

# ERCIM NEWS

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

Number 37

April 1999

## FRONT PAGE



**Christian Huitema,**  
chief scientist at the  
Telcordia (former  
Bellcore) Internet  
Architecture

Research Laboratory, head of the  
RODEO research team at INRIA  
from 1986 to 1996:  
"Next Generation Networks"

### SPECIAL:

Networking Technologies 8

## C O N T E N T S

|  |    |
|--|----|
| Joint ERCIM Actions  | 2  |
| The European Scene:<br>The Social Impact<br>of Networking Technologies | 8  |
| Special Theme:<br>Networking Technologies                              | 12 |
| Research and Development   | 38 |
| Technology Transfer  | 43 |
| Events   | 50 |
| In Brief   | 51 |

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



ERCIM will  
celebrate its 10th  
anniversary with  
a two days event  
in Amsterdam,  
4-5 November  
1999. See  
announcement  
on page 6.

### Next Issue:

Financial Mathematics

The last 15 years saw the development of the Internet, and more generally of data services. I usually try to explain this development by quoting two strong forces, Moore's law and Metcalfe's law. As predicted by Moore's law, the power of microprocessors has been regularly doubling about every 18 months for the past 25 years, the trend continues, computers are becoming much more powerful and affordable, and will pervade many more places. The networking effect often referred to as Metcalfe's law states that the value of a computer network grows as the square of the number of connected elements. This law resonates with one of the principles of the Internet, the trend to connect all computers on earth to a single network, and may well largely explain the success of the Net. In any case, we observe that data transmission already accounts for more volume than voice transmission, and grows much faster.

Today's telecommunication network have been designed to transmit voice, with data as an add-on service. But, at the current growth rate, in a few years, voice will only represent a small fraction of the global traffic. It will then not make sense to maintain a voice network, and we will see an inversion of the current situation: voice will run as an application within network primarily designed for data.

Next Generation Networks represent the convergence of multiple independent networks including voice, video and data into a single, unified, broadband network.

The broadband network is built using optical fibers, waveband digital multiplexing, ATM switches or IP routers. Cable modems or digital subscriber loops can provide high speed access to home. Voice gateways can be integrated in these modems. They provide a classic telephony interface to the residential user, but send the voice signals as packet over the data infrastructure. Network-based servers, the Call Agents, manage the establishment of phone calls. Other servers will manage the gateways to the existing phone network, including the translation between ATM or IP call set-up and the native signalling procedures of the phone network (SS7). Everything is designed to be 'transparent', at least in the first phase: the users will keep the same telephone handset, dial the same numbers, benefit from the same services such as call waiting or caller identification. There are already several such networks being planned, or built, in the United States and in Canada.

The deployment is motivated both by short term economies and by long term strategies. First, next generation networks can be deployed for a fraction of the cost of classic networks built of telephony switches and digital multiplexes. New entrants such as Level3 Communications or Qwest can thus compete without having to replicate the huge investments of the existing carriers. But, more important, the next generation networks are designed for the future. They will not only enable telephony and high speed access to the Web, e-mail and electronic commerce, but they will allow a next generation of applications.

My team at Bellcore, now Telcordia Technologies, has been working on the architecture of these new networks, designing for example the 'Media Gateway Control Protocol' that will be used by the Call Agents. But this is only the beginning of a new adventure. There are many new services that the ERCIM laboratories will help us invent!

**Christian Huitema**

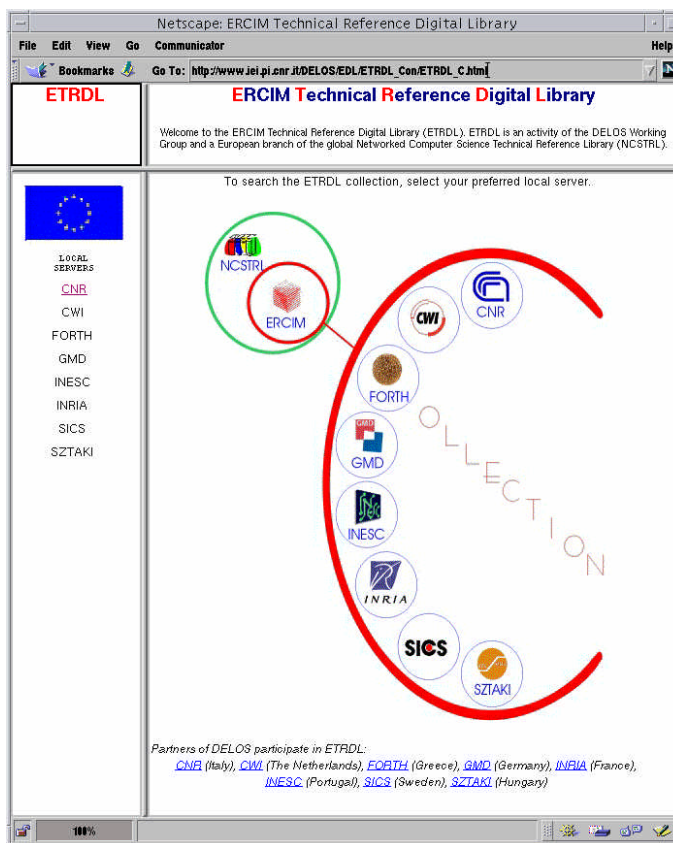
# The ERCIM Technical Reference Digital Library goes On-line

by Antonella Andreoni, Carlo Carlesi and Pasquale Pagano

One of the declared aims of ERCIM is to advance European research and development in the areas of information technology and applied mathematics. It is thus extremely important that results of the R&D activities carried out at the ERCIM member institutions can be made immediately available to the interested scientific community in an easily accessible form. For this reason, ERCIM has set up a digital library service for the technical reports and other similar reference material produced by its scientists: the ERCIM Technical Reference Digital Library (ETRD). An experimental service was first set up in 1998 – it is now publicly available.

The ERCIM Technical Reference Digital Library (ETRD) is a collaborative activity in which eight organisations (CNR, CWI, GMD, INRIA, SICS, INESC, SZTAKI, FORTH) currently participate. It is sponsored jointly by ERCIM, by the DELOS Working Group (ESPRIT LTR No. 21057), and by the participating institutions. The objectives of the digital library service were outlined in Ercim News, No.33, April 1998. We are happy to announce here that, after a one-year period of testing and refining the first prototype, ETRDL is now publicly accessible at <http://www.iei.pi.cnr.it/DELOS/ETRD>.

The service runs on an infrastructure based on the DIENST system (Version 4.1.9) developed by the US Networked Computer Science Technical Reference Library (NCSTRL) Consortium, which has been extended by the addition of new functionalities to meet the needs of the



The centralised Home Page.

European IT scientific community. These include the adoption of a common metadata description standard, the introduction of common classification schemes and methods to manage them, the implementation of interfaces for languages other than English. An author submission form has been included to facilitate the insertion of new documents by the users themselves.

ETRD is a distributed collection, consisting of the set of the local collections. These are maintained on the local servers of each partner institution. This has comported the implementation of two levels of Home Pages. A centralised access point has been provided to the system through the DELOS Web site, whereas a local home page is installed on each local server. The user who accesses the system through the centralised Home page is directed to select one of the local servers in order to search the ERCIM DL collection. Depending on the local server selected, he/she is also given a choice of language as most local servers will maintain interfaces in English and in the local language.

The Local Home Page interface caters simultaneously for two user classes:

information seekers and information providers by offering two main options: search/browse any collection; submit/withdraw a document to/from a local collection. From the local home pages, the search and browse functions can be activated over the entire NCSTRL collection, over the ERCIM collection, or over the collection(s) of the local institution. In each case, the user is not only accessing a different collection (or sub-collection), but is provided with a different perspective on the information, depending on the functions that have been implemented at that particular level. When searching on the ERCIM or the local collections, the user can switch between user interfaces in English or his/own language. On-line helps in both languages are available.

Adapting an existing system to meet new requirements is rarely easy. We had to address a number of problems when implementing the ETRDL. These included the implementation of a powerful search mechanism which allows the user to perform a simple or a fielded search throughout the entire library or on selected collections. Queries can include terms from controlled vocabularies (ACM and MSC

## ETRD Search form.

classification schemes) and free keywords. Constraints can be imposed on the type, date or language of the documents to be searched.

ETRD is now active. The next step will be to build up the local collections. All new technical documentation produced by the eight ERCIM institutions mentioned above (but hopefully others will join the activity soon) should be inserted in the library. Institutions may also consider including their most important past technical reports (eg from 1990 on); this implies classifying them according to the ETRDL rules.

In the mean-time, we are working on the development of the next version of the system (which will include modules to handle multilingual issues such as mechanisms for the correct search and visualization of all character sets, the maintenance of separate and linked indexes to permit the definition of the search language, semi-automatic procedures to assist authors in document classification, etc.).

In this stage, the feed-back we receive from users will be very important. So

please, try out the system, and send us your comments.

For more information on the ETRDL project and documentation on the system, see <http://www.iei.pi.cnr.it/DELOS/EDL/edl.htm>

### Please contact:

**Stefania Biagioni – IEI-CNR**  
**ETRD Co-ordinator**  
**Tel: +39 50 593481**  
**E-mail: [biagioni@iei.pi.cnr.it](mailto:biagioni@iei.pi.cnr.it)**

# EU-China High Performance Computing and Networking Initiative

by Eckart Bierdümpel  
and Thierry Priol

Following a request by the European Commission, ERCIM has set up an 'EU-China High Performance Computing and Networking (HPCN) Initiative' to foster business, industrial and technological co-operation between European and Chinese operators in this high-tech sector. HPCN is a crucial technology of strategic importance for European industrial and scientific competitiveness. It has the potential for great and far-reaching impact on productivity, industrial competitiveness, environmental management and many other aspects of economy and society. The use of HPCN systems for commercial applications is expected to be taken up vigorously in the second half of the decade, both in Europe and in emerging economies, such as China.

The first meeting of a Chinese Delegation with European experts and officials from the European Commission was held in Brussels. The meeting was followed by a familiarisation tour for the Chinese in Europe, organised by ERCIM and included visits at companies in Belgium, France, Germany, England and Italy. Two European workshops were organised by ERCIM in Pisa and Abingdon. There the organisers presented the initiative, while the Chinese partners gave an overview on HPCN activities and the potential of IT in general in China. As an example, the experience of setting up a high-tech joint venture between a European and a Chinese partner was presented, followed by a lively discussion on the obstacles and problems. Also the chances of setting

up such a joint venture and the initiative in general was discussed.

It was concluded that considerable interest was shown by the participating organisations in a co-operation with Chinese companies and institutions. Furthermore there was a noticeable interest by other companies who couldn't make it to the meetings. As a result of the European workshops, 75 companies registered as potentially interested in the initiative.

While ERCIM was waiting for the Chinese to organise a similar workshop in China, a preliminary report was released on the situation of HPCN in China. The report was based on the information gathered from Chinese web sites and newsletters; from people who visited China or already had experience in co-operations with Chinese institutions and also with some input from China itself.

The Chinese workshop was held in Beijing in October 1998. It started with presentations of the participating European companies, the Technology Transfer Network (TTN), a European



**Discussions between Chinese participants and representatives of European companies during the workshop in Beijing.**

initiative for technology transfer in the HPCN sector funded by the Esprit Programme of the European Union, and the ten selected Chinese institutions. After the presentations, the Chinese participants could discuss issues directly with the European representatives. An immediate positive result from this initiative was the contact established between different companies and institutions from Europe and China.

The Chinese side stated that the workshop was a useful experience and

that they found the TTN scheme very interesting. They were missing this kind of action in the part they learned from American and Japanese counterparts. At that time, China had established several HPC centres which mainly host HPCN



**During the workshop in Beijing.**

systems neglecting information dissemination and technology transfer. Therefore they were interested in setting up a remote node of the European TTN in China. In addition to the European TTN nodes, they would like to offer



access to HPCN platforms to Chinese end-users. Such platforms should be created from Chinese hardware with demonstration versions of European software including documentation and training. Such a node could serve as a European window in China to promote European technology and know-how plus enhanced co-operation between China and Europe. Another goal of this Chinese TTN node could be the summarised translation of documents from the European TTN in Chinese for dissemination in China. So the Chinese

participants said that they would recommend to the government to set up a similar scheme in China. The Chinese participants also showed interest in business contacts with the European companies represented at the meeting.

These companies were also quite satisfied with the feedback they got from the Chinese hosts.

In conclusion, the 4 goals of this initiative were mostly fulfilled:

- raising the awareness of the available technology and expertise for companies with potential interest in HPCN in China and in Europe
- informing the business community in this field about the market opportunities in China and Europe
- offering a business forum, where direct contacts between potential partners from Europe and China can be made
- trying to establish more co-operation between European and Chinese companies.

**Please contact:**

**Eckart Bierdümpel – GMD**  
**Tel: +49 2241 14 2256**  
**E-mail: eckart.bierduempel@gmd.de**

**Thierry Priol – INRIA**  
**Tel: +33 2 99 84 72 10**  
**E-mail: Thierry.Priol@irisa.fr**

## Conferences and Workshops sponsored by ERCIM

ERCIM sponsors up to five high quality internationally recognised conferences and six workshops or summer schools per year. The funding for both types of events is in the order of 2000 $\text{\textsterling}$ . In 1999, the bids of two conferences and three workshops were accepted.

Sponsorship guidelines for each category of events:

### Conferences

ERCIM invites sponsorship proposals from established conferences with international reputation with substantive overlap between the conference topic and ERCIM areas of activity. Typical cases would include annual conferences in Computer Science with international Programme Committees, substantial international participation, and proceedings published with established international scientific publishers. Conferences may receive funding under the following conditions:

- the ERCIM sponsorship has to be acknowledged in all printed PR material
- ERCIM information and promotion materials has to be distributed to the conference participants as a part of the participant's conference package
- an ERCIM information booth has to be arranged at the conference location
- conference organisers to provide free attendance to one ERCIM representative if ERCIM sends a representative to be present at the information booth
- the conference organisers have to include the ERCIM logo with a link to the ERCIM web site at the conference home page.

### Workshops and Summer Schools

ERCIM sponsors up to six workshops or summer schools under the following conditions:

- to be held at an ERCIM site
- to have some ERCIM partners having named individuals involved in the organisation committee
- participation from at least 2 non-ERCIM institutes on the Organising Committee underwriting by one ERCIM partner

The additional funding provided by ERCIM should be used to enhance the workshop by, for example, increasing the set of external speakers supported. This is not seen as a revenue gaining exercise for ERCIM so that, in case such events generated profits, the profit would remain with the external organisation. Publication of the proceedings is left to the organising committee but ERCIM would be prepared to help if appropriate.

### Events sponsored in 1999

Conferences:

- ICALP '99, Prague, 11-15 July 1999
- SOFSEM '99, Milovy, Czech Republic, 27 November - 4 December 1999

Workshops and Summer Schools:

- Fourth International Workshop on Formal Methods for Industrial Critical Systems, Trento, Italy, 11-12 July 1999.
- Summer School on Extended Database Technology, La Baule, France, 16-21 May 1999
- IWAN '99 - First International Workshop on Active Networks, Berlin, Germany, 30 June-2 July 1999 (see announcement on page 50).

Please contact:

ERCIM Office  
Tel: +33 1 3963 5040  
E-mail: office@ercim.org

### CALL FOR CANDIDATES

## Cor Baayen Award 1999

The Cor Baayen Award was created in 1995 to honour the first ERCIM President. Any young researcher from one of the fourteen 'ERCIM countries' (Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Norway, Slovakia, Sweden, Switzerland, The Netherlands, UK) can be nominated. The award consists of 5000 $\text{\textsterling}$  together with an award certificate. The selected fellow will be invited to the ERCIM meetings in autumn.

### Rules for nomination

- Nominations for each country are made by the corresponding ERCIM Executive Committee member. Those who wish that a particular candidate is nominated should therefore contact the ERCIM Executive Committee member for their country (list with contacts for each country at: <http://www.ercim.org/contacts/execom/execom.html>).
- Nominees must have carried out their work in one of the fourteen 'ERCIM countries'.
- On the date of the nomination, nominees must have been awarded their PhD (or equivalent) in the last two years.
- Each institute is allowed to nominate up to two persons from its country.
- A person can only be nominated once
- The selection of the Cor Baayen Fellow is the responsibility of the ERCIM Executive Committee.

Previous Cor Baayen Award winners:

- 1995: Christer Samuelsson (SICS)
- 1996: Dimitri Papadias (GMD)
- 1997: Kristina Höök (SICS)
- 1998: Matthias Grossglauber (INRIA)

See also : <http://www.ercim.org/activity/cor-baayen.html>

Please contact:

Frans Snijders – CWI  
Tel: +31 20 592 4171  
E-mail: Frans.Snijders@cwi.nl

**CALL FOR PARTICIPATION**

**5th ERCIM Workshop on User Interfaces for All**

**Dagstuhl, Germany,  
28 November - 1 December 1999**

The 5th ERCIM Workshop on 'User Interfaces for All' builds upon the results of the four previous Workshops held at Heraklion, Greece, 30-31 October 1995; Prague, Czech Republic, 7-8 November 1996; Obernai, France, 3-4 November 1997; and Stockholm, Sweden 19-21 October 1998. The emphasis of this year's Workshop, to be held at Dagstuhl, Germany, 28 November - 1 December 1999 (local organisers Prof. Alfred Kobsa and Dr Michael Pieper), is on theories, methodologies, techniques and tools which contribute to the development of User-Tailored Information Environments.

The vision of User Interfaces for All advocates the proactive realisation of the

design for all principle in the field of Human-Computer Interaction, and involves the development of user interfaces to interactive applications and telematic services, which provide universal access and quality in use to potentially all users. This user population includes people with different cultural, educational, training and employment background, novice and experienced computer users, the very young and the elderly, and people with different types of disabilities, in various interaction contexts and scenarios of use.

**Important Dates**

- 30 June 1999: deadline for electronic submission of all papers
- 15 August 1999: conditional notification of acceptance by e-mail for all papers (confirmation will be given upon registration)
- 25 August 1999: deadline for registration

- 10 October 1999: deadline for electronic submission of camera-ready papers

For more information, see:  
<http://zeus.gmd.de/5-UI4ALL-Workshop/call.html>  
 or  
<http://www.ics.forth.gr/ercim-wg-ui4all/index.html>.

**Please contact:**

**Constantine Stephanidis – ICS-FORTH  
 Chair, ERCIM Working Group "User  
 Interfaces for All"**  
**Tel: +30 81 391741**  
**Fax: +30 81 391740**  
**E-mail: cs@ics.forth.gr**

**Michael Pieper – GMD-FIT  
 Local Organiser of the Workshop**  
**Tel: +49 2241 14 2018**  
**Fax: +49 2241 14 2167**  
**Email: michael.pieper@gmd.de**



**ERCIM 10th Anniversary Event**

**Amsterdam, 4-5 November 1999**

ERCIM will celebrate its 10th anniversary with a two days event in Amsterdam, 4-5 November 1999. The first day will be an internal event for ERCIM-member personnel only, while the second day will be dedicated to Information Technology users in industry.

**ERCIM – a virtual laboratory for IT research in Europe, Amsterdam, Thursday 4 November 1999**

Under this slogan scientists of the ERCIM institutes will be given the opportunity to present their ideas on matters that are closely related to IT research. It is not research itself that should be targeted with these presentations but rather the issues that come up on a meta-level. To give some examples of the topics that might be tackled, we mention: 'The future of traditional scientific journals in the Internet Age 'Should Computer Science and Mathematics be kept separate?', 'Is freeware (cooperative development of software) a viable alternative to the traditional model?', 'Why is software less

reliable than a washing machine?', 'Why Framework Programs don't work', 'Alternatives to traditional scientific output metrics', 'Are European CS-curricula too theoretical?' Note that this is not an exhaustive list but just an indication. The presented ideas should be provocative and express opinions of individuals, not necessarily shared by their home institutes. ERCIM scientists interested to contribute to this event are invited to contact their Executive Committee representative.

**ERCIM – leveraging European R&D for Business and Society, Amsterdam, Friday 5 November 1999**

With the co-operation of world renowned European scientists and industry leaders, the November 5 event aims at taking stock of information technology, its advancement and its applications in business and society. Leaders from the fields of ICT, manufacturing, and information content will give their vision of the future of European R&D. The event will be concluded with a panel discussion. Besides the management of ERCIM's R&D members, attendance will consist of managers from industry, governmental authorities and European bodies.

**Please contact:**  
**ERCIM office**  
**Tel +33 1 3963 5303**  
**E-mail: office@ercim.org**



# The Social Impact of Networking Technologies

by Truls Gjestland

New information and communication technologies provide new ways of acquiring information, new ways of conducting business and new ways of interacting with people and organizations. The technology is thus including, as it offers more people a possibility to play an active role in the society. However, this technology is also potentially dangerous, as it may be excluding. Our society may be divided in those who have access and can use the new technology, and those who cannot. These aspects are being

addressed in many countries, and several research projects have been launched to study the social impact of ICT. It is also a growing trend to stress the slogan: design for everybody, so that the concept is more than a dinner speech item, but has become a reality, and a requirement in many new research projects: If feasible any new product or service should be available for as many as possible regardless of for instance personal handicaps.

## Organising Work as Telework – New Possibilities also to Disabled People?

by Eystein Fossum

The Research Council of Norway has set up a research programme to demonstrate how information networks can be used in the development of Norway as an information society. The programme supports the building of networks between industrial and governmental actors in common value chains in such a way that new knowledge about branches, services, processes and products is gained. In collaboration with organizations in different business areas there are set up several projects to demonstrate the possibilities.

One of the projects or group of projects, in this programme has been the telework and telecommuting project. Until now the main goal of this has been to develop telework as a good way of organising work in the Norwegian society.

During this project an idea came up to try out how telework could help organising work in a good way for disabled and handicapped people. Two separate projects now follow up this idea. These projects are co-ordinated with the main telework project. The reason for this is to get the benefit from using the telework

expertise established and to use parallel models for organising the projects.

The first project will gain experience of using telework in rehabilitation of disabled people. The plan is to set up a 'telework room' in a rehabilitation hospital. The project will find out which activities and to which extent patients can benefit from a telework-like situation, how to organise the collaboration between the patients employer and the hospital and how to "transfer" the telework situation to other institutions or the patients' own home after the hospital period. The present project is restricted to patients with motor handicap, which have been employed in a job where they have used dp-systems before they became disabled. These restrictions are chosen to limit the scope of the project. In the long term it is our hope to transfer the experiences from the project patients with other handicaps and with other job background.

The second project will focus on how telework can give new job possibilities to people with mental or physical handicap. There are many reasons for which handicapped people often have job problems. Some handicaps make it very laborious for the person to get to an office every day, and therefore a full time job may be impossible in the long run. Other persons may have some mental handicap as for example not being able to perform a job in common working environment with others. These are only two examples of handicaps that may be obstacles for getting into the work market and where it might be that telework can be a part of a solution. The project will try out and

gain experience with how telework may be used as a better way to organise jobs for some groups. The project, which is in its initial phase, will be organised in close collaboration with the working market authorities.

Please contact:

Eystein Fossum – Semco AS

Tel: +47 67 10 19 61

E-mail: eystein.fossum@semco.no

## The Seniornett Project

by Sidsel Bjørneby

The SENIORNETT project was about user aspects of Internet that are success criteria or create problems in relation to human factors aspects of elderly people. The project ran from 20 October till 31 December 1998 and was carried out by Human Factors Solutions in Oslo, Norway for the organisation Seniornett Norge. The users involved in the research were elderly Internet users at a library in Oslo.

The research was done as a pilot for a main project in year 1999 and 2000, where user interface will be developed as an iterative process, and where parallel research will identify effects on the communication patterns of elderly people on the Internet.

The rapidly growing number of elderly people, and the interest they show in

using the Internet is the background for starting the organisation Seniornett Norge in 1997, with the aim of increasing the participation of elderly people in the information society. After some user trials the organisation wanted to create their own gateway to world wide web through a web site where elderly people could feel more confident than when entering regular web sites. Human Factors Solutions was asked to help them with this, and www.seniornett.no was established early in 1998. Contents are created and links mostly identified by a group of 18 active, resourceful elderly Internet users, who are divided into five groups according to topics.

The aims of project:

- Identification of aspects of Internet as a baseline for developing interfaces and contents which are adapted to the functional needs of elderly people.
- Identification of types of networks where elderly can and want to communicate on the Internet.
- Identification on relevant web information on Internet interface for elderly and disabled users.
- Start the design of a web site that can be a good example of web sites adapted to the functional problems of elderly people.

Methods:

- Structured discussions with elderly Internet users about their opinions and experiences.
- Questionnaire to elderly Internet users.
- Internet search for relevant web sites
- Contact with organisations of handicapped people to find potential similarities between their use of Internet as compared to the Internet use of elderly people
- Contact Nordic and European web sites for elderly people to identify interest in common future projects.

### Identified User Aspects and Attitudes

Age related changes of relevance to Internet are mostly within visual, sensory and cognitive abilities. These are not identified as handicaps by the users themselves, but rather as normal changes that affect their use of the Internet, and

that requires specific adaptations to the web sites designed for them. These aspects need to be taken into account when developing web sites accessible to elderly people.

The overall experience and attitude of elderly people to the Internet is positive. They enjoy being able to take part in modern aspects of life, and to have the opportunity to collect information and to communicate on the Internet and with e-mail. They demonstrate a curiosity and courage to take the Internet into use which far extends the expectations of society at large.

### Contact with Senior Web Sites for Nordic and EU Co-operation

Positive contact is established with senior web sites in Norway, Denmark, Sweden, the Netherlands and Italy. There is a Norwegian group of elderly in Spain that are interested in contact/ co-operation.

As a result of the contact with organisations for handicapped people, it was agreed that it is possible with synergy effects on many points of interface design, but that the actual sites are better kept separate.

### Future activities

A main project is planned as a result of this little project, in order to achieve the aims of creating user interface and contents adapted to elderly Internet users. Firstly the project will be started in Norway, in close cooperation with elderly Internet users, but the aim is to include the work in further research in a project under Creating a User Friendly Information Society within EU's Fifth Framework Programme. This project will include two parallel activities:

- an iterative development process regarding user interfaces of Internet for the elderly
- researching the effects of Internet communication on elderly people.

The communication effects will be studied within different groups of elderly (geographically, more and less resourceful, within interest areas), between generations and between elderly people and society.

The research is planned to be carried out by Human Factors Solutions, in cooperation with elderly users of Internet, the Faculty of Psychology at the University of Oslo, The Norwegian research centre NOVA and the potential European partners. ■

Please contact:

Sidsel Bjørneby – Human Factors Solutions

Tel: +47 22 42 18 57

E-mail: sidsel@hfs.no

## The Role of the Internet Society

by Klaus Birkenbihl

**Growth rates of more than 100% per annum since 15 years. Though: Who does take care of the Internet? In 1992, a group around the Internet inventor Vint Cerf established the Internet Society (ISOC). In an environment increasingly influenced by market and competition, ISOC secures the conditions for the continuance and cohesion of the Internet. Since 1995, the Deutsche Interessengemeinschaft Internet (DIGI e.V.), which was also established in 1992, has been a German section of ISOC.**

For the normal users, the Internet is simply there. They pay the required charges to T-Online, AOL or another service provider and hope that the provider will guarantee that it works. However, the provider - even if it is called AOL or Telekom - maintains and controls only a minor part of what makes up the Internet. And this is not only technology. The Internet is an important part of the infrastructure for the global information society. The Internet is an economic and political factor whose further development and regulation will influence the live of everyone decisively.

Hardly visible in the stir caused by the conflict among the Internet giants Microsoft, Netscape, Oracle, IBM, CISCO and others, the Internet Society has undertaken the task to shape and



advance the technical, political and cultural future of the Internet. The Internet Society regards itself as the international organization which furthers the global cooperation and coordination for the Internet, its technologies and applications. Everyone interested in the further development of the Internet can become a member of the Internet Society. Conferences and technical and political bodies give ample opportunity to exert influence on international and national level. Individual members who want to participate in shaping the medium's future are equally welcome as companies which use or offer the Internet as modern technology for their purposes.

### Technology and administration

The standards for communication in the Internet are defined in the Internet Engineering Task Force (IETF). More than 80 working groups in 8 fields of work are concerned with specific technical subjects. IETF does its work during three big meetings of more than 1500 participants per year and by means of all the communication means provided by the Internet.

On behalf of ISOC and the Federal Network Council (FNC) of the USA, the central resources of the Internet, eg names, addresses and protocol parameters have been managed by the Internet Assigned Numbers Authority (IANA) under the direction of Ian Postel, who tragically died last summer. Currently ISOC is intensively involved in restructuring the management of these resources within the new organisation ICANN (Internet Corporation for Assigned Names and Numbers).

### Communication

ISOC carries out its mission to propagate and advance the idea of the Internet on various conferences whose biggest - the INET-Conference - is held alternately on different continents in summer. Pioneering efforts are put forth by specific workshops on these conferences. The workshop for developing and threshold countries, for example, contributed decisively to the fact that there is hardly any country in the world not having

access to the Internet. Further key subjects of INET are issues of political regulation and social impacts of technology in addition to technical innovations, new applications and areas of application.

The Internet Society clearly gives its members' view on the cryptography discussion, on the problem of admitted contents and on the free and secure flow of information on the Internet.

### Internationalization

The statement that the Internet is primarily a US undertaking is certainly not completely incorrect. But it is changing. The Internet meanwhile shows greater growth rates outside the USA than within the USA. Within the Internet Society, regional "chapters" give the opportunity to exert influence on the development on regional level. Many of these groups are still in the process of organization so that their effectiveness will be visible only in some years. Recently, a number of the European chapters formed a group called ISOC-ECC (ISOC European Coordination Council) to coordinate their activities and to better safeguard their interests towards the European Union.

### DIGI e.V. and ISOC.DE

In Germany, where the Internet has been poorly organized and hardly accepted till the early 1990's for different reasons, in 1992, simultaneously with the Internet Society - the Deutsche Interessengemeinschaft Internet (DIGI e.V.) was established to improve the conditions for propagating the Internet in Germany. The establishment of a national Network Information Center (DENIC), for managing addresses and names on national level, is equally a goal of this group as the technical and political discussion about the Internet in Germany in particular.

In 1995, DIGI e.V. was recognized by ISOC as "German Chapter" of the Internet Society. Since then, DIGI e.V. uses the abbreviation and the Internet domain ISOC.DE e.V. Like the global ISOC, ISOC.DE gives its view on organizational, technical and political

issues pursuing the further development of the Internet.

ISOC.DE clearly states its point of view on subjects such as safeguarding the DENIC which manages the .de domain, the cryptography discussion in Germany, Internet and censorship and structural issues of the Internet in Germany. The annual conference, Opennet, to be held this year in Bad Honnef from 15-17 November 1999 will give ample opportunity for discussions and first-hand information.

### Representative?

With hardly 7000 members worldwide and 250 in Germany, ISOC is certainly still in the initial states of its work. Therefore, the claim to speak for the Internet community cannot be backed by the number of members. Today, it is rather the expertise concentrated in ISOC, the efficient way of agreeing on standards and the absolute independence of any business interests which lend weight to ISOC. But this is changing. More and more companies discover the Internet as an important business sector where showing technical and organizational competence might be an asset. And the intention of these companies to swamp the Internet with proprietary standards is not to be overlooked. The political pressure to curb the Internet through regulation is also increasing. The Internet Society is therefore confronted with the urgent task to recruit a solid basis of members. This is the only way to attain the goal to exert a substantial influence on the development of an open Internet.

### Links

ISOC: <http://www.isoc.org>  
 ISOC.DE: <http://www.isoc.de>  
 IETF: <http://www.ietf.org>  
 INET: <http://www.isoc.org/inet98/>  
 DENIC: <http://www.denic.de/>  
 OPENNET:  
<http://www.isoc.de/opennet97/opennet.html>  
 W3C and ISOC.DE Offices:  
<http://www.gmd.de/w3c/welcome.html>

### Please contact:

**Klaus Birkenbihl – GMD**  
**Tel: +49 2241 14 2910**  
**E-mail: [Klaus.Birkenbihl@gmd.de](mailto:Klaus.Birkenbihl@gmd.de)**

# Introducing Intermodality for Social Tools

by Kent Saxin Hammarström and Kristina Höök

**The Convene project at the Swedish Institute of Computer Science deals with improving the interface to text-based virtual environments (MUD's), used for recreational and social purposes, for users with disabilities (a group that constitutes 10-20% of the world population). These systems are potentially useful 'social tools' for those restricted from full participation in social relationships and leisure, but have received little attention in terms of how to make them more accessible.**

MUD's are a type of networked systems, primarily used for recreational purposes, such as games and social interaction, since the early 80's. What distinguishes them from other forms of computer mediated communication is that they involve real-time interaction between multiple users and between users and a virtual world.

MUD systems could be excellent social tools for users with certain forms of disabilities that would open up both an accessible form of recreation and a social venue with lower barriers to interacting with others.

The aim of the project is to see what problems – if any – users with disabilities experience in these environments and how these could be alleviated. Our two studies done so far (one with a group of users with muscle dystrophy and one of blind users) indicate that one of the major, software-related problems is not qualitatively different from that of any user: information overload. This manifests itself in problems with navigation in the MUD worlds, problems with keeping the pace of the quite fast interaction between users, and finally, problems with the actual contents and social behaviours in these environments. We would like to

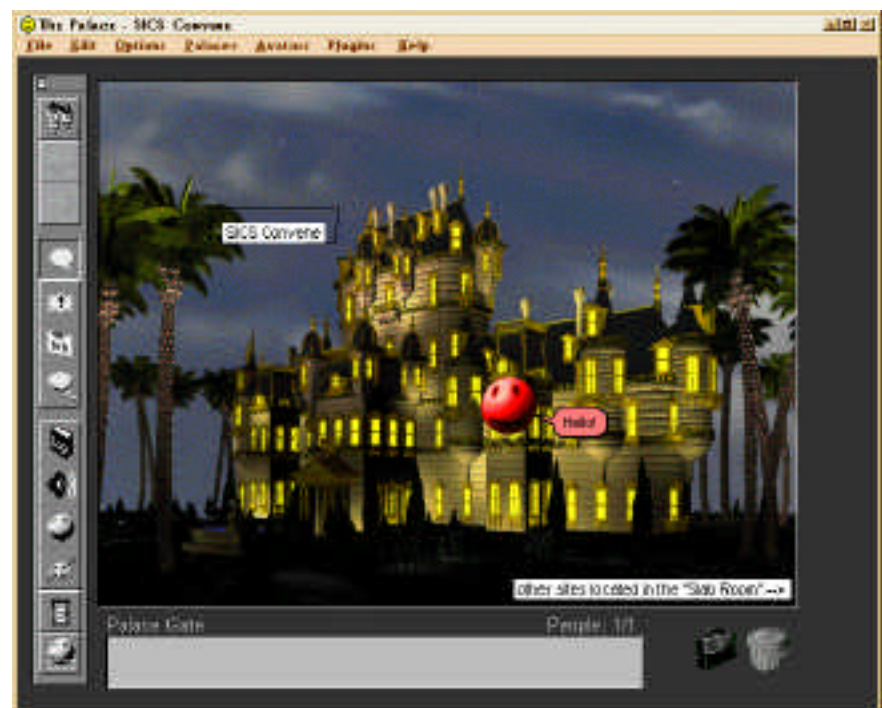
point out that we are all experiencing these problems to some degree; activity limitations (disabilities) simply exacerbate this.

We have chosen to use multimodality as the main thrust of how to improve these interfaces, in the form that we call 'intermodality', ie designing a modular software architecture that facilitates translating information from one representation (initially primarily textual) to one more suited to the user and the situation. We take textual input from the MUD world, such as, rooms descriptions, dialogues between users, events in the

handed back to the consumer. This approach has two facets:

- allowing compensation for functional impairments
- reducing the cognitive load by allowing the user to have information presented in complementary channels (eg utterances from other participants through speech synthesis and descriptions through a Braille display).

The project has a practical and problem-oriented, rather than technology-oriented, focus: Can we design a software architecture that allows creating programs



**Palace: A graphical MUD system.**

world, etc., and divide it into different channels: the room description channel, the dialogue channel, etc. The user can then choose which channel to present as audio, text, pictures, etc. For example, the room descriptions might be sent to the user as text, the dialogue with other users might be presented as speech output, and the structure of the world (which 'doors' a room contains) can be shown as a 'map'. Our fundamental vision is to solve the problem that it is the producer rather than the consumer determines the form of information - through dividing the information into different channels and allowing users to set these up as they please, the control is

that make existing MUD systems more accessible, without requiring additional, expensive equipment, using the resources the user would normally have at her disposal? The actual networking technologies used are thus the 'tried and true' in the form of existing MUD environments (predominantly using TCP/IP and Telnet over the Internet) rather than innovative ones, and the focus is on improving the software user interface to these technologies.

In summary, Convene aims to put the control over the form of information into the hands of the (disabled) users in order to enable them to access existing textual MUD environments.

This project is financed by the Swedish Institute of Computer Science (SICS), The Swedish Transport and Communications Research Board (KFB), and The Swedish National Board for Industrial and Technical Development (NUTEK). The Swedish Handicap Institute (HI) and the Swedish National Association for Disabled Children and Young People (RBU) contribute as advisors.

More information about the project can be found at:

<http://www.sics.se/~kia/convene/>

**Please contact:**

**Kent Saxin Hammarström – SICS**

**E-mail: [kent@sics.se](mailto:kent@sics.se)**

**Tel: +46 18 471 7041**

## Erasing Generation and Country Borders via Networking

by **Josef Koláři, Tomáš Macek,  
Bořena Mannová  
and Barbara Williams**

**Some IT courses delivered at the Department of Computer Science and Engineering, Czech Technical University (CTU) in Prague permit interesting educational experiments in which we aim both at technical and social aspects of teaching that uses modern Internet-based tools.**

The first experiment was conducted as a part of the Global Cooperation Project during the summer term of the 1997/98 school year with computer science classes at both the CTU and North Hennepin Community College, USA. Although the main aim both classes was programming, the students also obtained a basic instruction as how to use modern Internet-based communication tools. The project was conceived as a pilot project for teaching team cooperation on a global

scale. For three months the students worked in teams consisting of team members from remote places and different working and social background.

There were two videoconference sessions (at the beginning and at the end of the course) while most of the communication was done via e-mail, file transmission service, and WWW presentations. We tried to stress the social aspects of the projects so that the students were asked not to restrict their communication to professional issues. At the beginning, each student introduced him- or herself during a videoconference and a more complete presentation was done via personal WWW homepages they were asked to prepare.

For many participants the project was the first opportunity to have real working contact with a person from another country. They acquired new knowledge concerning educational systems in the two countries and learned a lot about differences in living styles. For the Czech students there was another positive outcome in the strengthening of their motivation to learn English. We noticed a remarkable improvement in their command of the language when comparing their initial and concluding presentations. All participating students got to know about current issues in the other country and also spoke about their personal interests.

Most of us who have ever worked in a different country realize that some of the local issues look different from elsewhere. The Global Cooperation Project helps the students to get this kind of different perspective at almost no extra cost as compared to international student exchange. It gives them a background that makes them more eligible for a stay abroad later on during their study. We anticipate that it will also result in a greater ability to establish personal contacts, which will greatly facilitate future business cooperation.

To the teachers the project confirmed that good motivation helps the students more than hours of lecture. We were surprised by the intricate presentations the students were able to prepare for their Web pages in spite of the fact that they were given

only a minimum instruction concerning the structure of HTML documents.

Another interesting teaching experiment is being developed in the first IT course for seniors, which is being taught at the department in the 1998/99 school year. It is designed as a bi-semester course stressing the practical aspect of computer usage and does not include any programming. The enrolment of this course is 35 people (10 men and 25 women) whose age is between 55 and 79 years (65 on average). Mostly, these people do not plan to use computers for work since they are pensioned already. An important motivation factor for them was to be able to understand their grandchildren and perhaps help them when playing and learning with computers.

It is a well-known fact that grown-up people who have not had any opportunity nor need to get acquainted with computers during their education or at the beginning of their professional career have both objective and subjective problems to start using them. Upon our meaning, the most important obstacles are subjective in nature. These people are frightened with the omnipresent buzz about computers considering these machines too complex and themselves too far behind to be able to understand them properly. This is not a problem at all for children ñ they have do not fear computers do not bothering what is inside but playing the games that attract them.

We are successfully trying to persuade our senior students that this is the way they should take it as well and Internet services (starting with WWW) is used instead of games. The potential of these services and the ease of using them are very convincing. Once the subjective barriers are removed, senior students advance very well in other topics making usage of the features of grown-ups: determination and dedication.

**Please contact:**

**Josef Koláři- CTU**

**Tel: +420 2 2435 7403**

**E-mail: [kolar@fel.cvut.cz](mailto:kolar@fel.cvut.cz)**

# Networking Technologies

by Yves Devillers

**The tremendous increase in the demand for and use of Internet, mobile communications and multimedia along with an improved or more exacting employment of existing facilities is creating new problems that call for innovative solutions. This special theme section on Networking Technologies discusses some of the ways in which ERCIM institutions are addressing these issues.**

The basic technologies such as Synchronous Data Hierarchy, Wavelength Multiplexing and optical switching needed to deploy high speed networks running at gigabits per second as envisioned by initiatives such as Internet 2, Canary, NGI and STARTAP are marginal to ERCIM's scope of interest. Attention in the ERCIM institutions is focussed on deploying testbeds to tame such speeds, to experiment new applications and to understand the impact of speed on applications when the end to end throughput is in the order of hundreds of megabit per second. This question is treated in first paper, by Wunderling and Hommes, on Gigabit Networking.

New architectures that include dedicated hardware and algorithms must be designed to handle routing and switching at such speeds since software alone is no longer sufficient. The unification of transport concepts will mean that one common architecture can handle several protocols (such as ATM, IP and PDH), see the papers from FORTH and VTT on pages 14 to 17.

The widely requested and long awaited introduction of multimedia on the network cannot be achieved without Quality of Service enforcement. Several tracks are relevant to this field such as resource reservation, service differentiation, simplified and integrated routing techniques, or even over dimensioning of the bandwidth resource,

with significant efforts in convergence toward one single protocol for the Internet. Cooper (page 22) summarises the problems to be expected in terms of scaling and of resource allocation and reservation. Queueing theory allows the identification of undesirable behaviour when mixing real time and best effort traffic or flowing long-trailed traffic (see Borst on page 20); the invited paper by Matthias Grossglauser (page 21) advocates monitoring traffic at different time scales in order to improve resource control mechanism.

In order to deploy simpler networks capable of handling high speeds and providing quality of service, emerging operators in North America usually rely upon IP absorbing ATM and SDH with the aim of running directly over

wavelength multiplexing and optical switches at 10 to 40 Gbps around 2002. As of today, ATM is still the preferred operational protocol for Telcos to enforce QoS over long distance networks at 155 Mbps (TEN 155 in Europe) and 622 Mbps (vBNS in the USA) with a few extensions and testbeds at 2.5Gbps. Multi Protocol Label Switching (MLPS) – IP routing protocols controlling ATM switches to support IP flows labelled according to requested QoS – is one example of convergence between IP and ATM; Chatzaki and Sartzetakis discuss this and related experiments, while Sartzetakis on page 24 addresses multi-vendor ATM management and Todrova and Brandt on page 26 warn about security issues raised by 'ATM to the desk' applications via public ATM networks.

## CONTENTS

|  |    |   |    |
|--|----|---|----|
| Networking Technologies<br>by Yves Devillers   | 12 | HOT – Home and Office Technologies<br>by Thomas Luckenbach  | 27 |
| Gigabit Networking is Reality –<br>2.4 Gbps via ATM Wide Area Network<br>by Peter Wunderling<br>and Ferdinand Hommes | 13 | Heterogeneous Inhouse Networking<br>Environment<br>by Dirk Elias  | 28 |
| Designing High-speed Communication<br>Systems<br>by Dimitrios N. Serpanos  | 14 | ARNO: Algorithms for Radio Network<br>Optimization<br>by Heinz Mühlenbein   | 29 |
| Wormhole IP over ATM<br>by Manolis Katevenis   | 15 | AMC - ATM based Wireless Mobile<br>Computing<br>by Enrico Gregori and Thomas<br>Luckenbach                                | 29 |
| High-Performance ATM Switching:<br>the ATLAS I Single Chip Switch<br>by Dionisios N. Pnevmatikatos                   | 16 | Internet Telephony merges with the GSM<br>Network<br>by Jori Paananen   | 31 |
| Integration of Multiple Switching<br>Disciplines<br>by Pertti Raatikainen  | 17 | SFT – Smart File Transfer<br>by Markus Laakso   | 31 |
| Multiservice Internet:<br>Service Network or Ham Technology?<br>by Chris Cooper                                      | 18 | Activities at SICS's Computer and Network<br>Architecture Laboratory<br>by Bengt Ahlgren                                  | 32 |
| Resource Allocation<br>in Integrated-Services Networks<br>by Sem Borst   | 20 | Monitoring and Displaying Traffic on the<br>World Wide Web<br>by Evangelos P. Markatos and<br>Athanasios E. Papathanasiou | 33 |
| Control of Network Resources<br>over Multiple Time-Scales<br>by Matthias Grossglauser                                | 21 | Investigation of Internet Object Cache<br>Performance Bottlenecks<br>by Bertold Kolics                                    | 34 |
| Internet Enhancements – Coexistence with<br>ATM<br>by Magda Chatzaki and Stelios<br>Sartzetakis                      | 22 | The Changing Face of Network<br>Management – A Hungarian Experience<br>by István Tétényi                                  | 35 |
| Technology Interoperation in ATM<br>Management<br>by Stelios Sartzetakis   | 24 | The new CNR Research Area in Pisa<br>chooses ATM Local Area Network: why<br>ATM?<br>by Marco Sommani                      | 36 |
| ATM Security Aspects<br>by Petia Todrova<br>and Hartmut Brandt   | 26 |   |    |

The ubiquitous Internet is now entering into our homes. Luckenbach (page 27) and Elias (page 28) tell us more about technologies and protocols of potential interest for hooking up the home and gives research news on in-house networking topics ('ATM QoS to the toaster').

'Everywhere' is anyway reachable by wireless local loops either via mobile phones (GSM, DECT) for Personal Data Assistant or even portable laptops, or for immobile LANs (LMDS) in the suburbs and industrial areas. Mühlenbein on page 29 talks about antenna dimensioning and placement together with frequency assignment problems while Gregori and Luckenbach describe a laptop technology based on DECT with - an expectable - sub-optimal TCP performance over non error-free media. GSM wireless and IP telephony are the topic of the papers from VTT on page 31: Paananen discusses the integration of a GSM network with the corporate intranet and Laakso talks about reliable file transfer over GSM to communicate with Finnish icebreakers. Research news from SICS (page 32) spans from network conscious applications to IP telephony and differentiated services.

Finally case studies tell us more about the monitoring and displaying of Web traffic (Markatos and Papathanasiou on page 33), about Web cache monitoring in Hungarnet and the day to day network management of this net (papers from SZTAKI on pages 34-35), and reasons behind the not so evident choice by CNR for an ATM rather than a Gigabit Ethernet solution for its new Pisa campus.

**Please contact**  
**Yves Devillers – INRIA**  
**Tel: +33 1 3963 5976**  
**E-mail: yves.devillers@inria.fr**

# Gigabit Networking is Reality – 2.4 Gbps via ATM Wide Area Network

by Peter Wunderling and Ferdinand Hommes

**One year after the successful start of the first German Gigabit Network Testbed (see ERCIM News Number 31, October 1997) the participants of the „Gigabit Testbed West” could demonstrate transmission speeds of 2.4 Gbps using ATM technology (Asynchronous Transfer Mode). The project is sponsored by the German Research Network Association ‘Verein zur Förderung eines Deutschen Forschungsnetzes, e.V. (DFN)’, and the Federal Ministry for Education, Science, Research and Technology of Germany with additional funding of the participants GMD and the Research Center Jülich (Forschungszentrum Jülich - FZJ).**

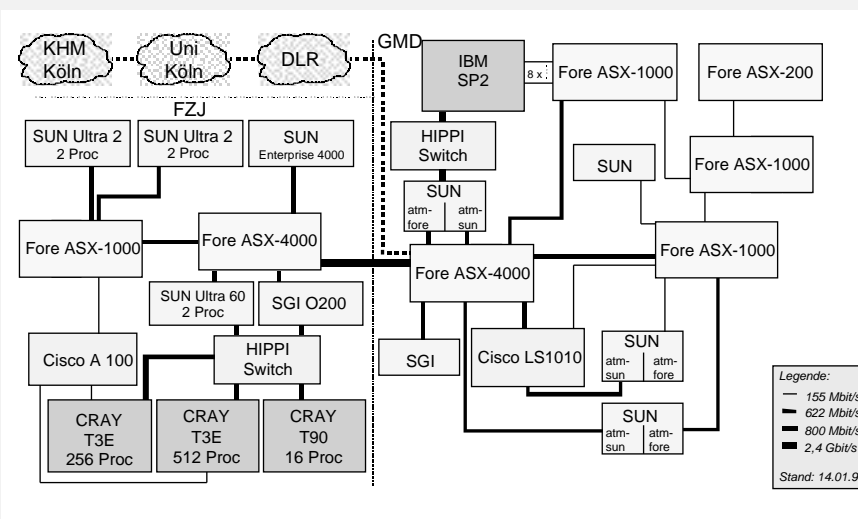
The goal of the testbed is to get experiences in high speed networking and to show the necessity for next generation applications. Experiences and knowledge of the testbeds are base of the

design and operation of the next generation network for the scientific community in Germany which is planned by DFN. The focus of the testbed is on metacomputing as well as on dedicated multimedia applications like 'Distributed virtual studio applications'. In this article only the networking aspects are shown.

The Gigabit Testbed West is one of two Gigabit Testbeds in Germany (the other one is Gigabit Testbed Süd) and started at August 1st 1997 with a 622 Mbps link between both research centers. The technique is Asynchronous Transfer Mode via SDH (Synchronous Data Hierarchy). The SDH equipment and the fiber optic link (about 100 km between the locations Sankt Augustin and Jülich) are provided by o.tel.o Service GmbH, a subsidiary company of a German energy supplier.

In the first step ASX-1000 switches (FORE Systems) connected the local ATM networks of both research centers running the link with 622 Mbps. The high performance computing (HPC) systems on both sides (CRAY T3E with 512 nodes in Jülich and IBM SP2 with 34 nodes at GMD) were connected to the ATM network via 155 Mbps interfaces. The first tests and measurements of the throughput between the HPC systems were disappointing - far below of 155 Mbps.

Since the announced 622 Mbps ports (ATM) for both systems are up to now not available, the following solution was



Configuration of the Gigabit Testbed West.

implemented to connect the CRAY T3E with the IBM SP2: FZJ and GMD are running HIPPI (High Performance Parallel Interface) to connect their HPC systems with transmission speeds of up to 800 Mbps. Using SUN and SGI systems with one or more 622 Mbps ports for the ATM network and interfaces for the HIPPI network it was possible to run these systems as HIPPI/ATM gateways.

The great success was at August 5th 1998. Both research centers could demonstrate the first 2.4 Gbps ATM WAN in Europe. Running ASX-4000 (the newest generation ATM switches of FORE Systems) with 2.4 Gbps interfaces on both sides the theoretical throughput of 2,17 Gbps for ATM data (in total: 2.488 Gbps capacity minus overhead for SDH and ATM) could nearly be reached. The tremendous data traffic was generated by several SUN systems each of them connected to the ASX-4000 with one or more 622 Mbps interfaces (SUN and FORE Systems).

The performance measurements of the ATM/HIPPI gateways showed acceptable results: 430 Mbps between CRAY T3E and H/A-gateway and 370 Mbps between IBM SP2 and H/A-gateway. Between the gateways a maximum throughput of 500 Mbps was measured. That means up to now both HPC systems can communicate with a maximum throughput of 370 Mbps.

The next steps in the Gigabit Testbed West are to get the link to Cologne operational and to implement the planned applications like 'Distributed TV-Productions' and '3V - Distributed Traffic Simulation and Visualization'.

For more information, see:  
<http://imk.gmd.de/hsn/>

**Please contact:**

**Peter Wunderling – GMD**  
Tel: +49 2241 14 2930  
E-mail: [peter.wunderling@gmd.de](mailto:peter.wunderling@gmd.de)

**Ferdinand Hommes – GMD**  
Tel: +49 2241 14 2932  
E-mail: [ferdinand.hommes@gmd.de](mailto:ferdinand.hommes@gmd.de)

## Designing High-speed Communication Systems

by Dimitrios N. Serpanos

**High-speed transmission technology experiences significant progress for over a decade. The development and deployment of high-speed links in the business environment and to the home has triggered the development of a wide range of networks and services. As a result, the design of communication systems faces significant performance challenges, especially in light of the requirements for delivery of real-time traffic (multimedia).**

The main problem in designing high-speed communication systems, ie adapters, bridges, routers, etc., is the throughput preservation problem: the problem to preserve the throughput of a high-speed link to applications through the protocol stack(s) implemented in a system. Solution to this problem requires the development of 'intelligent' communication system architectures, which implement high-speed data movement and protocol processing that meet the necessary performance requirements.

We are working on modular, scalable architectures for communication systems, because they are required to implement a wide range of functions and to meet the performance requirements of links as they scale to higher speeds. Considering the wide adoption and deployment of ATM technology, we focus on ATM systems, because they have demanding requirements, especially considering the scalability of the SDH/SONET physical layer. In our efforts, we evaluate and explore alternative technologies in order to obtain high performance systems at low cost.

An important module in high-speed communication systems is the memory manager, which manages logical data structures – typically queues – efficiently and enables high-speed data transfers

between system memory and link interfaces. We have developed a high-speed, scalable and re-usable Queue Manager, suitable for a wide range of ATM systems, such as workstation adapters, switches, routers, etc. The Queue Manager is a special-purpose processor that executes memory management instructions. To achieve re-usability, we have analyzed the requirements of the most common ATM functions (flow control, segmentation-and-re-assembly, etc.) and identified an instruction set that is sufficient to implement these functions. To achieve scalability, we have identified a minimal set of instructions as well as a minimal set of data structures to support. We have developed an architecture and several hardware implementations, which provide increasing performance at the cost of more complex hardware. A typical low cosqueues of t implementation supports 1024 8192 ATM cells and performs an Enqueue and a Dequeue operation in 132ns and 198ns, respectively. Such an implementation supports ATM systems with aggregate throughput close to 1 Gbps.

Considering the processing power that exists in some systems in the form of embedded processors, we have developed a software implementation for embedded systems as well, using the CYCLONE evaluation board with the Intel i960 processor at 40 MHz. The average delays of the Enqueue and Dequeue operations in the software implementation are 0.75 and 0.95 microseconds, respectively. We have evaluated the performance, cost and scalability of all implementations, so that one can choose the solution with the desired characteristics.

The Queue Manager provides the support for efficient queue management and fast data movement in an ATM system, but it is necessary to provide additional support in order to implement Quality-of-Service (QoS) in data transmission; typically this is achieved with adoption of priorities. Scheduling of outgoing traffic constitutes an additional challenge. In this direction, we have designed in our lab an efficient scheduler that uses weighted round-robin for QoS. The

scheduler operates in co-operation with the Queue Manager providing a complete solution in a prototype 4x1 ATM concentrator with 155 Mbps Sonet links.

As design methodologies and semiconductor technology progress, we can provide improved architectures and designs that exploit these advances. We are currently developing modules with more sophisticated and efficient architectures, targeting support of an increased number of functions and scalability to the higher speeds required by the emerging gigabit networks.

**Please contact:**

**Dimitrios Serpanos – ICS-FORTH**  
 Tel: +30 81 391663  
 E-mail: serpanos@ics.forth.gr

## Wormhole IP over ATM

by Manolis Katevenis


**Speeding up the Internet is of capital importance; this will need some hardware assistance. Hardware switches operate best on fixed-size quanta, while IP packets, on the other hand, have a variable size. It is not the first time, though, that hardware is called upon to route variable-size packets at high speed: Wormhole Routers did precisely that, in the eighties, for multiprocessor inter-connection networks. The same techniques are applicable today to the Internet: we introduced a novel, wormhole-like way to route IP over ATM, we simulated its performance, we built a first prototype, and we are now proceeding to testing it in real systems.**

IP is the uncontested protocol for data communications. At the same time, ATM technology finds widespread use, owing to its fixed-size “cells” that allow high-speed hardware switching, owing to the small size of these cells that allows fast preemption and hence low latency, and owing to its quality-of-service (QoS) architecture. IP can fruitfully run on top

of ATM, thus getting the best of both worlds. This has been done in software, on general-purpose computers; we do it in hardware, at a lower cost and with a lower latency.

ATM has some similarity with wormhole routing, the most popular multiprocessor interconnection network technique of the eighties. Just like virtual channels in wormhole routing carry packets segmented into “flits”, a number of hardware-managed VC’s in ATM can carry IP packets segmented into cells according to AAL-5; each VC is

ATM’s QoS capabilities; (iv) for IP traffic, the system operates equivalently to a network of low-latency gigabit IP routers, while being a lot less expensive; (v) packet routing delay is minimized owing to virtual-cut-through routing – segmentation and reassembly delays at intermediate routers are eliminated; (vi) packet routing delay is minimized for all packets – not just for the rest of the packets after a flow has been recognized, as in IP switching; (vii) the number of pre-established connections (labels) is small and fixed, and does not grow with the size of the network (as in tag



**2x155 Mbps Wormhole IP over ATM Router Prototype.**

dedicated to one packet for the duration of that packet, and is afterwards reassigned to another packet, in hardware. This was proposed by Barnett in 1997 and was named connectionless ATM. We, at ICS-FORTH in Heraklion, Crete, Greece, modified the Barnett proposal to make it applicable to existing ATM equipment: we proposed a novel single-input, single-output Wormhole IP Router, that functions as a VP/VC translation filter between ATM subnetworks.

Our wormhole IP routing filter has a number of advantages: (i) it works together with standard, existing ATM equipment; (ii) it allows the co-existence and integration of both IP and native ATM traffic in the same networks; (iii) the quality of service of native ATM traffic can stay unaffected by the added IP traffic, while IP can benefit from

switching), yet all packets are routed through pre-established connections. Based on actual internet traces, we showed by simulation that a few tens of hardware-managed VC’s per outgoing VP suffice for all but 0.01% or less of the packets.

We have built a first prototype of a bi-directional wormhole IP routing filter with two OC-3 ports (155 Mbps in each direction), which is shown in the photograph. Two 16-MByte DRAM SIMM’s plug into the sockets on the right; they hold the two-level IP routing table and the ATM connection table. The datapath, the control FSM, and the VCut free list (bit map) are contained in one FPGA. The routing delay is fixed, equal to about one cell time (a couple of microseconds), for all IP packets.

We are currently testing this prototype by connecting it to our ATM and IP network. Further development is needed, especially on the software side, and we are seeking partners for that and for commercially exploiting the technology.

For more information, see <http://archvlsi.ics.forth.gr/wormholeIP/>

**Please contact:**

**Manolis Katevenis – ICS-FORTH**  
Tel: +30 81 391664  
E-mail: [katevenis@ics.forth.gr](mailto:katevenis@ics.forth.gr)

## High-Performance ATM Switching: the ATLAS I Single Chip Switch

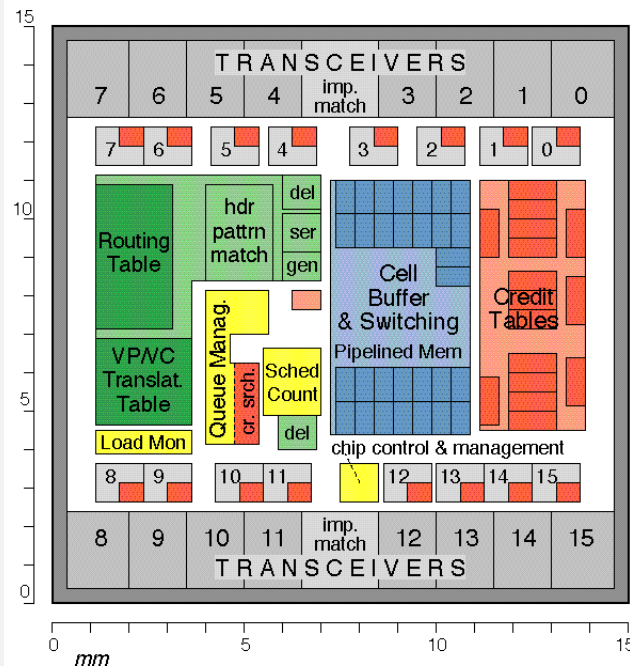
by Dionisios N. Pnevmatikatos

Asynchronous Transfer Mode (ATM) networks have become mainstream for wide area networks (WANs) and also in local area networks. They are based on packet switching, where 48 bytes of data together with a 5 header form the transfer unit called cell. This fixed packet size allows the parallel processing of cells inside the relay switches achieving high transmission throughput. ATM was originally designed to carry voice traffic, and as a consequence, ATM networks can efficiently handle voice, video and other types of traffic making them ideal for multimedia-type applications.

ATLAS I (ATm multi-Lane backpressure Switch One) is a high performance, single chip ATM switch designed in the Computer Architecture and VLSI Systems Division, of the Institute of Computer Science of FORTH. It is being developed within the ASICCOM Project, funded by the European Union ACTS Programme.

The ASICCOM Consortium includes two ERCIM members, FORTH and SINDEF. The other partners of ASICCOM are industrial partners (INTRACOM, Greece; ST Microelectronics, France and Italy; BULL, France), telecom operators (TELENOR, Norway; TELEFONICA, Spain), and research institutes (Poli. di Milano, Italy; Democritos, Greece). ATLAS I combines several desirable

credit-based flow control (backpressure). A cell in a backpressured service class (priority level) can only depart if it acquires both a buffer-pool credit for its outgoing link and a flow-group credit for its connection. A flow group is a set of connections that are flow-controlled together; ATLAS I supports up to 32 thousand flow groups. Credit-based flow control is useful in building switching



**ATLAS I floorplan. Sub-blocks are grouped by functionality.**

features such as a large number of links, high bandwidth links, advanced operating features, large buffer space and hooks for more efficient network management.

ATLAS I is a 6 million transistor chip designed in a 0.35 micron CMOS technology with five metal layers provided by ST Microelectronics. It features 16 input and 16 output ports (links) each operating at 622Mbits/s for an aggregate outgoing throughput of 10Gbits/s; the links use the IEEE 1355 'HIC/HS' protocol. The switch operates at 50 MHz and achieves a sub-microsecond cell cut-through latency. Buffer space for the ATM cells is provided in the form of a 256-entry shared buffer, which is split into multiple logical output queues; ATLAS I also supports three levels of priorities, multicasting and header translation.

The most distinctive and innovative feature of ATLAS I is its (optional)

fabrics with internal backpressure that provide the high performance of output queueing at the low cost of input buffering, and in making ATM sub-networks that never drop cells of data connections while fully and fairly utilizing all the available transmission capacity. An evaluation performed at FORTH shows that the credit protocol offers superior performance, especially in isolating well-behaved connections from bursty and hot-spot traffic.

ATLAS I also provides for efficient network load monitoring by means of accelerated Cell Loss Probability (CLP) measurement hardware. Very often in high-speed networks the probability of cell loss (for non-backpressured traffic) is very small, and measuring the CLP requires a very long observation period. The ATLAS I CLP hardware measures the cell loss rate of an emulated set of small buffers. Software can then use these measurements (over short periods of time)



to extrapolate the CLP of the actual traffic. ATLAS I also provides hooks for switch management through the network. Special management cells are addressed to an internal management port, which can execute read and write commands to any of the internal structures of the switch. The results of the operations are packed in a reply cell and returned to the sender. Security is achieved by only accepting management cells from designated input links and connections.

ATLAS I is a powerful building block for the construction of ATM networks. The large number of links of the switch minimizes the number of switches needed for the construction of a network, as well as the delay (measured in hops) through the network itself. In addition the smaller number of chips reduces the physical dimensions of the network and increases its reliability. The advanced flow-control features of ATLAS I allow the fair and efficient use of the network bandwidth in a mixture of services ranging from real-time, guaranteed quality-of-service to best-effort, bursty and flooding traffic.

Furthermore, the versatility in the management of the individual switches allows for the creation of powerful network management software.

ATLAS I was fabricated by ST Microelectronics, Crolles, France, in December 1998-January 1999, and is currently being packaged; we will proceed to testing by placing the chips in a demonstration system that was built by Intracom, Sindef and Telenor. Future interest includes the commercial exploitation of the chip, and further development work in ATLAS-to-SONET interfacing, and in flow-control and weighted round-robin scheduling architectures.

For more information on the ATLAS I switch see:  
<http://archvlsi.csi.forth.gr/atlasI/>

Please contact:

Dionisios Pnevmatikatos – ICS-FORTH  
Tel: +30 81 391657  
E-mail: pnevmati@csi.forth.gr

## Integration of Multiple Switching Disciplines

by Pertti Raatikainen

**Multimedia services presume that the underlying transport network supports a large variety of connections and guarantees required quality of service. A number of networking concepts has been tested, but none of them has gained ultimate acceptance. ATM technology has been the most favoured one, but other transporting solutions have also been considered and entirely new ones suggested, especially, for IP based communications.**

Development of a new unifying transport concept is a long process, and it can be envisaged that heterogeneous transport solutions will be used for multimedia service delivery for a lengthy period of transition. This implies that future network nodes, such as switches, have to interface to different kinds of networks, manage a number of parallel call control procedures and support switching of dissimilar data units.

As a solution to the emerging problem, VTT Information Technology and Helsinki University of Technology have launched SCOMS (Software

Configurable Multidiscipline Switch) project to develop a switching and routing solution capable of supporting multiple switching disciplines. The target is to develop a network node that integrates circuit, cell and packet switching into a single fabric. The project is funded by TEKES (Technology Development Centre of Finland) and six industrial partners. The project was started early 1998 and finishes by the end of year 2000.

### Experimental switch

A conventional telecommunications switch connects to a homogeneous transport network whereas a multi-discipline switch operates in a heterogeneous networking environment, ie, networks that utilise different transport technologies. Due to the different nature of the associated transport technologies, the switching equipment has to implement a set of different line interfaces and signalling procedures and, additionally, to support a varying set of connection types, quality of service measures and service types.

A straightforward way to implement such a switch is to logically divide the switch into separate 'virtual switches'. In this way, line cards, protocol stacks and associated call control functions of each virtual switch can work as if they were run in a homogeneous network. At the bottom layer, a common control block is needed to manage resources of the physical switch (eg switch configuration), convert application level

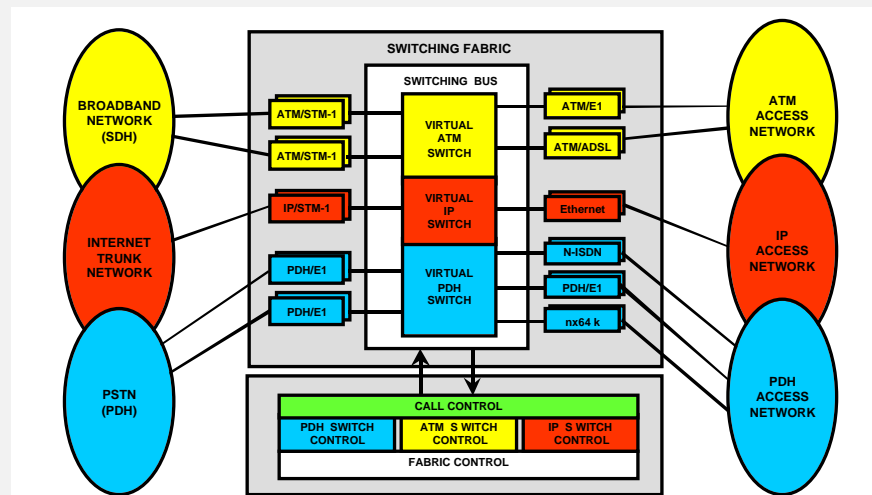


Figure 1: Block diagram of an experimental multidiscipline switch.

connection commands to physical connections and perform physical level routing of the by-passing data units.

The concrete goal of the SCOMS project is to develop an experimental switch integrating 64 kbit/s based circuit, ATM cell and IP packet switching (see Fig. 1). The physical switching platform is based on the Frame Synchronised Ring (FSR) concept, developed and patented by VTT Information Technology. This platform has characteristics necessary to implement the various fundamental features of a multidiscipline switch. These include, eg, inherent support for real-time multicasting, versatile addressing capabilities and support for simultaneous switching of different size data units. FSR's ring-shaped switching bus is dimensioned to perform effectively by implementing separate transport containers for circuit, cell and packet switched traffic.

Operation of the switching fabric is controlled by a workstation which implements fabric and call control functions. Signalling protocols of the three different transport networks (ATM, PDH and IP) are implemented separately, although, they use services of the common fabric control layer (see figure 2). Interworking between the different networks can be supported by a shared interworking call control module that is able to manage calls that originate and terminate at different networks (see figure 3).

The switch control module offers network specific application programming interfaces that implement functions necessary to reserve resources and manage connections through the switch. The control workstation is connected to the switch with an ATM/STM-1 link, and all protocols receive and transmit their signalling messages in reserved ATM virtual channels. Therefore, a non-ATM line interface must map the incoming signalling messages to a specific ATM virtual channel, assigned for transferring the signalling to the control workstation. ATM/STM-1 was chosen for the workstation connection, because it is a cost-effective way to have a high-speed signalling link, which offers a straightforward way to manage various control and signalling connections. An Ethernet interface is another possibility to connect the workstation to the switch.

The above configuration can be enhanced by introducing an interworking call model to route calls between the different networks (Fig. 3). In this configuration, the different protocol stacks are linked together via a shared call control layer. The call control must be able to decode, process and forward connection set-up requests from each protocol stack. The call control interworking function maps the connection set-up parameters between the different networks and maintains status information of all the established connections.

Currently, the SCOMS switch and its control architecture is in the implementation phase, and the first prototype, offering interworking between ATM and PDH networks, is expected by the end of year 1999. Further work will include interworking with IP networks, service architecture considerations as well as development of configuration and management applications. The inherent support for TCP/IP networking enables interesting views, for example, to develop distributed service control and web-based management.

**Please contact:**

**Pertti Raatikainen – VTT Information Technology**  
**Tel: +358 9 456 5644**  
**E-mail: pertti.raatikainen@vtt.fi**

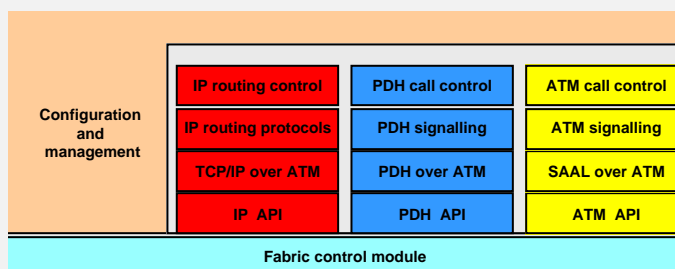
## Multiservice Internet: Service Network or Ham Technology?

by Chris Cooper

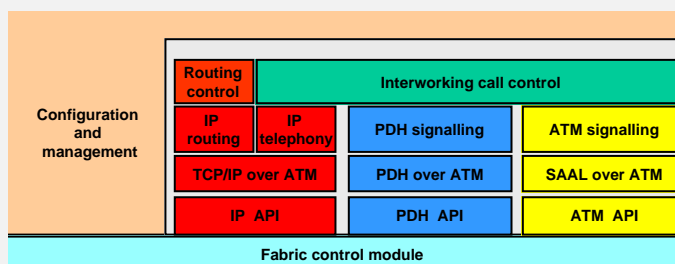
**As people collaborate over greater distances on a regular basis, so an increasingly significant, if as yet potential benefit of the global Internet is support of technical discussion, training, seminars, etc. – in fact the complete range of computer-supported co-operative working made accessible from each participants desktop.**

In a previous article (ERCIM News No.35, October 1998), David Duce described work undertaken in the European Framework 4 Telematics project MANICORAL to develop and begin initial trial assessment of a co-operative visualization system to support a European community of geodetic scientists and engineers. The experience gained with that project indicated that such desktop-based access to co-operative visualization had substantial potential: it also demonstrated how much still remains to be done to enable the

**Figure 2: Separate call control stacks for the different network architectures.**



**Figure 3: Call control stacks for interworking.**



Internet to support such styles of use in practice.

The demands of this sort of application are considerable when compared with, for example, a single user browsing world-wide web services. The latter demands support from the underlying network for only a single type of traffic: transport of traditional data, reliably, on a point-to-point basis-which the Internet has been providing on a best-effort basis for many years.

Collaborative visualization however requires not only traditional reliable data transport (for the visualization) but also continuous media transport for voice (and possibly video) communication. Moreover, both the data and continuous media support need to be on a multipoint basis in order to support (small) groups of more than two people. Experience in the MANICORAL project indicated that the current best-effort Internet fell far short of expectations in respect of what was required for routine collaboration, partly because multicast is not yet all-pervasive as a service, and partly because the quality of voice communication was too often so variable, including incomprehensibility or total loss, as to render a session abortive. There were a number of contributing causes to the project's specific experience. Here, a few of the underlying technological issues relating to the evolution of the Internet from a single-service (best-effort data) to multiservice are remarked upon.

The basis for multipoint transport in the Internet has been around for nearly a decade in the form of the multicast overlay network known as the Mbone. Great progress has been made in multicast routing, and this is now available in routers. Nevertheless, the transition from prototype multicast overlay to integrated multicast service still evidently has a long way to go to approach coverage comparable with point-to-point service. Even greater is the problem represented by the need to carry more than a single category of traffic. Continuous media traffic, such as speech, does not require that every bit is delivered correctly (though obviously most of it must be), but it does demand that voice

packets are delivered on a regular, timely basis. Packets arriving too late are as good as lost: too late means that no more samples are available to play out to the listener and there is a gap in the sound. Too many such gaps and speech becomes incomprehensible. In the case of conversational speech, the round-trip delay also needs to be short to prevent it interfering with normal conduct of the conversation.

The traditional approach to the problem of continuous media traffic transport in networks is to reserve resources for the lifetime of the activity, as is done in the telephone network during a call. This is also a part of the approach taken in seeking ways of integrating this traffic with traditional data traffic, whether the network technology is asynchronous transfer mode (ATM) for broadband ISDN, or integrated services Internet. There are two possibilities for reserving resources to handle traffic which cannot tolerate delay: the resources either need to be reserved dynamically by some sort of signalling or they must be permanently reserved through some service contract. In either case, some state has to be introduced into the switching elements of the network, the Internet routers. Resource reservation protocol, RSVP, has been developed as a means of letting a host signal its need for resources to be reserved dynamically in the network on a flow-by-flow basis. Considerable progress has also been made in the development of queue management scheduling disciplines to enable packets to be transmitted on an output link in such a way as to support the quality of service required for the different flows to which the packets belong.

So what's the problem? Well, two really: a scaling problem, and a basic problem related to resource allocation. The scaling problem arises from attempting to associate state with potentially every flow in the Internet. The memory space and packet processing associated with this may be manageable in more or less local or small regional networks, but it is not possible to support it in the core or backbone routers of the global Internet. The generally agreed approach to dealing with this problem is to aggregate flows

into a few types, to mark packets accordingly at the edges of the network, at the same time as operating admission control, so that the core need only perform very simple tests on a small field in each packet to determine which type of behaviour (scheduling) is appropriate for that packet, in the knowledge that the admission control will have effectively guaranteed that the resources are available in the network to handle the packet.

The other problem is that any form of resource reservation by one party implies that in times of high load or potential congestion some other party is denied service. Without additional criteria, there is no way of deciding who should be given service and who should be denied. One possible way out is to introduce pricing: different qualities of service attract different tariffs, and operators size their networks to provide service for amounts of different types of traffic according to observed traffic patterns and specific contracts. One effect of such a mode of operation is that there is a direct incentive to upgrade a network to cope with more traffic of a particular type as soon as traffic projections indicate approach to the region of denial of service: suitable investment might not only prevent subscribers moving to another provider but also offer the potential of further profit.

This article is based in part on a talk given by David Duce last summer at Multi-Service Networks'98, held each year a Cosener's House, Abingdon, Oxfordshire, UK. Details of this year's workshop will be available at <http://www.acu.rl.ac.uk/msn99/>

**Please contact:**

**Chris Cooper – CLRC**  
**Tel: +44 1235 44 6211**  
**E-mail: [Chris.Cooper@rl.ac.uk](mailto:Chris.Cooper@rl.ac.uk)**

# Resource Allocation in Integrated-Services Networks

by Sem Borst

**Service providers race to develop next-generation communication networks, integrating a wide variety of services, such as voice, video, and data, onto a common infrastructure. There are major obstacles to overcome, however, since different services not only have drastically different traffic characteristics, but also extremely diverse quality-of-service requirements. Performance evaluation research at CWI focuses on models and techniques from queueing theory for evaluating the relevant quality-of-service measures, such as delay performance. Important issues in integrated networks which have received particular attention include the impact of long-tailed traffic characteristics and the interaction between best-effort and real-time services.**

World-wide, the use of communication services is experiencing revolutionary growth. The growth is fueled not only by the expansion of conventional telephone services, but also the advance of data communications, the spectacular development of the Internet, and the proliferation of wireless communications. Driven by these demands, service providers rush to enhance their networks. A major trend is the integration of voice and data services.

Eventually, this is expected to result in the consolidation of a wide variety of services onto a common platform. While offering potential synergies, however, the integration of heterogeneous service classes also involves several fundamental problems. First of all, different services may have radically different traffic characteristics. To characterize traffic processes, it is convenient to adopt a

three-level hierarchy. The lowest level of granularity is the packet level. Packets are typically generated in bursty patterns occurring at a moderate rate, defining the burst level. Finally, at the connection level, connections are established and terminated.

Different service classes may also have extremely diverse quality-of-service requirements. Important measures are the amount of packet delay incurred during transmission, and the fraction of packet loss caused by buffer overflow. Voice traffic is rather sensitive to delay but can sustain some packet loss, whereas data traffic can tolerate some amount of delay but is quite vulnerable to packet loss.

The heterogeneity in quality-of-service requirements and traffic characteristics requires sophisticated allocation mechanisms to regulate the usage of network resources. The main resource allocation instruments include admission control, routing, scheduling, and flow control.

The design of efficient resource allocation algorithms requires techniques for evaluating the relevant quality-of-service measures, eg, delay performance, cell loss, call blocking. CWI has traditionally focused on techniques from queueing theory, which is concerned with the study of congestion phenomena in stochastic service systems. The basic queueing model consists of a group of servers, where customers arrive who require some kind of service. In the context of communication networks, the server usually represents a transmission link, or a buffer or port on a switch. The customers typically correspond to transmitted packets, bursts, or offered calls, in accordance with the three-level hierarchy described above.

Real-time traffic is traffic which is extremely sensitive to delay, such as voice and video-conferencing. In contrast, best-effort traffic is tolerant of some amount of delay, as long as the average capacity received over somewhat longer intervals is sufficient, file transfers being a typical example. Usually real-time traffic is given some form of priority over best-effort traffic in packet

scheduling. The best-effort traffic thus receives the capacity left-over by the real-time traffic. That remaining capacity may drastically fluctuate over time because of the bursty patterns in packet flows described earlier. The available capacity is then shared among the best-effort connections in certain fixed proportions. This motivated a detailed investigation at CWI of (discriminatory) processor-sharing models with time-varying service capacity. The results show that the average transmission time of best-effort files is asymptotically linear in the size of the files, which is commonly viewed as a desirable fairness property.

A second major research topic at CWI concerns the occurrence of long-tailed traffic characteristics. Measurements have provided increasing evidence that packet traffic exhibits long-range dependence and self-similarity over a wide range of time scales. Queueing models with long-tailed input traffic offer a useful approach to evaluating the impact of long-range dependence and self-similarity on network performance. It was shown for example that if the input traffic is long-tailed, then so are the buffer content and the delay (buffer content even 'one degree worse'). In integrated networks, where the transmission capacity is shared by several traffic streams, even a single long-tailed traffic stream may cause the buffer content to be long-tailed, unless the link rate is larger than the peak rate of that long-tailed connection plus the mean rate of the other flows. These results have important implications for admission control. Recently, the research has focused on the issue how scheduling algorithms may be used to neutralize the negative effects from long-tailed traffic phenomena. It was found that if traffic is not processed in order of arrival, but in processor-sharing fashion, then the tail behavior of the delay is no longer any worse than that of the input traffic. ■

**Please contact:**

**Sem Borst – CWI**  
**Tel: +31 20 592 4205**  
**E-mail: Sem.Borst@cwi.nl**

# Control of Network Resources over Multiple Time-Scales

by Matthias Grossglauser

**Networked multimedia applications require resource allocation because of their quality of service (QoS) requirements. On the other hand, network efficiency depends crucially on the degree of resource overbooking inside the network. A key problem in concurrently achieving both goals is caused by the fluctuation over multiple time-scales of the traffic load emitted by multimedia applications, because it makes it hard to predict resource requirements with sufficient accuracy. This in turn requires a careful design of control mechanisms so that they cover all time-scales.**

In our work, we examine resource control over three natural time-scales. On the packet time-scale, we evaluate the performance of traffic smoothing as a mechanism to accommodate bandwidth fluctuation. Our interest stems from the mounting experimental evidence that packet arrival processes exhibit ubiquitous properties of self-similarity and long range dependence (LRD). A random process exhibits long-range dependence if it has a non-summable autocorrelation function. Intuitively, this means that the process exhibits fluctuations over a wide range of time-scales. This property is of importance because it cannot be captured by Markovian traffic models, which have traditionally been the analytical tool of choice in the teletraffic community. However, we show that in the case of traffic smoothing, there exists a correlation horizon that separates relevant from irrelevant fluctuation time-scales for the purpose of performance prediction. This illustrates the general principle that the traffic, system, and performance metric time-scales together determine the set of candidate traffic models.



Matthias Grossglauser was the winner of the 1998 Cor Baayen Award competition (see announcement on page 5). The work for which Grossglauser – an EPFL-graduate – received the award was carried out at INRIA Sophia Antipolis under the guidance of Jean Bolot where he was a member of the RODEO team. He defended his thesis in spring 1998 on the topic 'Control of Network Resources over Multiple Time-Scales'. Grossglauser recently joined AT&T Laboratories in the US.

Per-flow smoothing is not effective in removing the longer-term traffic fluctuations. To achieve high utilization, we therefore need a mechanism to share the link bandwidth among multiple flows. We advocate renegotiation as an efficient mechanism to accommodate fluctuations over time-scales beyond the correlation horizon, which we call the burst time-scale. A new network service model called RCBR (Renegotiated Constant Bit Rate) combines network simplicity with desirable quality of service guarantees, while achieving much of the potential statistical multiplexing gain of bursty traffic. With RCBR, the network guarantees a constant bit rate to the application. The application can renegotiate this bit rate, but there is a small probability of renegotiation blocking. A network implementing RCBR is simple because there is no substantial buffering in the network, and therefore no need for elaborate buffer management and packet scheduling mechanisms. The quality of service is determined by the renegotiation blocking probability, which is kept small enough by limiting the number of flows in the system. This is achieved through admission control.

On the flow time-scale, we discuss measurement-based admission control (MBAC) as a means of relieving the application of the burden of a-priori traffic specification. The traditional approach to admission control requires an *a priori* traffic descriptor in terms of the parameters of a deterministic or stochastic model. However, it is generally hard or even impossible for the user or the application to come up with a tight traffic descriptor *before* establishing a flow. MBAC avoids this problem by shifting the task of traffic characterization from the user to the network, so that admission decisions are based on traffic measurements instead of

an explicit specification. This approach has several important advantages. First, the user-specified traffic descriptor can be trivially simple (eg, peak rate). Second, an overly conservative specification does not result in an overallocation of resources for the entire duration of the session.

Relying on measured quantities for admission control raises a number of issues that have to be understood in order to develop robust schemes.

## Estimation Error

There is the possibility of making errors associated with any estimation procedure. In the context of MBAC, the estimation errors can translate into erroneous flow admission decisions. The effect of these decision errors has to be carefully studied, because they add another level of uncertainty to the system, the first level being the stochastic nature of the traffic itself.

## Dynamics and Separation of Time-scale

A MBAC is a dynamical system, with flow arrivals and departures, and parameter estimates that vary with time. Since the estimation process measures the in-flow burst statistics, while the admission decisions are made for each arriving flow, MBAC inherently links the flow and burst time-scale dynamics. Thus, the question of impact of flow arrivals and departures on QoS arises. Intuitively, each flow arrival carries the potential of making a wrong decision, and each departing flow allows to recover from a past mistake.

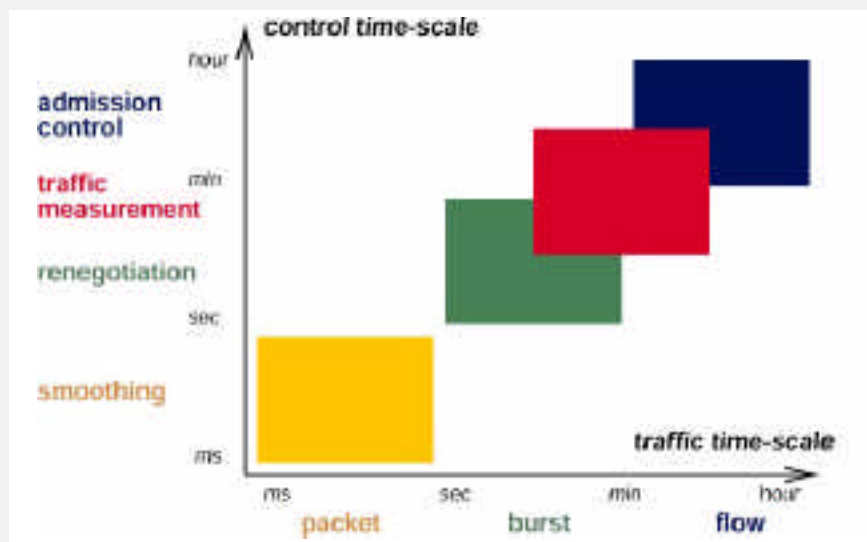
## Memory

The quality of the estimators can be improved by using more past information

about the flows present in the system. However, memory in the estimation process adds another component to the dynamics of a MBAC. Using too large a memory window will reduce the adaptability of MBAC to non-stationarities in the statistics. A key issue is therefore to determine an appropriate memory window size to use.

Using a simple model that captures the impact of measurement uncertainty and the interplay of burst and flow time-scale dynamics, we study all of the above

voice-over-IP, with more vendors providing its elements, and with an increasing number of operators competing for market share. The resulting system is of very large scale and complexity, to such an extent that one must assume *that there is always something wrong somewhere*. The resulting flood of alarm information makes it difficult for human network operators to detect, isolate, and repair faults efficiently. This stresses the importance of network management. Performance monitoring, fault



Resource control mechanisms have to cover all time-scales at which the load fluctuates. Traffic smoothing absorbs short time-scale fluctuations. Renegotiation in the RCBR service model accommodates traffic bursts and allows to share the link bandwidth efficiently. Admission control acts at the flow time-scale to avoid overload. Traffic measurement for the purpose of admission control overlaps with the burst and flow time-scales, because the quality of service depends on the burst time-scale, but admission decisions occur on the flow time-scale.

issues in a unified analytical framework. The goal is to shed insight on the design of robust MBAC schemes which can make QoS guarantees in the presence of measurement uncertainty, without requiring the tuning of external system parameters. The figure illustrates how the traffic time-scales and control mechanisms considered in our work relate to each other.

In our future work, we hope to address other problems that can affect the quality of service experienced by the user. The Internet grows rapidly in size and capacity, enabling more services such as virtual private networks (VPN) and

identification and localization, planning and resource provisioning, and configuration management are important and challenging future research topics, encompassing both architectural and performance issues.

More information about this project is available at:  
<http://www.research.att.com/~mgross/>

**Please contact:**

**Matthias Grossglauser – AT&T Labs - Research**  
**Tel: +1 973 360 7172**  
**E-mail: [mgross@research.att.com](mailto:mgross@research.att.com)**

## Internet Enhancements – Coexistence with ATM

by Magda Chatzaki and Stelios Sartzetakis

The IthACI project – Internet and the ATM: Experiments & Enhancements for Convergence and Integration – aims at contributing to the Multi-Protocol Layer Switching (MPLS) standardization efforts by developing and demonstrating enhanced features in the areas of multicast, QoS and resource management in IPSOFACTO and Tag-Switching, and mobility in a multicast environment. The project is setting up a testbed consisting of three distinct technology islands based on ATM switching technologies of different vendors, which incorporate different approaches for short-cutting IP traffic. NEC's 'IPSOFACTO', CISCO's 'Tag Switching' and ALCATEL's 'Yalsa technologies' are in operation in the test networks in three different countries in Europe, the enhanced features are under development on these technology islands, and their interconnection is planned over a Pan-European ATM wide area network.

Internet technology dominates communications today, with its TCP/IP protocols suite mostly appreciated because of its simplicity and flexibility. Traditional Internet supports only best effort service. However it transforms quickly into a commercial environment demanding support for Quality of Service (QoS). In order to provide means to guarantee QoS in the Internet environment we might increase its complexity by adding protocols for resource management, like RSVP (Resource reSerVation Protocol), so losing some of its original simplicity. Coming from the Telecom Operators world, ATM has been designed to support different QoS requirements providing simple and clean solutions. IP switching seems to be a promising new

approach to combine the speed of ATM switching with the simplicity and flexibility of IP internetworks. The basic idea of IP switching is to use IP network control on top of an ATM switch. MPLS (MultiProtocol Layer Switching) is the standard solution for IP switching proposed by the Internet Engineering Task Force (IETF). The MPLS standard is currently evolving to cover the areas of Quality of Service and Resource Management, Multicast support and Mobility for IP switching technologies.

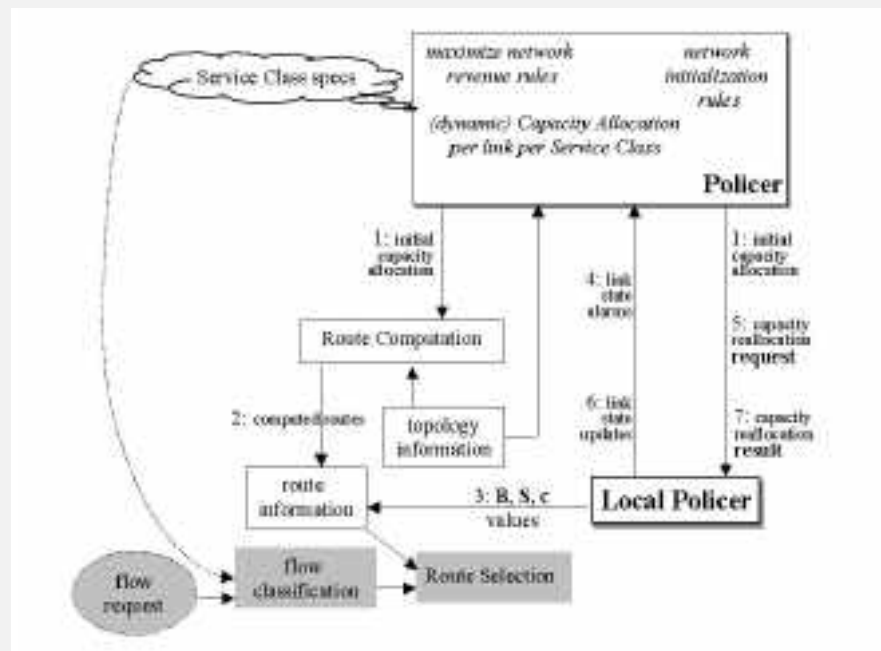
QoS and Resource Management Enhancements in Tag Switching aim at providing QoS and Resource Management functionality to the Tag Switching technology and consequently to MPLS. The enhancements to Tag Switching towards QoS and Resource Management are based on the idea of establishing an overall QoS provisioning framework from the viewpoint of a Service Provider. The QoS provisioning framework should provide differentiation between the services in terms of their transfer characteristics and above all differentiation among customers in terms of the services they may request. The framework is compatible with both IntServ and DiffServ approaches for QoS provisioning emerging in the Internet. It entails two aspects:

- definition of a set of discrete Service Classes (SCs) for requesting QoS-based services
- definition of a suitable architecture and protocols required for supporting the provisioning of the defined QoS-based services in the network.

The basic idea underlying this QoS framework is to create and dynamically manage a virtual topology induced to the physical one for allowing segregation of network resources (ie bandwidth) to the different SCs, cost-effectively and according to business policies. We name this virtual topology soft-network. The soft-network is used to configure and reserve the available network resources according to the characteristics of the supported SCs. Our architecture includes a mechanism that supports multiservice, multipath routing in Tag Switching (MPLS) capable networks. It supports initial allocation and dynamic reallocation

of network resources whenever necessary for multiple SCs. We consider a network that supports a number of SCs with different QoS guarantees. In our approach, network resources are initially

service classes treatment. The Resource Management enhancements first task is the development of a common prototype for the co-existence of a topology-driven unicast-oriented IP switching technique



**Architecture of a virtual network supporting QoS and Resource Management enhancements for Tag Switching.**

allocated to the different SCs according to the administrative policies of the network operator. This means that appropriate virtual networks are created to accommodate the traffic demands of the SCs according to their QoS characteristics. ICS-FORTH is mainly involved in this task and together with Algosystems S.A. develops the architecture shown in the figure in order to support the QoS and Resource Management enhancements for Tag Switching.

QoS and Resource Management Enhancements in IPSOFACTO include several tasks. The existing flow detection mechanisms in the IPSOFACTO architecture are being enhanced further with an real time protocol (RTP) flow detection mechanism. It should be noted here that RTP applications are very important, as they are actually real-time and therefore require a better service than traditional data applications. Further flow differentiation among the RTP flows is done using some clever assignment of flows with specific RTP profiles that will be mapped to the appropriate ATM switch

(MPLS) and a flow-driven IP switching technique (NECs IPSOFACTO). Interoperability and resource management issues involved with their co-existence will be studied. The second task is the design and implementation of a Flow Management Information Base (Flow MIB) to be used with NECs IPSOFACTO. The Flow MIB provides information on past and present flows.

For the Multicast Service enhancements we apply MPLS-style shortcut techniques to IP multicast in order to optimize the network without modifying the end-user environment and the existing IP multicast protocols. The shortcut point-to-multipoint ATM connections will follow dynamically the tree topology changes that are the result of IP hosts joining or leaving the group. Among various Multicast Routing Protocols being implemented and standardised IthACI selected the PIM-SM (Protocol Independent Multicast - Sparse Mode) because of the scalability limitations of the others. PIM-SM is deployed as the Multicast Routing Protocol in all three IthACI islands. However, each island

uses its own method to do the mapping of multicast streams onto layer 2 connections. Besides correct interworking of IP multicast and IP switching in the separate islands, the goal is to demonstrate the interoperability between Alcatels, NECs and Ciscos multicast solutions, either on layer 2 or layer 3.

The enhancements for mobility in a multicast environment follow the approaches supported by the IETF standards. The realisations in IthACI will conform to appropriate 'Requests for Comments' as long as they are available. Some ongoing work in IETF will also be considered and related Internet drafts are examined and applied accordingly. Two major schemes are possible to make the multicast service available to a mobile node. While away from home, a Mobile Node may be allowed to exercise its role either as if it were a logical entity in the home network or a logical entity in its current foreign network. In the first scheme, the Home Agent acts as the logical point of service. In the second scheme, the Mobile Node has to use a temporary co-located care-of-address to fulfil its role in the multicast scenarios.

The demand for higher throughput IP networks is huge. IthACI focuses on applying IP switching technologies to better satisfy this demand. Most of the major switch/router manufacturers have already announced MPLS support in their products, and will deliver by years end. IthACI also examines MPLS extensions and applicability to IP over DWDM (Dense Wave Division Multiplexing) layered networks.

IthACI project partly funded by ACTS started March 1998. The other consortium members are AlgoSystems S.A., University of Surrey, ABT, GMD-Fokus, CISCO Systems S.A., NEC Europe Ltd. and IMEC. Further information can be found at: <http://www.algo.com.gr/acts/ithaci/>

**Please contact:**

**Stelios Sartzetakis – ICS-FORTH**  
Tel: +30 81 39 1727  
E-mail: [stelios@ics.forth.gr](mailto:stelios@ics.forth.gr)

**Magda Chatzaki – ICS-FORTH**  
Tel: +30 81 39 1723  
E-mail: [magda@ics.forth.gr](mailto:magda@ics.forth.gr)

## Technology Interoperation in ATM Management

by Stelios Sartzetakis

**Providing cost-effective network survivability and availability; fast-responding, reliable fault detection and self-healing mechanisms, distributed dynamic routing functions with inherent load balancing capabilities, efficient VP layer design and dynamic network reconfiguration functions in an ATM network is a complicated task. Emerging technologies and standards are needed for the operation and management of ATM networks. The ACTS REFORM project used a number of these including: ATMF UNI and PNNI, OMG CORBA and Component Model and TINA NRA thus demonstrating the applicability, coexistence and interoperation of some of these technologies.**

Having the network provider's viewpoint in mind, who is offering QoS-based, switched, on-demand connectivity services, the REFORM project designed, implemented and tested a prototype system which provides the necessary means and functions for ensuring network resilience (availability & survivability) within acceptable levels, and in a cost effective manner. The REFORM system covers the configuration, performance and fault management functional areas. REFORM's control plane system (embedded in the network elements) hosts the required connection-oriented network layer functionality in addition to the REFORM-specific fast-responding fault detection, self-healing and QoS-based dynamic routing functions. The management plane system is concerned with the initial configuration and on-going dynamic management of the ATM VP layer. Specifically it hosts dynamic VPC bandwidth management, VP layer design and dynamic reconfiguration, fault management (filtering and correlation) functions as well as generic configuration

and network resource monitoring functions.

The entire REFORM system operates through a hierarchical model at different levels of time-scale and abstraction. The REFORM management system integrates the tasks of network planning and dimensioning with dynamic configuration, fault and performance management. The network planning and dimension functions lie at the heart of the REFORM system facilitating cost-effective network operation and enabling the implementation of the operator's business policy regarding service provisioning for the range of services being offered. Dynamic configuration, fault, and performance management are required for continuously optimising the performance of the network according to actual usage levels. Traditionally, these areas have been handled by disjoint systems for network planning, configuration management, performance management and fault management.

In our view, network management is much more than a data collection exercise for supporting configuration, fault and performance reports to be subsequently fed to network operators to determine the next course of action. REFORM considers network management functions as built-in, automated and intelligent facilities, which respond to changes in network conditions as and even before they happen. A human-orientated model of decision making may no longer be always viable for ensuring the cost-effective management of complex, multi-service networks. Network management functions need to actively interact with the network, exploiting and complementing the capabilities of the Network Elements (NEs), and therefore should be seen as an extension of embedded NE functionality.

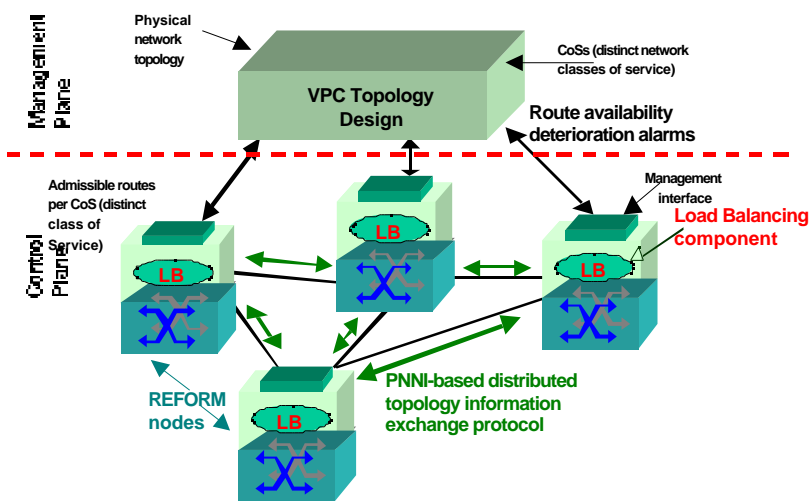
REFORM used networks built from multi-vendor NEs in its trials. As of today, implementations of the same signalling protocol by different vendors rarely fully inter-work and co-operate. Furthermore, the other REFORM control plane components were required to co-operate with the signalling components for collecting statistics and providing routing



information. This level of interaction with embedded capabilities is not feasible with today's commercially available NEs, as access to the internal functionality of the network equipment is very limited. To overcome these problems, the notion of

routing schemes. PNNI routing protocols may co-exist with other signalling systems, which support source node or hop-by-hop routing. However, some routing schemes, such as the one adopted by REFORM, may be of a hybrid nature combining the merits

CORBA was used as the ultimate integration means that allows transparent communication not only between distributed objects, but also between objects located in the same machine; eg between two components in the same network node. In the same way that CORBA insulates distributed applications from network details it also abstracts implementations from operating system peculiarities. We are in no doubt that future management frameworks will be based on distributed object technologies, with CORBA being a prime candidate. It is also likely that CORBA can be used in the control plane to support open, object-oriented network layer functionality, though this requires lightweight operations and mappings of its protocols over specific network technologies such as ATM AAL5. The use of the OMG Component Model in the REFORM control plane was a step in this direction.



**Routing in the REFORM system; a hybrid scheme involving centralised and distributed logic.**

the REFORM Node was introduced. The REFORM Node incorporates the control plane functions of the REFORM system, allowing vendor-independent interaction with the NEs and uniform signalling and routing functions throughout the network. This is achieved by isolating vendor-specific technology from the system components through an intermediate adaptation layer to the specific NE. An existing ATM Forum UNI 3.0 signalling stack implementation was modified in order to offer uniform signalling capabilities at both UNI and NNI reference points of the REFORM Node. However it should be noted that any NNI protocol eg B-ISUP, PNNI could have been used. Appropriate modifications according to the REFORM model for QoS provisioning had to be made.

The PNNI specification can be seen as containing two parts: the routing protocol and the signalling protocol. The PNNI routing protocol was used as the basis for dynamic routing in REFORM. The REFORM QoS-based routing algorithm runs over the PNNI topology information distribution mechanisms. The PNNI routing protocol specifications constitute a rich platform for building intelligent

of centralised and distributed routing. The PNNI specifications should not be seen as dictating the particular routing scheme to be applied, but rather as a framework enabling distributed exchange of topology information.

In the REFORM system, a hybrid routing scheme is adopted shown in the figure. The centralised part (residing in the management plane) defines the admissible routes per source-destination and CoS (distinct Class of Service) so that the QoS requirements of the CoSs and certain network-wide cost-effective criteria are met. The distributed part is responsible for influencing the routing decisions (from the many possible pre-defined admissible routes), according to actual traffic conditions, with the purpose to drive the network towards a load-balanced state. This hybrid routing scheme maintains the merits of dynamic routing and at the same time harnesses routing dynamicity to operate within the overall network operational policies with respect to QoS provisioning. Experimentation has shown that the distributed PNNI routing protocols can interoperate with overlying network management systems for QoS-based routing.

The REFORM system is one of the first attempts to validate the TINA NRA (Telecommunications Information Networking Architecture Network Resource Architecture) in a practical network management application encompassing configuration, performance and fault management aspects. REFORM field trials were conducted on several testbeds: the ACTS EXPERT testbed in Basel, Switzerland, and on ATM networks in Norway, Greece and Japan.

The REFORM system realisation showed that the evolution towards flexible telecommunications systems built from off-the-shelf intelligent software and hardware components is feasible in the near future. The project REFORM was part of the ACTS programme. The other partners were NTT, Telenor, OTE, Alcatel, UCL, NTUA, Skelton, IONA, with Algosystems as the co-ordinator. Further information can be found at: <http://olympus.algo.com.gr/acts/reform/public.htm>

**Please contact:**

**Stelios Sartzetakis – ICS-FORTH**  
**Tel: +30 81 39 1727**  
**E-mail: [stelios@ics.forth.gr](mailto:stelios@ics.forth.gr)**

# ATM Security Aspects

by Petia Todorova  
and Hartmut Brandt

**Due to the increased usage of ATM (Asynchronous Transfer Mode) for broadband backbones in local area and wide area networks and the introduction of public ATM networks, existing security concepts and protection mechanisms have to be revised. This article gives an overview of some aspects which are currently being studied at GMD Institute for Open Communication Systems.**

Today ATM is used, on one hand, to provide bandwidth for high quality multimedia applications (ATM to the desk) and, on the other hand, to serve as backbone for local and wide area IP (Internet Protocol) networks. With the advent of signalled public ATM networks it is now possible for companies with local networks and telephone systems in different locations to integrate voice and data communication and to connect their local systems through the public ATM. The complexity of ATM and the variety of protocols and approaches involves the danger of bypassing traditional protection mechanisms and exposes the networks to new threats.

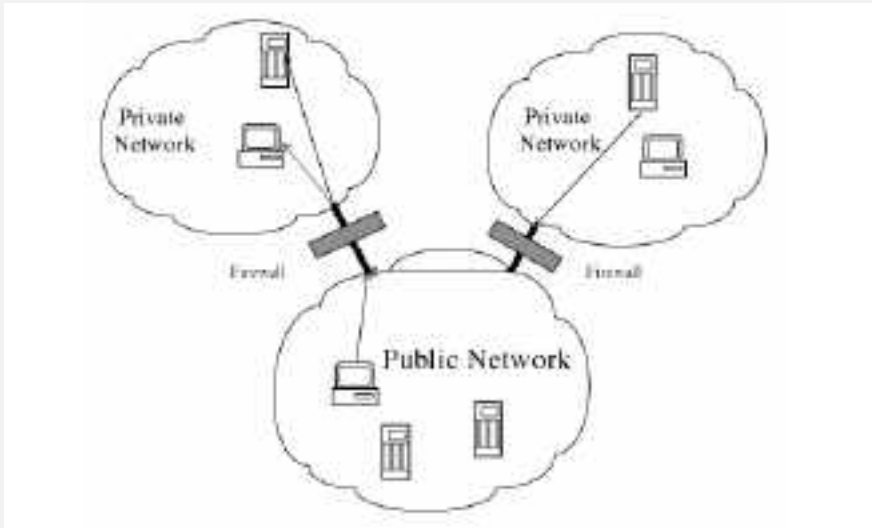
## ATM to the desk

High quality multimedia applications like video conferencing and video retrieval demand very high bandwidth and quality guarantees from the network. ATM fulfills these requirements, but makes it necessary to use native ATM applications, eg applications that build directly on ATM and ATM signalling. Connecting these systems to a public switched ATM network (which only makes video conferencing useful) opens these systems to threats from third parties. Malicious attackers are able to establish connections to the local ATM systems to get access to sensible information or to provide 'denial of

service' attacks. They may monitor and analyze traffic at intermediate nodes.

The ATM-Forum as one of the big ATM standardisation bodies is currently proposing an entire set of security mechanisms for ATM. These mechanisms require changes in switch

Additionally, the inherent complexity of ATM protocols makes it hard to predict possible threats. So, for example, it is possible to establish a direct ATM connection to CLIP clients, if its address is known. Because this connection bypasses any IP firewalls, other



**Interconnecting ATM islands through the public ATM Network.**

software and are partly based on hardware (for encryption and checksumming). For this reason these features will probably not become available in the near future from the major ATM vendors.

## ATM backbones

The main use of ATM in private companies is as high speed backbone for IP networks. Companies which have facilities in different locations may need to interconnect their local networks through the public network (see figure). With the advent of public switched ATM networks it is now possible to connect local CLIP (Classical IP over ATM), LANE (LAN emulation) or MPOA (Multiprotocol over ATM) islands more economically than by leased lines. Routing can be optimised and bandwidth resources need not to be pre-allocated.

Connecting a local IP network to a public ATM backbone involves an entire set of new threats to the IP network. Traditional level 3 protection mechanisms (IP firewalls) can be bypassed by attackers.

mechanisms are needed to prevent these kinds of attacks.

## Conclusion

The ATM security studies performed in the Center of Competence for Advanced Network Technologies and Systems (CATS) at GMD Institute for Open Communication Systems indicate that traditional IP security mechanisms do not apply to ATM networks. A long term solution is the implementation of the ATM-Forum security standard, a short term solution is the implementation of an ATM-Firewall, which is currently done. Important topics for future work include the implementation of user plane authentication, data integrity and security audit and alarm reporting functions.

For more information on CATS, see <http://www.fokus.gmd.de/research/cc/cats>

### Please contact:

**Petia Todorova – GMD**  
Tel: +49 30 3463 7251  
E-mail: [todorova@fokus.gmd.de](mailto:todorova@fokus.gmd.de)

**Hartmut Brandt – GMD**  
Tel: +49 30 3463 7352  
E-mail: [brandt@fokus.gmd.de](mailto:brandt@fokus.gmd.de)

# HOT – Home and Office Technologies

by Thomas Luckenbach

The realization of networks for the interconnection of sub-computer devices like light switches, sensors or actors can be done by means of a variety of networking technologies. The article will give an overview about some of these technologies for the SOHO (Small Office Home Office) market which are currently being studied in the I3C (Integration of Infra-/Intra-/Internet Communications) project at GMD Institute for Open Communication Systems.

The area of microelectronics is one of the fastest developing technical areas in the world. Microelectronics together with modern software technologies will become one of the driving forces behind the envisaged age of the information and knowledge society.

While there are 'only' about 200 million personal computers worldwide there are already about 12 billion microcontrollers installed within nearly all kind of electronic equipment. In the home and office environment these microcontrollers can be found within stoves, washers and dish washers, coffee machines, refrigerators, TV-sets/VCR's, video beamers, etc.

The interconnection of these devices by means of suitable network technologies will lead to a new type of networks termed 'Infranetworks'. These Infranetworks can be used for a variety of applications and services like facility management systems for home and office automation, theft and security control or personalized environments/person and location aware systems based on active badges and suitable sensors within the Infranet.

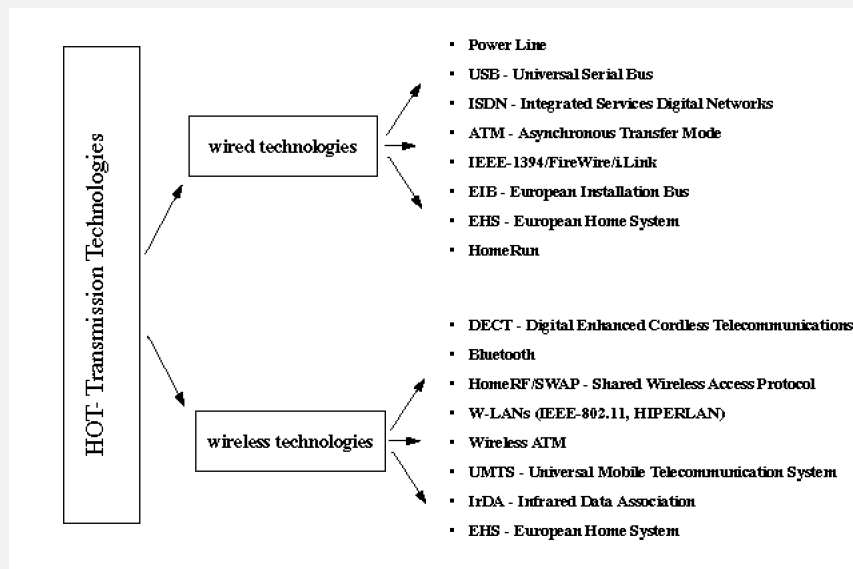
The realization of these Infranetworks can be done by means of different networking technologies, which can be classified into wired and wireless Infranet technologies (see figure).

## Wired Technologies

Power line systems are using the existing inhouse cabling installed for power distribution for the exchange of control information with a maximum bandwidth of a few Kbit/s. The Universal Serial Bus is designed for the interconnection of low

The European Home System EHS is designed as a 'home communication system' supporting different physical media like power line, coax, twisted pair and wireless infra-red. Most of the current systems are based on power line with 2.4Kbit/s.

HomeRun is currently being developed with the goal of providing a simple, high-speed Ethernet technology for the home environment based on existing telephone wires. It is initially planned for 1Mbit/s with upgrade plans to 10Mbit/s and



## Home and Office Networking Technologies.

bandwidth devices like keyboard, mouse, joystick or low resolution cameras. It operates over distances of up to 5 meter with a total bandwidth of 12Mbit/s.

IEEE-1394 is designed to interconnect high-end digital audio/video equipment like camcorders, digital VCR's or video printers as well as to interconnect personal computers (1394 home PC network). It supports data rates of up to 400Mbit/s at a distance of up to 4.5 meters. The predecessor of 1394 is FireWire, an initiative by Apple Comp. while Sony has started a branding initiative for 1394, called i.Link.

The European Installation Bus EIB is a free topology, decentralized system for home and building automation. It is based on a CSMA/CA technique via a 4-wire bus and provides 2.4Kbit/s bandwidth.

beyond and operates at distances of up to 150m.

The Local Operating Network (LON) is mainly based on the LonTalk communication protocol and the Neuron-Chip. LON currently supports different media like twisted pair, free topology networks, power line and network interfaces like PC cards.

## Wireless Technologies

The DECT standard (Digital European Cordless Telecommunication Systems) has been developed by ETSI (European Telecommunication Standards Institute) for in-building voice communications. It operates around 1.9 MHz with a total bandwidth of 1.152 Mbit/s (2 Mbit/s option available) and allows distances of

up to 300m between end systems and base station.

The Bluetooth proposal is targeted towards 'a global specification for Wireless Connectivity' and will operate initially with a bandwidth of 1Mb/s in the 2.4GHz ISM band.

The Shared Wireless Access Protocol by the HomeRF working group is based on a combination of DECT and W-LAN (802.11) technologies. It uses Time Division Multiple Access (TDMA) and Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) in the 2.4GHz band and provides 1Mbit/s or 2Mbit/s depending on the modulation scheme.

The IrDA (Infrared Data Association) specifications currently include protocols for point-to-point as well as multipoint data transfer via short distances (<10m) with 115Kbit/s, 4Mbit/s and even 16Mbit/s.

### Conclusion

I3C is a project running since the beginning of 1998 at GMD Institute for Open Communication Systems. The goal of the I3C project is the provision of a seamless networking infrastructure as a basis for a ubiquitous communication environment. In order to reach that goal current activities concentrate on the integration of LON (Local Operating Network), power line, IrDA, Inhouse ISDN (Integrated Services Digital Network), Ethernet, ATM (Asynchronous Transfer Mode) and ADSL (Asymmetric Digital Subscriber Line). First results indicate the suitability of standard Web-browsers to monitor and control Infranet devices via the Internet.

More information at:

<http://www.fokus.gmd.de/research/cc/cats/projects>.

Please contact:

**Thomas Luckenbach – GMD**

**Tel: +49 30 3463 7245**

**E-mail: [luckenbach@fokus.gmd.de](mailto:luckenbach@fokus.gmd.de)**

## Heterogeneous Inhouse Networking Environment

by Dirk Elias

**Starting with a kick-off meeting on 13 January 1999, the EURESCOM project P915 has been established.**

**Main goal of the project is to establish a pan-european demonstrator for a heterogeneous inhouse networking environment. Although the demonstrator will be distributed over Europe, it will be accessible over the EURESCOM web server (<http://www.eurescom.de>). The availability of new broadband access network technologies forces the European telcos to think about existing and future inhouse networks and communication applications for residential customers to meet their demands and to provide new services.**

Partners of the project are British Telecom, CSELT, Deutsche Telekom Berkom, France Telecom, Swisscom, Telefonica and Telenor. GMD Institute for Open Communication Systems is participating in the project for Deutsche Telekom Berkom.

Due to the increasing use of the ATM technology (Asynchronous Transfer Mode) as well as major enhancements of various forms of the Digital Subscriber Line (xDSL), the commercial introduction of integrated broadband services for the home and office environment is about to start. Although there is a clear view of the architecture of and supported services within access networks (AN) today, the integration of heterogeneous customer premises networks (CPN) and the related spectrum of application scenarios has not been researched sufficiently until now. A minor part of the work proposed by this Project is currently being discussed under the term SOHO (Small Office Home Office) networking technology.

The project HINE - Heterogeneous Inhouse Networking Environment - aims to investigate the basis for the use of a broad range of applications in a completely heterogeneous in-house networking environment as well as their integration, based on a CORBA platform (Common Object Request Broker Architecture), into public networks (Internet) via broadband access networks. Considered in-house networks will be Narrowband Integrated Services Digital Network (NISDN), Ethernet and ATM (Intranet) as well as networks for entertainment, automation and control applications (Infranet). Especially these Infranetworks may lead to a paradigm shift from traditional end systems like personal computers, telephones or fax systems as the mayor source of traffic (and business) to stand-alone or integrated micro-/nano-communicators in a home or office environment. To these networks various transmission technologies can be applied. Besides existing copper (telephony, power line, coax) infrastructures, new transmission technologies (wired & wireless, optical & electrical) have to be analysed for their suitability.

The direction of the HINE research activities is the extension of the ATM technology from a mainly public network operators' oriented exploitation towards new areas of operation and the seamless integration of already existing infrastructure for the provision of a global, ubiquitous communication environment.

Target environments today are mainly in the in-house, business and home area with wireline and wireless systems, including those, which have generally been out of the scope of global communication scenarios until today. In addition to 'classical' customer premises equipment (telephone, fax, personal computers, etc.), today one is surrounded by a huge amount of microelectronic controlled equipment (stereo/video products, home theatre, white goods, light control, burglar alarm, person identification/location systems, etc.). Today most of these systems are stand-alone by the means of forming small, not interconnected and proprietary islands.

Including these devices in an integrated communication environment using existing or easy-to-install networks (power-line, wireless, plastic optical fiber, etc.) will allow major enhancements to existing and the creation of completely new applications and services. There already appears to be a number of chip design companies and manufacturers developing single chip low cost solutions for connecting consumer devices to the Internet. The driving force for this development is unclear, but it may come from several directions. One example could be that competition amongst manufacturers may become so great that keeping contact with their products and customers becomes imperative.

**Please contact:**

Dirk Elias – GMD  
Tel: +49 30 3463 7342  
E-mail: elias@fokus.gmd.de

## ARNO: Algorithms for Radio Network Optimization

by Heinz Mühlenbein

**In the application-oriented project Algorithms for Radio Network Optimization (ARNO) the Breeder Genetic Algorithms are used to solve optimization problems in the area of digital mobile telecommunication. The optimization problems include the Site Positioning Problem (SPP) and the Frequency Assignment Problem (FAP).**

The SPP consists of the task to place antennas in the area where the mobile telecommunication network is developed. It is a very complex task, since many different factors influence the positioning of the antennas.

First of all, the developer must have information of the number and the distribution of the potential mobile

services users. Also, the area has to be known from a geographic point of view. Usually, this information is available in the form of digital terrain data. Antennas have to be placed in such a way that users from the whole area may use the mobile services. Another task of the SPP is to predict the number of needed frequencies for each antenna to allow for a maximum number of users to be served simultaneously.

After the SPP task has been completed, the next step is finding an assignment of frequencies to the antennas (the FAP problem). The number of possible frequencies is limited and always much smaller than the total number of frequencies required in such a network.

### Frequency Assignment Problem

The assignment of frequencies to the antennas must fulfill several requirements which reflect the following electromagnetic compatibility constraints:

- the co-site constraint: frequencies which are assigned to the same antenna have to respect a minimum distance with regard to the frequency spectrum
- the co-channel constraint: identical frequencies may be assigned but only to antennas which are a minimum distance from each other; otherwise, the frequencies would interfere with each other, making their use impossible
- the adjacent-channel constraint: same as the previous constraint, but for adjacent frequencies regarding the frequency spectrum.

Generally speaking, the FAP may be viewed as the task of assigning a number of available frequencies to a set of requesters, subject to a set of given constraints. The main goal of a FAP solution is to find an assignment that violates as few as possible constraints. Other usual goals are to find an assignment which uses a minimum number of different frequencies, or an assignment whose span (the difference between the smallest and the highest assigned frequency) is at a minimum.

The FAP is a NP-hard problem. The well known Graph Coloring Problem (GCP) is a special case of the FAP.

### Cooperation partners

Cooperation partners are France Telecom - Centre National d'Etudes des Telecommunications, University of Wales Cardiff - Department of Computer Science, GMD - German National Research Center for Information Technology, Laboratoire de Genie Informatique et d'Ingenierie de Production and Etude et Conseil en Techniques Informatiques Avancees.

**Please contact:**

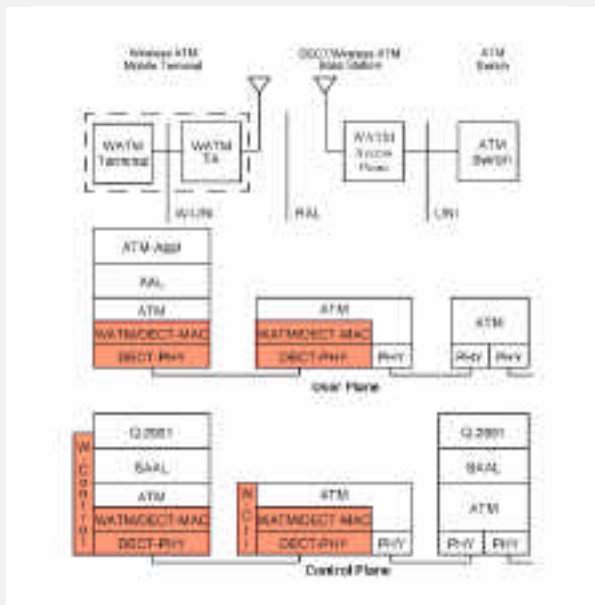
Heinz Mühlenbein – GMD  
Tel: +49 2241 14 2405  
E-mail: muehlenbein@gmd.de

## AMC - ATM based Wireless Mobile Computing

by Enrico Gregori  
and Thomas Luckenbach

**Within the framework of a scientific co-operation between GMD and CNR, at the end of 1997, GMD Institute for Open Communication Systems in Berlin and CNUCE-CNR in Pisa began to work on developing a wireless access technology to be integrated in the ATM technology.**

The objective of a collaboration now under way between FOKUS-GMD and CNUCE-CNR in the wireless mobile computing area is to implement an efficient connection of mobile devices capable of wireless communication with the World Wide Web. A wireless access technology (based on the DECT - Digital Enhanced Cordless Telecommunications - interface) has been deployed to connect mobile devices to backbone ATM networks. Up to now, the emphasis has been on mobility support in ATM infrastructure networks. However, besides the wireless ATM technology, the user also needs an efficient connection to Internet. Thus, we have also studied integration between the WATM protocol



**Figure 1: WATM/DECT protocol stack.**

stack and TCP/IP protocols. In the first stage of the project, we focused primarily on performance problems caused by the use of TCP as an end-to-end transport protocol when one end of the connection is a portable computer connected to our WATM network.

The project activity was distributed between the partners: FOKUS-GMD has been primarily involved in the prototype development; CNUCE-CNR has studied the integration of the WATM technology in Internet.

The WATM prototype developed by FOKUS consists of a base station controller and a terminal adaptor for Sun workstations. The base station and the terminal adaptor communicate in the 1.880 MHz radio frequency band as defined by ETSI within the DECT standard. The maximum available bandwidth on the DECT physical layer

is 1.152Kbit/s based on a TDMA (Time Division Multiple Access) transmission scheme.

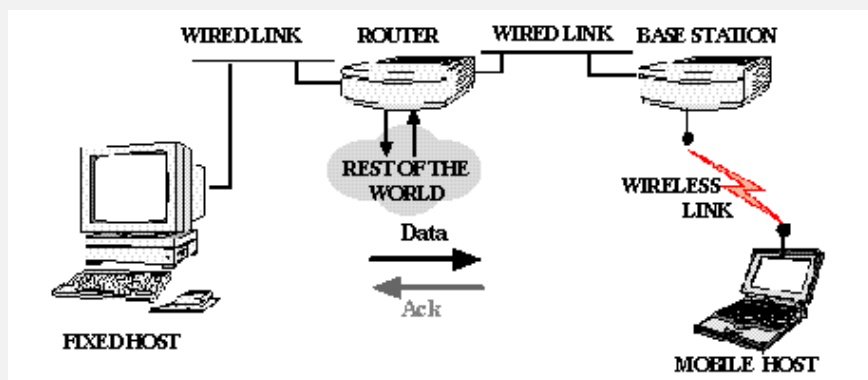
In the WATM prototype the access to the TDMA physical layer is achieved by means of a WATM/DECT Media Access Control (MAC) layer developed for the prototype. Instead of allocating fixed parts of the available bandwidth for the duration of individual connections - as is done by the DECT higher layer protocols - the MAC layer of the prototype supports the flexible allocation of bandwidth for different ATM connections on the basis of a 'request/grant' protocol between the base station and the end systems. Using this MAC protocol, it is possible to redistribute the actual available bandwidth every 10ms (which is the duration of a DECT time frame) on an on-demand basis.

The MAC protocol has been designed and implemented in order to support real-time and non-real-time voice/video and data services as defined by the ATM Forum. In the WATM/DECT prototype the protocol runs on a TMS320C40 processor with parts of the software realized within Xilinx PFGA's (Field Programmable Gate Arrays).

As far as the integration of the WATM technology into Internet is concerned CNUCE has developed a simulator of our environment to identify and eliminate sources of inefficiency. The environment under study is shown in Figure 1: a transmitter (the fixed host) sends data to a receiver (the mobile host) via Internet. The data transmitted by the fixed host cross the fixed network and arrive at the base station, and from there they are sent to the mobile host via a wireless connection. The TCP connections are modelled faithfully on the TCP protocol in accordance with the specifications of the RFCs and, where such specifications depend on the implementation, reference is made to the TCP/IP version supplied by Berkeley: 4.4 BSD-Lite. We also modelled the new protocol mechanism included in TCP Tahoe, Reno, SACK and New Reno.

Our simulator was used for an extensive performance evaluation of the TCP/WATM environment. The target of our analysis on TCP throughput was twofold: i) a study of the impact of user mobility, and, ii) a study of the impact of wireless links transmission error rates. Results obtained so far have shown that the extension of the TCP protocols to support nomadic computing are far from being optimal. We are currently investigating the use of a mechanism to split the TCP Connection (ie, Indirect TCP techniques) in order to prevent the wireless link from introducing severe throughput limitations.

**Please contact:**  
 Enrico Gregori – CNUCE-CNR  
 Tel.: +39 050 59 32 50  
 E-mail: e.gregori@cnuce.cnr.it  
 Thomas Luckenbach – GMD  
 Tel.: +49 30 3463 7245  
 E-mail: luckenbach@fokus.gmd.de




**Figure 2: Simulator environment.**

## Internet Telephony merges with the GSM Network

by Jori Paananen

VTT Information Technology has played a significant role in the development of the Nokia GSM Intranet Office product, which was launched at the CeBIT'98 exhibition. This product is the first in the world to combine Internet Protocol (IP) telephone systems and the worldwide GSM standard using normal GSM phones as terminal devices. By merging the corporate intranet and the GSM network, the Nokia product makes internal company calls extremely cheap or even totally free of charge.



**Nokia GSM Intranet Office.**

The Nokia GSM Intranet Office product uses standard GSM phones as terminal devices and GSM base stations connected to the corporate intranet. The system utilises the corporate intranet to route internal calls to other GSM phones or workstations (H.323 terminals) equipped with software for IP telephony. Outside the office the calls are routed normally through the public GSM network. The integration of the corporate intranet and the GSM network represents a quantum leap towards total mobility and reachability of company employees at a reasonable cost. On the other hand, new

cost-effective services will enable telecom operators to improve customer loyalty and increase the use of mobile phones.

The GSM Intranet Office prototype system demonstrated at CeBIT'98 was equipped with a Gatekeeper and an Intranet Location Register (ILR) developed by VTT. For the GSM phones registered to the system, the Gatekeeper is equivalent to a Mobile Switching Centre (MSC). The other system network elements are connected to the Gatekeeper over an IP-based intranet. The Gatekeeper checks the ILR to determine whether the phones involved in the call are authorised to use the system and whether they are located in the GSM Intranet Office area. If they are, it routes the GSM signalling and traffic inside the system. If either one is located outside the area, or is not registered as a user, it routes the call through the A-interface Gateway to the public GSM network.

The development of the GSM Intranet Office system towards the final product has continued throughout 1998. VTT Information Technology has participated in the development of various parts of the system.

For more information, see:  
<http://www.nokia.com/networks/16/gio/index.html>

**Please contact:**

**Jori Paananen – VTT Information Technology**  
**Tel: +358 9 456 5646**  
**E-mail: [jori.paananen@vtt.fi](mailto:jori.paananen@vtt.fi)**

## SFT – Smart File Transfer

by Markus Laakso

SFT, Smart File Transfer, is an application developed by VTT Information Technology for automatic transfer and distribution of files over mobile networks. The need for this emerged from the Finnish Maritime Administration as the Finnish icebreaker fleet had problems in transferring data to and from the icebreakers. Reliable data transfer is essential for optimal fleet management. Using SFT, the nine Finnish icebreakers get satellite images and ice charts daily and use this information in combination with information about the ship traffic to co-ordinate the assistance of ships through ice. The most cost-effective network in the areas of operation is still the NMT-450 analogue cellular network. In the Gulf of Finland GSM-data has been used and the newest multi-purpose icebreakers are equipped with satellite communication facilities.

Using SFT, the information is automatically distributed to the icebreakers. Controlled by its parameters, SFT scans predefined input directories and transfers any files found in these directories to the remote nodes and vice versa. As the connection is unreliable, and may break in the middle of a file transfer, SFT uses an acknowledging protocol with resendings if needed, to guarantee that a file, once queued for sending, is received once and only once at the receiving node.

SFT uses standard transfer protocols on top of which additional layers have been built. In the present implementation, RCP (Remote Copy) has been used for file transfer, but could be substituted by FTP or some secure transfer protocol if needed. In this way standard routers, with dial-on-demand features, can be used for establishing the connection.

The applications using the services provided by SFT, may have different types of transfer requirements: small files related to operational information should be transferred with as little delay as possible, but maps and satellite images are large and do not have as stringent delay requirements. SFT implements two types of queues: a priority queue and an ordinary queue. Files in the priority queue are sent before files in the ordinary queue. In addition, new priority files suspend the transfer of an ordinary file. This multiplexing makes it possible to send even large files as background transfer without blocking the channel for time sensitive priority files.

#### Automatic file compression

To save transfer time, the files are compressed before sending using lossless compression (Lempel - Ziv) and expanded at the reception. This enables simple formats to be used by different applications without the need of separate compression prior to sending. To guarantee that the file is correctly reassembled an MD5 checksum algorithm is used to check that the original file has been transferred unchanged.

Because SFT is used unattended 24 h/day 7days/week, special logfiles are automatically created, which gives information about the performance of the system as well as trace information in case of unexpected errors.

To summarise, SFT atomises the information exchange between the coordinating centre and the icebreakers providing reliable unattended file transfer over unreliable connections. Although the use of standard protocols does not necessarily maximise the net transfer rate, the robustness and modularity of this approach are considered to be more valuable for operational use than maximised channel capacity utilisation. ■

#### Please contact:

**Markus Laakso – VTT Information Technology**  
**Tel: +358 9 456 4568**  
**E-mail: Markus.Laakso@vtt.fi**

## Activities at SICS's Computer and Network Architecture Laboratory

by Bengt Ahlgren

**The focus of the Computer and Network Architecture (CNA) laboratory is in the networking technology and the computing platform for advanced distributed applications on the Internet. The networking technology ranges from the design of hardware for IP routers over core Internet protocols in the network and transport layers to how distributed applications use the network service.**

The CNA lab participated in the EU Esprit LTR project HIPPARCH together with researchers from INRIA Sophia-Antipolis, University College London, Uppsala University, Dassault Electronique (now part of Thompson-CSF) and University of Technology, Sydney. The project was finished at the end of August 1998 with excellent reviews. The project therefore received additional funding for exploiting the results further.

In HIPPARCH we have been investigating novel architectures for high performance communication protocols based on the Application Level Framing (ALF) and Integrated Layer Processing (ILP) concepts. The project also adopted the notion of "network conscious applications", that is, applications that are aware of varying network conditions and that adapt in one way or another to these conditions in order to optimize application performance. This notion is based on the belief that it is only the application itself that can make the best adaptation decision, for example, change audio encoding algorithm.

The results from the project include a protocol compiler supporting the development of ALF based applications, control algorithms for network consciousness, a multicast based reliable distribution protocol, a QoS model for network conscious applications and a network bandwidth probing tool.

We are also participating in the EU Esprit project Pegasus II. The other partners are University of Cambridge, University of Twente, University of Glasgow and Citrix Ltd (Cambridge, UK, formerly APM Ltd). The Pegasus II project develops an operating system, Nemesis, which is a generic platform for multimedia applications featuring effective user control over resource allocation policy.

A mobile ad-hoc network, or "MANET", is a self-configuring multihop wireless packet network, where all nodes participate in forwarding packets to their destinations. All nodes are also mobile, and thus free to move arbitrarily. Ad-hoc networking is a hot research area in the area of mobile communication. The research issues we are working on in this area are power management, routing protocols and addressing to provide Internet connectivity to ad-hoc networks.

In a previous project we have developed a dimensioning model for IP telephony. We are continuing this work with simulations and measurements in experimental networks to validate the model and to evaluate if an IP network with differentiated services can provide the necessary network quality of service for telephony.

For more information on the Computer and Network Architecture Laboratory, see <http://www.sics.se/cna/cna.html> ■

#### Please contact:

**Bengt Ahlgren – SICS**  
**Tel: +46 8 633 1562**  
**E-mail: bengta@sics.se**



# Monitoring and Displaying Traffic on the World Wide Web

by Evangelos P. Markatos and Athanasios E. Papathanasiou

Developed by the Computer Architecture and VLSI Group at the Institute of Computer Science - FORTH, Palantir is an application that can be used to display the origin, volume and type of the incoming requests of a web server. A good knowledge of the geographic distribution and access patterns of the clients creating these requests may indicate a more efficient way of serving them.

World-Wide Web traffic continues to increase at impressive rates. Busy web servers may get as many as several millions of hits (accesses) in a day. Accesses may originate from all over the world and may result in a 'rush hour' that lasts 24-hours per day. Web traffic will probably continue to increase as more people gain access and new applications (including commercial ones) are emerging. To meet the demands of this ever-increasing traffic, webmasters should design their web servers in such a way as to disseminate information (and sell or advertise products) effectively and reliably. A first step towards effective information dissemination is understanding a web server's client base, and reaching out to it. Palantir is a web traffic monitor and visualization tool that can be used to display the origin, volume and type of the incoming requests of a web server. A good knowledge of the geographic distribution and access patterns of the clients creating these requests may indicate a more efficient way of serving them.

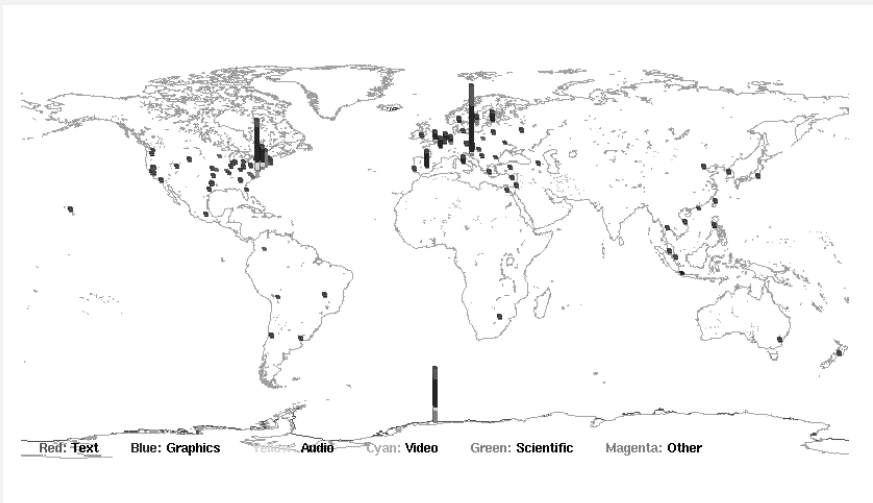
Palantir reads the log files of a web server and after categorizing the requests into six major types (text, graphics, audio, video, scientific, and other - that is any other kind of request) displays the origin

of the requests on a geographical landscape (in pictorial form). The geographical origin of the request is found by comparing its IP address to databases of information about geographical locations of the address. In the case of addresses registered in the US, the addresses are compared to registry databases and then that information is compared against a zipcode database to determine the geographic origination of the address. The type and magnitude of requests that

exceeded their time life (old requests) are deleted. In this way a visualization of the way that network traffic varies with time is given! Moreover, Palantir checks continually the log file to display new incoming requests at real-time.

To facilitate its use, Palantir provides an graphical interface with several functions most important of which are:

- aggregation: requests originating from several predefined (or user-defined)



**Palantir's Traffic Viewer Window. The picture shows the requests received by the Web Server of the Computer Architecture and VLSI Systems Division of ICS-FORTH of Crete at Heraklion during two consecutive days.**

originate from each region are shown in the map either stacked bars or as concentric circles. Different types of requests are easily identified by being displayed with different colors.

Palantir can animate the web traffic in static mode or in dynamic mode. In the first case, the requests, which have occurred during a specific period of time and are contained in the selected log files, are animated in the viewer. Each request remains displayed until the end of the simulation (it has an unlimited time life). Thus, the stacked bars (or the concentric circles) present the total amount of requests cumulatively (summary of traffic over a specified period). In the Dynamic mode, Palantir's viewer tries to capture the instant traffic of requests. Each request, contained in the log file, is considered to have a limited time life. As time passes, new requests are displayed on the viewer, while those that have

regions are aggregated into a single stacked bar (or concentric cycle)

- zooming: Palantir may zoom in a specific location in order to study more effectively the traffic that originates from a particular geographic region.
- filtering: Palantir provides two kinds of filters: Domain Filter and Request Filter. The Domain Filter is used to display only requests that come from a specified domain, while the Request Filter displays only those requests that ask for a specific kind of files.

Palantir is completely written in Java, mainly for portability reasons and may be accessed though its Home Page at ICS-FORTH:

<http://archvlsi.ics.forth.gr/OS/www.html>

A Palantir server is currently running at <http://sappho.ics.forth.gr:9000>, which may be freely used for downloading Log

files and viewing Network Traffic Visualizations.

Current plans for Palantir include visualizing a domain's outgoing requests. A simple way for achieving this is by monitoring outgoing TCP packets that have as a destination port the 80 port (default port for http requests) of a host. ■

#### Please contact

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

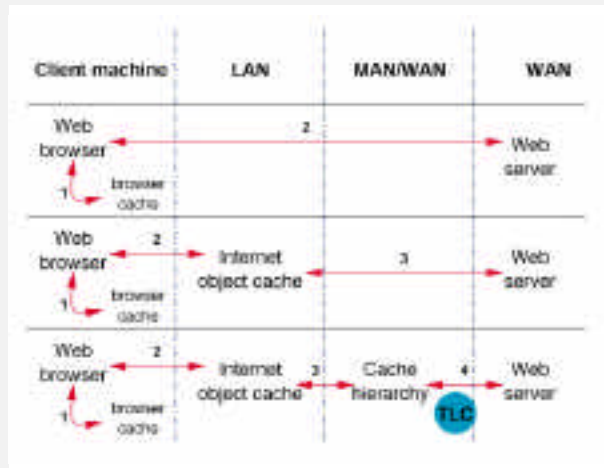
Athanasios E. Papathanasiou – ICS-FORTH  
Tel: +30 81 391437  
E-mail: papathan@ics.forth.gr

## Investigation of Internet Object Cache Performance Bottlenecks

by Bertold Kolics

Internet object caching is one of the biggest driving forces in the Internet related development area, because it saves bandwidth and reduces access latency for popular objects. Since 1995 about 50 cache servers have been installed in Hungarnet, the local academic and higher education network in Hungary. These servers form a loose, at most three-level hierarchy with one server at the top. It is in the nature of the currently used hierarchy that the topmost server perceives the largest load. Accordingly, the investigation focuses primarily on the performance of the top-level server. The primary goal of the project conducted at SZTAKI is to develop tools for identifying performance bottlenecks, by using extensive monitoring and log analysis at the investigated host. As the number of requests increases day after day, the prediction of performance bottlenecks in the future is also of interest.

Different levels of Internet object caching are illustrated in Fig. 1. In the case of Hungarnet a loose cache hierarchy has been deployed, with only one server at the top. TLC in a circle represents the place of the investigated host - the top-level cache. This server acts as a parent cache, ie, its child caches request objects from it if the requested URL cannot be found in the local cache. (In this context



Different levels of Internet object caching (the numbers on the arrows represent the order of communication, the columns show the type of the communication network).

an object is an entity identified by a unique URL.) Furthermore, the top-level server also changes information with other caches, ie, it can only retrieve an object from such a 'neighbour' cache if the requested URL can be found at the neighbour. Therefore, in this kind of hierarchy the top-level server perceives the largest load.

In most cases Squid is used as a proxy cache software and it follows that the host operating systems are UNIX-clones. Thus, the performance problems can have network-, Squid-, hardware- or UNIX-specific sources.

The monitoring of the caching host is performed in several ways: First, an SNMP client retrieves status data regularly from the caching software and the result is displayed in a graphical output produced by MRTG. Variables such as the number of requests, service time of hits (objects serviced from the cache) and misses (objects serviced using external sources), traffic volumes are tracked. Trends, unusual events can be easily identified with this type of monitoring. Second, a summary of the cache activities is created (with

Calamaris) on a daily basis by using the request log file of Squid. This log file includes, among other things, the timestamp of the request, the service time, the client's IP address, the requested URL, the size of the object. From this summary the cache manager can identify what the cache is primarily used for (eg which sites/domains are popular), how the cache behaves towards

different clients, what the gains of the cache hierarchy are, how much savings the cache achieved on the given day.

Third, the underlying OS is monitored. The top-level server runs on a recent Solaris operating system. OS specific tools as iostat, vmstat can be used for monitoring the performance of the system, but at this server we use 'The SE Performance Toolkit'. The toolkit makes use of the standard libraries shipped with Solaris to read kernel values. It evaluates these values in every 30 second and logs the unusual or unwanted system performance specific events (eg too much load on network interface card or on a specific disk). 30 second sampling time is long enough to eliminate bursts in the system load and small enough not to lose valuable information by taking averages for long periods of time.

By evaluating the information gained by this three-level cache analysis, the cache manager can identify possible bottlenecks in the cache infrastructure more accurately in an easier way.

The project itself is Hungarnet-specific, but the results and experiences of the

investigations can be useful for other cache operators as well. A cache management system with performance measuring/evaluation module is to be built in the near future.

Currently, two institutions from Hungary are involved in the project: SZTAKI, Network Department and Systems & Control Laboratory and the Technical University of Budapest, Department of Control Engineering and Information Technology. The only international partner of the project is the TF-CACHE project of TERENA. Co-operation with other ERCIM members in this field would be welcome.

**Please contact:**

**Bertold Kolics – SZTAKI**  
Tel: +36 1 349 7532  
E-mail: Bertold.Kolics@sztaki.hu

## The Changing Face of Network Management – A Hungarian Experience

by István Tétényi

**TEN-155 happened the other day. TEN-34 became history after 18 months. Research networking is still fashionable. People, organisations have accustomed to the results of the Internet revolution. We take ubiquitous, instant access granted. Is it not true? Just the opposite. The situation is only seemingly under control. Networking is like teenagers with long arms and legs and short clothes. Network management is much like the caring parent that tries to meet the ‘demands’. This article tries to summarise the results that we have achieved to manage the largest slice of the Hungarian Internet of several hundred institutions.**

The Network Services Department of SZTAKI (NSD) is a small ISP without

the commercial chore of selling the services on the open market. NSD has been managing Hungarnet’s IP network for many years. (Hungarnet is the Hungarian Academic and Research Network Organisation, <http://www.hungarnet.hu>) Some people define the task that we do as network operation. Well, it is true on the one hand but a bit less than the full image. It is true that we operate a network infrastructure. On the other hand, this is ever expanding in capacity and has got a wide range of technologies from ISDN to SDH.

Working in a conservative and under funded environment sometimes helps. You do not have to worry about the management approval of requested funds and the bureaucratic procedures of procurement. Although you will not end-up with glossy leaflets and boxes that are a bit outdated in the day of delivery. On the other hand, you might collect and develop a set of tools running on different platforms and providing the very services needed. Traditional network management meant that there were a few dedicated desktops from where the operators and the real techies could check the network status. The Internet revolution and the web explosion affected network management. We are somewhere halfway through. The big names (Microsoft, Cisco) are about to invade this territory. Till it happens, there is room for home-made solutions. Below you will find our approach and the results.

As in every modern organisation, the Intranet is the basic infrastructure component that supports groupwork. We (NSD) do have a loosely coupled set of tools put together in the recent years in order to manage the networks we operate. Naturally, the tools alone - without the people and procedures - cannot solve the tasks. What we use at the moment is a Network Management Intranet that was developed over the years. It has been refined and its functionality extended many times.

Network Management Intranet = several functions (tools) that are equally available for a group of people solving common tasks together.

A cornerstone of our approach is: there should not be any dedicated console for dedicated network management. We have to have tools that can be relocated in case of failures; we have to have a flexible environment where at least a group of people can share the same view of the network:

- Basic core tools: monitoring network reachability (pinger), monitoring service availability (servers alive), network events administration with ticketing (home-made ticketing software)
- Statistical tools: collecting interface traffic statistics (mrtg, monster), collection individual IP statistics (home-made collection tool)
- Extended monitoring tool: line database (home-made), interface database (home-made), network database (home-made), web based console for monitoring network reachability (home-made), service level monitoring (home-made)
- Reporting - monthly reporting: service level reporting (home-made), line overloads (home-made), traffic monitoring (home-made), ticket statistics (home-made)
- Special tools: monitoring special events (collisions, temperature, VC statistics), backup handling (home-made), paging (SMS) regularly and in case of special network events.

### Expected future developments

We plan to introduce netflow statistics collection this year to have even more precise information about traffic characteristics. Application monitoring will be extended. We try to strengthen the database background for the tools we use. Security issues have to be refined. We plan to migrate interface statistics collection and visualisation from monster to rrdtool (mrtg, cricket). We are looking forward to directory enhanced network management. We are open to co-operate with anyone who truly believes in vendor independent tools and network management.

**Please contact:**

**István Tétényi – SZTAKI**  
Tel: +36 1 349 7532  
E-mail: tetenyi@sztaki.hu

# The New CNR Research Area in Pisa chooses an ATM Local Area Network

by Marco Sommani

Starting in late spring of 1999, the 15 CNR institutes in Pisa will move to a new campus, known as the 'Area della Ricerca di San Cataldo'. The backbone of the local area network of the new campus has been designed using ATM LAN Emulation (LANE) and Multiprotocol over ATM (MPOA). The benefits of an ATM based solution appeared clear as recently as two years ago. Now, however, with the appearance on the market of powerful layer-3 switches with gigabit Ethernet interfaces, the advantages of ATM based LANs have become much less evident, to the point that their deployment is discouraged by many experts. This article claims that, even if layer-3 switching techniques over gigabit Ethernet can minimise the price/throughput ratio, an ATM-based LAN remains the best solution if multiple broadcast domains are scattered over the same physical infrastructure.

## Design Goals and Constraints

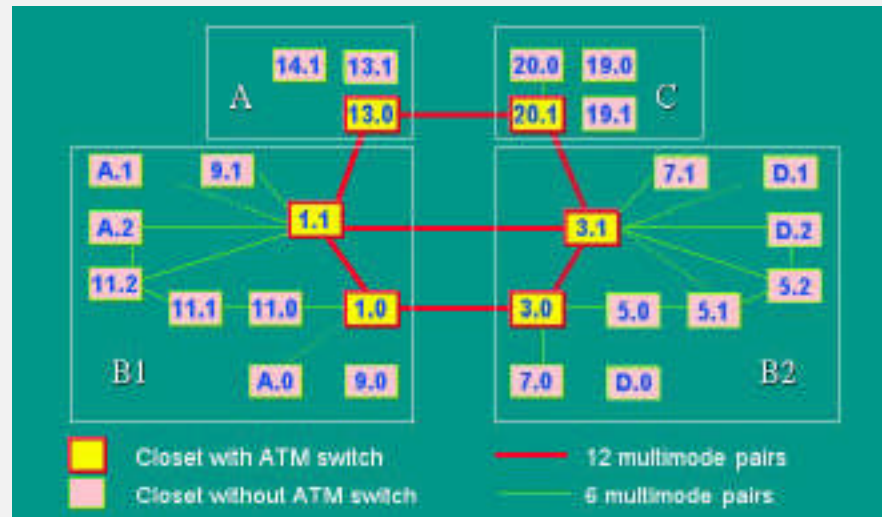
Currently the 15 CNR institutes in Pisa are scattered throughout the city. Typically, each institute has its own internal LAN and communicates with the rest of the world via a slow connection (64 Kb/s or basic rate ISDN in most cases, with a few 2 Mb/s exceptions). The router with the external Internet connection is located in a block hosting four institutes: CNUCE, IAT, IEI and IMC. The internal WAN topology is roughly a star centred at the site with the Internet connection.

Inter-institute traffic is quite low. Most communications are either internal within an institute or directed to remote Internet

destinations. Only IP is used for inter-institute traffic. Intra-institute traffic is also based on other protocols, most commonly AppleTalk, IPX and NETBEUI.

The LAN of the new campus must be able to meet two conflicting requirements: it must keep a clear separation between the broadcast domains of different institutes and enable

The ATM switches (IBM 8260s) are located in the wiring closets of the central hexagon (1.1 – 1.0 – 3.0 – 3.1 – 20.1 – 13.0). LAN emulation, MPOA, routing and bridging services are provided by equipment (IBM MSS) co-located with two ATM switches. All other wiring closets are just one fibre hop away from the nearest ATM switch. Every wiring closet contains Ethernet switches (IBM 8271s) with fibre ATM uplinks to the



The fiber optic backbone.

high-bandwidth and low-delay inter-institute communications. This is because (i) it has been considered undesirable to force network users to change their habits suddenly, and (ii) inter-institute co-operation is expected to become increasingly important and some institutes may be requested to host servers providing services to the whole campus.

## The Solution

A graphical description of the new LAN is available at <http://soi.cnr.it/~mario/piante/Default.html>.

The figure gives a simplified schema representing the fibre connections between wiring closets. The four big rectangles represent the main buildings in the campus (A, B1, B2 and C). Smaller rectangles represent the 27 wiring closets concentrating all copper cable coming from the rooms.

nearest ATM switch. Fibre connections between peripheral wiring closets shown in the figure are currently not used.

Typically one institute spans from two to six wiring closets and shares some closets with other institutes. Ethernet ports on the Ethernet switches are statically assigned to broadcast domains; every broadcast domain is associated with an emulated LAN (ELAN) in the ATM cloud. The Ethernet switches now in place are not MPOA-capable. As a result, all traffic between different broadcast domains must transit through one of the two IBM MSS acting as a router, resulting in an obvious bottleneck for all inter-institute traffic. This will not be a real problem in the beginning, when inter-institute traffic is expected to remain low; however, as soon as inter-institute traffic starts to increase, it will become necessary to add MPOA-capable edge devices such as the IBM 8371 to the infrastructure.

**Comparison with a Gigabit Ethernet Solution**

The alternative to the ATM solution is to place layer-2 and/or layer-3 switches in every closet and to interconnect them via point-to-point Ethernet links at appropriate speeds (100 Mb/s or 1000 Mb/s).

The non-ATM solution is appropriate when there is a close correspondence between broadcast domains and the physical topology of the network. In all other cases, the answer in a non-ATM world is to use Virtual LANs (VLANs), ie to use bridging to exchange traffic between hosts of the same broadcast domain attached to different switches and to "tag" frames sent over inter-switch links with their VLAN identifier.

With complex VLAN topologies, management may be a real nightmare. In fact, if all inter-switch links are enabled for transporting traffic for all the VLANs, then all broadcasts and Multicasts of the VLANs are propagated to all switches; otherwise, if inter-switch links are configured just for those VLANs that they should interconnect,

network management becomes too complex.

With complex VLAN topologies, it is also very difficult to configure inter-VLAN routing in a way that minimises the number of hops over inter-switch links. This problem is only partially reduced by special features that are common on many layer-3 switches, such as IP auto-learn.

**Conclusions**

When the LAN structure is simple, gigabit Ethernet and layer-3 switches can offer the best price/performance ratio. In all other cases, ATM LAN emulation, with its two-level addressing structure, provides a much better solution, because it optimises traffic patterns and simplifies network management.

In the last two years, there has been a tendency in the academic community to under-estimate the advantages that can be obtained with ATM. This probably depends on the fact that typically broadcast domains inside universities can easily be mapped over the physical infrastructure.

Strangely enough, ATM is also quite uncommon on commercial sites. A possible explanation is the fact that many network consultants prefer to promote what they already know (Ethernet) rather than spending time to learn something new (LAN emulation and MPOA).

If this trend continues, prices of Ethernet-based solutions will continue to fall and ATM prices will remain stable. Even worse, switch manufacturers might decide to stop investments in ATM products.

In order to stop this trend, every effort should be made to clarify ATM advantages to as wide an audience as possible.

**Please contact:**

**Marco Sommani – IAT-CNR**  
**Tel: +39 050 593 313**  
**E-mail: Marco.Sommani@iat.cnr.it**

**Order Form**

*If you wish to subscribe to ERCIM News or if you know of a colleague who would like to receive regular copies of ERCIM News, please fill in this form and we will add you/him/her to the mailing list. send or fax this form to:*

**ERCIM OFFICE, Domaine de Voluceau,  
 Rocquencourt  
 BP 105  
 F-78153 Le Chesnay Cedex  
 Fax: +33 1 3963 5052  
 E-mail: office@ercim.org**

(Data from this form will be held on a computer database. By giving your email address, you allow ERCIM to send you email. For subscribers from the UK: The Council for the Central Laboratory of the Research Councils is registered under the Data Protection Act and is authorised to hold details from this form on a computer database for any use which complies with the requirements of the Act.

Name: \_\_\_\_\_

Organisation/Company: \_\_\_\_\_

Address: \_\_\_\_\_

\_\_\_\_\_

Post Code: \_\_\_\_\_

City: \_\_\_\_\_

Country: \_\_\_\_\_

E-mail: \_\_\_\_\_

You can also subscribe to ERCIM News and order back copies by filling out the form at the ERCIM web site at [http://www.ercim.org/publication/Ercim\\_News/](http://www.ercim.org/publication/Ercim_News/)

# Software: Objective Zero Defect

by Maurice Mashaal

**To write programs that are at the same time complex and without logical faults or bugs, that is a feat, or more likely a utopian dream. In a day and age when computers are everywhere, one cannot help but be a little anxious at moments: isn't there any risk that the people who deal with my credit card will fleece me without my noticing it? Are all the electronic circuits of the plane I am about to board entirely compatible with one another? In general, the software at work has been validated and tested beforehand, all the more so if its malfunctioning can have grave consequences. Validation and testing take time and effort, however, which cost money. Furthermore, how can we be sure that everything was verified?**

This is the interest of research that aims at systematizing and automating such procedures. The problem arises foremost for distributed or parallel systems that comprise several communicating or interacting elements. For example, communication protocols must be certified to be locking free to avoid situations in which two different programs are waiting for each other to supply an information that is necessary for them to keep executing or in which several programs need to write simultaneously at the same memory address.

Such research is being carried out by several INRIA projects, with a significant theoretical aspect. In effect, one of the approaches consists in analyzing the semantics of the programming language to obtain formal models of the software under consideration. This step facilitates the design of the actual software verification tools. The verification can in fact be effected by trying to formally prove properties and good functioning criteria. However, most often verification

entails a recourse to brute force, that is to say to observe the behavior of the system for each set of data. This procedure is equivalent to exploring all the states that the system comes to occupy, which are modeled by a 'states graph'.

The problem is that the number of configurations to be explored is immense, even if the elements of the system can only occupy a finite number of states, which is often the case in practice. The difficulty thus resides in carrying out this exploration in an 'intelligent' way, by restricting attention to significant cases or to sub-systems and so on. This can avoid a comprehensive verification that would be way too costly in terms of time and memory. In addition, the goal is to design tools that can be used by non-specialists. This is far from being a negligible aspect if one has industrial applications in mind.

## MEIJE

One of the projects that are working in this direction is MEIJE, which is interested in particular in the Esterel language. Esterel is a programming language developed within MEIJE that is adapted to command systems that must react in real time to an external environment and function in a 'synchronous' fashion between parallel components. The applications considered by MEIJE thus naturally concern embedded systems for aircraft, automobiles, or material circuits among others. This is an area where software reliability is crucial. MEIJE is thus working in collaboration with Dassault Aviation on the modeling of embedded systems in aircraft for their verification. Similarly in micro-electronics, MEIJE is experimenting on the use of Esterel techniques for the synthesis and design of integrated circuits. From a formal viewpoint, such devices are equivalent to concurrent software systems.

## The VASY Initiative

The VASY initiative is contributing to the development of a compilation and verification toolbox called CADP and already distributed to more than 172 sites

worldwide. This environment is constructed around LOTOS, a language that makes it possible to describe protocols for 'asynchronous' distributed and parallel systems. LOTOS was chosen because it is the only language of this type that has the status of an international standard and whose semantics are rigorously defined.

Applications of the CADP toolbox range from databases to communication networks including cryptographic security protocols for electronic commerce. CADP is being currently used by Bull, in the framework of the Bull/INRIA Dyade GIE, to validate the multiprocessor architectures that the company is devising for their future high end servers. CADP was also used to detect an error in the 'Firewire' bus (IEEE standard 1394), a high speed serial bus for micro-computers that was adopted by the main computer manufacturers, software publishers and audiovisual equipment manufacturers. This error was practically impossible to catch by hand since it only occurs after a very definite sequence of some fifty operations!

## PAMPA

The PAMPA project is developing validation tools in the framework of the UML (Unified Modeling Language) object-oriented method, among others. This method is widely used in industry to write programs and model them. One of PAMPA's strong points is the joint development with CNRS laboratory Vérimag of a software called TGV (Test Generation with Verification) that automatically produces compliance tests. With this tool, PAMPA for example is collaborating with the European project Modistarc whose goal is to set up a testing methodology for embedded communication and management protocols in automobiles. A large part of the applications of verification tools however concerns the field of telecommunications, in which systems are distributed in essence. Together with four other INRIA teams, PAMPA is collaborating with Alcatel in the design of tools for formal operations within a

telecommunication software development line defined by Alcatel.

For more information see:

MEIJE: <http://www-sop.inria.fr/meije/meije-eng.html>

VASY: <http://www.inrialpes.fr/vasy>

PAMPA: <http://www.irisa.fr/pampa>

**Please contact:**

**Robert de Simone – INRIA (MEIJE)**

**Tel: + 33 6 92 38 7941**

**E-mail: Robert.De-Simone@inria.fr**

**Hubert Garavel – INRIA (VASY)**

**Tel: + 33 4 76 61 52 24**

**E-mail: Hubert.Garavel@inria.fr**

**Claude Jard – CNRS (PAMPA)**

**Tel: +33 2 99 84 71 93**

**E-mail: Claude.Jard@inria.fr**

## Critical Path in Fuzzy Network Analysis

by Milan Mareš

**The project ‘Fuzzy Set Theoretical Models of Cooperative Behaviour of Economic Subjects’, was solved in 1996-1998 at the Institute of Information Theory and Automation of Sciences of the Czech Republic. The project team has also included specialists from the Faculty of Civil Engineering of the Slovak Technical University. The widely conceived orientation of the grant has included, beside others, the construction and investigation of the co-ordination of activities in complex production and building processes which is usually described by means of network analysis and critical path method (CPM).**

The deterministic version of the CPM algorithms is well known for many years already and an effective software is commercially available. Anyhow, the determinism of the model can be questioned in many practical situations in which non-standard procedures are used in the modelled sequence of activities. The main goal of this part of research was to generalize the well known model and include uncertain and

vague phenomena. Namely, it was supposed that the assumed durations of particular activities are not exactly known in advance, and that the exclusive character of some of them and objectively existing uncertainty of the technological and economic environment influencing some others lead to only uncertain idea about the time needed for their realization. Moreover, this uncertainty is not based on statistical dispersion of possible values and, consequently, application of probabilistic methods is not effective. Such situation may appear, eg, if some activities are realized by new technologies going beyond the stabilized experience, if a non-standard object (atypical building, satellite or tanker) is produced, or if the complete production process modelled by the referred method is to last so long that some exactly unpredictable changes of technologies cannot be excluded.

In such case a fuzzy set theoretical model of durations of activities was used, fuzzy durations of paths and fuzzy float values were derived, and it was also shown that the concept of critical path itself is fuzzy. The described mathematical model of network analysis with vague components is based on the paradigm due to which the uncertainty of input data (durations of activities) generates also uncertainty of some properties (criticism and sub-criticism of paths, ordering of paths with respect to their durations) and output data (floats of paths and activities). The derivation of the characteristic of vagueness of the output data from the vagueness of the input ones is based on the processing of fuzzy numbers.

From the point of interpretation the most interesting output is the concept of fuzzy floats which indicates possible risks of delays with numerically structured possibilities. The fact that being critical becomes a fuzzy property of paths means that the critical paths form a fuzzy subset of the set of all paths to which any path belongs with some possibility. Also the vague durations of paths mean that for any pair of paths each of them can last longer than the other one with some possibility and it is meaningful to compute the floats regarding all possibly critical paths. The possibility with which

these floats reach negative values with respect to all other paths indicates the possibility with which a delay in realization can jeopardize the punctual fulfilment of given time-limits. These possibilities of negative floats and their structure offers an interesting information about the certainties and risks being related to the modelled production process and, in this sense, the fuzzy set theoretical analysis of paths and their durations offers a more relief and more finely structured information than the deterministic model. The model is based on some former works of the researchers and it probably opens the possibility of development regarding further elements of the network analysis.

The fuzzy set theoretical analysis of the critical path has illustrated one methodological discrepancy of the operations with fuzzy quantities based on the so called extension principle, namely the enormous increasing of uncertainty extent if the algebraic operation of summation is used. The paths in the network analysis which usually consist of numerous activities show this problem in an illustrative way. Such rapid increasing of uncertainty, moreover, does not fully correspond with the everyday practical experience. It appears to be useful to look for some alternative approach to the arithmetical processing of fuzzy numbers. Such approach could be based on the separation of the quantitative and fuzzy semantic component of a vague number and on their separate processing by arithmetical and fuzzy logical methods, respectively. One of the affiliated outcomes of the referred research is a suggestion of an alternative model of fuzzy quantities and their processing, respecting the heuristic principle formulated above.

Further information (reprints of publications) is available on request by e-mail.

**Please contact:**

**Milan Mareš – CRCIM**

**Tel: +420 2 688 4669**

**E-mail: mares@utia.cas.cz**

# EUROgatherer – Personalised Information Gathering System

by Costantino Thanos

The explosive growth of Internet and the World Wide Web makes an enormous amount of information widely available; however, this abundance of information can easily become a mixed blessing without methods to filter and control the potentially unlimited flux of information from sources to their receiving end-users (the well-known information overload syndrome). The EUROgatherer project aims at alleviating this problem by making information available to users in the appropriate form, amount and level of detail, and at the right time.

The EUROgatherer project aims at designing and implementing a system which provides a personalised information gathering service on the Internet. In particular, the EUROgatherer system will provide functionalities:

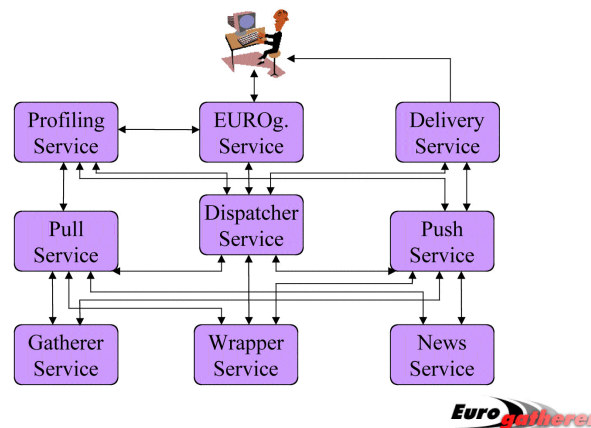
- to access a variety of information sources
- to create meaningful abstracts of the retrieved documents and classify them appropriately
- to acquire and retain user profiles and act upon one or more goals based on such profiles
- to support a relevance feedback mechanism.

A service-based architecture has been designed to meet the project requirements and is now under development. The set of EUROgatherer services (UserProfiling, Delivery, Pull, Push, and Wrapper services) provided by the system are visible to the user and accessible both separately or as a member of a cluster of services (for example, Pull + Profiling + Delivery or Push + Profiling + Delivery, etc.). One service can access

another creating a cluster of interoperable services. Thus, an Internet user can access a single EUROgatherer service or a composite cluster of services. Such a cluster should provide a more complete service (for example a personalized push or pull service).

The openness of the architecture is guaranteed by the fact that all the system components:

- share a common user profile model which complies with the standard P3P (see the Profiling service below)



A federated view of the EUROgatherer services.

- share a common category taxonomy
- use standard communication mechanisms/protocols.

The EUROgatherer system is composed of the following services:

- User Service - the main entry point to the system services. It provides the users with a global view of system functionality and supports them when accessing the system components.
- Profiling Service – responsible for the management (storage, maintenance and retrieval) of the user profiles.
- Dispatcher Service – responsible for dispatching information requests to other system services.
- Delivery Service – responsible for delivering the results of information requests to the users according to specified delivery modes.
- Pull Service – responsible for collecting information meeting user needs expressed both as long term needs (specified in a user profile) as well as short term needs (ad hoc queries, results on demand).

- Push Service – mainly responsible for filtering an information flow with respect to user long term information needs.
- Wrapper Service – responsible for retrieving HTML pages by querying online Web databases and search engines and transmitting them to the Pull and Push services.
- Gatherer Service – responsible for retrieving HTML pages by spidering the Web and transmitting them to the Pull and Push services.

- News Service – responsible for collecting continuous information flows transmitted by News Agencies or Usenet news discussion groups using several transmission means (satellite, internet, teletext, videotext) and transmitting them to the Push service.

The project adopts an advanced federated approach to the Internet services domain (see the figure). It will build on existing standards and contribute to the development of new ones.

The users of the EUROgatherer system can be divided in two categories:

- Casual users – casual users are generic Internet users who use the EUROgatherer system like any other Web search engine. They have at their disposal the AdHocQuery functionality, provided by the EUROgatherer User service, to issue queries to be executed by the system. This functionality satisfies the short term information needs of a user.
- Registered users – registered users are users who have specified their long term



information needs through the definition of a profile. This profile is stored in a database managed by the Profiling service. Retrieval of information specified in a user profile in pull/push mode can be triggered either directly by the user or by the dispatcher service on behalf of the user. Two interactive modes are available: 'Result on Demand' and 'Scheduled Queries'.

The EUROgatherer project is funded by the Telematics Information Engineering programme. It is conducted by a consortium composed of the following partners: Italia On Line, Italy; Xerox Research Centre Europe, France; CINET, Spain; Eurospider Information Technology, Switzerland; Consiglio Nazionale delle Ricerche – Istituto di Elaborazione dell'Informazione, Italy; University of Dortmund, Germany; Dublin City University, Ireland.

**Please contact:**

**Costantino Thanos – IEI-CNR  
EUROgatherer Project Coordinator  
Tel: +39 050 593 492  
E-mail: thanos@iei.pi.cnr.it**

## Cross-Language Information System Evaluation

by Carol Peters

**The development of methodologies and measures for system evaluation is an important topic in the information retrieval field. In recent years, the US National Institute of Standards and Technology (NIST) has done a great service to the IR community with the organisation of the Text Retrieval Conference (TREC) series in which various types of IR systems are evaluated by assessing their performance against set tasks (<http://trec.nist.gov>). Since 1997, TREC has included a track for cross-language information retrieval (CLIR) system evaluation.**

The interest shown in the cross-language track since its introduction into the TREC conferences in 1997 (TREC-6) shows the importance of this emerging area. There are many applications where information should be accessible to users regardless of language. In the global information society, situations when a user is faced with the task of querying a multilingual document collection are becoming increasingly common. Many users have some foreign language knowledge, but their proficiency may not be good enough to formulate queries to appropriately express their information needs. Such users will benefit enormously if they can enter their queries in their native language, because they are able to examine relevant documents in other languages even if they are untranslated. Monolingual users, on the other hand, can use translations aids to help them understand their search results. There is thus a growing demand for efficient cross-language query systems; the aim of the CLIR evaluation effort is to stimulate research in this area by providing a forum for the exchange of ideas and communication of results and to help create the necessary test collections for effective evaluation.

### The CLIR track at TREC-8

The main task of this year's CLIR track is to search documents in a multilingual pool containing news documents in four different languages: English, German, French, and Italian. The goal is to retrieve documents from all languages, rather than just a given pair. This is to encourage groups to work with as many languages as possible. For each topic (ie, query), a list of document identifiers is to be submitted which are ranked in decreasing order of the estimated relevance of the documents; the list will usually contain documents in all four languages. The main evaluation will be based on this list, although a simplified task will be considered with only English and a second language. In addition, there will be a special subtask. This task consists of a second data collection from the social science field. The rationale of this subtask is to study CLIR in a vertical domain (ie social science) where a German/English thesaurus is available.

The experience of the first CLIR track showed that it is difficult to produce topics in several languages at a single site. Thus, since TREC-7, a distributed approach to topic creation and results assessment has been chosen. There are currently four different sites, each located in an area where one of the topic languages is spoken natively:

- English: NIST, Gaithersburg, MD, USA
- French: University of Zurich, Switzerland
- German: Social Science Information Centre, Bonn/University of Koblenz, Germany
- Italian: IEI-CNR, Pisa, Italy

The track coordinator is Peter Schäuble, EIT, Zurich, Switzerland.

Each site is responsible for creating a certain number of topics and for translating the topics created by the other sites into the local language. In this way, an equivalent set of topics is created for each of the four languages. Sites are also responsible for the evaluation of results against the local language document collection.

Although the CLIR track is coordinated in Europe, the results are presented in the United States, at the TREC Conferences in November. It is thus probably not surprising that, so far, the participation has been dominated by US groups. At TREC-7, there were nine CLIR participants: 6 from North America - 5 US and 1 Canada - and just 3 from Europe (1 each from France, Switzerland and the Netherlands). This is clearly not representative of the European situation where multilingual issues are an everyday reality. We very much hope that this imbalance will be at least partially redressed this year at TREC-8.

### CLIR System Evaluation in 2000

Together with the TREC coordinators at NIST, we are now looking towards the future. It seems evident that issues regarding multilingual information retrieval system effectiveness will gain in importance in the next few years. We are thus seriously considering the possibility of setting up a European

forum specifically for cross-language and multilingual system assessment. In this way, we could extend the scope of the current CLIR evaluation activity to cover other important, related issues, such as cross-language evaluation for multi-modal systems, monolingual evaluation for non-English IR systems, etc. Such a European forum would continue to participate with NIST but all coordination activities would be centred in Europe, including the maintenance of a repository of tools, training data and evaluation suites for use by the scientific community for system evaluation activities.

For further information and instructions on how to participate in TREC-8, see <http://www-ir.inf.ethz.ch/TREC8/index.html>

**Please contact:**

**Martin Braschler** – EIT, Zurich  
Tel: +41 1 365 3052  
E-mail: [braschler@euospider.ch](mailto:braschler@euospider.ch)

## Presenting Swedish Language Technology

by Björn Gambäck

**The Information and Language Engineering group at SICS is currently setting up a web-based information service for human language technology in Sweden. The aim is to present the work done in the field, both by researchers and by industry, to the Swedish public.**

At the same time SICS will act as a national coordination point for the academic and industrial language processing communities, that is to channel information between the different groups working in Computational Linguistics and Speech in the country. The work on the web-site ([www.svenska.sics.se](http://www.svenska.sics.se)) is funded by the Swedish government through NUTEK,

the National Board for Industrial and Technical Development. Since the main target groups are within Sweden, the web-site will mainly be in Swedish ('svenska' actually means 'Swedish' in Swedish!); however, there will also be a part targeted at other European groups looking for Swedish research partners.

SICS has been chosen as the coordinator for this project since SICS's language group has a long-established reputation as being one of the leading groups in the field in the country. In fact, there are already two projects at SICS with somewhat similar goals: in one project SICS is building a platform for Swedish language processing and collecting resources in the field from the various other research groups, freely distributing the results to the research community. In another project SICS is creating a digital library by collecting language technology related research material from all the Nordic countries.

The work on the information service will mainly consist of two parts: on the one hand, administrating the actual information site; and on the other, the editorial work with creating and collecting the material which will appear on it. The latter part of the work will be carried out by an editorial board consisting of representatives from several different research groups, adding both geographical and competence spread. The information content must reflect a number of research areas, and include and interest commercial partners as well as the Swedish public. Thus, no single person will be able to keep the information accurate and up to date. The role of the editor in chief will be to co-ordinate a distributed, collective effort. The editorial board will be headed by Björn Gambäck and Jussi Karlgren, both from SICS.

For further details, please see <http://www.svenska.sics.se/>

**Please contact:**

**Björn Gambäck** – SICS  
Tel: +46 8 633 1535  
E-mail: [gamback@sics.se](mailto:gamback@sics.se)

## Mathematical Morphology and Image Processing

by Henk Heijmans

**In the processing and analysis of images it is important to be able to extract features, describe shapes and recognize patterns. Such tasks refer to geometrical concepts such as size, shape, and orientation. Mathematical morphology uses concepts from set theory, geometry and topology to analyze geometrical structures in an image. A substantial part of CWI's research theme Signals and Images is connected with multiresolution methods, based on the application of fractals, wavelets and morphology. One line of research explores the relationships between wavelets and morphological methods, aiming at a unified approach.**

The word 'morphology', stems from the Greek words morph and logos, meaning 'the study of forms'. The term is encountered in a number of scientific disciplines including biology and geography. In the context of image processing it is the name of a specific methodology designed for the analysis of the geometrical structure in an image. Mathematical morphology was invented in the early 1960s by Georges Matheron and Jean Serra who worked on the automatic analysis of images occurring in mineralogy and petrography. Meanwhile the method has found applications also in several other fields, including medical diagnostics, histology, industrial inspection, computer vision, and character recognition.

Mathematical morphology examines the geometrical structure of an image by probing it with small patterns, called 'structuring elements', of varying size and shape, just the way a blind man explores the world with his fingers or a stick. This procedure results in nonlinear image operators which are well-suited to

exploring geometrical and topological structures. A succession of such operators is applied to an image in order to make certain features apparent, distinguishing meaningful information from irrelevant distortions, by reducing it to a sort of caricature (skeletonization). For example, in optical character recognition one may transform the digital image of a symbol by reducing each connected component

level. From a frequency point of view, the resulting difference signals (known as detail signals) form a signal decomposition in terms of bandpass-filtered copies of the original signal. There is neurophysiological evidence that the human visual system indeed uses a similar kind of decomposition. This tool has been one of the most popular multiresolution schemes used in image

sequential filters, are essentially multiresolution techniques. Bearing this in mind, one may wonder what are the relationships between the existing linear (wavelets) and nonlinear (morphological) multiresolution approaches. Is it possible to unify both approaches into one mathematical framework? Does there exist such a thing as a 'morphological wavelet'? The somewhat ambitious goal of CWI's research effort is to answer such questions. First promising steps have been made by the author in collaboration with J. Goutsias (Johns Hopkins University).



**Decomposition of a grey-scale image by a morphological wavelet (max-lifting scheme). The original image is shown at the left and its decomposition into four subbands at the right.**

to a one-pixel-thick skeleton retaining the symbol's shape. Such a skeleton suffices for recognition and can be handled much more economically than the full symbol.

Experience in image processing and computer vision has shown that for a comprehensive understanding of a scene analysis at a broad range of resolution levels is necessary. The resulting multiresolution techniques (quadrees, pyramids, fractal imaging, scale-spaces, etc.) all have their merits and limitations. For example, fractals have been exploited with great success in image compression but to a much lesser extent for segmentation problems.

In the earliest multiresolution approaches to signal and image processing, the most popular way was to obtain a coarse level signal by subsampling a fine resolution signal, after linear smoothing, in order to remove high frequencies. A 'detail pyramid' can then be derived by subtracting from each level an interpolated version of the next coarser

processing and computer vision. CWI is currently investigating a general axiomatic pyramid scheme encompassing most existing linear as well as nonlinear (morphological) pyramids.

The emergence of wavelet techniques has considerably boosted the multiresolution approach. Unfortunately, application of wavelets to problems in image processing and computer vision is sometimes hindered by its linearity. Coarsening an image by means of linear operators may not be compatible with a natural coarsening of some image attribute of interest (shape of object, for example), and hence use of linear procedures may be inconsistent in such applications.

Mathematical morphology (nonlinear) is complementary to wavelets (linear) in that it considers images as geometrical objects rather than as elements of a linear (Hilbert) space. Many of the existing morphological techniques, such as granulometries, skeletons, and alternating

**Please contact:**

**Henk Heijmans – CWI**  
**Tel: +31 20 592 4057**  
**E-mail: Henk.Heijmans@cwi.nl**

## Virtual Environments support Interactive Engineering

by Gerold Wesche and Jürgen Wind

**Special-purpose Virtual Environments, designed to fit the needs of 3D visualization and simulation applications, enter the research and development departments of the car industry. They do not only support engineers in their daily work, they also provide a useful platform for presentations to managers, to decision makers or to the public. GMD pioneered this work introducing a Virtual Environment called 'Responsive Workbench™', a horizontal table-sized stereo projection system. Together with a tailor-made interface, it is the ideal workplace for 3D visualization and interactive steering applications.**

The CyberStage, a CAVE™-like (CAVE Automatic Virtual Environment) Virtual Environment with an integrated eight-

channel sound system and a sound floor, supports interior design studies, especially for car cabins. Interface design, ergonomics and inter-activity are the main research topics in this field of the GMD and its project partners or collaborators Daimler-Benz, BMW and Volkswagen. The visualization of the new Mercedes-Benz S-Class car cabin for the board of Daimler-Benz and the VW Sharan presentation in the

Workbench. The user specifies visualization parameters - such as insertion points for streamlines - using six degrees of freedom input devices.

These parameters are continuously sent from the Silicon Graphics computer to the computational server. In turn, the server rapidly computes high resolution visualization primitives, like streamlines or iso surfaces, and sends these back to

vector objects are computed in real time and emanate from the plane in the direction of the flow.

The system is running on a Responsive Workbench installation at Daimler-Benz, and it has been used for demonstrating the new Mercedes-Benz S-Class to members of the board of Daimler-Benz.

### Interactive Steering of crash tests

Simulations of structural mechanics like crash tests benefit from the coupling of parallel numerical computation power and the rendering capabilities of a high-end graphics workstation. Adapting the interface or the Responsive Workbench to this application, a system has been built that integrates numerical simulation, 3D data exploration and interactive steering together with Virtual Reality techniques.

The major objective of this project is to evaluate the potentials and obstacles of interactive simulations. A parallel version of the PAM-CRASH simulation package from the European Software Institute, evaluating BMW's crash tests, is running on a 12-processor SGI Onyx with InfiniteReality Graphics.

### Interior design studies


The CyberStage as an immersive Virtual Environment is very suitable for the exploration of interior car designs. Easy orientation in space, head tracking and interactivity give a nearly perfect illusion of depth and three-dimensionality.

In a presentation for Volkswagen, the interior design of the VW Sharan model was a convincing experience for our collaborators. At their lab they set up their own version of a CyberStage installation. ■

#### Please contact:

**Gerold Wesche – GMD**  
Tel: +49 2241 14 2708  
E-mail: gerold.wesche@gmd.de

**Jürgen Wind – GMD**  
Tel: +49 2241 14 2290  
E-mail: juergen.wind@gmd.de



VW Sharan (Data courtesy of Volkswagen).

CyberStage for Volkswagen are a proof-of-concept for the approaches developed to date.

3D visualization of Computational Fluid Dynamics, interactive steering of crash tests and interior design studies of car cabins are the main applications together that have been developed together with the car industry.

### Visualization of Computational Fluid Dynamics

A 3D visualization application for fluid dynamics in Virtual Environments has been developed by GMD and Daimler-Benz AG. It supports the evaluation and exploration of huge simulation data sets in a highly interactive way. Visualization of the simulation data and interaction occur through the Responsive

the Silicon Graphics, where they are displayed at high frame rates.

Scalar volumes such as temperature or pressure fields, can be visualized using interactive cutting planes controlled by hand. The color coded temperature visualization is updated every frame without any noticeable latency. The intersection computation of the plane with a 3D texture is implemented in hardware. This visualization technique is even usable for unsteady fields. However, during the animation of the sequence the textures have to be reloaded into texture memory because the majority of the memory is filled with the data produced by one timestep. This may reduce the smoothness of the animations. In a 3D environment, interactive cutting planes are also applicable for the visualization of velocity vectors. Small

# Virtual Reality for Oil and Gas Industry

by Bernd Fröhlich and Martin Göbel

**An International consortium has been established to demonstrate the impact of Virtual Reality in the Oil and Gas industry. This effort involves the development of prototype Virtual Reality applications specific to Oil and Gas activities.**

The Virtual Reality applications will enable consortia members in Oil and Gas industry to evaluate how working in Virtual Environments can most effectively be used in Oil and Gas industry. Consortia members are Arco, Amoco, BHP, EXXON, Landmark, Mobil, Saga, Schlumberger, Shell, Smedvig, and Statoil. These members represent both Oil and Gas companies and the software developers of Oil and Gas applications.

The consortium has selected the Virtual Environment Technology Lab of the University of Houston as the American research and development partner and GMD as the European Research and Development partner. Silicon Graphics is supporting this consortium.



**GeoScience data on the 2-sided Responsive Workbench.**

Both research partners have many years of experience operating projection based virtual environments such as the CAVE (CAVE Automatic Virtual Environment), a four side room size rear projection system and Responsive Workbenches table type rear projection systems using high end SGI Onyx2s.

GMD and the University of Houston will develop the demonstrator for visualization, sonification and interactive exploration within Virtual Environments. As a proof of concept the demonstrator will cover aspects like device independence, intuitive interaction with geoscience data, telecommunication, distributed and collaborative visualization as well as cross technology applications.


### Goals

The overall goals of the demonstrator development are to apply and evaluate Virtual Environment technology for Reservoir discovery, characterization and management to enable multi-disciplinary collaboration, to generate synoptic views of data from multiple sources, and to support well planning by multi disciplinary teams.


### Status and Future Work at GMD

A first demonstrator is available since late 1998. Important features of our system are the visualization of multiple data types, support for well planning, combined visualization and sonification of well log data, and multi-resolution techniques. For navigation and interaction with the system we engineered a new input device tailored to geo scientific data. This new device allows users to focus on their exploration task rather than on operating the computer.

The highest priorities for the next two years are support for remote collaboration, exploration of other navigation tools and integration of SGI's Volumizer volume rendering Application Program Interface. In particular, we are going to extend our current demonstrator to run under our distributed Avocado Virtual Reality



**The CUBE: A new input device for GeoScience applications.**



**Visualization of seismic data and horizons.**

system and we will support video conferencing capabilities within the distributed Virtual Environment application. Additionally we are looking into the integration of a large SensAble PHANToM force-feedback device to facilitate more intuitive interaction. ■

### Please contact:

**Bernd Fröhlich – GMD**  
 Tel: +49 2241 14 2038  
 E-mail: bernd.froehlich@gmd.de

**Martin Göbel – GMD**  
 Tel: +49 2241 14 2367  
 E-mail: martin.goebel@gmd.de

# Virtutech commercializes SimICS

by Peter Magnusson

**The SimICS simulation-based development tool was commercialized in 1998, with rights of development and exploitation being licensed to Virtutech AB, the first formal spin-off company in SICS' history.**

Software and hardware designers of high-end computer systems that play a key role in emerging complex information systems, ranging from database engines to telecom/datacom transaction-oriented systems, lack adequate debugging tools both as far as functionality and performance issues are concerned. Such tools would drastically reduce time-to-market thus improving competitiveness in a fundamental way.

Developed in the computer architecture group at SICS over a period of several years, SimICS is a "system level instruction set simulator". It is a specialized development tool for computer architects and operating system developers that models a target computer at the level of individual instructions. It offers unique benefits in high-end computer system design work beyond traditional tools and methods. For example, SimICS can run a realistic commercial workload in a fully virtualized system, allowing perfectly repeatable measurements that can track bugs that violate specifications regarding function as well as performance.

Virtutech specializes in high-end full-system simulation technology. In the brief period since opening its doors in the summer of 1998, Virtutech has signed a series of key strategic customer agreements, including with Sun Microsystems, Ericsson, and Hewlett-Packard.

"The simulator group that founded Virtutech is one of the world's foremost research groups in system level

simulation", according to David W. Yen, vice president and general manager of enterprise server products for Sun Microsystems, Inc. "Sun has worked with them for many years, and it is good to see they have discovered the commercial potential of their work and are beginning to capitalize on its worth."

SimICS is now marketed under the label Virtutech SimICS, and is the first product sold by Virtutech, in addition to specialized consulting services. Virtutech expects to have a profitable first year of operation, and 1999 is already "sold out" with customer projects.

"The tools and technology that we provide customers are unique in the commercial space," says Virtutech CEO Peter Magnusson. "Historically, this class of technology has been restricted to in-house proprietary efforts, deemed too sensitive and strategic to be outsourced. We have pioneered the concept of a specialized, independent company in the field. The technical and economical benefits of our specialization are significant. In today's competitive world, no stone remains unturned in improving quality and time to market. Our first customers are exemplary leading edge companies, early adopters of new technologies and processes."

## It began with Architecture and a Misunderstanding

The work leading up to SimICS originated in computer architecture and operating systems research at SICS in the late 1980s and early 1990s. In projects led by Seif Haridi, Erik Hagersten, and Andrzej Ciepilewski, it was observed that the existing simulation techniques failed to capture significant characteristics of 'real' commercial systems. A tradition of focusing on scientific workloads had led to a general feeling in the field of computer architecture that available benchmarks and techniques had systematic methodological weaknesses.

A computer architect uses simulation to provide input into his or her architecture models. Much of computer architecture involves careful optimization of common cases, perhaps to the detriment of less

common events. The relative frequencies of events in real uses of computers are thus important both to guide design choices, and to predict their relative impact. Because of limitations in existing simulation techniques in the late 1980s, computer architects in academia were restricted to models driven by (a) scientific, rather than commercial, programs, and (b) only so-called 'user-level' events, as opposed to operating system involvement. By contrast, there were indications that typical commercial workloads, such as large databases, behaved rather differently. In later years, this was confirmed by detailed studies performed by IBM, among others.

The first directions in the work leading to SimICS was based on a misunderstanding. An existing, ground breaking simulation environment developed by Robert Bedichek at the University of Washington was extended to support a multiprocessor system and to mimic real devices of a prototype research architecture, the SICS DDM. The first programmer on the project developed simulator device models sufficiently accurate to boot and run completely unmodified operating systems, including commercial on-board firmware. This strategy was chosen on the - incorrect - assumption that it was common practice. The strategy proved cumbersome, and eventually it became evident that SICS was the first open research facility to seriously undertake such a programming task. The early programmers on the project were confident that the task was not only feasible, but would take around six months or so. The obvious benefits of running a real, commercial workload, led a series of research managers at SICS to support continued work in that direction.

Some six calendar years, twenty man years, and several hundred thousand lines of code later, in 1997, the simulation group in the Computer and Network Architectures (CNA) group at SICS finally succeeded in the original goal: booting a commercial operating system (Solaris 2.6) on a simulated Sun Microsystems server (sun4m architecture). This was the first known occasion of an academic group running

an unmodified commercial operating system in a fully simulated environment.

In the meantime, the architecture research at SICS achieved great success, and significantly influenced commercial designs. SimICS has followed in the wake of SICS computer architecture results, and has now also taken the step into the commercial field.

The simulation group at SICS eventually grew to five people, all of whom became founding employees of Virtutech: Magnus Christensson, Fredrik Larsson, Peter Magnusson, Andreas Moestedt, and Bengt Werner. The board of directors of Virtutech include Bo Hedfors, an earlier chairman of the SICS Foundation, and Stig Larsson, the current chairman, as well as Per Stenström, professor of computer architecture at the Chalmers University of Technology.

Other than those already mentioned, significant contributions to the success of simulation work at SICS over the years were made by Torbjörn "tege" Granlund, Anders Landin, and David Samuelsson.

Virtutech is located in downtown Stockholm. For more information, see <http://www.virtutech.se> and <http://www.simics.com>

**Please contact:**

**Peter Magnusson – Virtutech AB**  
**Tel +46 8 690 0720**  
**E-mail: [info@virtutech.se](mailto:info@virtutech.se)**

## Gene-IT – A new INRIA Spin-off Company

**Jean-Jacques Codani**

**Gene-IT SA is a new start up of INRIA, active in the BioInformatics sector. Gene-IT's core technology is based on LASSAP (Large Scale Sequence comparison Package), a multiple algorithms, intra- and inter-databases high performance search engine, running on Unix platforms. It is a modern open system, based on a new approach to sequence comparison (genomic, proteic, EXTs, patterns,...). LASSAP provides biologists tools and methods that rigorously test their understanding of Genomics.**

Designed to fit the needs of large-scale sequence analysis, LASSAP has unique capabilities to analyse very large data sets of sequence and pattern. By handling annotation and sequence information simultaneously, LASSAP allows the scientist to ask complex queries on databases or clusters of databases.

Various post-analysis techniques can be performed on the results, from simple selections on alignment properties up to clustering analysis, allowing complex line of thought theories to be tested and answered immediately. The versatility of LASSAP allows different tasks to be performed by the scientist from a simple web based client environment, for example:

- improvements of sequence quality and automatic analysis within sequencing centres
- complex scanning on nucleic, proteic, ESTs and motifs databases
- building protein families
- reducing database redundancy, checking for patents by using proprietary algorithms.

From an application point of view, LASSAP allows to build or to find

families of proteins, among entire genomes and/or specific sets. This process relies on an in-depth study of gene duplications, which lead to the disclosure of a number of unexpected and singular regularities, which have a strong predictive value and can be viewed therefore as laws of Genomics.

As a consequence, this straightforward and rigorous methodology is a powerful automatic annotation tool, which favourably impacts on any Drug Discovery process. Moreover, once a family of protein has been chosen, LASSAP can perform complex tasks, mixing public domain and private databases (such as Geneseq and LifeSeq) to reduce the set of ESTs candidates, relative to this family. This methodology reduces drastically the amount of work usually needed to select EST clones as a genetic material for candidates.

LASSAP is progressive software, which does not pretend to provide the unique solution to the complex process of Drug Discovery. It has been devised to let scientist's design new strategies. They benefit from using a tool which lets them think in terms of science, and not in terms of informatics. LASSAP is not a fat genomic data management system designed to simplify navigation through 'mouse clicks', nor a single-use database: it is an automatic annotation tool, a data mining engine which provides new understanding and novel candidates.

Through its command line interface, LASSAP can be used by the BioInformatics team to implement sophisticated tasks. A Java front and back end is available which allows scientists to use the same functionality's using a Graphical User Interface. Its client/server model allows rapid deployment with minimum training to both BioInformatics and scientific personnel via the Intranet.

For more information, see:  
<http://www.gene-it.com/>

**Please contact:**

**Jean-Jacques Codani – Gene-IT S.A.**  
**Tel: +33 1 39 63 52 48**  
**E-Mail: [gene-it@inria.fr](mailto:gene-it@inria.fr)**

# Eidetica – A new CWI Spin-off Company

by Henk Nieland

Eidetica is the most recent of CWI's spin-off companies (seven were established in the last five years). It was created by researchers in the Interactive Information Engineering theme, combining expertise in linguistics and mathematical clustering of dependency networks for textual domains. Eidetica provides software solutions for information managers, responsible for, eg, an intranet, a trend watching service, or a library, enabling them to easily set up high quality browsing, searching and analytic services, thus increasing both their productivity and variety of competence.

Eidetica is a Twinning venture – a concept introduced by the ministry of Economic Affairs to support emerging ICT companies with excellent office facilities, business coaching and financial support in the form of seed or venture capital. Both CWI and Twinning have a share in Eidetica. The company maintains intensive contacts with the CWI research group, which focuses on automatic creation of user friendly access and control environments for large data and information collections based on content analysis.

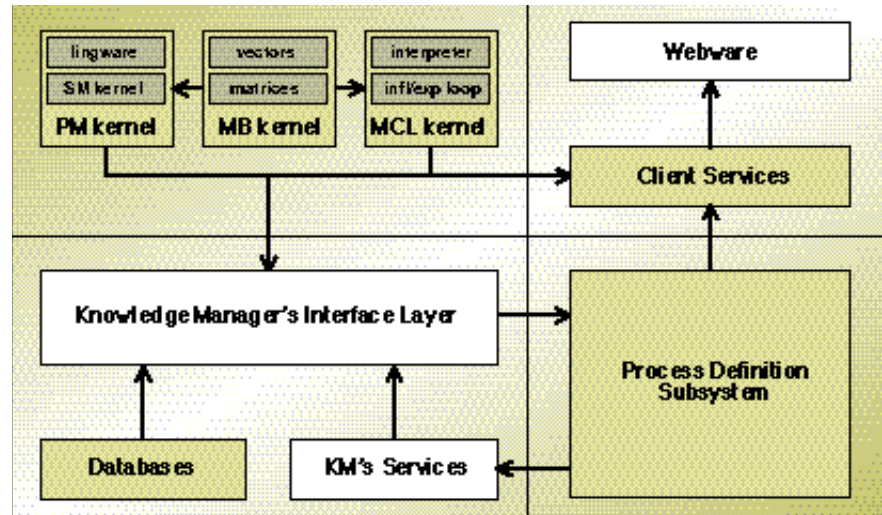
The unique quality of Eidetica's software architecture is the ability to extract quality controlled subject maps from large collections of documents. (The term 'eidetica' refers to the ability to recollect an earlier observation as a realistic image in one's mind.) Using these expert supervised subject maps as a basis, Eidetica's workbenches automatically generate intranet browsers and document delivery systems that reliably and efficiently guide their users through large data collections so as to obtain all – and only – the latest relevant information.

Eidetica's products range from a full text subject indexer through the metadata agent

architecture to the knowledge manager's workbench, each containing the document support browser technology.

In order to do more with electronic information collections than search for a few words and wrestle one's way through endless lists of not very relevant results,

with such questions and with outside contacts, routing the continuous intake of documents and information, deciding on who should read what, and archiving and arranging collections and providing access mechanisms for clients. It also offers expanding possibilities in the analysis of large document collections,



Architecture of the Eidetica system.

Eidetica's software enables browsing through document collections at a conceptual level. This is achieved through the following steps:

- merge the complete, miscellaneous document collection into a uniform format;
- read all documents to extract a dictionary of subjects
- create various 'maps' of the collection: which documents address which subjects, what authors write about what subjects, which subjects are connected to other subjects
- quality control: visualize the constructed maps and give the information manager tools to refine them
- use the subject maps to build browsing and querying interfaces that guide the user through the collection to find precisely the right information.

Eidetica's software helps improve efficiency and comfort in performing the information and knowledge manager's responsibilities, such as: collecting publications responding to a specific question, matching company expertise

such as personnel activity matching in a large organization, and trend watching in strategic literature.

At the core of Eidetica's system is a proprietary clustering method that can efficiently find groups of related information in collections about 1000 times larger than conventional methods, and advanced methods to extract subject keywords inside documents and titles.

In a nutshell, Eidetica's benefits are:

- simple and intuitive interfaces
- structured screen information
- better support of existing functionality
- more efficient performance of traditional tasks
- analysis methods
- open architecture.

For more information, see: <http://www.eidetica.com>

Please contact:

**Annius V. Groenink – CEO Eidetica**  
**Tel :+31 20 888 4126**  
**E-mail: avg@eidetica.com**




# The Computer makes Music in Pisa

by Leonello Tarabella

The keywords characterising the activity of the Computer Music laboratory of CNUCE-CNR, Pisa, are man-machine interaction and real time gesture control. The computer interacts with the artist to produce music in many different ways. In order to achieve this, a special language called Real-Time Concurrent Pascal Music (RTCPM) has been implemented. Using this language, a musical performance can be defined in terms of many procedures running simultaneously and interacting with the composition program during execution. As an evolution of RTCPM, a new language named GALileo (Graphic/ALgorithmic) which includes both visual programming and signal processing "facilities" is now under development.

In order to create and control interactive live computer music/graphics performances, the CNUCE Computer Music lab employs infra-red beams and video captured image processing technologies for the remote sensing of moving objects handled by performers or of gestures made by them. The most important devices that have been developed using this approach are:

- TwinTowers (in collaboration with IEL-CNR) – this system receives information from the movement of the performer's hands. It consists of two sets of four infra-red sensing devices that create two zones of the space (the vertical edges of two square-based parallelepipedon, or towers) in which an object can be detected in terms of distance and front and side rotations with respect to a reference frame. The system thus implements a kind of double aerial tri-dimensional joy-stick.
- Aerial-Painting-Hands – this system, developed by Giuseppe Scapellato and normally used by the artist Marco Cardini (see Ercim News No.35),



The Aerial Painting Hands system set up.



Leonello Tarabella at the Twin Towers.

detects the positions and movements of a performer's hands; the performer wears a pair of different coloured gloves and moves his/her hands in a CCD camera video range. The signal is sent to a computer which processes the images and recognises the x-y positions of both hands and the open/close status of the right hand. The computer then synthesizes in real-time coloured images projected on a large video screen according to the movements of the performer's hands. While painting in the air, the performer is also creating computer generated music in real-time. An advanced version of this system is now under development. The new system will recognise the posture, rotation and position in space of the hands, enabling the transmission of more sophisticated information and thus a greater expressiveness during the performance.

- UV-stick – this system consists of a short straight stick (lit by a UV lamp) and moved in the visual range of a CCD camera connected to a computer which



Massimo Magrini at the UV-stick.

recognises the 2-D position and 3-D rotations of the stick. The data acquired is used to manage sound synthesis algorithms and create music in real-time.

- Imaginary Piano – where the hands of a performer play in the air with no real keyboard in front of a camera, connected to a video digitizer card. A special application running on the computer recognizes the player's gestures, interacts with a RTCPM program and executes an external sampled piano sound.

The Computer Music Laboratory also participated in the EC ESPRIT project CATS (Computer Aided Theatrical Score) for theatre and cinema direction simulation.

For more information, see:  
<http://spcons.cnuce.cnr.it/music/cmd.html>

Please contact:

Leonello Tarabella – CNUCE-CNR  
 Tel: +39 050 593 276  
 E-mail: l.tarabella@cnuce.cnr.it

**SPONSORED BY ERCIM****SOFSEM'98 - The Successful Silver Jubilee**

by **Gabriela Andrejkova**  
and **Jan Vinar**

**SOFSEM (SOftware SEMinar) is a full seven-day international conference devoted to the theory and practice of software systems. Its aim is to foster cooperation among computer science professionals from academia and practice. Last year's Silver Jubilee Session was organized by the Slovak and Czech Societies for Computer Science and SRCIM (Slovak Research Consortium for Informatics and Mathematics).**

It was held from 21-27 November 1998 and attracted 137 participants from fifteen countries, about two thirds from Bohemia and Slovakia.

Sponsors played a very important role, the main sponsors were the Andersen Consulting, IBM, Digital Slovakia, SAP Slovensko, with other companies and organizations (ERCIM, Microsoft Slovakia, ORACLE Slovensko, VSZ Informatika) providing important contributions. Thanks to the support, equipment, and staff of Telenor Slovakia the participants enjoyed the Internet connection for the first time in the history of SOFSEM.

The scientific program of the seminar was divided into four principal areas: Distributed and Parallel Systems, Electronic Commerce, Electronic Documents and Digital Libraries, and Trends in Algorithms. There was also a Silver Jubilee Track devoted to major achievements of computer science. There were twenty invited lectures and eighteen contributed papers which were selected from among 48 contributions by the Program Committee chaired by Branislav Rován of the Comenius University of Bratislava.

**Invited Lectures**

Invited lectures were given under the following tracks:

- Major Achievements (The Silver Jubilee track)
- Parallel and Distributed Computing
- Electronic Commerce
- Electronic Documents and Digital Libraries
- Trends in Algorithms.

The Proceedings of the conference appeared as volume 1521 of the Lecture Notes in Computer Science of the Springer-Verlag.

The next conference, SOFSEM'99, will be held at Milovy in the Bohemian Highlands (approximately two hours by car from Brno). We hope that it will be as successful as it was on numerous previous occasions in the same venue.

For more information, see:  
SOFSEM '88 web page:  
<http://www.dcs.fmph.uniba.sk/~sofsem98>  
SOFSEM'99 web page:  
<http://www.ics.muni.cz/sofsem/sofsem.html>

**Please contact:**

**Gabriela Andrejkova and Jan Vinar – SRCIM**  
Tel: +421 95 6221128  
E-mail: [andrejk@kosice.upjs.sk](mailto:andrejk@kosice.upjs.sk)

**SPONSORED BY ERCIM****IWAN '99 – First International Working Conference on Active Networks**

**Berlin, 30 June-2 July 1999**

The First International Working Conference on Active Networks will take place in Berlin, 30 June - 2 July, 1999. It is being organised by the GMD Institute for Open Communication Systems and Hitachi Ltd, Japan. The International Federation for Information Processing (IFIP) is sponsoring the conference.

**Scope**

Recent advances in mobile code and agent technology allow for a new approach to the control and management of broadband networks and distributed

applications. Rather than running network elements with fixed functionality, protocol engines and applications can be dynamically deployed at routers, switches, base stations, and end systems. Enhanced functionality for better resource allocation and control of quality of service can be embedded into the network elements. This opens a new dimension in networking paving the way to future Active Networks.

**Tutorials (30 June)**

The first day is set aside for two tutorials. Prof. Jonathan Smith of the University of Pennsylvania will speak on Security:

- The Threat Model
- Techniques for Addressing Threats
- Architectural Implications.

The second speaker is yet to be confirmed.

**Conference (1-2 July)**

On the following two days topics such as:

- Foes or Allies: Active Networks and Open Interfaces
- Emerging Standards
- IP Services Creation

shall be presented and discussed. Jonathan Smith will hold a keynote lecture on 'Active Networks: Myths and Measurements'. Further Keynote addresses are to be determined.

A social evening is planned to allow for more informal exchanges among noted researchers from both academia and industry.

For details please refer to the website <http://www.fokus.gmd.de/iwan99>

**Please contact:**

**Stefan Covaci – GMD**  
Tel: +49 30 3463 7171  
E-mail: [iwan99@fokus.gmd.de](mailto:iwan99@fokus.gmd.de)

## SPONSORED BY ERCIM

## ICALP'99 – 26-th International Colloquium on Automata, Languages, and Programming

Prague, 11-15 July 1999

With respect to the positive experience from previous conferences, ICALP'99 is going to keep and to enhance the idea of a two-track conference. Similarly as in the case of two tracks of the journal Theoretical Computer Science, Track A of the conference will correspond to Algorithms, Automata, Complexity, and Games, while Track B to Logic, Semantics, and Theory of Programming.

For further information, see the conference website at:  
<http://www.uivt.cas.cz/ICALP99/>

**Please contact:**

**Jiri Wiedermann (Conference Chair) – Institute of Computer Science, Czech Academy of Sciences**  
Tel: +420 2 6605 2083  
E-mail: [wieder@uivt.cas.cz](mailto:wieder@uivt.cas.cz)

## CALL FOR CONTRIBUTIONS

## Tenth DELOS Workshop on Audio-Visual Digital Libraries

Santorini, Greece,  
24-25 June 1999

The digital library of the future will increasingly be multimedial and multilingual. The collaborative interaction of image, sound, speech and natural language understanding technology is needed in order to successfully populate, segment, index, and search digital audio-visual collections (such as film, music, or photo archives) with satisfactory recall and precision.

The Tenth Delos Workshop will provide a forum to bring together researchers and content providers interested in this topic. Invited speakers will present their views on the pressing issues in this area. The rest of the two-day meeting will be dedicated to the presentation of position papers and the description of on-going work. Ample time will be allowed for discussion of the issues raised.

**Topics**

Possible topics include: image/video/sound object representation; integrated models; metadata (models and management); storage management for continuous data; efficient indexing algorithms; cross-domain searching; content-based retrieval; query languages; automated summarization; multilingual access; human computer interfaces; concurrent handling of heterogeneous media.

**Submissions**

We invite contributions in the form of position papers (2 pages) to be sent to [carol@iei.pi.cnr.it](mailto:carol@iei.pi.cnr.it) by 15 May 1999. Notification of acceptance will be e-mailed by 1 June 1999. Each presenter will be asked to provide a technical note (5-6 pages) for the Proceedings by 15 June.

For more information, see  
<http://www.iei.pi.cnr.it/DELOS/WORKSHOP/workshop.htm>

**Please contact:**

**Tarina Ayazi – IEI-CNR**  
Tel: +39 050 593 401  
E-mail: [tarina@iei.pi.cnr.it](mailto:tarina@iei.pi.cnr.it)

## IN BRIEF

**CNR – The Italian government has approved a decree for the reorganisation of the National Research Council.** One of the main objectives is to bring the Italian scientific community into line with the rest of Europe by providing the single research institutions with greater financial and decisional autonomy and by setting up international groups of experts for each discipline in order to obtain a more objective evaluation of research programmes and results. The interaction between research and industry will be stimulated by encouraging the participation in private consortia and the formation of spin-off companies. The result should be a more efficient organisational structure with a greater impact on Italian competitiveness in the world's markets.


**INRIA – UDLR (Uni Directional Link Routing), a new protocol for satellite communication** has been set up by INRIA's RODEO research team. UDLR is a new routing protocol which solves seamlessly the problem of integrating in the Internet, at the IP routing level, unidirectional satellite links. UDLR uses a terrestrial link to transmit the routing informations from the receiver to the emitter. A consortium led by Alcatel, in partnership with INRIA, Eutelsat, Polycom and Softway has started experimenting this technology for parallel broadcasting in video-conferencing, videosevers, mirroring and feeding cache servers. Applications in connecting schools are also studied by another consortium to provide tools for cooperative work and tele-training. UDLR is being studied by the IETF relevant Working Group as a candidate

for an Internet standard on the subject, and might be definitely agreed upon in March 1999.

**INRIA – INRIA has launched INRIA-Transfert and I-Source Gestion**, two complementary subsidiaries to promote the creation of IT start-up companies. The mission of the INRIA-Transfer, a 100% INRIA subsidiary, is to accompany future entrepreneurs in the preliminary phases of company creation. I-Source Gestion (in partnership with other financial partners) has a budget of 100MF (19 Mio€) to invest as seed-money in start-up companies. Since 1984, INRIA has spun off 35 companies which generated a global turnover of almost 110 Mio€ in 1998. It has acquired thus an experience in this field, worth capitalizing upon.

**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. ERCIM News has a circulation of 7,000 copies.

ERCIM Central Office:  
 Domaine de Voluceau  
 Rocquencourt  
 B.P. 105  
 F-78153 Le Chesnay Cedex  
 FRANCE  
 E-mail: office@ercim.org

Central Editor:  
 Peter Kunz

E-mail:  
 peter.kunz@ercim.org

Telephone:  
 +33 1 3963 5040

Local Editors:  
 Gabriela Andrejkova (SRCIM)  
 Erzsébet Csuhaj-Varjú (SZTAKI)  
 Truls Gjestland (SINTEF)  
 Michal Haindl (CRCIM)  
 Kersti Hedman (SICS)  
 Bernard Hidoine (INRIA)  
 Pia-Maria Linden-Linna (VTT)  
 Siegfried Münch (GMD)  
 Henk Nieland (CWI)  
 Flemming Nielson (DANIT)  
 Carol Peters (CNR)  
 Martin Prime (CLRC)  
 Constantine Stephanidis (FORTH)

andrejk@kosice.upjs.sk  
 csuhaj@sztaki.hu  
 truls.gjestland@informatics.sintef.no  
 haindl@utia.cas.cz  
 kersti@sics.se  
 bernard.hidoine@inria.fr  
 pia-maria.linden-linna@vtt.fi  
 siegfried.muench@gmd.de  
 henkn@cwi.nl  
 fn@daimi.au.dk  
 carol@iei.pi.cnr.it  
 martin@inf.rl.ac.uk  
 cs@csi.forth.gr

**You can subscribe to ERCIM News free of charge by: sending e-mail to your local editor; posting paper mail to the address above or filling out the form at the ERCIM web site at <http://www.ercim.org/>**



**Central Laboratory of the Research Councils**

Rutherford Appleton Laboratory  
 Chilton, Didcot  
 GB-Oxon OX11 0QX

Tel: +44 123582 1900  
 Fax: +44 1235 44 5385  
<http://www.cclrc.ac.uk/>



**Centrum voor Wiskunde en Informatica**

Kruislaan 413  
 NL-1098 SJ  
 Amsterdam

Tel: +31 205929333  
 Fax: +31 20 592 4199  
<http://www.cwi.nl/>



**Consiglio Nazionale delle Ricerche**

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

Tel: +39 050 593 433  
 Fax: +39 050 554 342  
<http://www.iei.pi.cnr.it/>



**Czech Research Consortium for Informatics and Mathematics**

FI MU  
 Botanicka 68a  
 CZ-602 00 Brno

Tel: +420 2 6884669  
 Fax: +420 2 6884903  
<http://www.utia.cas.cz/>  
<http://ercim/home.html>



**Danish Consortium for Information Technology**

DANIT co/CIT  
 Aabogade 34  
 DK - 8200 Aarhus N

Tel: +45 8942 2440  
 Fax: +45 8942 2443  
<http://www.cit.dk/ERCIM/>



**Foundation for Research and Technology - Hellas**

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

Tel: +30 81 39 16 00  
 Fax: +30 81 39 16 01  
<http://www.ics.forth.gr/>



**GMD - Forschungszentrum Informationstechnik GmbH**

Schloß Birlinghoven  
 D-53754 Sankt Augustin

Tel: +49 2241 14 0  
 Fax: +49 2241 14 2889  
<http://www.gmd.de/>



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

B.P. 105  
 F-78153 Le Chesnay

Tel: +33 1 39 63 5511  
 Fax: +33 1 39 63 5330  
<http://www.inria.fr/>



**Swedish Institute of Computer Science**

Box 1263  
 S-164 29 Kista

Tel: +46 8 633 1500  
 Fax: +46 8 751 7230  
<http://www.sics.se/>



**Swiss Association for Research in Information Technology**

Dept. Informatik  
 ETH-Zentrum  
 CH-8092 Zürich

Tel: +41 1 632 72 41  
 Fax: +41 1 632 11 72  
<http://www-dbs.inf.ethz.ch/sarit/>



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

SINTEF Telecom & Informatics  
 N-7034 Trondheim

Tel: +47 73 59 30 00  
 Fax: +47 73 59 43 02  
<http://www.informatics.sintef.no/>



**Slovak Research Consortium for Informatics and Mathematics**

Dept. of Computer Science, Comenius University  
 Mlynska Dolina M  
 SK-84215 Bratislava

Tel: +421 7 726635  
 Fax: +421 7 727041



**Magyar Tudományok Akadémia - Számítástechnikai és Automatizálási Kutató Intézete**

P.O. Box 63  
 H-1518 Budapest

Tel: +36 1 4665644  
 Fax: +36 1 466 7503  
<http://www.sztaki.hu/>



**Technical Research Centre of Finland**

VTT Information Technology  
 P.O. Box 1200  
 FIN-02044 VTT

Tel: +358 9 456 6041  
 Fax: +358 9 456 6027  
<http://www.vtt.fi/>