconds after task T1") whereas order dependencies simply define the ordering between tasks.

Currently such business processes are either "hardcoded" in several applications written on paper, or they are just in the minds of company workers. No advanced support is available for specification of the processes. Rationalizing, maintaining, reusing and developing the processes is troublesome.

From the reusability point of view, workflow systems can be understood as an integrator, or "glue", between different, heterogeneous, office tools, applications and even systems. The integration as a goal puts high requirements on the architecture of workflow systems. Since workflows can span heterogeneous platforms and systems, or even organizational borders, there must be a homogeneous area where uniformly described specifications exist. A solution envisaged is the high level specification language which is to be developed in the TransCoop project, based on the LOTOS and TM modeling languages. Also, the interfaces of the workflow components different underlying existing services (like data communication, databases, and existing applications) need to be defined in a homogenized way.

An essential part of the architecture is the functionality of the transaction-oriented workflow management system. Clearly, the traditional (database) transaction management approach (or, the traditional transaction model) will not be suitable for transactional workflows, although traditional transactions can be used as atomic building blocks or basic services. This is because workflows may be nested, might have a long duration, and might have real-time constraints or dependencies. To support these properties, an extended transaction model has to be developed and formalized. This is one of the most challenging tasks of the project.

The other partners in TransCoop are GMD-IPSI (Germany, prime), and the University of Twente (The Netherlands). For further details see ERCIM News, February 1994, pp. 12, 25.

Please contact:

Jari Veijalainen or Aija Palomäki –VTT
Tel: + 358 0 456-6014, -6016
E-mail: Jari.Veijalainen@vtt.fi
or Aija.Palomaki@vtt.fi

Automatic Page Layout - Applications in Catalogues, Newspapers, and Electronic Publications

by Juha Ylä-Jääski, Hannu Ahonen, Teppo Veijonen, Juhana Niku-Paavola and Virpi Leikola

VTT develops automated page layout techniques, which have been applied in catalogues and newspaper advertisements for some years now. A further extension is the automatic production of editorial pages from an article database.

The aim of page layout is to produce aesthetically attractive pages, with a clear structure, containing the essential information in an easily readable form. Obviously the page layout has to be configured differently, depending on the application, page size and format, and the media used to communicate the information to the user. In pages with advertisements, minimization of the amount of empty space is usually the main criterion for page layout, while the layout of individual articles, including the choice of the proper typography and placement of graphics and pictures, plays a major role in editorial newspaper pages.

Despite the introduction of computer-aided techniques, the layout of newspaper pages and complicated catalogue pages still involves a lot of manual work. VTT first began to apply automated page layout techniques in the production of yellow-page telephone directories. The software for this application is already used to produce all yellow-page telephone directories in Finland and in the United Kingdom.

Yellow-page directories

In the United Kingdom, the yellow-page prepress production at Pindar Infotek in Scarborough consists of some 80,000 original pages a year; over 2 million line ads and some 350,000 display ads; and a directory with 1,000 pages approximately every fourth working day.

The page layout system operates in a fully automated fashion, producing the layout for a directory with 1,000 pages in less than one hour, using one standard Macintosh or Unix workstation. The automatic process includes the placement of all material onto the pages: line and display ads, classification headings, fillers, and continuation lines. Pindar has estimated that for a fully manual page layout process they would require a staff of some 330 persons, compared with only 2 running the automated page layout software today.

In 1992 when the page layout system was introduced into production at Pindar, the software came first in the UK Trade and

Industry's Manufacturing Intelligence Award. Over 400 companies showed interest in the competition.

Newspapers

In newspapers, the major goal for advertisement page layout is to automate and integrate the placement of display and classified ads. Several special requests restricting the placement of individual ads have to be taken into account, such as the colour requirements for an ad, predefined positions with respect to a specific page or topic section, and possible rejection of two competing advertisers on the same page. The first installations of the newspaper ad system for daily production are in progress for two major European newspapers: Nederlandse Dagbladunie (publisher of Algemeen Dagblad and NRC Handelsblad) in the Netherlands and Dagens Nyheter in Sweden.

Electronic news

As a further extension, the page layout techniques have been applied to produce editorial pages in a fully automatic fashion. The editorial page layout software is configurable for applications ranging from electronic news distribution for computer screens to broadsheet newspapers. This makes it possible to use the

same software with merely different parameter settings, independent of the particular application or the media used to view the result.

The editorial page layout process can be carried out on-line, according to an interest profile of the user, and directly from a database, which includes all the articles currently available for the publication. The processing time for the layout of a page is only a few seconds. An editorial article naturally contains several additional attributes, compared with the rigid size of an advertisement. Attributes for an article include alternatives for the size and contents of the title and the actual article text, possible references to images or graphics, and a priority level within a proper topic section.

Pilot projects to apply the editorial page layout system in a production environment for electronic news distribution are in progress. We also intend to apply the tools of the system to support editorial page layout in daily newspapers.

Please contact:
Juha Ylä-Jääski – VTT
Tel: +358- 0 456 4541
E-mail: Juha.Yla-Jaaski@vtt.fi

Abductive Information Retrieval of Multimedia Data

by Adrian Müller and Ulrich Thiel

The problem of automatic query reformulation is studied in the context of a logic-based information retrieval system that employs - in contrast to approaches based on deductive reasoning - an abductive inference engine. Given a query, the abduction process yields a set of possible expansions to the query. An architecture for an interactive retrieval system based on abduction is proposed comprising a schema-level representation of the documents' contents and structure, an abductive retrieval engine, and a user interface which allows to control the inference process. The retrieval engine was tested on a collection of SGML-structured texts which contains multimedia data.

An increasing demand for end-user oriented retrieval systems motivated a lot of research on methods for automatic query expansion in information retrieval. In most cases, the approaches are based on vector space or probabilistic information retrieval models. In these approaches, the contents of the documents is only represented to a rather limited extent. As this may restrict the enhancement that is achievable by automatic query expansion, we claim that the construction of elaborated queries should be accompanied by a corresponding enhancement of the document representation in the system, and, of course, by the provision of retrieval procedures that exploit them.

The logic-based approach to information retrieval, as put forward by C.J. van Rijsbergen in his seminal papers, defines the retrieval problem as follows: Given a query Q , we can determine for each document D the probability that D implies Q . The documents with high probability estimates are regarded to be sufficiently relevant to the query. Most proposals to use this notion of information retrieval in pro-



VTT's automated knowledge based page layout system.

typical systems are mainly based on deductive inference in more or less restricted forms of first-order or probabilistic logic. However, since the consequence (the query) is known, and we want to know the set of potential premises (the documents), we might ask for an inferential process, which allows us to find those premises directly. Now, this is exactly, what abductive reasoning has to offer: It generates a set of hypotheses, which imply the consequence (the query). They may then be checked by accessing the database. The hypotheses may be regarded as system-generated interpretations of the user's information need. Since they may be rather different, they can be combined as multiple sources of evidence.

For enhanced retrieval operations we employ a unified approach to semantic modeling and the representation of document structures. Using a formal representation of complex objects we can define syntactic and semantic relationships, which are used in an object-oriented model of multimedia documents. As a consequence, the abductive reasoning process yields not a query expansion on the level of search terms, as e.g. in thesaurus-based systems, but also produces different possible "readings" of the query which may

differ in their meaning on both the semantic and structural level.

Figure 1 illustrates the process of query expansion which is embedded in the human-system dialogue. A query is entered by the user in a formbased or direct manipulative interaction, and represented in the dialogue model. The abductive retrieval engine treats this formal representation as a proposition and employs several knowledge sources to generate hypotheses which allow to infer this given proposition. Semantic rules from a domain model as well as structural information from the document model and entries from a data dictionary are used during this reasoning process. Constraints resulting from the current dialogue situation, or a model of the user, prevent the system from producing inappropriate query interpretations. The expanded query, or different possible expansions, can then be validated by the user, or used to retrieve alternative answer sets from the database, which are offered to the user.

Please contact:
Adrian Mueller, Ulrich Thiel – GMD
Tel: +49 6151 869 855
E-mail: adrian.mueller@gmd.de
or ulrich.thiel@gmd.de

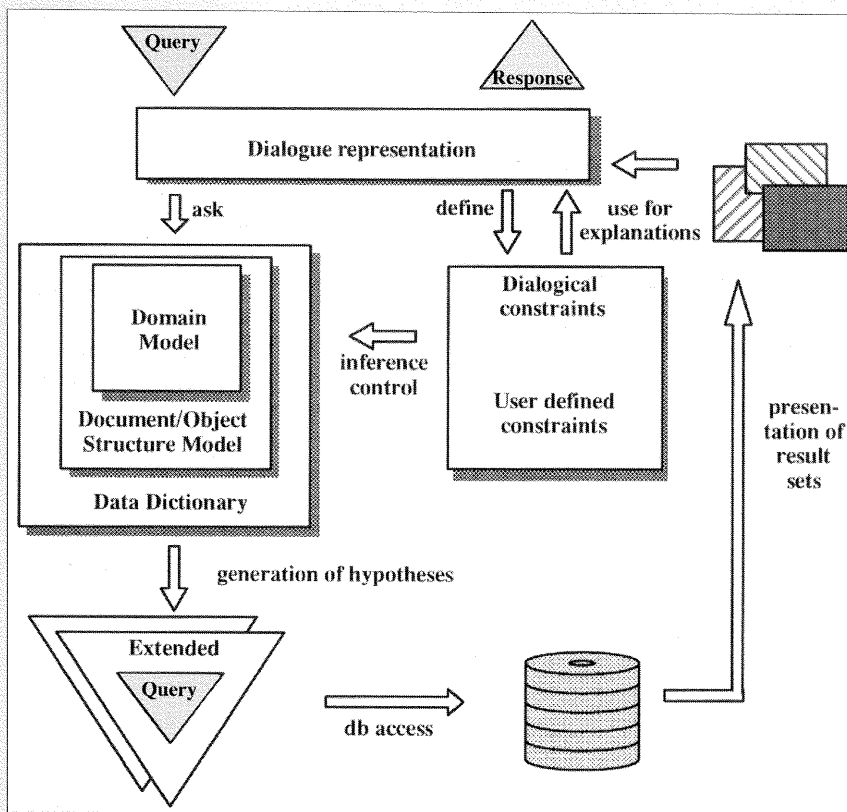
Knowledge-Based Access to Multimedia Document Bases

by Carlo Meghini, Fabrizio Sebastiani and Costantino Thanos

The recent, impressive achievements of hardware technology have opened the way to a new generation of information systems: multimedia document bases (MDBs). An MDB is a repository of complex objects that integrate components of various kinds: text, images, graphics, video, and audio.

MDBs of realistic size are expected to contain thousands of multimedia documents, if not more. Among the many challenges posed by the management of this type of information system, the problem of multimedia document retrieval (aka Multimedia Information Retrieval - MIR) does not appear to be close to a solution, because it involves the content of the documents, which is inherently difficult to understand and model, let alone handle algorithmically. The answer to the problem entails the design of a multimedia information retrieval (MIR) model, that is, of: (1) a language for abstractly representing a) multimedia documents and b) user information needs, and (2) a matching function relating (a) and (b) in order to establish the relevance of documents to given user information needs. Our MIR modelling strategy attempts to combine the logical approach to IR and the conceptual modelling approach to information systems.

In our recent work, we have proposed a logical model of IR based on Terminological Logics (TL). In fact, TLs are particularly suitable for modelling MIR along the guidelines of the conceptual modelling approach to in-



The process of query expansion

formation systems, as they can be employed in representing:

- documents under a variety of aspects (e.g. structural, layout, semantic content)
- user information needs
- domain knowledge.

This proposal, however, does not yet take into account the imprecision that is intrinsic in the MIR enterprise: given that queries are inherently imperfect representations of user information needs, and that document representations are inherently imperfect representations of documents, we cannot expect relevance judgements on the part of the system to coincide with relevance judgements on the part of the user, and we can thus reasonably expect to compute only the probability that a document is relevant to user information needs. In a further development of our work we have thus proposed a Probabilistic TL for modelling MIR, i.e. a logic allowing the expression of real-valued terms representing probability values and possibly involving expressions of a TL. Two different views of probability, i.e. probability as statistical information and probability as information about degrees of belief, can be accounted for in our approach. We have presented a formal syntax and a denotational (possible-worlds) semantics for our logics.

The research outlined in the preceding sections opens up a number of issues, such as the relationships between statistical information and information about degrees of belief, and the role of user feedback. We plan to investigate these issues in the near future.

Please contact:

Carlo Meghini – IEI-CNR

Tel: +39 50 593405

or Fabrizio Sebastiani – IEI-CNR

Tel: +39 50 593407

or Costantino Thanos – IEI-CNR

Tel: +39 50 593492

E-mail:

{meghini,fabrizio,thanos}@iei.pi.cnr.it

Using Domain Knowledge in a Multimedia Toolkit

by Wernher Behrendt

MIPS is a client environment that homogeneously presents multimedia information stored in remote, distributed databases. An Embedded Knowledge Based System supports the mapping of single user queries to multiple remote database query languages, and later optimises presentation by consolidating retrieved information into a single dynamically growing hyperdocument. The system is aimed at the Value Added Reseller market: a general platform can be tailored to the needs of a general application domain, or a specific application document, and thereafter to the specific needs of individual user sites and individual users.

Designing for flexibility at several levels poses some difficulties because the levels are in fact overlapping knowledge domains which interact in various ways. For example, at the user task level (e.g. a design studio), distinctions could be made between photos, paintings, and drawings. At the level of digital multimedia information, further knowledge is required as we need to know how the image is encoded (JPEG, GIF, TIFF, EPS, ...). At the level of a particular business site, we need to know how the encoded image can be displayed on the specific hardware and software platform of that business: this requires a general specification of presentation tools and their capabilities of rendering different encoding formats. Furthermore, a firm may have a very constrained set of display tools but a good range of conversion filters - thus, accessing prima facie, incompatible data sources becomes a planning matter whether a conversion chain can be found from one format to another. A multimedia toolkit must enable the user to do their work without

worrying too much about hardware and software details. It must enable information designers to create applications which are structured according to principles pertaining to the domain of the application (e.g. tourism, or aircraft maintenance). And yet, these principles must be flexible enough to enable automated updating without compromising the integrity of the information presented. MIPS hyperdocuments could be regarded as knowledge assisted information retrieval filters.

The embedded KBS which is being developed at RAL aims at making the interactions between knowledge about multimedia system capabilities and application domain specific concepts transparent. That way, the system becomes capable of interpreting a user's query not only in terms of what information is required, but also in terms of how this information is best put together and presented, given the constraints of the domain, business, and hardware/software platform. Separate knowledge modules exist for describing:

- the schemata of remote databases and their access details,
- the semantic constraints on queries by allowing partial or full domain models to assist in query interpretation
- user preferences and security classes
- domain dependent clarification dialogues
- domain dependent templates for hyperdocuments which get filled in through user interaction with the system

A demonstrator is currently being built to assist travel agents in Greece with multimedia presentations of travel and holiday offers - the interactive holiday brochure.

Please contact:

Michael Wilson – DRAL

Tel: +44 235 44 66 19

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

A Software Test Pilot for Performance Assessment Support

by Martin Kersten

CWI coordinates the ESPRIT project Pythagoras and develops a Software Testpilot to assess the quality of database management systems.

In Pythagoras, the objective is to develop the means of predicting, assessing and tuning the performance of Advanced Information Servers (AIS). To achieve this goal, the project seeks to develop a tool-kit and methodology for performance quality evaluation and prediction of the parallel database servers which are intended to support an AIS. Pythagoras delivers a performance suite aimed at quantifying the effectiveness of AIS in geographic and OLTP business applica-

tions. These exemplars are complemented by in-depth studies to improve the knowledge of the key software technology required, such as concurrency control, query optimization, and data placement. Finally, the project delivers several engineering tools for designers and users of a parallel database server, including an Adaptive Performance Evaluator and an Oracle-base performance predictor.

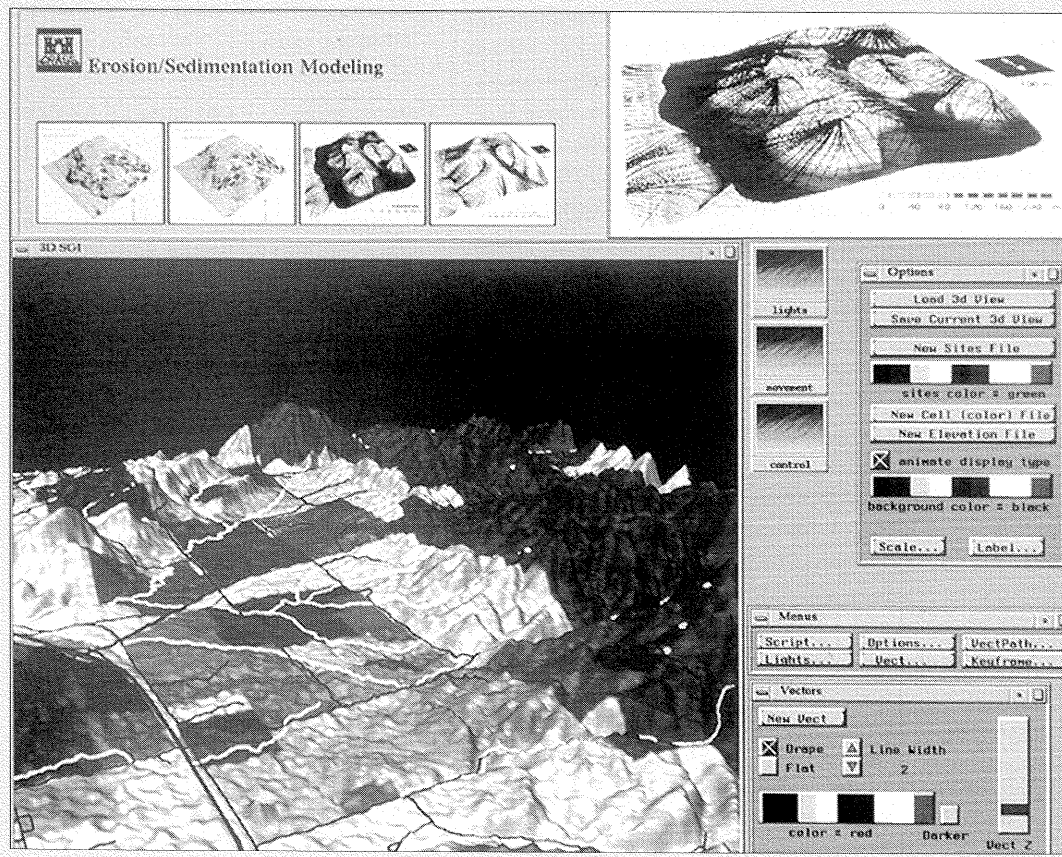
The Software Testpilot, under development at CWI, assists both database designers and users to assess the robustness and the performance of applications running on a (prototype) system.

The common approach to obtaining performance characteristics of a DBMS is benchmarking. Benchmarks may be domain-specific, i.e., they specify a synthetic workload characterizing a typical application problem domain. The benchmarks are often complemented by test-suites geared at exposure of weaknesses in specific system implementations.

A major drawback of the benchmarks and test-suites is that they only represent a few points in the workload search

space. Hence, a DBMS engineer or user with a workload characteristic which is slightly off those measured, may find his system performing badly. Likewise, test-suites may be biased by isolated implementation challenges, e.g. novel data structures or query optimization techniques, and thus neglect possible side-effects. We concluded that DBMS engineers and users require a performance assessment method which goes beyond benchmarking. What they need are intelligent tools to explore a large workload search space, quickly seeking out the slope, top, and knees of performance figures.

In our approach the performance assessment support system (PASS) experimenter specifies an abstract workload search space, a small interface library with the target system, and a description of the expected behaviour within a test-suite. Thereafter, it is up to the PASS to select the actual workload parameter values and to execute the corresponding target system transactions, such that the performance characteristics and quality weaknesses are determined at minimal cost (=time). A novel aspect of this ap-



Geographical information, stored in a database, is used, e.g., to model terrain erosion and sedimentation. Here efficiency and reliability of the database management system are of crucial importance. CWI coordinates the ESPRIT project Pythagoras, which is developing a Software Test Pilot to assess the quality of new database management Systems.

proach lies in using a behavioral model of the target system to drive the process.

The requirements on a PASS are derived using the user's interaction and the target system demands shown in Figure 1.

From the perspective of user interaction, a PASS system should provide:

- flexible language for experimental space specification
- an extendible module to encapsulate algorithms for exploring a large workload search space
- a graphical user interface enabling the user to control and direct execution of the experiments
- statistical validation and fitting of user supplied and system generated hypotheses

to adapt the behavioral model to the target system

- persistent storage of measurement data for post-session analysis by domain specific tools
- documentation of experiments, in the form of a short narrative description of the objectives and boundary conditions of the experiment.

The target system under test should be manageable in isolation from its environment and its state be deterministic, it should be machine controllable, crash recoverable and permit state inspection.

The Software Testpilot is being used on a daily basis both to assess performance of parallel DBMS and to tune database cost models. The approach taken is to

provide a declarative language, called TSL, to specify a large performance test space. Finding an efficient exploration is left to the Software Testpilot. The result is a highly flexible language where the user's prime responsibility is to define actions objects, their relationships, and their interface with the DBMS. A clear separation between test space exploration and target system interaction ensures a highly generalizable tool applicable to a broad application domain.

The Software Testpilot software can be obtained from the author.

Please contact
Martin Kersten – CWI
Tel: +31 20 592 4066
E-mail: mk@cwi.nl

RESEARCH ACTIVITIES

ATM Systems

by Bjarne Helvig

Currently, telecommunication services are handled by dedicated networks, leased lines and the emerging 64 kbit/s ISDN. It is expected that in the future all services will be carried by a single high capacity network, the Broadband-ISDN. This has a number of advantages and will, for instance, enable cost-efficient integration of sound, video, text, graphics etc. into real-time multimedia communication.

The information transfer will take place in an "asynchronous transfer mode" (ATM), where the information stream is split into a number of small, separately transferred "packages", named cells. The ATM technology is about to enter both in the public network and in LAN domains. One of the "features" of ATM is that cells may be transferred as needed at any time. If the maximum traffic carrying capacity of the network

is exceeded, the service will be impaired. The probability for this to happen should, however, be extremely small.

At DELAB, theoretical and experimental work is carried out to determine the traffic/load carrying capacity of ATM networks. During a RACE project (PARASOL R1083), the theoretical basis for an efficient generation of a realistic synthetic traffic load were derived, which enables emulation of a large number of independent sources. Based on this, an advanced traffic generator, the Synthesized Traffic Generator (STG), was developed. The STG enables performance and quality of service measurements of ATM equipment and networks, under a wide range of foreseen and currently unforeseen load conditions. The STG hardware was developed in cooperation with Norwegian Telecom Research and KTAS (DK).

Currently, we are working on a new version of the control software and man-machine interface for an industrial pilot series. This development is a cooperation

between Wandel & Golterman (D) and Norwegian Telecom. The equipment is about to be put into operation for evaluation of ATM pilot systems on a European basis.

Other work at DELAB in this area encompass development of a highly reconfigurable ATM switch equipped for traffic measurements, special techniques for measurement of rarely occurring events like cell losses, modelling of traffic sources, traffic flow management related to service primitives, and analysis of end-to-end ATM traffic processes.

Most of the work in this area is done under contracts with Norwegian Telecom Research.

Please contact:
Bjarne Helvig – SINTEF DELAB
Tel: +47 73 59 26 67
E-mail: bjarne.helvig@delab.sintef.no

Image Retrieval by Pictorial Content in Medical Image Databases

by Catherine Chronaki and Stelios Orphanoudakis

The large number of medical images currently generated by various diagnostic modalities have made not only the interpretation of such images by humans very difficult, but also their management. The use of computer methods in the analysis and interpretation of medical images, as well as in the proper integration of multimodal imaging information, is one of the major trends shaping the future of medical imaging. Similarly, computer methods are currently being used for the efficient transmission, storage, and retrieval of medical images and image related data.

The introduction of Picture Archiving and Communication Systems (PACS) for the purpose of digital image analysis, management and communication in the hospital environment provides easy access to images by physicians and other user groups in the hospital, and supports medical training, clinical research, computer aided diagnosis, radiotherapy planning, surgery planning, etc. In this environment, image indexing and retrieval by pictorial content will give added value to existing PACS by providing considerable clinical decision support, increasing the effectiveness of diagnostic image interpretations, and a valuable tool for medical training and research. However, one must first provide appropriate descriptions of image content and corresponding measures of similarity, in a way that medical images can be compared based on their morphological appearance, their clinical content, or a combination of the two. In other words, the information system should capture and possibly extend the knowledge of medical experts.

"I²C" is an information system for the indexing and retrieval of medical images by pictorial content, which has been developed by ICS-FORTH in collaboration with KNOSSOS Technologies S.A., through their participation in the project EurIPACS (European Integrated PACS), which is funded in part by the AIM programme of the European Union. "I²C" has evolved out of the integration of a set of different tools for the extraction, indexing and storage of image descriptions. Such tools include noise reduction, segmentation, and line approximation algorithms, as well as a contour editor and a database. Some of these tools are interactive, others are automatic. The system is open in the sense that new tools may be added to the system with minimal effort, in some cases without even disrupting the operation of the system.

Image content description and clinical similarity of medical images are captured in the core concepts of the "I²C" design:

Image classes.

Images are classified when entering the system. Images originating from the same diagnostic imaging modality and corresponding to the same part of the anatomy are assigned to the same class, while subclasses can also be created for images sharing other common characteristics.

This allows for content based queries to be directed to an appropriate class of images, thus reducing the search space.

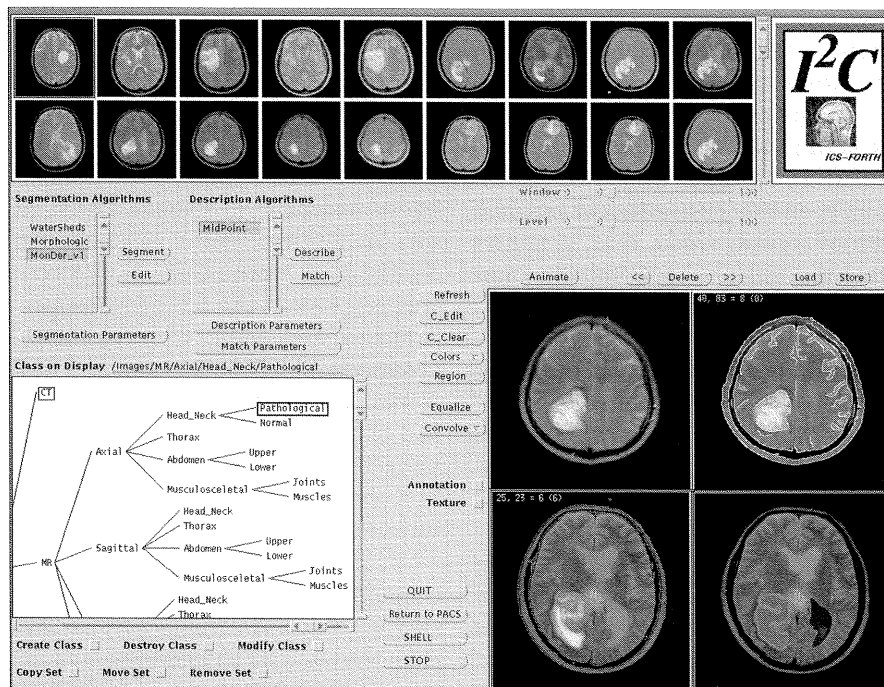
Algorithms customized for specific image classes.

Algorithms and tools used in the extraction, indexing and storage of image descriptions are all attached to classes. Thus, "I²C" may accommodate tools created for a specific class, which have and use built-in knowledge of the specific characteristics of that class. General-purpose tools, which may be used in any class of images, are also parameterized to allow them to be tuned to a specific class.

Image Description types.

An image description type encapsulates image content extraction, indexing, storage, and retrieval methods and implements a particular approach to image content description. An image description type developed for MR knee examinations may include knowledge of knee anatomy and characteristic malfunctions as depicted by magnetic resonance. It may also assume knowledge of the noise characteristics of the particular modality and use it to improve the processing of the image.

The aim of "I²C" is to capture the expert knowledge in the context of medical image classes with the appropriate de-



The User Interface of the I²C Prototype

scription types. This is an evolving activity. Researchers in image description are using the system to develop new methods and improve old ones. One of these activities involves the development of *AttributeMatch*, a general and sophisticated indexing and retrieval by content scheme that attempts to capture clinical similarity of medical images. For each image, the description type maintains a hierarchical image description based on the output of segmentation. This description consists of contours, segments and their properties and points. At each description level, a set of attributes is defined. Some attributes are precomputed and stored, while the rest are computed on demand at retrieval time to maintain a compact image description. In addition to stored attributes, each description components associated with an importance factor which denotes the contribution of each component to the description at higher levels of abstraction. A query by pictorial content involves the specification of a grey scale image, a grey scale medical image with a region of interest specified on it, or even a sketch drawn on the contour editor. The users are allowed to direct the search via a set of search parameters, such as relations among objects, relative position, texture, orientation, border type, etc. Using such options, a user may specify a complex clinical similarity criterion.

For example, the user may consider the presence of objects of a given area and texture, with a sharp border, irrespective of location as sufficient to yield a good match.

Current efforts are concentrated on improving the system's user interface and refining the image description types used, as well as implementing new ones, in order to better capture the knowledge of medical experts. The final "I²C" prototype is expected to provide added value to integrated picture archiving and communication systems, and to support computer aided diagnosis, medical training, and research. The system will be integrated and evaluated in a clinical PACS environment in terms of its usability and acceptance by the medical community. ■

Please contact:
Stelios Orphanoudakis – FORTH
Tel: +30 81 391600
E-mail: orphanou@ics.forth.gr

Architecture for an Interactive Multimedia Retrieval System

by Ulrich Thiel

The GMD-Institute for Integrated Publication and Information Systems is developing an information retrieval system which exploits the power of abductive reasoning for accessing relevant data, providing a convenient interface that supports unexperienced users. Following van Rijsbergen's suggestion, GMD-scientists regard retrieval as an inferential process that determines those documents which imply the query. Abductive reasoning generates a set of hypotheses which allow to conclude that a document is relevant with respect to the query. The hypotheses are checked by accessing the database.

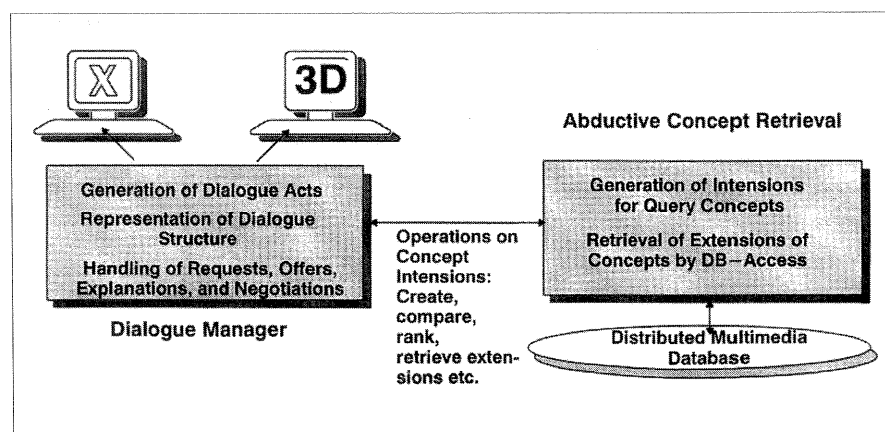
Traditionally, researchers in information retrieval (as well as in the database field) assumed that it is sufficient to look at the current query and to identify the alleged relevant documents by a procedure that computed those documents matching the query according to a given retrieval

model. In our experiments, we found that the abductive reasoning can produce unanticipated hypotheses, that may be surprising or inadequate. Thus, a full-fledged interface to an abductive retrieval system should be capable to explain the hypotheses, and if necessary, to negotiate them.

Here, we employ an appropriate formal model of the dialogue, which is interpreted as a multimodal conversation. This includes, as in our previous prototype MERIT, an abstract representation of the interface and the user's actions. Thus, we have at hand the means to capture the tactics of single dialogue contributions (in terms of speech act theory: their illocutionary parts) as well as the strategic plan the user is assumed to pursue (e.g. inferred by a case-based reasoning process).

The system should not only retrieve relevant data, but also provide an appropriate layout, and the software tools (e.g. for decompression or visualization) needed to access the data in the current setting. For enhanced retrieval operations, e.g. concept retrieval, a semantic modelling of the domain as well as an explicit representation of the contents of the document will be necessary (including the information about dynamic layout and the required software environment as well).

We regard a multimedia document appearing on the screen as an interactive object combining three aspects: The raw data, i.e. the set of texts, pictures, videos



The Interactive Multimedia Retrieval System

etc., which are part of the document (called underlying objects) are to be determined by appropriate retrieval methods according to their type. The layout of the interactive object is inherited from a specific layout object specifying the temporal and spatial relationships of the document components. Finally, the media control components needed to realize the visualization are regarded as media objects.

Please contact:
Ulrich Thiel – GMD
Tel.: +49 6151 869 855
E-mail: thiel@ darmstadt.gmd.de

Conversational Roles Model

by Adelheit Stein

The design of flexible and cooperative user interfaces to information systems, such as databases, requires an elaborate model of the interaction process. Scientists at the GMD-Institute for Integrated Publication and Information Systems propose a conversational approach to modelling information seeking interactions between two partners - human and computer. The "Conversational Roles" (COR) model was used to design and implement several prototypical interfaces, e.g., the systems MERIT and CORINNA.

COR is a speech act oriented dialogue model which allows to describe the formal structure of a dialogue in an abstract, modality-free way. Thus, it allows to interpret linguistic as well as non-linguistic (graphical, deictic) actions of user and system as conversational contributions to the ongoing dialogue, i.e., as dialogue acts or communicative acts. The complete interaction is referred to as multimodal conversation. The model was influenced by the "Conversation for Action"

(CfA) model proposed by Winograd and Flores, where discourses are interpreted as "negotiations". We extended their model for the situation of information-seeking dialogues, applying concepts of Systemic Linguistic approaches to discourse modelling and Rhetorical Structure Theory (RST). Basically, COR defines the generic dialogue acts available (e.g., requesting or offering information, rejecting offers, withdrawing commitments) and describes their functional interrelations. For instance, an information offer or request for information can be accepted or rejected by the addressee, which would then require different possible responses; or the decision of how to proceed is postponed by inserting a clarifying subdialogue where the conditions for this decision are negotiated or backing information is exchanged. COR can be represented as a recursive state-transition network (see figure 1) where arcs signify 'dialogue acts' and nodes mark 'dialogue states'.

Note that the network is recursive, i.e., the arcs are also state-transition networks which may contain - among other transitions - dialogues of the basic schema. Additionally, dialogue acts define both the role assignment between information seeker and information provider and the changes of those roles that may occur in certain dialogue situations, e.g., in clarification dialogues.

Please contact:
Adelheit Stein – GMD
Tel.: +49 6151 869 841
E-mail: stein@darmstadt.gmd.de

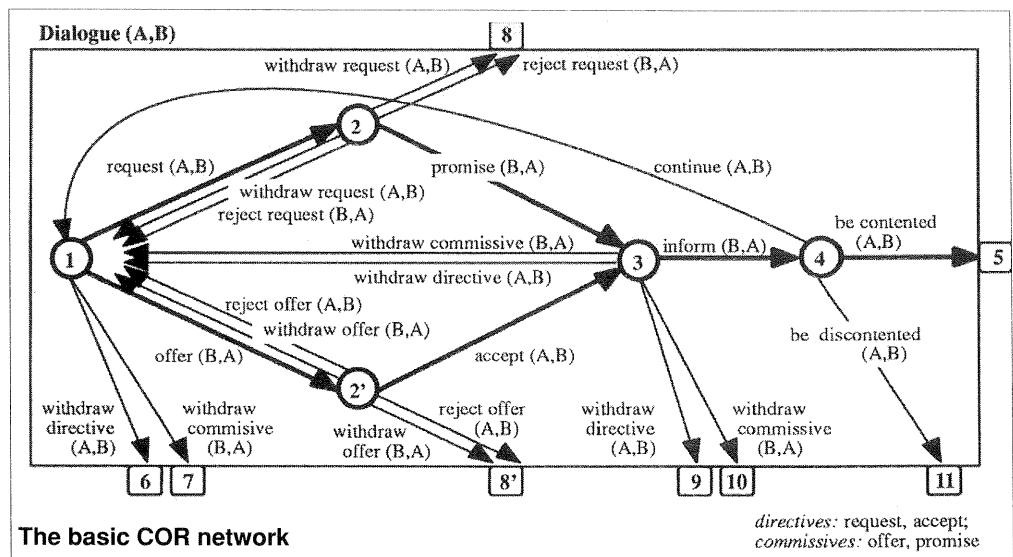
Real-time Programming Methodology

by Rolv Braek

DELAB is the leading partner in SISU, a national technology transfer programme introducing improved methods and tools for real-time software development into the Norwegian software industry. SISU is a cooperative programme, and DELAB's participation involves a very close cooperation with a variety of industrial firms.

According to the initial plan, the project beginning in 1988, was due to finish in 1993. The participating companies found the cooperation so rewarding that it is now extended for another four year period under the name SISU II.

A central part of the collaboration has been to develop and document a common systems engineering methodology. Based on our long running experience from industrial developments, DELAB has been able to contribute substantially in this work. Our staff has co-authored an internationally published textbook on the methodology. This has contributed to the methodology becoming internationally recognized and accepted, and stimulated an increasing number of Norwegian companies to take it into use.



DELAB has also developed methods for configuration management and tools for automatic program generation that are now being used by the SISU project participants in their product development projects.

Please contact:
Rolv Braek – SINTEF DELAB
Tel: +47 73 59 26 70
E-mail: rolv.braek@delab.sintef.no

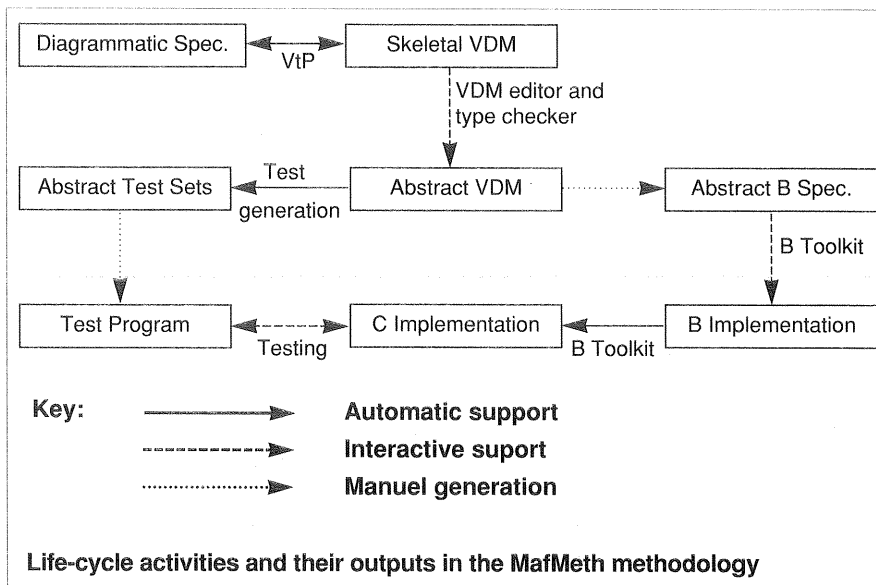
Application Measurement and Dissemination of a Mathematically Formal Methodology

by **Jeremy Dick** and **Juan Bicarregui**

The MaFMeth project is an application experiment funded under the EC ESSI programme. The project will develop part of a high-integrity system by applying mathematically formal techniques to all phases of development. It is a collaboration between the Bull development centre (Hemel Hempstead), Bull S.A. (Paris), B-Core Limited (Oxford) and EPSRC-DRAL (Oxon).

The software to be developed is a component of a distributed transaction processing monitor deployed in financially critical banking environments. This software allows applications running on heterogeneous platforms to coordinate their transactional activities via message passing. The monitor provides message storage and transformation services.

The toolkit "VDM through Pictures" will be used to create an initial abstract specification. The specification structures are initially described in a diagrammatic form which is then used to generate a skeletal



VDM-SL specification to which further detail is added. A set of functional "black-box" tests cases is then generated automatically from the specification.

The abstract specification will then be translated into the B Abstract Machine notation and the B Toolkit used to refine the design towards code. The B to C code generation facilities available in the toolkit will be used to generate code from the refined specification.

The C code generated from the concrete B specifications will be tested using the test cases generated from the abstract VDM specification thus providing the formal link between the VDM and B models of the system.

The process and product developed using the formal methodology will be compared with other system components developed using the established development process. The measurements will be taken in terms of faults found during unit test, validation test and customer use.

The project started in April 1994 and will run for 14 months.

The Bull development centre at Hemel Hempstead is the prime user, Bull S.A. and B-Core Limited will provide the two toolkits and EPSRC-DRAL will provide expertise in the two formalisms.

Please contact:
Juan Bicarregui – DRAL
Tel: +44 235 44 6380
E-mail: jcb@inf.rl.ac.uk

The ESPRIT Special Action for Greece in the Multimedia field

by **Yannis Vassiliou**

The ESPRIT Special Action for Greece in the Multimedia field is a concerted initiative of the European Commission in cooperation with FORTH to support and promote the introduction and use of Multimedia technologies and methods in the Greek industry

Two of the main fields of applications are:

- Preservation of Cultural Heritage
- Computer Integrated Manufacturing and Engineering

Multimedia in preservation of cultural heritage

The objectives of Hel-MME on Cultural Heritage are to develop an economically viable prototype application, able to stimulate the dissemination and preservation of the European cultural heritage; to propose appropriate standards (data format) able to harmonise similar applications in Europe and world-wide; to

generate business opportunities for hardware and software providers and for the electronic publishing industry. Two projects are under Hel-MME on Cultural Heritage:

- Multimedia application based on the Greek cultural heritage of northern Greece (ALEXANDER project)

The aim of the "ALEXANDER" project is to introduce and establish advance research facilities in representative areas of the cultural heritage based on treasures of Northern Greece. Facilities for capturing and recording objects of art will be established in Dion and Vergina. To create databases of Art objects, a database mechanism will be established for achieving the retrieval of images and text. A prototype CD-I title referring to historical aspects of ancient Macedonia will be made. This title will demonstrate the capabilities of the new technologies towards the introduction of electronic publishing methods in the rich area of the Greek cultural heritage. The project will develop two demonstrators: one in the premises of the Byzantine Museum in Athens and another in the University of Thessaloniki.

The project participants are: ATC, Intrakom, Byzantine Museum of Athens, Aristotle University of Thessaloniki, Lambrakis Research Foundation, New Media Productions, and the National Technical University of Athens.

- Interactive Multimedia applications on the on the history and cultural heritage of northern Greece with emphasis on Macedonia (IMAGINA project)

IMAGINA aims to apply the Multimedia technology in the important field of the history and culture of Macedonia. The history and culture of Macedonia - not only the political history but also political institutions, social structures, economy, language, religion, art and intellectual achievements of the people of Macedonia as interlinked and organically interwoven manifestations, from Pre-history to the present-day - will be presented using Multimedia technology.

The CD-ROM that is going to be published, will take advantage of the state-of-the-art of the authoring of Multimedia applications using the latest technological platforms and standards. The Multimedia material will be collected and retrieved by the most well known scientists and artists of Greece.

The project participants are: IDEAL, Ekdotike Athinon, MUSIC - Technical University of Crete, KERA, Exodus, and Multimedia Systems Center.

Computer Integrated Manufacturing and Engineering (CIME)

Three projects have been created to support and promote the introduction and use of CIME technologies and methods in the Greek industry:

- IASON

The IASON project is designed to provide the means of diffusion of CIME Technology in the Greek Industrial Environment. The technological baseline is focused on selected ESPRIT projects addressing the domains of architectural modelling (CIM-OSA/VOICEI & II and GRAI technique), industrial networking (CNMA & DINAS- DOS), novel production management architectures (COSIMA & IMPACS), production management in the process industry (PROFIT & EPIC) as well as technological developments available in the consortium (CIMBAN prototype/product and IDEAL-CIM).


The IASON project revolves around the technologies of production management in the shop-floor, utilising and establishing distributed architecture based on the concept of coordinating autonomous units (intelligent agents). The manufacturing domain covered by IASON covers a wide range of industries, and in particular, process industry, wood processing, semi-process sector, electronics, construction-site and food processing.

The project participants are: Intrasoft, Intarkom, Heltec, Prism, Zenon, Shelman, Ergo Medical, Econ Electronics, Elval, University of Patras, and the National Technical University of Athens,

- DELTA-CIME

The objective of the new generation of Computer Integrated Manufacturing and Engineering (DELTA-CIME project) is to plan and control the manufacturing resources on business and financial impacts, using the latest technological tools.

The industrial controls, quality assurance and MRP II Systems will be applied in the Industry DELTA S.A. The monitoring system for quality control will be applied in ETMA S.A. The rest of the participant industries will take part in a special dissemination programme where the problems and the benefits deriving from the application of the developed system will be analysed. For the development of the planning production and monitoring systems, innovative approaches will be attempted.



Opening screen of the "Alexander" Project CD-I

The project participants are: ITCC, Delta S.A., Etma, S.A., Elais S.A., Akridas S.A., Bekka S.A., Biokarpet S.A., CTI, University of Thessaloniki, Etat S.A., Technical University of Crete and Athens Innovation Center.

• CIMEIG

The CIMEIG project (Computer Integrated Manufacturing in Electronic Industry in Greece) focuses on assisting Greece with manufacturing management. The CIMEIG seeks to provide and develop a user friendly production and quality control system which will be implemented in large Greek manufacturing organisations.

The CIM system that is going to be developed consists of locally distributed independent stations connected to a central unit. Each station is running a specific application based on its own computing power. All stations are connected to each other through the central unit, in order to exchange data. The system will require minimum installation and customisation and will enable user companies to improve their production methods. The development and the implementation of such a system could be in parts and in phases. The proposed research project aims at the development of some CIM components and its implementation in an integrated application, as the basic structure for future development and integration. More specifically the CIMEIG project develops the following sub-systems:

- Production Planning System
- Production Control System
- Statistical Quality Control System (SQC systems)

The project participants are: Hourdakis S.A., National Technical University of Athens, and Intelltech.

Please contact:
Yannis Vassiliou – FORTH-ICS
Tel: +30 1 7245324
E-mail: yannis@ics.forth.gr

An Integrated System for Feature-Based Modelling

by Bianca Falcidieno

A system which aims at integrating two different approaches towards feature-based modelling: design-by-features and feature recognition, is now under development. In this system, the user will be able to generate a feature-based description by using both features and geometric primitives and to map feature-based descriptions between different application contexts through a feature conversion mechanism. The system is being implemented in a collaboration between the Istituto per la Matematica Applicata (IMA), CNR, Genoa, Italy, and the Fraunhofer Institut für Graphische Datenverarbeitung, Darmstadt, Germany.

A new generation of more intelligent CAD systems is now being employed to exploit the potential of features, considered as communication links between the application and the geometry worlds, since they assign a functional meaning to geometric and topological information within a product description. A feature-based model can be created by following two different approaches: design-by-features and feature recognition. Using the first strategy, the user creates the object model directly with its relevant features whereas, in the recognition approach, feature information is extracted from the geometric model of the object. The main limits of these two methods are that while, on the one hand, design features can differ from the features of interest for a given application, on the other, no feature recognition algorithm can capture the designer's intent. As a consequence, the solution for efficient feature-based modelling seems to be a combination of both strategies.

Our research project has achieved integration through the definition of an intermediate representation which is shared by the recognition and the design systems and from which a feature conversion mechanism, based on recognition techniques, can derive different application dependent representations.

IMA's role in this project is to supply the feature recognition system and to define the intermediate model. A feature recognition system which was developed at IMA in the past has been partly re-designed and extended in order to be integrated in the global system. Thus, the intermediate model has been defined taking into account the requirements of the design, recognition and conversion mechanisms. The model belongs to the hybrid model class and aims at capturing the advantages of the most used solid representations (boundary and volumetric) and has a hierarchical graph structure in which feature relationships are defined. From the application viewpoint, manipulations on the model are performed by the feature converter, which maps the intermediate model into different context-dependent feature-based models. These manipulations basically operate on the graph structure of the intermediate model as feature composition and decomposition operators.

The system is being implemented in C++ on top of the commercial solid modeler ACIS (provided by Spatial Technology) which supports non manifold geometry. The project leaders are Prof. José Encarnacao, in Germany, and Dr Bianca Falcidieno, in Italy.

This program is partially supported by a HCM CEE project. People interested in participating may submit a curriculum to the following address:

Prof J. Encarnacao
 FHG Wilhelminenstrasse 7
 D-6100 Darmstadt
 Fax +49-6151-155199
 Code ERB4050PL930918-MATH

Please contact:
Bianca Falcidieno – IMA-CNR
Tel: +39 10 6475 667
E-mail: falcidieno@image.ge.cnr.it

ARIS Architectures for Information Security

by Winfried E. Kühnhauser

In the last decade, progress in computer systems architecture was strongly influenced by the evolution of open, distributed systems that provided redundancy, parallelism, and modularity. Nevertheless, new technologies generally also bring forth new risks. In the field of computer security these risks become specifically obvious in critical applications like hospital information systems, juridical document processing, and electronic banking. In these areas, information technology will only be successful and socially acceptable if we manage to establish trust in the security of these systems.

Traditional computer security concepts generally assume an isolated, monolithic system base. Modern system architecture whose very basic concepts include openness and distribution, makes traditional security concepts fairly obsolete. Thus while providing a significant progress with respect to availability, performance, and flexibility, modern system architectures are a real threat to system security. The virulence in this situation becomes clear by an analogy: no car manufacturer would sell a new and much faster car without providing more powerful brakes at the same time.

A second major impact on a modern security architecture is imposed by today's sophisticated application systems. Traditional concepts for security architectures assume a single, fixed set of rules (the security policy) describing the behaviour of the system with respect to security. In traditional security architectures, the security policy is an unseparable part of the security kernel, an encapsulated system component that implements every security-related feature of the system. According to the goals of the original initiators of computer security research, this

security policy had only a single goal: the protection of the confidentiality of (military) documents.

Today, the need for computer security goes far beyond military applications. The use of information technology in hospitals, law authorities and banking doesn't allow any longer such a narrow view. As an example, the availability of information in hospital information systems is in many situations more vital than confidentiality, and electronic cash systems regard integrity as a major security goal.

As a fundamental difference to early security architectures, the traditional way to maintain a single systemwide security policy within a security kernel is obsolete. The security goals in today's systems are defined by the application and no longer by an underlying operating system or security architecture. Each application has its own individual security goals, and a basic functionality of a modern security architecture is the binding of an application to its application-specific security policy.

In the last decade, several security mechanisms have been developed to cope with these two problems. Unfortunately, these mechanisms only cure local diseases (e.g. using encryption to guard integrity and confidentiality of distributed systems communication) and fail to provide a more general solution that revives the ideal conditions of isolated monolithic systems with a single burnt-in security policy.

As a first step in this direction, ARIS is developing a new security architecture for distributed systems that is a framework for the cooperation of applications, application-specific security policies and security mechanisms.

In contrast to the traditional security kernel concept, the security architecture supports individual application-specific security policies. To cope with future requirements, the architecture allows to dynamically integrate new security policies.

As engineering requirements, specification and implementation of a security policy is a most expensive software devel-

opment (usually requiring formal verification), reusability of security policies is a major concern. The architecture supports a separation of application and security policy, together with a dynamic binding mechanism that submits an application to a security policy. This property allows to view application and security policy as orthogonal system components that permit to run the same application in different environments under different security policies.

As part of the dynamic integration of new security policies, the architecture supports the integration of new security mechanisms that may be needed to enforce new policies. The long term goal of ARIS is to demonstrate and evaluate the postulated properties of our security architecture. Several major components of the architecture will be built, and we plan to develop sample applications that will make use of the architecture's special features ■

Please contact:
Winfried Kühnhauser – GMD
Tel: +49 2241 14 2480
E-mail: winfried.kuehnhauser@gmd.de

CORINNA – Cooperative Retrieval Dialogues

by Ulrich Thiel and Markus Fischer

Unfamiliar user interfaces often impede users who want to access to an information system. The natural language as a familiar medium enables even naive users to interact with a system. The conversational model underlying such a human computer interaction regards user and system as cooperative partners who share mutual goals. A comprehensive conversational model has been developed at GMD-Institute for Integrated Publication and Information Systems and was successfully used in the prototypical interface CORINNA. The model - called COR

(COversational Roles) provides a framework for the formal description of information seeking dialogues.

The CORINNA system (COoperative Retrieval Interface based on Natural language Acts) employs the COR model - augmented by elements of the Rhetorical Structure Theory (RST) - as a means to describe the structure of dialogue situations and the dialogue history. During a dialogue the system is able to provide hints for a goaloriented reformulation of queries according to appropriate retrieval tactics.

The realisation of dialogue contributions is based on composable sentence plans that represent the illocutionary as well as the propositional aspects of the utterances. These plans are processed by the Penman system which generates the English text. (The employed version of the Penman system was developed in a cooperation between the KOMET department of GMD-Institute for Integrated Publication and Information Systems and the Information Sciences Institute). The formulation of natural language queries is supported by a system of menus which allows the user to compose a sentence by menu selections.

The system was programmed in Common Lisp. The demo version provides access to the relational database containing a download from the CORDIS database offering information on research

programmes and projects funded by the European Community.

Please contact:
Ulrich Thiel - GMD
Tel.: +49 6151 869 855

The screenshot displays three main components of the CORINNA interface:

- dialogue history:** A list of interactions between the User and System.
 - User: Show me all programs.
 - System: Can you tell me a term that is more specific than program?
 - User: Why?
 - System: I am asking you because program is too general.
 - User: OK
 - User: Environmental - protection - program
 - System: Environmental - protection - program is OK.
 - System: The table contains all relevant entities.
- menu for user acts:** A menu for user choice.
 - OK
 - This is not sufficient.
 - I do not like that. Let's stop the dialogue.
- retrieval results:** A table listing research programs.

ACRONYM	TITLE
ENVIRO 4C	Multinational R&D programmes (EED) in the field of the environment, 1986-1990
ENVPROT 4C	Research programme (Euratom, EEC) to be implemented by the Joint Research Centre - environmental protection -, 1984-1987
ENVPROT 3C	Sectoral research and development programme (EEC) in the field of environment - environmental protection - (indirect and concerted actions), 1981-1985
ENVPROT 4C	Multinational research and development programmes (EEC) in the field of the environment - protection of the environment -, 1986-1990
ENVPROTJ 5C	Specific research programmes (EEC) to be implemented by the Joint Research Centre - environmental protection -, 1988-1991
ENVCLI C	Sectoral research and development programme (EEC) in the field of environment: Environmental protection and climatology (indirect and concerted actions), 1981-1985
ENVIREG	Community initiative (EEC) concerning the environment, 1990-1993
ENV C	Specific research and technological development programme (EEC) in the field of the environment, 1990-1994

CORINNA describes the structure of dialogue situations and the dialogue history

TECHNOLOGY TRANSFER

Verification for Distributed and Concurrent Systems

by Henk Nieland

The use of computers in safety-critical applications is increasing rapidly, as is interest in the theoretical foundations for the design of reliable systems. Such systems are used in embedded applications and in interconnected networks. CWI is looking into formal techniques to specify and ver-

ify these systems. These techniques were recently applied to prove the correctness of part of a communication protocol used in the Philips 900 audio system.

The first step in finding an automated solution to a problem is to specify the problem, that is to say, describe it in general terms, for example without referring to a particular type of programming language or computer. A specification provides the criteria for a program's design; if these criteria are met the program is called correct. The criteria are formulated in a specification language. Given the increasing size and complexity of the problems to be automated, attention is focusing nowadays ever more frequently on formal methods.

In line with the specification is the verification of the resultant software, in other words, proving that a program possesses a particular property. Of course, ideally, proof of correctness is supplied along with the construction of the program. However, this is still proving to be extremely troublesome in practice.

Specification, design and verification of software have existed since programming began. With the arrival of distributed computer systems and parallel processing, the need arose to develop new techniques for concurrent programs. That very soon proved to be much more difficult than with sequential programs. For example, one of the problems is to prove that nothing will go wrong with the synchronization of concurrent processes. The use of process algebra is a highly promising approach. CWI has made important

contributions to its development. With this approach, concurrent processes can be described in formal algebraic terms. Recently, CWI has obtained a strong interest in "hybrid systems". These are reactive systems consisting of both digital and analog components. The digital components are typically computers or microprocessors controlled by programs, whereas the analog components will be continuously changing environment variables such as the voltage on a wire and the speed of a train.

Much of CWI research in this field takes place within the framework of European programmes. The recently completed RACE project, SPECS (Specification and Programming Environment for Communications Software), investigated how to automate and optimize as far as possible the design of telecommunication software. CWI played a key role in developing the m-CRL specification language, based on process algebra. In the RACE project, BOOST (Broadband Object-Oriented Service Technology), a software environment is being designed for the rapidly expanding network services (for

example airline company flight schedules). Furthermore, research into calculi and algebra's for concurrent processes is being carried out in the ESPRIT project, CONCUR. While the European approach mainly focuses on algebraic techniques, the work being carried out in the important cooperation with MIT (USA) leans more towards assertional methods and simulation relations. This includes work with "invariants": statements which remain valid throughout every stage of a program, in which they perform a function comparable with conservation laws in physics.

An important development, especially for industrial applications, has been the extension of the specification and verification techniques developed so far to real-time and hybrid systems. These applications range from audio equipment and fighter planes to process controllers. CWI research currently includes the development of a process algebra for real-time systems, and the use of general purpose theorem provers to assist in finding and validating invariant assertions for hybrid systems.

Recently, CWI has analyzed a key fragment of a protocol developed by Philips for the exchange of control information between the various devices of an audio system (tuner, CD player,...). The protocol runs on microprocessors that also have to perform several other tasks, such as positioning the laser of a CD-player and updating a display. Because these other tasks frequently have higher priority, the key difficulty in the protocol is a significant uncertainty in the timing of events: sometimes when the protocol wants to send a message it may not be allowed to do this right away, because a microprocessor is busy with some other task. The analysis of the Philips protocol has been carried out using a model of hybrid systems and proof methods that have been developed in collaboration with MIT. ■

Please contact:
Frits Vaandrager – CWI
Tel: +31 20 592 4232
E-mail: fritsv@cwi.nl

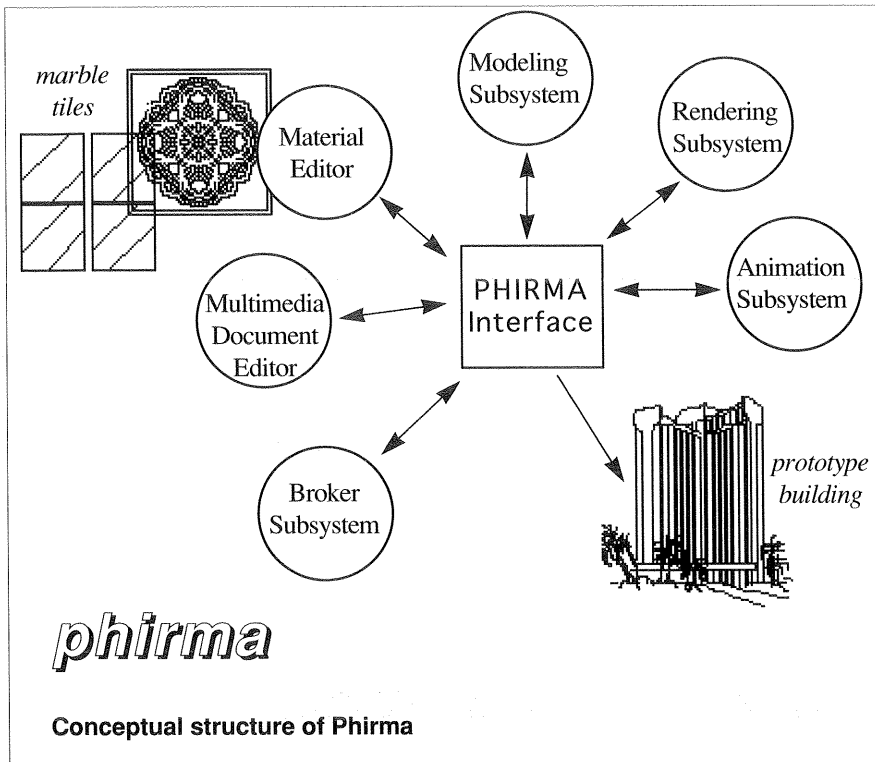
PHIRMA – Photorealistic Image Rendering for Manufactured Articles

by Patrizia Palamidese

The main objective of the PHIRMA project (Photorealistic Image Rendering for Manufactured Articles) is the development of a high-quality and competitive cost system for the production of photorealistic synthetic images of manufactured objects and their integration in an electronic broker system. The project is funded by the CEU Craft program (a subprogram of BriteEuram), which is directed towards small/medium enterprises (SMEs).



Fully fledged computer networks are standard features in today's consumer electronics, like this Philips 900 audio system. CWI has proved the correctness of part of the Easylink real-time protocol used for the system's inter-communication. Photo courtesy n.v. Philips Industrial Activities Leuven.



The motivating observation behind the setting up of PHIRMA is that computer-graphic simulation of industrial products becomes a competitive technology for many small/medium enterprises (SMEs) whenever it can replace, either partially or totally, some of the production phases of a product. The project is being initially promoted by the marble industry of Carrara (Italy) and by some medium sized furniture industries in the Pisa area, in collaboration with researchers from CNUCE-CNR. In the marble industry, the availability of low-cost but high-quality graphics systems would allow the designer to create buildings or interiors and test the appearance of large marble surfaces before proceeding with execution. Similarly, the furniture industry would see a significant improvement in its productivity, in terms of cost and time, if only it could get round the need to set up real, furnished rooms as prototypes, in order to test the harmony between styles, colors and fabrics and to create illustrated catalogues of the final products.

One crucial requirement, stressed chiefly by the partners working with marble, is the need to simulate realistic surfaces of large buildings by reproducing the dominant characteristics (in terms of colour, tone and pattern structures) of the mar-

ble slices cut from a particular marble block. Hence, a basic objective is the study of techniques for generating materials and textures from real samples effectively enough to overcome the drawbacks of current texturing methods (repetitivity of patterns on large surfaces, unrealistic effects of algorithmically-generated textures).

In addition, realistic images of the final product can also be used in the preparation and editing of electronic sales catalogues, for producing animated advertisements, and for defining contract conditions. This information, together with geometric models, textual descriptions and cost data, can easily be interchanged via networks between parts producers, retailers and clients. The market share of software tools for telematic product diffusion and sales in the coming years is expected to be very high.

PHIRMA demonstrates that advanced techniques can be smoothly integrated into existing user environments, while adding to the value in comparison with existing products. The basic idea is to bring together up-to-date solutions (see figure) in user interface design, computer graphics, image-processing and multimedia-document-editing disciplines. PHIRMA is an open system,

based on an object-oriented client-server design which guarantees extensibility and independence from existing packages such as modelers and renderers.

Please contact:

Patrizia Palamidese – CNUCE-CNR

Tel: +39 50 593226

E-mail: patrizia@vm.cnuce.cnr.it

or Roberto Gagliardi – CNUCE-CNR

Tel: +39 50 593259

E-mail: robert@vm.cnuce.cnr.it

or Primo Coltelli – CNUCE-CNR

Tel: +39 50 593308

E-mail: coltelli@vm.cnuce.cnr.it

General Design — A new Company for User Interface Design

by Lon Barfield

General Design is an Amsterdam based user interface company set up by three former CWI employees. The company is committed to quality design of the user aspects of interactive technology and offers a wide range of services to industry and to academic groups seeking industrial collaboration.

The services include:

- Evaluation:

The evaluation of the usability of existing or proposed systems.

- Design:

The design of user interfaces for new or existing systems, from the underlying concepts and metaphors through to detailed visual and audio design.

- Prototyping:

On-screen prototyping and simulation of user interfaces, for evaluation purposes or for demonstrations and other publicity activities.

- **Implementation:**

General Design's computing background means that they can deliver not just designs but also implementations of user interfaces.

- **Lecturing:**

Further to these activities General Design is also involved in the delivery of courses on user interface design and multi-media to both academia and industry.

The majority of today's interactive systems is not properly designed for ease of use. This is typified by the fact that each new product is accompanied by huge incomprehensible user manuals,

telephone help lines, text books on how to get the best out of it and long and expensive training courses on how to use it. General Design is working in a fast growing area tackling these usability problems and thinking about tomorrow's. On the horizon are multi-media, interactive TV, virtual reality and the interactive world promised with the advent of data superhighways. Each of these new steps will bring its own unique usability problems to the world of interactive technology. Solving these problems is what General Design is all about.

Among the first commissions were the simulation of an electronic payment system and the design of a User Guide for a Software Test Pilot, both under development in the ESPRIT projects CAFE (Conditional Access for Europe, see also ERCIM News 16) and PYTHAGORAS (see elsewhere in this issue), respectively, and both coordinated by CWI.

Please contact:

Lon Barfield – General Design
Tel: +31 20 592 4140
E-mail: lon@cw.nl

INESC – International Relations en route to the River of Pearls

by Isabel Caetano

INESC has been improving its contacts and associations with institutions operating in the Territory of Macao as well as in different Chinese Provinces since 1993. We have submitted several projects involving cooperation under the terms of the protocols signed not only between Portugal and the Peoples' Republic of China but also between the Economic Community and Macao. Professor José Manuel Tribolet, the Director of INESC has taken part in two missions to China, Macao and Hong Kong and was present at a mission to China from 1-14 May.

The mission was promoted and supported by the European Communities DG XIII via the Agreement for Trade

and Cooperation signed between the European Union and Macao. Contacts have been established with universities, companies, government bodies and institutions operating in the telecommunication and information technology sectors. The basic aim of the mission was to define and explore ways and methods of cooperation with China and Macao within the framework of the existing cooperation agreements between these States and the European Union.

The strategic guidelines for cooperation between INESC, Macao and China encompass the following basic objectives:

- to promote national expertise and technologies developed at INESC and similar institutions;
- to further activities with a scientific and technological base of economic interest at the institutional and individual partners from China, Macao and Portugal, involved.
- to internationalise marketable products with possible application in the Macao Chinese marketplace (communications systems, control systems etc...)
- the involvement of INESC in the computerisation and modernisation of the organisational structure of public administration in Macao.

In the field of support for the training of human resources, INESC has selected six recent graduate scholarship winners from the University of Macao, giving them practical training, including aspects involved with the day-to-day running of a Research and Development Institute. This project was financed by various institutions including the European Communities Commission, JNICT (National Scientific and Technological Research Board) and the Macao Foundation.

INESC is also taking part in a Portuguese – Chinese – English automatic translation project in cooperation with the University of Macao and Chinese Universities and was represented at a meeting organised by the University of Macao which took place 16-18 March this year.

Please contact:

Isabel Caetano – INESC
Tel.: +351 1 1 3100367
E-mail: isc@eniac.inesc.pt

CALL FOR PARTICIPATION

WTC '94**Cernobbio, Italy, 5-7 September 1994**

The World Transputer Congress '94 is being held in the Villa Erba, Cernobbio, Lake Como, Italy from 5 to 7 September 1994. The wide ranging Tutorial Programme will be held on the 3 and 4 September at the same location.

The Conference, Tutorials and Exhibition will offer delegates an opportunity to listen to and interact with leading individuals in the field of transputer-based parallel processing and see the latest parallel processing products, not exclusively transputer-based from the leading suppliers.

The Keynote and Invited Speakers this year include Professor David May FRS (Inmos), Professor G S Stiles (Utah State University), Professor Marzano (University of Rome) and Professor S Noguchi (Tohoku University, Japan). 66 submitted papers from authors worldwide, covering applications, systems and theory, will be presented in a number of parallel streams.

The event is sponsored by The Transputer Consortium, SGS-Thomson, the Commission of the European Union (CEU) and the Italian Transputer User Group (ItTUG).

The CEU sponsorship is in the form of an award through the Human Capital and Mobility (Euroconferences) Programme. This award will allow us to support the attendance of some 40 young European researchers at WTC '94.

Please note there are significant discounts available on delegate fees for both the Conference and Tutorials if you register by 7 July.

To obtain Draft Conference and Tutorial Programmes, Registration and Accommodation Booking Forms, and further information on WTC '94 :

Please contact:
Mike Jane, TTC Coordinator, - DRAL
Tel: 0235 445408
Fax: 0235 445893
E-mail: mrj@inf.rl.ac.uk

CALL FOR PAPERS

WTC '95**Harrogate International Centre, England, 4-6 September, 1995**

The Transputer Consortium (TTC) is pleased to announce that the WORLD TRANSPUTER CONGRESS 1995 (WTC '95) will be held on 4 - 6 September 1995 at the Harrogate International Centre, West Yorkshire, England.

WTC '95 is also sponsored by the Commission of European Union.

The local partner for WTC '95 is the World occam and Transputer User Group (WoTUG).

Exhibition

Companies and other organisations wishing to exhibit transputer-based and other relevant products at WETC 95 should contact us at the address given below. There is also space for posters on research projects.

Instructions to Authors

Four copies of submitted papers (not exceeding 16 pages, single-spaced, A4 or US 'letter') must be submitted by 1 March 1995. Authors will be notified of acceptance by 24 May 1995. Camera-ready copy must be delivered by 23 June 1995, to ensure inclusion in the proceedings.

Call for Tutorials and Workshops

Tutorials will be held on 2-3 September 1995 in the Royal Baths Assembly Rooms. Suggestions or proposals must be submitted by 1 March 1995.

Registration of Interest

To register your interest in WTC '95, please provide your name, address, e-mail, tel, fax, details and please indicate what you wish to receive:

All information on WTC '95 Submitting a paper to WTC '95 Attending WTC '95

Exhibiting at WTC '95 Submitting a Tutorial/Workshop Proposal.

Please contact:
Susan C Hilton - DRAL
Tel: + 44 235 6154
Fax: +44 235 44 5893
E-mail: sch@inf.rl.ac.uk

CALL FOR PAPERS

IMACS - SAS '95

5th International Symposium on System Analysis and Simulation

Berlin, Germany, 26-30 June, 1995

In the tradition of the past four conferences held in 1982, 1985, 1988 and 1992, the objective of the symposium is to provide a forum for exchange of research results in the area of computer simulation, especially the mathematical and computational aspects of systems analysis as well as a platform for discussions of fruitful applications. The conference in 1995, sponsored by the International Association for Mathematics and Computers in Simulation (IMACS), will be organized by the GMD-Research Institute for Computer Architecture and Software Technology, Berlin.

In the last decade computer simulation became a key technology in industry, economics, and natural sciences. The request to manage and control difficult processes has forced the development of very complex dynamic systems and extended simulation tools. Now the increasing availability of parallel computer systems offers a new dimension for real-time simulation of such complex dynamic systems. Beside papers that deal with aspects of parallel processing, advanced

tools and applications in important fields of systems analysis and simulation are of particular interest.

Computational and Mathematical Aspects of High Performance Scientific Computing (HPSC) and Parallel Applications

- Numerical Algorithms and Approximations
- Computer Methods for Partial Differential Equations
- Computer Methods for Ordinary Differential Equations
- Parallel and Supercomputing
- Real World Computing
- Virtual Reality

Analysis of Complex Dynamic and Nonlinear Systems

- Systems Analysis and Mathematical Modelling
- Continuous and Discrete Systems
- Control Theory and Robotics
- Fuzzy Systems
- Large Scale Control Systems
- Stochastic Systems
- Optimization and Decision Support
- Neural Nets
- Petri Net Simulation
- Bond Graph Modelling

Challenges and Applications in Industry, Natural Sciences, Economics, Medicine, Informatics, Engineering and Environment

- Management and Production Systems
- Computer Integrated Manufacturing
- Environment and Ecology
- Circuit Simulation
- Semiconductor Devices
- Electrical Machines
- Traffic and Vehicles
- Computational Fluid Dynamics
- Socio-Economic Systems
- Biological and Medical Systems

Simulation and Programming Tools

- Hardware and Software Environments
- Scientific Visualization
- Simulation Languages

Deadlines:

- Abstracts (3 copies, 1-2 pages) before December 15, 1994
- Notification of acceptance before February 1, 1995
- Camera-ready paper (2 copies, 4-6 pages) before April 1, 1995

Please contact:

Achim Sydow – GMD
Tel.: +49 30 6392 1814/1800
E-mail: sas95@first.gmd.de

CALL FOR PAPERS

ISADS '95 2nd International Symposium on Autonomous Decentralized Systems

**Phoenix, Arizona, USA,
 25-27 April, 1995**

Due to the continuing growth in the power, intelligence and openness of computer, communication and control technologies, highly efficient and dependable business and control systems are realized. Dynamically changing social and economic situations demand next-generation systems based on emerging technologies and applications. Such systems are expected to have the characteristics of living systems composed of largely autonomous and decentralized components. Such systems are called Autonomous Decentralized Systems (ADS).

After the successful first International Symposium on Autonomous Decentralized Systems held in 1993 in Japan, the second ISADS will be held in Phoenix, Arizona, USA on April 25-27, 1995. ISADS invites papers and panel proposals that will foster increased interactions among researchers and practitioners in computer, communication, control and

other related fields from academia, industry and government.

The scope of discussions on ADS shall include, but not be limited to:

- Distributed and parallel computer architectures for ADS
- Computer-supported cooperative work
- Intelligent distributed computing/intelligent networks
- Heterogeneous distributed information systems
- Software development and maintenance for ADS
- Fault-tolerance, on-line expansion and maintenance for ADS
- Component integration in ADS
- Interactive communication systems/multimedia systems.

Applications of ADS:

- Manufacturing systems
- Office automation
- Robotic systems
- Transportation (air/land/sea) systems
- Intelligent Buildings
- Home information/entertainment systems
- National health care systems;
- Intelligent vehicle/highway systems.

Deadlines:

- 1. August 1994: All paper and panel proposals due.
- 5. September 1994: Panel organizers notified of acceptance.
- 27 September 1994: Organizers of accepted panel proposals provide final information on session chairperson and panelists.
- October 23, 1994: Authors notified of acceptance.
- December 29, 1994: Camera-ready copies of accepted papers and panelists' position papers due.

Please contact:

Radu Popescu-Zeletin
Tel.: +49 30 25499 206
E-mail: zeletin@fokus.gmd.d400.de
or Joseph Urban - Arizona State Univ.
Tel.: +1 602 965 2774
E-mail: jurban@asuvox.eas.asu.edu

CALL FOR PARTICIPATION

ESA '94 – 2nd Annual European Symposium on Algorithms

Utrecht, The Netherlands,
26 - 28 September 1994

The symposium covers all research on algorithms and their analysis relative to the fields of theoretical computer science, discrete applied mathematics, and all other areas of algorithm-oriented research and its applications.

Topics:

Papers in the following areas will be particularly discussed:

- Graph and Network Problems
- Computational Geometry
- Algebraic Problems and Symbolic Computation
- Combinatorial Optimization
- Neural and Genetic Computing
- Pattern Matching
- Cryptography.

Applications to concrete problems in science and industry, as well as the implementation of algorithms in real-world problems, will also be addressed. The programme committee is chaired by Jan van Leeuwen (Department of Computer Science, Utrecht University, e-mail: jan@cs.ruu.nl). The proceedings will appear in the series Lecture Notes in Computer Science of Springer-Verlag. ESA'94 is supported by the European Association for Theoretical Computer Science.

Please contact:

Marko de Groot - Utrecht University
Tel: +31 30 53 4095/1454
E-mail: marko@cs.ruu.nl

CALL FOR PAPERS

EPIA '95 – 7th Portuguese Conference on Artificial Intelligence

Funchal, Madeira Island,
Portugal, 3-6 October, 1995.

As in previous issues, EPIA'95 will be run as an international conference, English being the official language. The scientific program encompasses tutorials, invited lectures, demonstrations, and paper presentations. The conference is devoted to all areas of Artificial Intelligence and will cover both theoretical and foundational issues, and applications as well.

Parallel workshops on Expert Systems, Fuzzy Logic and Neural Networks, and Applications of A.I. to Robotics and Vision systems will run simultaneously.

The following researchers have already confirmed their participation, as guest speakers:

- Marvin Minsky (MIT, USA)
- Manuela Veloso (CMU, USA)
- Borges de Almeida (IST, Portugal)
- Rodney Brooks (MIT, USA)

Submission of Papers

Authors must submit five complete printed copies of their papers. They are requested to select 1-3 appropriate keywords from the list below: Applications, agent-oriented programming, automated reasoning, belief revision, case-based reasoning, common sense reasoning, constraint satisfaction, distributed AI, expert systems, genetic algorithms, knowledge representation, logic programming, machine learning, natural language understanding, nonmonotonic reasoning, planning, qualitative reasoning, real-time systems, robotics, spatial reasoning, theorem proving, theory of computation, tutoring systems.

Deadlines

- 20 March 1995: Submission of Papers
- 15 May 1995: Notification of Acceptance
- 12 June 1995 Papers returned

Publication

The proceedings will be published by Springer-Verlag (Lecture Notes in A.I. series). Authors will be required to transfer copyright of their papers to Springer-Verlag.

Parallel Workshops

In the framework of the conference three workshops will be organized: Applications of Expert Systems, Fuzzy Logic and Neural Networks in Engineering, and Applications of Artificial Intelligence to Robotics and Vision Systems. Real word applications, running systems, and demos are welcome. ■

Please contact:

EPIA95 – INESC
Fax: 351-1-525843
E-mail: epia95@inesc.pt

Eurographics '95

Maastricht, The Netherlands,
28 August-1 September 1995

The annual Eurographics Conference will have its 1995 edition in Maastricht, The Netherlands. The event will be organized by CWI and will be held in Maastricht's conference center MECC.

The conference will cover all research in the area of computer graphics, multimedia and virtual reality. Apart from paper sessions, EG'95 will feature state-of-the-art sessions, tutorial lectures, a slide/video/film competition and an exhibition. A call for papers will be issued this summer. ■

Please contact:

Eurographics '95 – CWI
Tel: +31 20 592 4171
E-mail: eg95@cwi.nl

CALL FOR PARTICIPATION

ECAI '94 – 11th European Conference on Artificial Intelligence

Amsterdam, The Netherlands
8-12 August 1994

The aim of the conference is to cover all aspects of AI research and to bring together basic and applied research. Of special interest this year are papers which address applied AI.

Main topics:

Applications, Automated Reasoning, Cognitive Modelling, Connectionism and PDP, Distributed AI, Enabling Technology & Systems, Integrated Systems, Knowledge Representation, Machine Learning, Natural Language; Philosophical Foundations, Planning, Scheduling and Actions; Reasoning about Physical Systems, Robotics; Social, Economic and other Implications; Standardisation, User Interfaces, Vision and Signal Understanding; Verification, Validation and Testing.

Full details of the technical sessions, invited and survey speakers can be obtained by e-mailing ecai94-programme@scs.leeds.ac.uk or by anonymous ftp from [agora.leeds.ac.uk](ftp://agora.leeds.ac.uk), directory: ECAI94. Extended workshop information can be obtained by anonymous FTP from [cs.vu.nl](ftp://cs.vu.nl), directory/pub/ecai94 or via electronic mail to ecai94-workshops@cs.vu.nl. Extended tutorial information can be obtained by anonymous FTP from [swi.psy.uva.nl](ftp://swi.psy.uva.nl), directory pub/ecai'94.

Please contact:

Mirjam de Leeuw – Erasmus University
Tel: +31 10 4082302,
Fax: +31 10 4530784,
E-mail: M.M.deLeeuw@apv.oos.eur.nl

CALL FOR PAPERS

RCS '95 – 4th International Workshop on Responsive Computer Systems

Berlin, Germany, 29-30 March
1995

The disciplines of parallel, distributed computing, fault-tolerant computing and real-time systems still are, to a large extent, evolving separately. As computer users demand timeliness and dependability to be ensured when using arbitrarily complex system architectures, there is an urgent need to integrate theory and practice of real-time systems, fault-tolerant computing and parallel, distributed computing. We call such integrated systems responsive, as they have to respond to internal guiding programs or external inputs in a timely, dependable and predictable manner. Responsive systems raise a significant number of fundamental issues of theoretical and practical relevance.

RCS'95 will focus on concepts, methods, algorithms and tools for such systems. RCS'95 will pay particular attention to new results in the areas of algorithmic proofs and formal approaches addressing issues falling in at least two of the three disciplines.

Papers that are purely descriptive are discouraged. We view RCS'95 as an opportunity for participants to gain a clear picture of most recent original results in "difficult" research areas. For example, how to express finite upper bounds on response times in the absence of advance knowledge regarding future operational conditions, is the synchronous model appropriate to solve the 10*-9/hour problem, what would be a complete design

methodology that permits to formally PROVE that a pair <design+implementation> satisfies user requirements specifications in the case of responsive systems?

Topics:

- Formal methods for specification, design and verification of responsive systems
- Proofs of logical or/and timing properties
- Algorithms and communication protocols for achieving time-dependent task scheduling, concurrency control and fault tolerance in an integrated manner
- Language and tool support for responsive systems
- Modeling, analysis, performance evaluation
- Application to real systems

Participation

Active researchers in the areas of distributed and parallel systems, fault-tolerant computing and real-time systems are invited to participate. To maximize vigorous exchange of ideas, the workshop attendance will be limited. All prospective attendees who wish to submit a paper are invited to send five copies of a full manuscript not exceeding 5000 words to the Program Chairman by 21 October, 1994. Preprints of full papers will be distributed at the workshop. A digest of papers will be published.

Deadlines

- Submission of papers: 21 October, 1994
- Notification of Acceptance: 20 January, 1995
- Camera-ready Copy: 28 February, 1995

Please contact:

Gerard Le Lann – INRIA
Tel: + 33 1 39 63 53 64
Fax: + 33 1 39 63 53 30
E-mail: Gerard.Le_Lann@inria.fr

RAL and Daresbury Laboratory merge to form DRAL - In May 1993 the UK Government published its Science and Technology White Paper 'Realising our Potential'. This signalled a major change in the structure and organisation of British science, including a restructuring of the Research Councils. On 31 March 1994 the Science and Engineering Research Council (SERC) ceased to exist and responsibility for its programmes was shared between two new bodies, the Engineering and Physical Sciences Research Council (EPSRC) and the Particle Physics and Astronomy Research Council (PPARC), and two that continue, the Natural Environment (NERC) and the Biotechnology and Biological Sciences (BBSRC) Research Councils. - DRAL has become part of DRAL (Daresbury and Rutherford Appleton Laboratories) by merging with its sister laboratory in the North of England, the Daresbury Laboratory. This combined operation is expected to become more independent from any one research council in due course but is temporarily under the aegis of the EPSRC.



Prof. Dr.-Ing. Wolfgang K. Giloi
(Photo München)

GMD - Prof. Dr.-Ing. Wolfgang K. Giloi, Professor of computer science at the Technical University Berlin and director of GMD's Institute for Computer Systems and Software Technology has been elected to the Berlin-Brandenburg Academy of Sciences as regular member. In summer 1992, the Academy was founded as successor of the former Prussian Academy of Sciences.

CWI - Dr. ir. Gerard van Oortmerssen, previously managing director, has been appointed general director of CWI from the 1st of May 1994. CWI's scientific director, prof.dr. Cor Baayen, who served as ERCIM President from April 1991 until May 1994, will take early retirement from the 1st of April 1995.

CNR - Prof. Franco Denoth (IEI-CNR) has been elected as President of the Committee for Information Sciences and Technology of CNR. The Committee is responsible for approving and supervising CNR supported research and technological activities in the IT area, including many of the projects under way in the universities. Most of these projects are conducted in collaboration with industrial partners. Among the first tasks of the new Committee will be the improvement of the national scientific networking facilities and the alignment of Italian research activities with the objectives of the Fourth Framework Programme.

INRIA has signed in April a **framework agreement with Thomson-CSF**. Numerous collaborations are at work presently and this agreement is intended to coordinate the latter. The fields are cognitive sciences and image processing, distributed architectures, numerical computing and signal processing. New fields are planned to be explored in the near future.


GMD - Horst Schwichtenberg, researcher at GMD's Institute for Algorithms and Scientific Computing, has been elected to the "Quality Team of October 1993" together with members of IBM Germany and IBM Kingston. They have got the IBM award for the successful first installation of the "IBM Scalable PowerParallel System 9076-SP1" in GMD. According to IBM "this project is the first successful Europe-wide implementation of parallel computing in a complex environment."

INRIA - Prof. Alain Bensoussan has been appointed general director of INRIA for a new three-year period. Alain Bensoussan is directing INRIA since 1984.

CNR - Italy's access to EUROPANET will be managed in Pisa by CNUCE-CNR. It will provide a 2Mbit IP and a 64K X.25 international access to the Italian research network (GARR). CNUCE-CNR was the Italian participant in the EUREKA 8 COSINE (Cooperation for OSI Networking in Europe) project, which finished late 1993. COSINE developed the EUROPANET infrastructure and the DANTE service management system in order to provide data communication services to users from academic, governmental and industrial research organisations throughout Europe.


INRIA - To experiment interoperability of various information systems architectures directed towards end-users, INRIA has launched in June 1994 together with the "Banque Nationale de Paris" (BNP), "Electricité de France" (EDF) and France Télécom/CNET (Centre National d'Etudes en Télécommunication) the three year initiative "PREVISIA". The first experiments will address EDI, X500 home servers and distributed transactions based on OSI-TP.

GMD - Dr. Gerhard Paaß, researcher at GMD's Institute for Applied Information Technology, has been appointed adjunct professor by the Queensland University of Technology in Brisbane, Australia. As adjunct professor, Gerhard Paaß will give lectures on neurocomputing at the University's computers science department. His first stay will begin in August 1994.



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: Peter Kunz	Local Editors: Werner Behrendt (RAL) Lars Bergman (SISU) João Bilhim (INESC) Truls Gjestland (SINTEF) Bernard Hidoine Päivi Johansson (VTT) Siegfried Münch (GMD) Henk Nieland (CWI) Carol Peters (CNR) Constantine Stephanidis (FORTH) Miguel Delgado (AEDIMA)	E-mail: kunz@nuri.inria.fr wb@inf.rl.ac.uk lars@sisu.se jafb@eniac.inesc.pt truls.gjestland@delab.sintef.no bernard.hidoine@inria.fr paivi.johansson@vtt.fi siegfried.muench@gmd.de henkn@cw.nl carol@vm.iei.pi.cnr.it cs@csi.forth.gr mdelgado@ugr.es	Telephone: +33 1 39 63 50 40 +44 235 44 5281 +46 8 752 1613 +351 1 572 895 +47 73 59 26 45 +33 1 39 63 5484 +358 0 456 5936 +49 2241 14 2303 +31 20 592 4092 +39 50 593 429 +30 81 39 17 41 +34 58 24 40 94
Lay-out Peter Kunz				

AEDIMA				
Asociacion Española de Informatica y Matematica Aplicada UPC Av. Gregorio Marañon, s/n E-08028 Barcelona	Centrum voor Wiskunde en Informatica Kruislaan 413 NL-1098 SJ Amsterdam	Consiglio Nazionale delle Ricerche IEI-CNR Via S. Maria, 46 I-56126 Pisa	Foundation of Research and Technology – Hellas Institute of Computer Science P.O. Box 1385 GR-71110 Heraklion, Crete	Gesellschaft für Mathematik und Datenverarbeitung mbH 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 Fax: +31 20 592 4199	Tel: +39 50 593 433 Fax: +39 50 554 342	Tel: +30 81 39 16 00 Fax: +30 81 39 16 01	Tel: +49 2241 14 0 Fax: +49 2241 14 2889

					
Institut National de Recherche en Informatique et en Automatique B.P. 105 F-78153 Le Chesnay	Instituto de Engenharia de Sistemas e Computadores Rua Alves Redol 9 Apartado 13069 P-1000 Lisboa	Rutherford Appleton Laboratory Chilton DIDCOT GB-Oxon OX11 0QX	Swedish Institute of Computer Science Box 1263 S-164 28 Kista	Stiftelsen for Industriell og Teknisk Forskning ved Norges Tekniske Høgskole SINTEF DELAB N-7034 Trondheim	Technical Research Centre of Finland VTT Information Technology P.O. Box 1200 FIN-02044 VTT
Tel: +33 1 39 63 5511 Fax: +33 1 39 63 5330	Tel: +351 1 310 0217 Fax: +351 1 572 897	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	Tel: +358 0 456 6041 Fax: +358 0 456 6027