ERCIM NEWS

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

Number 54, July 2003

Special: Applications and Service Platforms for the Mobile User

Contents

Cover photo
courtesy of VTT

Next issue:
October 2003

Special theme:
Machine Perception

8th Workshop on Formal Methods for Industrial Critical Systems

by Thomas Arts and Wan Fokkink

The 8th Workshop on Formal Methods for Industrial Critical Systems (FMICS'03) was co-located with the ERCIM meeting in Røros in Norway, June 5-7, 2003. The FMICS workshops aim to strengthen research on formal methods, and to promote the use of formal methods within industry. This year the workshop attracted over thirty participants.

The fourteen regular contributions to the workshop, which were selected from 25 submissions, dealt in large part with the application of formal methods in industrial case studies. Some notable application areas were smart cards, Java, MPEG and avionic systems. Two invited speakers were present: Reiner Hähnle from Chalmers University of Technology presented a paper entitled 'Integration of Informal and Formal Development of Object-Oriented Safety-Critical Software: A Case Study with the KeY System', and Werner Damm from the University of Oldenburg talked about 'Recent Trends in the Development of Safety-Critical Systems'.

The FMICS workshop series is organised by the ERCIM Working Group on Formal Methods for Industrial Critical Systems. On the occasion of the co-location with the ERCIM spring meetings, ten representatives from ERCIM institutes within this Working Group were invited to present an overview of current research on formal methods in their own research groups.

ERCIM President Gerard van Oortmerssen presents the working group award 2002 to Stefania Gnesi, coordinator of the Working Group on Formal Methods for Industrial Critical Systems.

Simona Orzan proudly shows her best paper award.

As a gesture of involvement, the European Association for the Study of Science and Technology handed out two best paper awards, one being given to Paul Ziemann and Martin Gogolla from the University of Bremen, and one to Stefan Blom and Simona Orzan from CWI.

The FMICS Working Group has managed to achieve a broad public visibility. This good interaction with the wider scientific community was recognised by the ERCIM board of directors, and as a result FMICS was, during the workshop dinner, rewarded with the ERCIM award for the most successful Working Group of 2002.

Link:
FMICS Working Group:
http://www.inrialpes.fr/vasy/fmics/index.html

ERCIM PhD Fellowship Programme

ERCIM offers 18-month fellowships in leading European information technology research centres. Fellowships are available for PhD-holders from all over the world. Next Deadline for application: 30 September 2003

Fellowships are of 18 months duration, generally spent in two ERCIM member institutes. The fellow will receive a monthly allowance which may vary depending on the country. In order to encourage the mobility, a member institute will not be eligible to host a candidate of the same nationality.

A preliminary list of topics includes but is not restricted to the scientific fields covered by the ERCIM Working Groups:

- Applications of Numerical Mathematics in Science
- Constraints Technology and Applications
- Control and System Theory
- Dependable Software-Intensive Systems
- Digital Libraries
- E-Learning
- Environmental Modelling
- Formal Methods
- Health and Information Technology
- Image and Video Understanding
- Matrix Computations and Statistics
- Soft Computing
- User Interfaces for All.

Conditions

Applicants must:

- have a PhD degree or be in the last year of the thesis work
- be fluent in English
- be discharged or get deferment from military service.
- start the grant before October 2004
- have completed their PhD before starting the grant.

How to apply

Applications must be submitted online.
A detailed description and the application form are available at:
http://www.ercim.org/fellowship/

GRIDs Expert Group

The European Commission's Information Society Directorate (INFSO), headed by Commissioner Erkii Liikanen and with Fabio Colosanti as Director-General, has a new Unit (F2) led by Wolfgang Boch concerned with Grids for Complex Problem Solving.

The F2 unit arose from the work done on GRIDs within Wolfgang Boch's previous Unit concerned with IST and the environment, including natural and man-made hazards and long-term environmental monitoring.

A workshop was called in January 2003 with some key speakers. It attracted more than 200 participants. Among the key speakers, three formed a 'core group of experts' to moderate discussions and to assist the Commission (in particular Roman Tirler who is the officer in charge of this activity) in pushing forward the research priorities and work programme. These three are ERCIM members: Thiery Priol (INRIA), Domenico Laforenza (CNR) and Keith Jeffery (CCLRC). Subsequently names were suggested to grow the group to the current 14 experts. This larger group includes Seif Haridi (SICS) and Ludek Matyska (CRCIM).

The group is working very productively. It has convened in two meetings to provide an expert's view on the five to seven year research priorities for the European context. The final report entitled 'Next Generation Grid(s) – European Grid Research 2005 - 2010' is now available on the CORDIS website at:
http://www.cordis.lu/ist/grids/index.htm
under the heading 'GRIDs Highlights'.

Expert Group report for downloading:
ftp://ftp.cordis.lu/pub/ist/docs/ngg_eg_final.pdf

Euro-Legal

News about legal information relating to Information Technology from European directives, and pan-European legal requirements and regulations.

Unsolicited Mail

Most of us have received at some time or other unsolicited emails advertising products and services. A number of Preference Services now operate within the Communication Sector, which enable consumers to withdraw their names and addresses from marketing circulation lists. Member organisations undertake to remove the addresses of people from their mailing lists who want this done. However, unscrupulous companies tend to trade names and addresses with other scam promoters, and those mailing lists may never be updated. The Stop Now Orders (EC Directive) Regulations 2001, which is a Code of Practice enforced by the Advertising Standards Authority and the Independent Committee for the Supervision of Standards of Telephone Information Services (ICSTIS), enables the Office of Fair Trading to bring an injunction against traders breaching regulations on misleading advertising. Action can now be taken across EU borders so scams operating across Europe can be pursued.

But even worse than the marketing scams themselves is the sheer volume of unsolicited commercial email (UCE), more commonly known as "spam", that can shift the cost of advertising from the advertiser to the recipient when with one email account a junk emailer can send a message to millions of recipients, turning that message into hundreds or possibly thousands of megabytes of data.

Vint Cerf, Senior Vice President, MCI is quoted as saying "Spamming is the scourge of electronic-mail and newsgroups on the Internet. It can seriously interfere with the operation of public services, to say nothing of the effect it may have on any individual's e-mail system.....Spammers are, in effect, taking resources away from users and service suppliers without compensation and without authorization." Whilst we would all agree that spamming is annoying, it is true to say that it also constitutes a theft of service.

One way of dealing with spam is to block the sites from which these UBE's (Unsolicited Bulk Emails) come. Pressure can be put on spammers by blocking access to websites that persistently advertise by UBE. However, spam is almost always forged. The address from which the message appears to have come is made to look like a real or innocent person.

Organisations such as EuroCAUCE (The European Coalition Against Unsolicited Commercial Email) have been set up to promote legislation which would outlaw UCE. European Directive 2002/58/EC on Privacy and Electronic Communication, first published in July 2002, will be implemented in Members States by 31 October 2003. The Directive states "For such forms of unsolicited communications for direct marketing, it is justified to require that prior explicit consent of the recipients is obtained before such communications are addressed to them". "When contact details are obtained, the customer should be informed about their further use for direct marketing in a clear and distinct manner, and be given the opportunity to refuse such usage". "Certain electronic mail systems allow subscribers to view the sender and subject line of an electronic mail, and also to delete the message without having to download the rest of the electronic mail's content or any attachments, thereby reducing costs which could arise from downloading unsolicited electronic mails or attachments. These modalities may continue to be useful in certain cases as an additional tool to the general obligations established in this Directive." For more details on the Directive see:
http://www.euro.cauce.org/en/amendments1a.html.

by Heather Weaver, CCLRC
Heather Weaver regrets that she is unable to reply personally to emails or telephone calls seeking legal advice.

W3C Approves Patent Policy

The World Wide Web Consortium has approved the W3C Patent Policy based on review by the W3C Advisory Committee and the public on 21 May 2003. Written by the Patent Policy Working Group, the policy received more support from the W3C members than any Recommendation in recent history. The W3C Patent Policy encourages royalty-free Web standards and aims to solve a specific problem – to reduce the threat of blocking patents on key components of Web infrastructure.

"On April 30, 1993, CERN decided to make the Web work available to any users, without charging a licensing fee. Ten years later, W3C's members have taken the lead and affirmed that, even after revolutionary changes in the world of information technology, the path to new markets is more smoothly paved with open standards," declared Tim Berners-Lee, W3C Director. "W3C now sets a benchmark for the pragmatic way to successfully develop royalty-free Web standards in the current patent environment."

W3C's Membership takes Lead in Supporting Royalty-Free Standards for the Web's Second Decade

Based on overwhelming support of the W3C Membership, consensus in the Patent Policy Working Group and support from interested members of the public, Tim Berners Lee has determined that the proposed Royalty-Free Patent Policy should become the Patent Policy for W3C. "The policy affirms and strengthens the basic business model that has driven innovation on the Web from its inception. The availability of an interoperable, unencumbered Web infrastructure provides an expanding foundation for innovative applications, profitable commerce, and the free flow of information and ideas on a commercial and non-commercial basis," he declared.

Beyond establishing a commitment to royalty-free standards, the policy provides W3C with:

- a stable, practical patent policy
- a clear licensing framework
- consistent disclosure obligations
- an exception handling process when problems arise.

After Three Years of Work, Diverse Parties Create a Common Path

The W3C Patent Policy Working Group was launched in October 1999, after a patent claim against P3P derailed the development of that technology. Based on a legal analysis of the claim, the threat was removed, and work successfully resumed. This and other experiences raised awareness of patent issues in the W3C membership. The Working Group was created in part to make more concrete how W3C Working Groups could successfully work on Web standards in the evolving patent climate.

"The Patent Policy represents what may be the most thorough effort to date in defining a basic patent policy for standard-setting," continued Berners-Lee.

The primary goal of the W3C Patent Policy is to enable W3C Recommendations to be implemented on a royalty-free basis. The policy also requires patent disclosure by W3C members when they are aware of patents that may be essential to the implementation of W3C Recommendations.

In simple terms, the Patent Policy provides that:

- All who participate in the development of a W3C Recommendation must agree to license essential claims (that is, patents that block interoperability) on a royalty-free basis.
- Under certain circumstances, Working Group participants may exclude specifically identified patent claims from the royalty-free commitment. These exclusions are required shortly after publication of the first public Working Draft, reducing the likelihood that surprise patents will jeopardize collective Working Group efforts.
- Patent disclosures are required from W3C Members and requested of anyone else who sees the technical drafts and has actual knowledge of patents that may be essential.
- Patent claims not available with terms consistent with the W3C Patent Policy will be addressed by a exception handling process.

Links:
W3C Patent Policy:
http://www.w3.org/Consortium/Patent-Policy-20030520.html

Director's Decision, W3C Patent Policy:
http://www.w3.org/2003/05/12-director-patent-decision-public.html

The W3C Semantic Tour

The World Wide Web Consortium (W3C) held a series of five one-day events across Europe. The tour, supported by the European Commission, through the Question-How project, started on 10 June in Rome, 12 June in London, 17 June in Munich, 19 June in Athens, and ended on 24 June in Brussels. The goal of this W3C Semantic Tour was to promote W3C technologies bringing to the Web the idea of having data defined and linked in a way that it can be used for more effective discovery, automation, integration, and reuse across various applications. The events have been organized by ERCIM, the European Host of W3C, and five of the European W3C Offices (Benelux, Germany and Austria, Greece, Italy, and United Kingdom and Ireland).

For the Web to reach its full potential, it must evolve into a Semantic Web, providing a universally accessible platform that allows data to be shared and processed by automated tools as well as by people. The Semantic Web is an initiative of the World Wide Web Consortium, and has been established to serve a leadership role, in both the design of specifications and the open, collaborative development of technology. ERCIM News has dedicated a special issue to the Semantic Web in October 2002 (No. 51).

Link:
http://www.w3.org/2003/03/semantic-tour.html

Continued Focus on the Mobile Web

W3C continues to focus on the mobile web. Mobile versions of the vector grahpics format SVG (Basic and Tiny) were issued as W3C Recommendations early this year. Moreover, the W3C Device Independence Working Group has just released the first public Working Draft of Core Presentation Characteristics. The group is building a basis for adapting content to device presentation capabilities. This draft provides a common set of property or attribute definitions that may be reused in future vocabularies. Finally, the Multimodal Interaction Working Group released an update to the W3C Multimodal Interaction Framework W3C Note. The framework identifies the major components for multimodal systems. The group is writing specifications to extend the Web user interface to offer input and output choices 'anywhere, on any device, anytime.'

Link:
http://www.w3.org/2001/di/Activity

Web Services at W3C

The World Wide Web is more and more used for application-to-application communication. The programmatic interfaces made available are referred to as 'Web services'. W3C's goal is to build an open, extensible technical foundation for the Web, and Web Services is one part of this foundation. W3C has five Working Groups/Task Forces focusing on Web services technologies, and at least another five which have produced the foundation for all other Web services work. With the recent completion of the SOAP 1.2 W3C Recommendation, over 19 technical specifications are under development at W3C within the Web Services Working Groups and Task Forces.

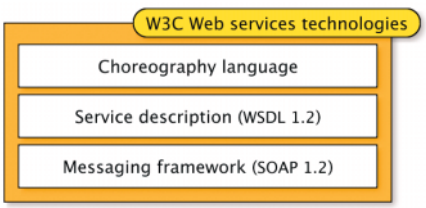

The first building block for Web services, important to all organizations interested in Web Services, is being produced by W3C's XML Protocol Working Group – SOAP Version 1.2. SOAP is a central piece of the Web services architecture, enabling complex application-to-application communication. SOAP Version 1.2, advanced to W3C Recommendation status on 24 June 2003, is a robust and extensible XML-based messaging framework based on SOAP/1.1, with full support for W3C Recommendations, support for multiple protocols (HTTP/1.1 as well as others).

The Web Services Activity is composed of three other Working Groups, addressing the overall architecture of Web services, their description and their composition. The work in Web Services Activity rests in turn on other work done in W3C. The XML Activity is responsible for the XML specification, XML Schema, XPath, XQuery, and other relevant specifications. Work in our Technology and Society Domain has defined specifications for encrypting and digitally signing XML documents or parts of XML documents, which answer important security needs in the Web services area. Other groups at W3C are connected to these groups, either through requisite coordination, or through task forces. For example, the W3C Internationalization Web Services Task Force, which looks at internationalization issues and how they impact Web services applications and infrastructure. These technologies serve as the basis for many other specifications, such as the WS-Security Specification first produced by IBM and Microsoft and now being worked on further in OASIS.

Links:
Homepage for W3C's WS Activity:
http://www.w3.org/2002/ws/

Web Services published Technical Reports:
http://www.w3.org/TR/tr-activity.html#WebServicesActivity

SOAP 1.2 FAQ
http://www.w3.org/2003/06/soap12faq.html

From SOAP/1.1 to SOAP Version 1.2 in 9 points
http://www.w3.org/2003/06/soap11-soap12.html

The Working Groups and Task Forces at a glance:
XML Protocol Working Group
http://www.w3.org/2000/xp/Group/
All work on SOAP 1.2 happens here.

Web Services Architecture Working Group
http://www.w3.org/2002/ws/arch/
This group identifies the overall architecture, identifies missing pieces, and determines whether new work is needed or existing work (inside or outside of W3C) meets needs.

Web Services Description Working Group.
http://www.w3.org/2002/ws/desc/
This group is at work on WSDL 1.2, the Web Services Description Language, based on WSDL 1.1..

Web Services Choreography Working Group.
http://www.w3.org/2002/ws/chor/
This group is developing a language describing the composition of Web services.

Internationalization Web Services Task Force
http://www.w3.org/International/ws/
Internationalization of Web services allows applications to function across national, linguistic, or cultural boundaries.

Last W3C Recommendations

- XPointer element() Scheme
- XPointer Framework
- XPointer xmlns() Scheme
- SOAP Version 1.2 Part 0: Primer
- SOAP Version 1.2 Part 1: Messaging Framework
- SOAP Version 1.2 Part 2: Adjuncts
- SOAP Version 1.2 Specification Assertions and Test Collection

An exhaustive list of all W3C Technical Reports:
http://www.w3.org/TR/

ERCIM is the European host of W3C.

Applications and Service Platforms for the Mobile User — Introduction

by John Krogstie

The most recent wave in information systems development is the introduction of networked palmtop computers, made possible by the continued miniaturisation of electronic circuits and the enormous resources spent by industrialised countries on deploying a ubiquitously available broadband wireless communication infrastructure.

Because of the mobility of users of these devices and the characteristics of wireless communication, the operating environment for any accompanying services will be much more dynamic than is typical for traditional distributed systems. In order to retain system quality and usability under such circumstances, R&D problems within a number of areas must be solved.

User-orientation and Personalisation

Mobile information systems often address a wide user group, which means that user interfaces should feature prominently and early in the design process and often need to be very simple. The resulting user interfaces cannot assume a prior acquaintance with computers, and input and output facilities may be severely restricted (no keyboard, small screen-size etc) or based on new modalities (speech recognition and synthesis etc). This means that individualisation of mobile information systems becomes increasingly important, both at the individual level, where user-interface details such as commands and screen layout are tailored to personal preferences and hardware, and at the work level, where functions are tailored to fit the user's preferred work processes. The latter case can also include the support of novel workplace designs and methods of work organisation that enable collaboration of multi-location and mobile workers.

Individualisation means information systems that are able both to automatically adapt themselves to the preferences of the user and to be explicitly tailored by users through a specific user interface. The main goal is to achieve usability of the applications on all possible interfaces, based on adaptation to the different physical devices. This calls for intelligent, adaptive and self-configuring services that enable automatic context-sensitivity, user profiling and personalisation in a trusted and secure environment, as well as multilingual and multi-cultural presentation, and multiple modes of interaction.

Technological Aspects including Convergence and Multi-Channel Support

Mobile devices have severely limited processing, memory and communication capacities compared to other kinds of computers. Performance considerations therefore become increasingly important in the first design steps. Analytical predictive methods are necessary in order to assess a large number of alternatives during the design of mobile information systems. Mobile information systems also pose new challenges to achieving information systems dependability. The drive toward mobile devices has led to the integration and convergence of various technologies into a wide range of innovative mobile and multi-modal applications. Mobile and other new technologies provide many different ways to offer the same or similar services to customers. Novel approaches are therefore needed for the

development and evolution of applications on and across different mobile and traditional platforms.

Methodology for Development to ensure Organisational Return

Mobile information systems are radical, and will therefore reward an increased focus on idea generation early in the design process. Understanding and modeling mobile users' requirements for new services is consequently of great importance. One needs both to be able to develop these systems and to address the major hurdles for the deployment of applications and services for the mobile user. Another effect of the introduction of mobile information systems is the spawning of further initiatives for changing other information systems. Overall, it is important to focus on the interoperability of services and the ability to roam across heterogeneous networks and service environments including, for example, working, billing, payment, ticketing and accounting services, as well as having seamless access to corporate and government resources.

This special issue highlights some of the existing European R&D in Applications and Service Platforms for the Mobile User. Although far from being a complete collection, it gives a good impression of where the European research community is putting its effort today, and where one can expect, or at least hope for, more results in the future.

Please contact:
John Krogstie, NTNU and SINTEF
Tel +47 22067425
E-mail: John.Krogstie@sintef.no
http://www.idi.ntnu.no/~krogstie

ARTICLES IN THIS SECTION

In Search of Knowledge about Mobile Users

by Ee-Peng Lim, Yida Wang, Kok-Leong Ong and San-Yih Hwang

At the Center for Advanced Information Systems at Nanyang Technological University, Singapore, research is carried out on using knowledge discovery techniques such as frequent pattern mining, classification, and clustering to create new mobile applications or to enhance existing mobile applications.

Mobile phones and other mobile devices are fast becoming indispensible in our modern society. According to a recent survey by Frank N. Magid Associates and Upoc.com, 59 percent of Americans age 12 and over (about 140 millions of them) own mobile phones, and that almost a quarter of non-owners plan to buy a mobile phone in the near future. The sales of mobile phones worldwide was 385 millions in 2001 and it has been predicted to reach 675 million in 2006. In tandem to this growth trend, we also witness the emergence of many new applications and businesses that exploit mobile phone technologies in different ways. Before other wearable computing gadgets become more feasible and popular, mobile phones are likely to remain as the more dominant wearable devices in the coming years.

Mobile phones, unlike computers connected to wired networks, are highly personalizable. While it is common for a user to own a few mobile phones, it is very unlikely for different users to share a mobile phone. Also unlike other personalized accessories such as watches, walkmans, etc., many of the mobile phones are trackable. They are trackable because they have to maintain regular contacts with the mobile telecommunication networks in order to receive and make calls. With these trackability and personalization features, one can conceive many unique and interesting mobile applications for end users.

At the Center for Advanced Information Systems, we conduct extensive research on using knowledge discovery techniques such as frequent pattern mining, classification, and clustering to create new mobile applications or to enhance existing mobile applications. Examples of such mobile applications include e-commerce systems, databases, email systems, search engines and web browsers. We investigate both the functionalities and operational efficiencies of mobile applications, and determine the kinds of knowledge required for enhancing them. We also develop algorithms for discovering knowledge from mobile user data and evaluate their performance. We further study how the discovered knowledge and mining algorithms can be integrated with the operational systems turning the knowledge into actions.

In the e-commerce application domain, we envisage the importance of using mobile user movement data to derive knowledge about mobile users. As users' purchase behaviors are often highly correlated to their group affiliations, knowing the latter well will allow e-commerce vendors to develop group-specific pricing models and marketing strategies to better meet the buying needs of the user groups. While there have been several existing methods developed to address the problem of discovering user groups, they are usually based on transaction histories and user profiles, and the accuracies may not be satisfactory.

In our research, we introduce a new approach to mine user groups, known as *group pattern mining*. In group pattern mining, we determine the customer grouping information based on the spatio-temporal distances among the customers. We assume that the user movement data are first collected by logging location data emitted from the mobile phones and similar devices. Mobile users that are always close to one another for significant amount of time are modeled as a group pattern represented by 4-tuple $<G, max\_dist, min\_dur>$ where G denotes the set of mobile users, *max_dist* denotes the maximum distance between any pair of users for them to be considered close, and *min_dur* denotes the minimum duration in which the users in G are close to one another. For example, for {Tom, Mary} to be a valid pattern, Tom and Mary must be less than max_dist apart for at least min_dur continuous time period.

Compared to traditional frequent pattern mining, group pattern mining incurs much more computation and storage. To discover group patterns from a set of mobile users, we have proposed a few efficient algorithms that are based on frequent pattern mining. As part of this research, we also investigate new data summarization techniques and data structures to reduce the computational and storage overheads. In brief, our research framework for mobile data mining is divided into three core areas, which we

Further Reading

- Upoc.com. "More Mobile Owners Turning to Text Messaging", http://www.upoc.com/corp/news/news-emarketer.html, Feb. 2003.
- Reed Electronics Research, "RER-The Mobile Phone Industry - A Strategic Overview," Oct. 2002.
- U. Varshney, R. Vetter and R. Kalakota, "Mobile commerce: A New Frontier," IEEE Computer: Special Issue on E-commerce, Oct. 2000.
- Ee-Peng Lim, Keng Siau, Advances in Mobile Commerce Technologies, Idea Group Publishing, Jan. 2003.
- MIT Media Lab. The Context Aware Cell Phone Project. (under the MIThril wearable computing initiative). http://www.media.mit.edu/wearables/mithril/phone.html
- Yida Wang, Ee-Peng Lim, San-Yih Hwang, "On Mining Group Patterns of Mobile Users," 14th International Conference on Database and Expert Systems Applications (DEXA2003), Prague, Czech Republic, September, 2003.

shall highlight some of the issues that should be further investigated. They are:

- Infrastructure for mobile data mining – we explore the design issues involving a warehouse for mobile data. This may include the algorithms for efficient aggregation and transformation of mobile data. The challenge in this case is to deal with the heterogeneous data format from different mobile devices at different locations with different bandwidth and computing resources.
- Algorithms for mobile data mining – once data is available, the next challenge is to make sense of the data. Hence, algorithms are needed to find knowledge to improve the efficiency of mobile applications/queries, and to enhance the user experiences of the phone.
- Incorporating mobile mining results into operational systems – knowledge obtained from data mining must be integrated with the operational systems. The challenge is to develop algorithms that evaluate which are the 'actionable' data mining results, and then apply them in a timely fashion to ensure the effectiveness of the mobile data mining system.

Mobile data mining, as a very new area of research, has created a wide range of opportunities for researchers, engineers and developers to create new interesting applications for both the end users and businesses. We believe that our group pattern mining work is only one of these many examples. As many more knowledge about mobile users can be mined in the near future, we therefore expect the upcoming applications and systems to be able to adapt more seamlessly into our daily lives.

Please contact:
Ee-Peng Lim, Yida Wang
and Kok-Leong Ong
Nanyang Technological University,
Singapore
E-mail: aseplim@ntu.edu.sg,
ongkl@pmail.ntu.edu.sg,
wyd66@pmail.ntu.edu.sg

San-Yih Hwang
National Sun Yat-Sen University, Taiwan
E-mail: syhwang@mis.nsysu.edu.tw

The Missing Link – User-Experience and Incremental Revenue Generation on the Mobile Internet

by Barry Smyth

Intelligent navigation and personalization technology developed by ChangingWorlds and University College Dublin's SmartMedia Institute has proven to significantly enhance the mobile user experience, drive usage and take-up, and so increase data-services revenues for mobile operators.

Paradoxically, the mobile Internet represents both a dramatic step forward and a significant step backward from an information access standpoint. While it offers users greater access to valuable information and services 'on the move', mobile handsets are hardly the ideal access device in terms of their screen-size and input capabilities. As a result, Mobile Internet users are often frustrated by how difficult it is to quickly access the right information at the right time. In our research we have pioneered the use of artificial intelligence and personalization techniques to provide practical solutions to these access problems and our technology is now deployed with leading European operators. Specifically, the ClixSmart Navigator™ product-suite, developed by ChangingWorlds Ltd. (www.changingworlds.com), provides mobile operators with a complete multi-access, portal platform that includes patented intelligent navigation services that automatically learn about user preferences in order to automatically adapt the structure of a portal for the individual user. ClixSmart Navigator can reduce content access times by 50%, significantly improving usability, and so leading to direct increases in mobile usage.

The Mobile Navigation Problem

The core usability problem with Mobile Internet services (such as WAP portals) is that users spend a significant time navigating to content through a series of menus (see Figure). Recent research highlights the scale of this problem and the mismatch between user expectations and realities. One study claims that while the average user expects to be able to access content within 30 seconds, the reality is closer to 150 seconds. WAP navigation effort can be usefully modelled as click-distance – the number of menu selections and scrolls needed to locate a content item. A recent analysis of 20 European mobile portals reports an average click-distance in excess of 16. In other words, a typical European mobile user can expect to have to make 16 clicks (scrolls and selects) to navigate from their portal home page to a typical content target. Moreover, on average, European portals are organised such that less than 30% of content sites are within 10-12 clicks of the portal home page; 10-12 clicks corresponds to a navigation time of about 30 seconds, which is the tolerable limit for most users. In other words, more than 70% of mobile portal content is essentially invisible to users because of its positioning within its parent portal.

The Personalization Solution

Large click-distances are a fundamental feature of a 'one size fits all' approach to portal design and even optimising a

menu structure for the needs of some imaginary 'average user' is unlikely to benefit individuals. However, instead of presenting the same portal to each and every user it is possible to use personalization techniques to learn about the preferences of individual users and to strategically adapt the structure of the portal on a user-by-user basis.

ClixSmart Navigator adopts a collaborative, probabilistic approach to personalization. Individual user accesses are tracked and used to generate comprehensive user profiles that reflect the user's portal access patterns (the services they access, when, and how often), their access-device characteristics, and various temporal data. These profiles are used to develop a probabilistic access model for each user that is capable of estimating the probability that a user will access an option, o (which could be another portal menu or content page), given that they are currently viewing portal menu, m. This in turn allows ClixSmart Navigator to personalize static portal menus by reordering their default links based on the target user's access probabilities and by promoting new links that are likely destination targets for this user.

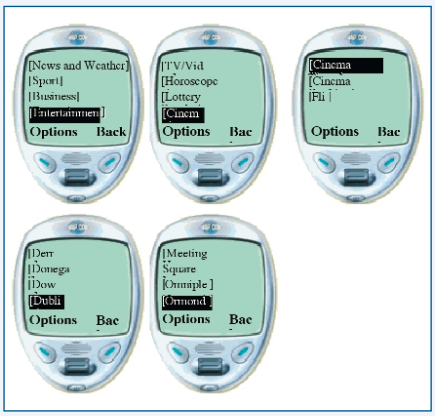

A sample navigation sequence on a WAP portal: accessing the Ormonde Cinema site takes 16 clicks (scrolls and selects) from the portal home page.

Over time the portal structure gradually adapts to the access patterns of an individual user. For example, a regular movie goer may find that, at the weekend, their local cinema has been promoted to their portal home page. Of course, it is vitally important that any restructuring takes place in a fashion that makes sense to the user. Rapid changes in portal structure can confuse the user rather than assist them, and as such ClixSmart Navigator contains a range of features to control the speed and scope of perosnalization in order to maximise the benefits to the end-user.

The Benefits

ClixSmart Navigator is currently deployed on the Vodafone (Ireland) and O2 (Germany) networks and is under evaluation by a range of European mobile operators. To date the benefits to the end users, mobile operators, and content providers are clear and compelling. For example, extensive evaluations have proven that ClixSmart Navigator's intelligent navigation can reduce navigation effort by up to 50% within 4 weeks. Moreover, for every one second of navigation time that is saved, the average user engages in additional 3 seconds of content time. This translates into significant increases in overall airtime (20%+), page impressions (30%+) and data download volumes (30%+) when compared to control-group usage.

Link:
http://www.changingworlds.com/

Please contact:
Barry Smyth
ChangingWorlds, Dublin
and University College Dublin
E-mail: barry.smyth@ucd.ie

User-Centric Peer-to-Peer Service Environment for Interaction with Networked Appliances

by Daniel Pakkala, Pekka Välitalo, Pekka Pääkkönen and Juhani Latvakoski

The focus of research at VTT has been to design a service platform for the future wireless world. Special attention has been paid to user centricity in the design of future service platforms.

The mobile Internet is now being realised and launched on the market. It enables different kinds of mobile terminal to be connected to the Internet anywhere and at anytime. However, due to the massive amount of information content and the number of devices connected to the Internet, users' lives are becoming more and more complicated. Consequently, in the future, user centricity requirements will become increasingly important, making this an essential research challenge.

User centricity and the flexibility to deal with multiple use scenarios have been starting points for the development of the service platform. We have chosen the future residential home environment for the validation of solutions because it is a good example of a future intelligent environment using many networked appliances with different capabilities. In a residential home environment, some networked appliances can be very simple and it would not be cost effective to implement the control channels of these devices based on Internet Protocol (IP). Instead, these simple appliances - like X.10 light switches, for example - can be controlled via a service gateway from the IP networks. For devices like alarm systems or remotely controllable devices with more processing capabilities and functionalities, it is reasonable to implement the control channels based on IP. This variety of appliances creates the need for flexibility and transport layer

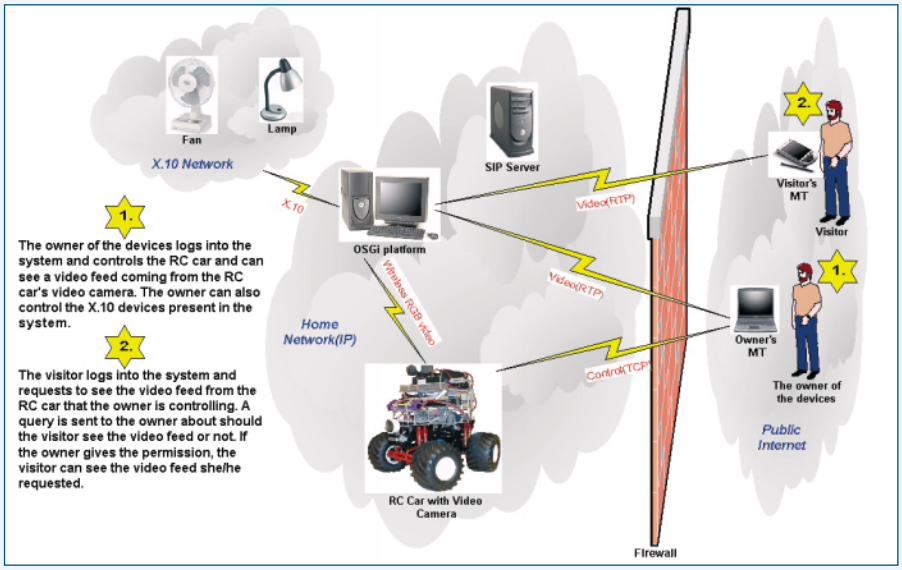

Figure 1: Approaches for controlling networked appliances.

protocol independence for the service platform.

The problem to be solved by user centricity in service platform design can be compared to the situation some of us already have in our living rooms. For example, owning a TV, VCR, hi-fi system and DVD means there are at least four different remote controls to manage. Even with this relatively small number of devices, the separate remote controls present an annoying configuration from the user perspective. In future homes there could be many more appliances that need to be controlled and managed by the user. The purpose of user centricity in the service platform is to make the control and management of future networked appliances (NAs) as easy as possible for the end user. For example, a single button press on a mobile terminal (MT) can display a floor plan of the home that shows all the appliances that the user is able to control. When a user selects, for example, a living room lamp, the user interface (UI) of the living room lamp appears on the screen and the user can then remotely control it.

In Figure 1, two different approaches for controlling networked appliances are presented. On the left is the legacy independent UI approach, and on the right, the applied user centric peer-to-peer approach that we have used in the developed service platform. In the legacy approach, the user interfaces of the IP-capable appliances are located at the MT, whilst those of the non IP-capable appliances are located at the service gateway. This kind of approach to controlling system design will lead to a non user-friendly situation in the future, when the number of NAs will grow fast. In the applied user centric peer-to-peer approach, all the devices should be represented in the service gateway by their agents or control software, and the UI components of all the NAs present in the system should be available at the service gateway. The final UI that is uploaded to the user's mobile terminal is generated from these UI components. The UI components contain the address information of the NAs providing the service in question. After the final UI is activated at the mobile terminal, it automatically establishes the control connections needed and becomes visible to the user.

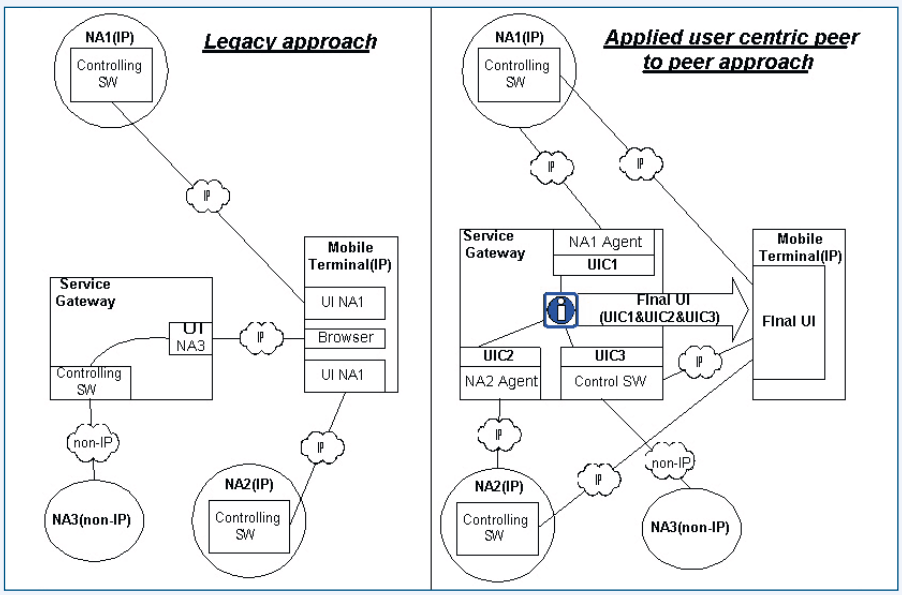

Figure 2: Demonstration system.

In our latest demonstration we introduced the concept of the User Centric Peer-to-Peer Service Environment (see Figure 2). In the demonstration we controlled two X.10 devices and a remotely controllable car with a video camera through the service platform. The X.10 devices represented simple non IP-capable devices and the car represented a more complicated device with IP capabilities. The configuration of the demonstration, final data channels and use cases are presented in Figure 2.

The work described in this article has been carried out in a strategic project OmniPresent rEference seRvice Architecture - OPERA at VTT Electronics during year 2002.

Link:
http://opera.vtt.fi/

Please contact:
Daniel Pakkala, VTT Electronics
E-mail: Daniel.Pakkala@vtt.fi

Juhani Latvakoski, VTT Electronics
E-mail: Juhani.Latvakoski@vtt.fi

Pekka Välitalo, VTT Electronics
E-mail: Pekka.Valitalo@vtt.fi

Pekka Pääkkönen, VTT Electronics
E-mail: Pekka.Paakkonen@vtt.fi

Kontti — Context-aware Mobile Portal

by Juha Kolari

For many, a mobile device is a constant companion during daily routines. Information and service needs vary according to the user's immediate situation. A context-aware service responds to this by offering contextually relevant information. By identifying these contexts and the varying information needs users have in them, mobile services can be made more topical, personal and easily available.

In order to study the technology and concept of context-awareness, VTT has developed a context-aware mobile portal. The contents of the service are adapted to the user's current context and device. The contexts can be identified automatically according to time or location. They can also be manually defined and activated by the user. The contents of the portal can be existing mobile services and Web pages as well as personal notes and files.

There will be two field trials to evaluate the prototype and the concept. Over 70 users will be involved as the prototype is tested in the field. Conceptually, the evaluations approach the subject from three different angles, namely, everyday contexts, context-aware services and social surroundings.

Everyday Contexts

Everyday use plays a visible role in the evaluation. The goal is a natural integration of the service with what the users actually do. As 'context' itself is a highly personal concept, so context-aware services need to be flexible enough to conform to the varied situations in users' lives. The system allows users to create their own personal contexts, such as 'at home', 'at work', 'out partying' and 'feeling blue'. The users can link personal notes and existing services to the contexts relevant to them. They can also inform others of their current personal context. The challenge is in achieving this high personal conformity without requiring too much effort from the users.

Context-Aware Services

Ready-made context-aware services may not adapt to the user's everyday life but can provide content with little need for further personalisation. The portal allows the delivery of a context-aware service directly to the user. Visitors to a festival could, for instance, subscribe to the festival schedule. The schedule and upcoming events can be presented according to time and the user's location. We will study the service from this angle at the week-long Tampere Theatre Festival in August 2003. Another case will be a location-based historical tour in the city of Tampere, which will be evaluated in the summer of 2003.

Context-awareness can bring about a more personal and spontaneous way of using mobile services.

Social Surroundings

One motivation for creating and maintaining everyday contexts is to convey the user's own status to others; eg whether you are available for contact or what your plans for the evening are. The user can specify which other users are allowed to view his/her current status. Even then, the user can define how the status will appear to others, and the location or context itself is not necessarily conveyed in this description.

User-generated material naturally means the inclusion of context-based messages. The system allows users to send messages which are delivered only when the chosen target context is active, ie when the user arrives at work or home. Targeting the message allows, for instance, a work-related message to be sent to work on a weekend. The message will be delivered once the recipient enters the workplace.

Appropriate technological matches have been chosen for these concepts. Use in everyday life requires a mobile device that is available all the time and is not limited to a certain area. Consequently, the prototype has been designed to work in WAP-enabled mobile phones. It is optimised for use in phones with colour displays and multimedia capabilities. Both WML and XHTML versions are available. The system is open with regard to positioning methods, and the current system supports both cell-based and proximity-based positioning. A trial on a WLAN platform will complement the evaluation both conceptually and technologically. The WLAN trial will be used to study the promising concept of context-aware services more thoroughly.

The development and studies are being done under the project Kontti (Context-Aware Services for Mobile Users). The two-year project began in 2002 and will run through 2003. The project is funded in part by Tekes and its NETS programme. Other collaborative partners in the project are Nokia, Radiolinja and Teamware. In addition, the Finnish Broadcasting Company (YLE) provides content in the trials.

Links:
http://www.vtt.fi/tte/projects/kontti/

Please contact:
Juha Kolari, VTT (user studies)
Tel:+358 3 316 3311
E-mail: Juha.Kolari@vtt.fi

Timo Laakko, VTT (project manager)
E-mail Timo.Laakko@vtt.fi
Tel: +358 9 456 4505

Contextual Mediation enables Appropriate Data Selection

by Dan Chalmers, Morris Sloman and Naranker Dulay

Mobility produces a wide range of context, which has a significant impact on users' experience of computing. Some solutions exist for issues such as small screens and low bandwidth, although these generally offer rather indiscriminate modification of data and support only a small range of devices. A more sensitive approach has been investigated at Imperial College London: causing a selection of data in response to a wide range of contexts.

Context aware computing refers to the treatment of the user's environment in system behaviour. Aspects of context are wide-ranging: location, speed of travel, noise levels, device in use, task engaged in etc, and mobile computing exposes users to the effects of these. Context has many uses in mobile computing including: description of context to the user, eg location notification; context-sensitive resource use, eg location aware printing; contextual annotation of data, eg noting other people present as meta data to meeting minutes. At Imperial College we have been developing techniques to specify contextual mediation. This is the use of context to modify the data selected amongst the data offered, in order to make the best presentation of information to the user given for context of use. Techniques include: scaling graphics and summarising text to fit documents within a screen; limiting the quality of images to reduce data size when using a slow network; including or highlighting data describing required facilities such as induction loops for the deaf; omitting attractions which are closed when presenting a map of a tourist site. Some of these techniques have been deployed in isolation, often hard-coded into applications and targeted towards a few specific devices – the effect often being unsubtle and failing to capture all the user's needs. Our research provides a general way to associate a description of needs and limitations with the context of use – and use this to select the best available combination of data.

Our work initially focused on maps. Map data may be large in size, causing long delays on slow networks and it may include a level of detail, which cannot be interpreted on low-resolution displays. Vector map data lends itself to partial presentation and processing of different elements. We were able to use this to allow the prioritisation of data according to its semantic types, eg roads, rivers, buildings, distance from location of interest; and its properties, eg different scale representations of the same feature. In this way mode of transport affects the priority of representations of roads, railways, footpaths etc. Speed affects inclusion of distant or small features. Task (work, leisure etc.) affects inclusion of tourist symbols. We also describe limits over managed resources in order to constrain the selection. Screen resolution and available bandwidth affect the combination of data selected, in order to meet specified goals: time to download may depend on speed, drawing complexity may depend on vibration. Note that the selection of data to display considers the combined effect of all the data in the selection, rather than processing data on an element-by-element basis.

Contextual mediation is illustrated in the Figure. The left-hand map shows the complexity resulting from no mediation. The mediated maps are for navigating in a car at 30 to 40mph, hence the concentration on major roads and the omission of minor roads away from the immediate vicinity (the centre). The data was loaded within a 25s deadline over a simulated GPRS network. The difference between the tourist and worker's

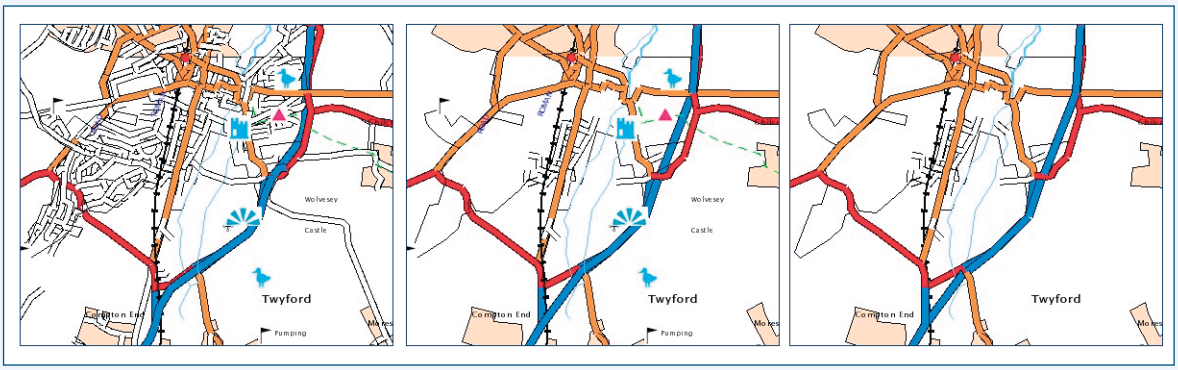

Figure (left to right):
Unmediated map, all features included.
Mediated for tourist in a car, less peripheral detail.
Mediated for worker in a car, note omission of tourist sites, golf courses, footpaths etc.

maps is subtle, but helps the worker to scan the map without distraction. Where features have hyperlinks from them, eg describing tourist attractions or access to delivery points, the ease of selecting links also improves. In tests we found the resulting maps to be more predictable in their download time than unmediated maps. Predictable delays have been shown to be a key factor in user satisfaction in the web. The omission of unnecessary detail was shown to provide improved clarity and faster navigation in user tests.

Our approach is applicable to other semantically rich data intensive applications. We are now investigating applying this approach to collaborative tools such as message services, shared whiteboards, engineering and maintenance plans and web based documents. In collaboration with the University of Southampton we have been investigating the mediation of hyperlinks and the use of mediation in pervasive information systems. There is also related work starting on Ubiquitous Computing for Healthcare in the Community, which will be focusing on monitoring of patients with clinical conditions as they go about their normal activity.

Link:
http://www-dse.doc.ic.ac.uk/Research/ubicomp.html

Please contact:
Morris Sloman,
Imperial College London, UK
E-mail: m.sloman@imperial.ac.uk

mBlog: a Mobile Information Service for All

by Emmanuel Frécon

Easy-to-use Web sites, called 'weblogs' or 'blogs' are the latest massive Internet phenomenon. mBlog, developed by SICS in collaboration with Ericsson, takes blogs a step further by adding a mobility component.

Weblogs, or blogs for short, are the latest Internet phenomenon. Having started as an underground culture in the late 90s, the movement has now reached such an extent that these Web sites have become an important medium for the general spread of opinion and information on the Internet. Thousands of new blogs are created every day by ordinary people or groups, and cover themes such as personal diaries, political or professional commentaries, fanzines around culture and sports, group diaries around fishing, excursions, etc.

Technically speaking, blogs are an easy way to produce a Web page by filling out a form. The text entered is displayed on a page organised in reverse chronological order: the most recent entry is at the top of the page and older entries follow. Blogs can provide an archive mechanism for all entries, and a search feature to look for specific information within older content. Typically, a blog will have a small number of authors and any number of readers. A commenting system will allow readers to react to the opinions expressed by the authors on the Web page.

The simplicity of this formula makes blogs available to the mass of Internet users, relieving them from the complexity of HTML editing. Apart from this simplicity, blogs are interesting both for their sociological aspect and for their content or topic area. From the latter point of view, blogs have two main purposes: they help maintain a community around any given topic (from a class trip to war in Iraq), and they allow one to make oneself publicly heard.

mBlog, developed by SICS in collaboration with Ericsson, takes blogs a step further through the addition of a mobility component. With mBlog, people can maintain any number of blogs simply by sending images and texts by MMS, SMS, e-mail or by filling out a form in a Web browser. Furthermore, each piece of information can be geographically positioned.

Mobility: the Future of Blogs

The very nature of blogs and their usage lead us to consider them as 'real-time' entities. Indeed, for a blogger, no matter what the topic of the blog, it is necessary that information be published as soon as possible. For example, if a blog was to support a local football team, being able to blog the main events of a match in real time is a way to keep supporters who could not join the game aware of its progress, as well as an effective method of keeping a record of the match. Consequently, the immediacy of information production and consumption naturally lends itself to extending the concept of blogs towards mobile environments.

The idea behind mBlog is the ability to post from a mobile phone (MMS, SMS, WAP input, e-mail), and the ability to read entries at any time using a number of mobile outputs (WAP browser, speech synthesis, etc). mBlog also introduces a new dimension to weblogs, namely space. Each blog can have a geographical dimension, with each entry being positioned within this space. This offers novel ways to access the information.

Business and Sociological Results

The genesis of mobile blogs will support the building of communities and social networks of many kinds, both at the professional and personal levels. For example, we foresee professional usage within mobile fleets of salesmen.

Being ubiquitous by nature, mobile blogs will generate traffic within several parallel distribution channels (GPRS, 3G, Internet), and will therefore benefit

mBlog takes blogs a step further through the addition of a mobility component.

the telecom industry. In particular, mobile blogs will drive MMS, SMS, WAP and data traffic, as well as increasing the sales and usage of multimedia phones. Through the provision of a service that is anchored in our social lives and behaviours, mobile blogs could act as one of the kick-off applications for MMS and 3G services.

How does it work?

Users send SMS, MMS or e-mail messages to a dedicated number or e-mail address. mBlog then creates a new blog entry for each message. This entry is placed at the top of the Web page while previous entries are moved down. Upon posting, entries are time-stamped and automatically given a geographical position, ie that of the device. This position can later be changed if necessary.

Alternatively, users can direct their browser to a specific Web page and fill in a simple form. Upon completion, an entry is created as above. The resulting blog can be viewed from any Web browser or WAP phone. A geographical map shows the location of the entries. It is useful to note that new multimedia phones now include fully-fledged Web browsers.

Technology behind the Scene

Blog systems are based on XML technology, which makes their content easy to manage, arrange and exchange. Lately, a number of standard APIs have emerged, meaning that blog systems have now become fully-fledged Web services. mBlog makes extensive use of standard Internet protocols and is implemented in a scalable and distributed manner:

Each mBlog service can be run separately on any machine with Internet access. Upon reception, the content of MMS and SMS is analysed and routed to the appropriate blog through a communication standard called XMLRPC.

WAP output can be applied to any existing blog, and is triggered through continuous surveillance of any Web page (in this case a blog). Upon modification, blog content is fetched through XMLRPC, converted to WML and posted back through FTP.

Entries are automatically geographically positioned using a positioning server made at SICS. Currently, we only support accurate WiFi positioning, but we are hoping to support GPS and GSM in the near future. The geographical map of the blog is continuously updated as new entries are posted.

Future Work

mBlog provides a simple means for reading and writing blogs on the run. Compared to other 'moblogging' attempts, mBlog takes a unified approach through the support of most existing mobile devices. Additionally, mBlog adds the dimension of space to blogs, to enhance the association of information with places.

We are currently looking at extending the system by building it on top of a scalable P2P network infrastructure. Information can be replicated along this network as needed. This has two main advantages: it guarantees its very existence against external ‘attackers’ of all sorts, and access to consumers is widened.

Link:
http://ice.sics.se/projects/mblog.html

Please contact:
Emmanuel Frécon, SICS
Tel +46 8 633 15 34
E-mail: emmanuel@sics.se
http://www.sics.se/~emmanuel/index.shtml

A Mobile Schema Browser for Integration Specialists

by Gerald O'Connor and Mark Roantree

Middleware developed at Dublin City University allows a mobile user to query metadata on an object-relational database and automatically display its structure. This is particularly useful for example for displaying and browsing complex schema information on a Portable Digital Assistant (PDA).

Integration Engineers are often faced with a requirement to display and analyse the complex schemas of information systems to be merged. As these systems can be dispersed over a wide geographic area, a Portable Digital Assistant (PDA) provides a flexible means of viewing and displaying schema information. However the browsing process, which is often complex and problematic on a workstation screen, becomes more difficult on the smaller PDA. Using our metadata middleware, a mobile user can query metadata on an object-relational database and automatically display its structure. To demonstrate this, an application was developed which exploits our interface to the extended object Object-Relational (O-R) schema repository, to manipulate complex metadata information.

The O-R model is a complex structure combining features of relational and object models. The relational model consists of tables, triggers, constraints, views and procedures, while the object model consists of types, associations, aggregations, encapsulation, inheritance and other complex structures. This combination of features provides a powerful environment for representing data, yet if this data is not fully specified and manageable through a metadata interface, its reuse and sharing becomes a time consuming and expensive task. The Efficient Global Transactions for Video media research project (EGTV) has defined a complete metadata interface to O-R metadata and developed a tool to demonstrate and evaluate its effectiveness.

Mobile Layer

The Figure illustrates the deployment architecture for O-R metadata access. In the Mobile Layer, a PDA uses the metadata middleware in a specific application. The Schema Browser queries the schema using the Extended-SQL metadata query language, developed by the Interoperable Systems Group (ISG). This includes metadata query options for any O-R schema, role views and multimedia objects. The extended-SQL queries are wrapped in XML to provide a robust, non-proprietary, persistent and verifiable file format for the storage and transmission of data. Using the XML version of the language, the Schema Browser automatically queries and displays the complex metadata required to display the O-R schema model.

XML Translation Layer

The XML Translation Layer resides at the database server. Its purpose is to provide a standard interface to the Extended-SQL query language. When receiving a query it is unwrapped to form an Extended-SQL query, which is subsequently passed to the SQL translation layer. After execution, results received from the SQL Translation Layer are XML-wrapped using a basic rule set, and then returned to the application.

SQL Translation Layer

The SQL Translation Layer is where most of the metadata processing takes place. Current approaches to interfacing metadata for O-R databases were examined before this layer was specified. It was found that relationships between types very often complex and difficult to retrieve using SQL:99. In fact, SQL:99 is not powerful enough to extract some metadata (nested table associations) and programming language extensions were specified to overcome this. The SQL Translation Layer accepts an Extended-SQL query, which is parsed to invoke a sequence of actions against the schema repository. The results may be comprised of conventional O-R metadata, role metadata and/or multimedia metadata. After the results are restructured to an O-R format, they are passed back to the XML Translation Layer.

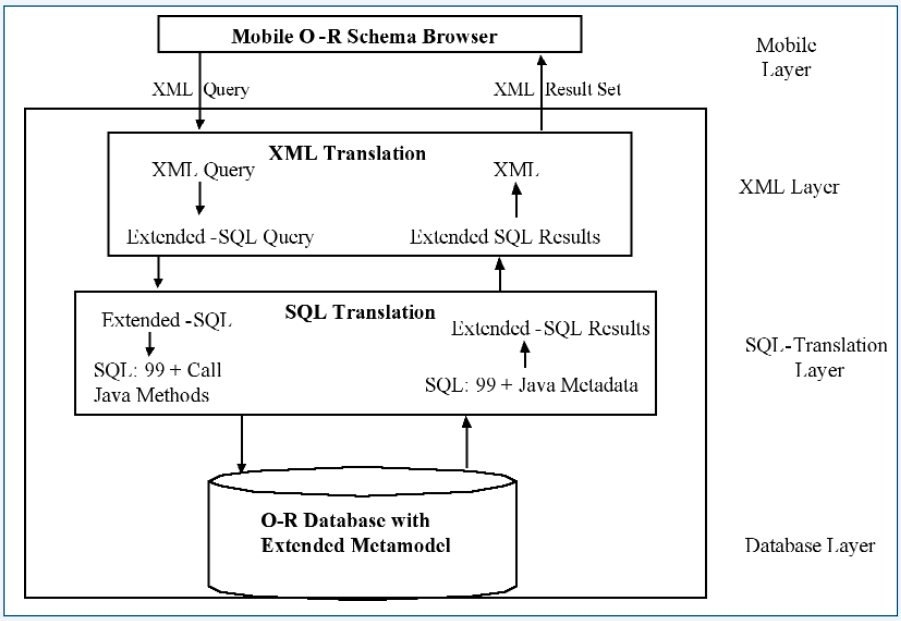

Mobile Object-Relational Schema Browser.

Database Layer

The O-R schema repository was extended by other EGTV researchers to provide new interfaces to role and multimedia metadata. Role metadata was added because it adds to the expressiveness of O-R databases. Roles provide temporal aspects to entities, a feature that is missing in conventional models. Without roles, a new object must be created each time the structure of an object evolves and many complex issues are involved in maintaining such an operation. The current popular O-R repositories fail to manage multimedia data adequately, as it is often represented as a BLOB with no further metadata to explore its application or use. This was addressed with the metamodel extensions. All of this metadata can be accessed using this framework and the Extended-SQL:99 language. Thus, it becomes accessible to integration engineers using mobile devices.

Conclusions

The mobile display tool can automatically display schema, view, role, and multimedia information in just a few selection buttons. This provides the mobile integration engineer with a fast completion of the first phase in the integration process. It also illustrates that the mobile device can offer a viable platform for some CASE tools. Current research is focused on providing update facilities through the mobile device, in order to make structural (metadata) changes to underlying O-R schemas, and to define new role views.

Link:
http://www.computing.dcu.ie/~isg/

Please contact:
Gerald O'Connor and Mark Roantree
Dublin City University, Ireland
E-mail: {goconnor, mark}@Computing.DCU.ie

A Peer-to-Peer Middleware for Mobile TeamWork

by Carlo Ghezzi, Gianpaolo Cugola, and Gian Pietro Picco

The advantages of a peer-to-peer architecture go well beyond the realm of Internet file sharing, becoming crucial in supporting business processes and especially collaborative work involving mobile users. To support this view, we designed and experimented with PEERWARE, a core communication middleware for TeamWork applications.

Collaborative work is intrinsically peer-to-peer in nature. Members of a team typically interact directly with each other, each responsible for a given set of documents and carrying with them the subset relevant for discussion. On the other hand, most of the currently available tools supporting collaboration exploit a rigid client-server architecture.

This results in an 'architectural mismatch' between the external view provided by the application and its internal software architecture. The effect of this mismatch is a lack of flexibility in carrying out the interactions, which must all be funneled through the server. This limitation is even more evident when mobility becomes part of the picture. People need to communicate and collaborate even while in movement, and independently of their location. However, in similar situations, server access is often prevented by technical or administrative barriers.

We argue that a peer-to-peer approach holds significant advantages over traditional client-server architectures. When a peer-to-peer architecture is adopted, data and services are no longer gathered in a single point of accumulation. Instead, they are spread across all the nodes of the distributed system. Users may directly host the resources they want to share with others, with no need to publish them on a particular server.

Interestingly, these features are relevant not only in mobile scenarios but also in fixed ones, where the decentralized nature of a peer-to-peer architecture naturally encompasses the case of multi-site or multicompany projects, whose cooperation infrastructure must span administrative boundaries, and is subject to security concerns.

Unfortunately, most of the peer-to-peer applications developed in recent years started from premises that are rather different from those outlined thus far. They target the Internet and aim at providing peer-to-peer computing over millions of nodes, with file sharing as their main application concern. The difference in perspective from the domain of collaborative work is made evident by their search capabilities, which typically do not guarantee to capture information about all matching files. In most cases they do not take into consideration features like security or the ability to support reactive interactions, which are crucial in cooperative business applications. Moreover, they bring peer-to-peer to an extreme, where the logical network of peers is totally fluid, and none can be assumed to be fixed and contributing to the definition of a permanent infrastructure. This radical view prevents access to resources exported by non-connected peers, which is unacceptable in the business world, where critical data is often required to be always available, independently of its owner.

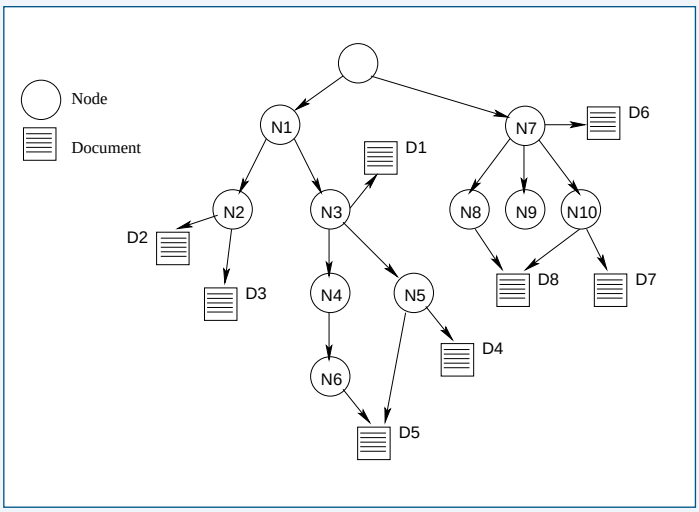

Figure 1: The data structure provided by PEERWARE.

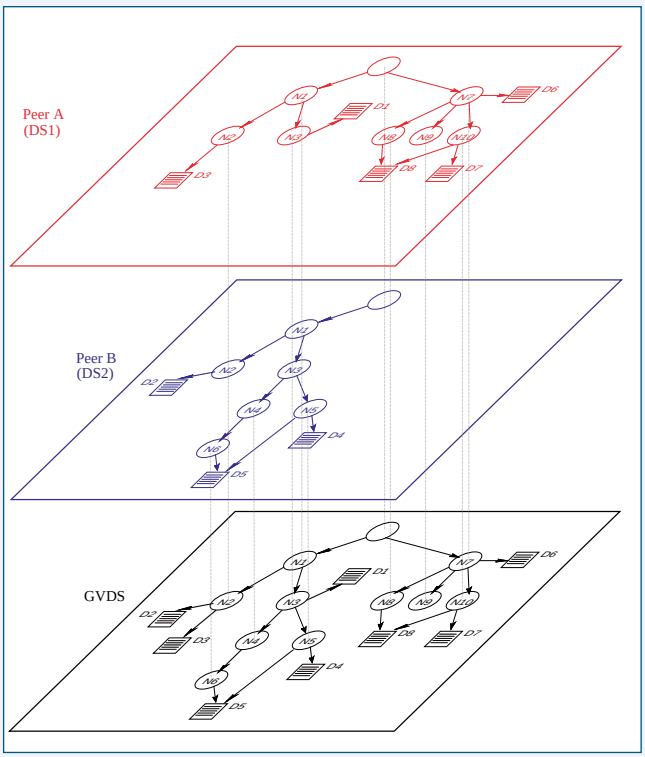

Figure 2: Building the GVDS in PEERWARE.

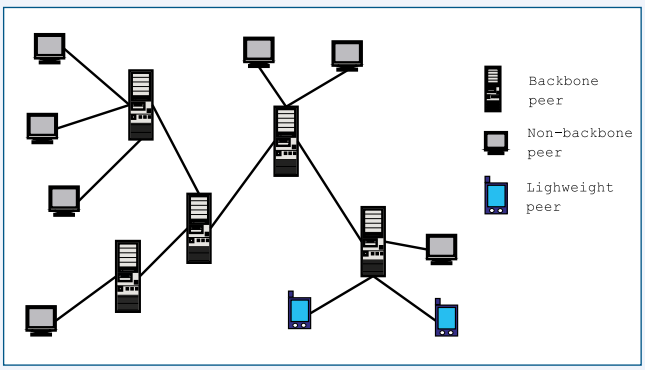

Figure 3: The PEERWARE run-time architecture.

PEERWARE

On the basis of the above considerations, we have developed PEERWARE: a peer-to-peer middleware for teamwork support specifically geared towards the enterprise domain.

PEERWARE is both a model and an incarnation of this model in a middleware. In developing both, our first concerns were minimality and flexibility.

The Model

The PEERWARE coordination model exploits the notion of a global virtual data structure (GVDS), which is a generalization of the LIME coordination model. Coordination among units is enabled through a data space that is transiently shared and dynamically built out of the data spaces provided by each accessible unit.

The data structure managed by PEERWARE is a hierarchy of nodes containing documents, where a document may actually be accessible from multiple nodes, as shown in Figure 1. This structure resembles a standard file system, where directories play the role of nodes, files are the documents, and Unix-like hard links are allowed only on documents.

When a peer is isolated, it is only given access to its own tree (stored locally) of items (ie, nodes and documents). However, when connectivity with other peers is established, the peer has access to the virtual tree constructed by superimposing the trees contributed by all the peers in the system, as illustrated by Figure 2.

In search of minimality, PEERWARE provides only three main operations to operate on the GVDS:

- the execute operation allows peers to execute an arbitrary piece of code on a selected set of items held by connected peers. The results are collected and returned to the caller
- the subscribe operation allows peers to subscribe to events occurring on a selected set of items, while
- the publish operation allows peers to notify the occurrence of events.

By exploiting these primitives, peers can query the GVDS and also subscribe to events and receive the corresponding notifications. The hierarchical structure of the GVDS provides a natural scoping mechanism, thus leading to an efficient implementation of searches.

The Middleware

Currently, the PEERWARE model has been implemented in two middleware: one, developed in Java, has been used as the core of the MOTION platform, the other, developed in C# under the Microsoft .Net infrastructure, is the core of the PeerVerSy configuration management tool.

Both implementations are tailored to the business domain and distinguish between a set of permanently available backbone peers, and a fringe of mobile peers, which are allowed to connect and disconnect as required. To optimize routing, these peers are connected to

form an acyclic graph in which the mobile peers represent the leaves, as shown in Figure 3.

Access control and security are critical issues in the domain we target and they are addressed by two separate modules. One provides mechanisms to establish encrypted channels among peers and to manage the security information necessary to authenticate a peer. The other embeds the actual security policy that determines the capabilities of a given peer.

To increase flexibility, the functionalities provided by the security modules and also by the repositories holding local documents are sharply decoupled from the specific implementation provided for these functionalities. Thus, the security protocols, as well as the format of the security information used to perform authentication, and the repository effectively used can be changed easily, eg, to adapt them to the common practice of a specific business environment.

Conclusions

Collaboration defines a scenario where interaction is intrinsically peer-to-peer. We exploited this idea by developing PEERWARE, a peer-to-peer middleware explicitly tailored for collaboration. The model of collaboration adopted by PEERWARE is minimal but expressive enough to support complex collaboration schemes, including both reactive and proactive interactions. Current incarnations of this model in a middleware have been tailored to the domain of enterprise-wide collaboration. We are now working on a different implementation oriented toward more dynamic scenarios, including ad-hoc networks.

Links:
PEERWARE:
http://peerware.sourceforge.net
MOTION project:
http://www.motion.softeco.it/pages/
PeerVerSy website:
http://sourceforge.net/projects/peerversy/

Please contact:
Gianpaolo Cugola, Politecnico di Milano, Italy
Tel.: +39 022 399 3493
E-mail: cugola@elet.polimi.it

Fluid Computing

by Marcel Graf

It has long been desirable to be able to let an application 'flow' from one user interface to another, depending on the user's situation and the capability of the interface being used. The IBM Zurich Research Laboratory has developed middleware which realises this wish. The service provided is called Fluid Computing.

Fluid Computing is a sub-area of Pervasive Computing. It denotes the replication and real-time synchronisation of application states on several devices. Thus the application state flows as a 'fluid' between devices.

There are three main application areas:

1. Multi-device applications: several devices may be temporarily coupled to behave as one single device, for example, a mobile and a stationary device.

2. Mitigation of the effects of variable connectivity: applications on ubiquitous devices can exploit full or intermittent connectivity, or operate in disconnected operation - all in a seamless fashion.

3. Collaboration: multi-person applications enable several users to collaborate on a shared document in synchronous mode (at the same time) or in asynchronous mode (at different times) without the user having to switch the application.

The Fluid Computing middleware replicates data on multiple devices and achieves coordination of these devices through synchronisation. Each device has a replica of the application state, which allows it to operate autonomously. A synchronisation protocol tries to keep the replicas consistent, depending on the quality of the network connectivity available (weak consistency). The Fluid Computing synchronisation protocol operates in two modes. Batch mode replication occurs when connectivity is regained after disconnection, exchanging all updates that have accumulated during the disconnection. Trickle mode replication occurs as long as there is some connectivity, propagating updates in real-time as soon as they are generated at a replica.

The weak consistency approach enables disconnected operation: we want to allow the user to continue working with a device and making updates even when the device has no connectivity. Replicas are therefore allowed to become inconsistent when disconnected. The propagation of updates for synchronisation has two special properties, explained below.

Epidemic update propagation: A user who carries several devices (mobile phone, PDA, watch etc) will often be in a situation in which the devices have connectivity between one another via a wireless proximity network (eg Bluetooth), but no wide-area connectivity. We want to be able to use the proximity connectivity to make the devices consistent among themselves. Updates can thus be exchanged with any

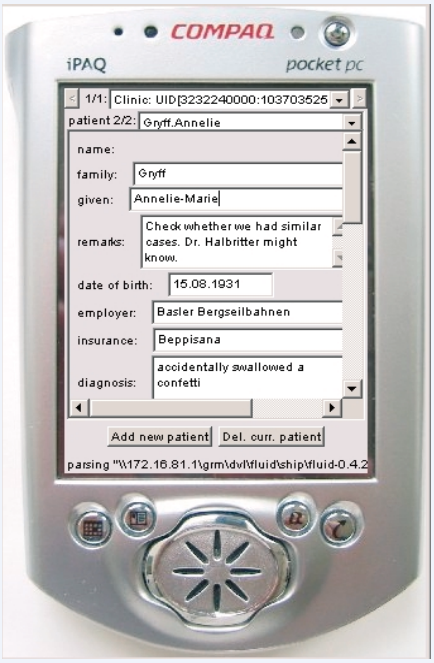

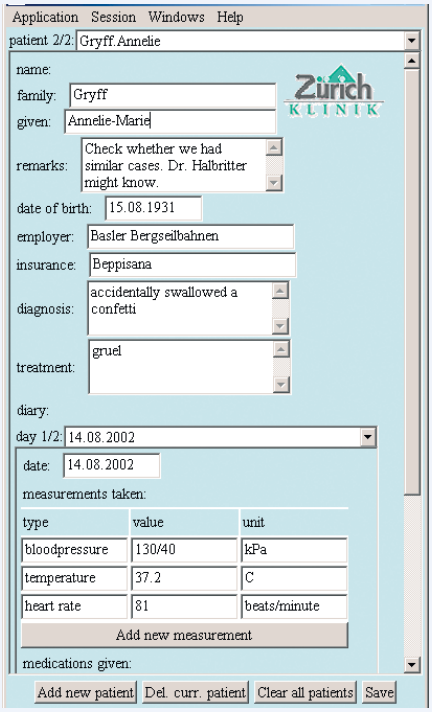

A Fluid Computing application is no longer tied to a particular device. It may run simultaneously on several devices which the Fluid middleware coordinates, for example a small mobile PDA (left) and a big stationary PC display (right). When wireless network connectivity is available synchronization occurs keystroke-by-keystroke, otherwise updates are queued until connectivity is regained.

other replica, may even propagate through several other replicas, and are not required to go through a central server.

Real-time update propagation: Mobile devices are limited in their human/computer interaction capabilities because of their small size. This problem can be overcome by temporarily coupling several devices tightly via a wireless proximity network. Tight coupling means that a user action (eg a keystroke) on one device is immediately reflected on the other devices, such that the user experience is that of a single device. This enables the combination of devices to be optimised for a certain input or output modality, for example, the combination of a mobile phone with a large stationary display. Therefore when there is connectivity between a set of replicas, and the user performs an update, this update is sent immediately to the other replicas. The replicated application state allows a user to decouple devices at any time without breaking the application.

The Fluid Computing middleware is a lightweight Java library that runs on PDAs such as Sharp Zaurus™ or PocketPC™ devices. The data structure of the replicated data is that of an XML document (other data structures are supported by the middleware as plugins). The API used by applications to manipulate that data structure is the Document Object Model (DOM) API of the W3C. Application developers interact only with the XML data structure and are shielded from the complexities of the replication mechanism.

A Sample Scenario

Imagine yourself to be a doctor at a renowned Zurich Clinic on the shores of the Lake of Zurich, where the rich and famous regain their health. All doctors and nurses are equipped with Personal Digital Assistants (PDAs) that give them access to their patients' data wherever they are. A wireless LAN provides connectivity to the PDAs throughout the clinic's premises.

You are walking along an aisle while reading the data record of Annelie Gryff on your Compaq iPAQ *. You have to decide on her exercise program. First you prescribe some swimming. Because the display is small you soon find yourself scrolling back and forth. You need a larger screen! You walk to a PC. Gryff's record appears on the PC, including the swimming exercises you have already entered. You continue listing prescriptions using the PC. Every keystroke is immediately reflected on your PDA. In an emergency you could immediately leave the PC - it would not be necessary to press a 'sync button' to synchronise the two devices.

For lunch you decide to go to Hiltl's restaurant, a vegetarian restaurant near Bahnhofstrasse. Obviously there is no longer a wireless connection to the clinic's intranet, but you still have Gryff's patient record on your PDA. You decide to prescribe a vegetarian diet - maybe that will do some good for her weight. As soon as you return to the premises of the clinic, this prescription propagates to the central patient database.

Link:
http://www.zurich.ibm.com/csc/mobile/

Please contact:
Marcel Graf, IBM Zurich Research Lab
Switzerland
Tel: +41 1 724 8318
E-mail: grm@zurich.ibm.com

Designing Usable Multi-Platform Interactive Applications

by Fabio Paternò and Carmen Santoro

Scientists at CNR propose a solution for the design and development of nomadic interactive applications, accessible through heterogeneous platforms and from different locations. This solution is supported by TERESA (Transformation Environment for inteRactivE Systems representAtions), a mixed-initiative tool supporting multiple transformations able to adapt the interactive part of an application according to the platforms employed.

A wide variety of new interactive platforms are now being offered on the market, ranging from small devices such as interactive watches to very large flat displays. This is changing the nature of many interactive applications, converting them to nomadic applications which need to support user access through different interactive devices, from various locations. Consequently, one fundamental issue is how to support software designers and developers in building such applications. In particular, there is a need for innovative methods and tools able to support the development of interactive software systems that can adapt to different targets while preserving usability.

In current practice the design of multi-platform applications is obtained through the development of several versions of the same application, one for each platform considered. Such versions can at most exchange data. This solution is somewhat limited because it implies high implementation and maintenance costs. Another solution proposed is the use of style sheets. Each platform is associated with a different set of stylesheets. Thus, the same elements are presented differently according to the type of platform available. This can be useful, but it is still insufficient; style sheets only help by trying to better support the same tasks through different platforms. However, this is not always meaningful because users often want to do different things according to the type of platform, and there may be dependencies among tasks performed through different platforms.

We have developed a solution for these thorny issues. Our approach can be summarised in four words: One Model, Many Interfaces. This means that we start with an abstract description of the activities to be supported and obtain different user interfaces for each available platform in order to support these activities. In particular, we start with a task model of a nomadic application, which describes the activities that should be supported in order to reach the user's goals. Designers are able to create effective user interfaces for the various platforms considered through a number of transformations supported by our new tool TERESA (Transformation Environment for inteRactivE Systems representAtion).

This work is being carried out in the IST European CAMELEON Project. The objective of the project is to build methods and environments supporting the design and development of highly usable context-sensitive interactive software systems by providing the means to express context-dependent information and developing tools that support the use of the information contained in multiple logical representations.

Within the context of this project, the HIIS lab of ISTI-CNR has been developing TERESA, a publicly available tool. The aim is to support the design and development of nomadic applications. The method underlying TERESA is composed of a number of steps (see Figure) that allow designers to start with an idea of the overall task model of a nomadic application and then derive concrete and effective user interfaces for multiple devices through multiple levels

The main phases of the method underlying TERESA.

of abstractions and related in-between transformations.

A number of requirements have driven the design and development of TERESA:

- *Mixed initiative*, the tool is able to support different levels of automation ranging from completely automatic to highly interactive solutions where designers can customise or even radically change the solutions proposed by the tool. This is important to satisfy a variety of needs. Situations where the time available is short, the application domain is narrow, or the designer has no expertise call for completely automatic solutions. When designers are expert or the application domain is broad or has specific aspects, then more interactive environments are useful because they allow the designer to make important design decisions directly.
- *Model-based*, because the variety of platforms increasingly available can be better handled through multiple abstractions that provide designers with a logical view of the activities to support.
- *XML-based*, XML-based languages have been proposed for every type of domain. In the field of interactive systems there have been a few proposals that partially capture the key aspects to be addressed.
- *Top-down*, this approach is an example of forward engineering. Various abstraction levels are considered, and we support cases where designers have to start from scratch, first creating logical descriptions, and then moving on to more concrete representations until they reach the final system.
- *Different entry-points*, our approach aims at being comprehensive and supporting the entire task/platform taxonomy. However, there can be cases where only a part of the task or platform needs to be supported. For example, when only different devices (such as different brands of mobile phones) referring to the same type of platform are considered. In this case, there is no need for a nomadic task model, given that only one type of platform is involved.
- *Web-oriented*, due to the spreading use of the web, we decided that Web applications had to be our first target. However, the approach is also valid for generating user interfaces for other types of software environments, such as Java applications.

This approach has been applied to the development of a number of nomadic applications and has emerged as a promising and effective solution when addressing the problem of generating UIs for heterogeneous devices, capable of providing usable multimedia interfaces for a broader set of mobile devices, including vocal interaction techniques.

Links:
CAMELEON Project
http://giove.cnuce.cnr.it/cameleon.html.
HIIS lab, ISTI-CNR
http://www.isti.cnr.it/ResearchUnits/Labs/hiis-lab/
TERESA:
http://giove.cnuce.cnr.it/teresa.html

Please contact:
Fabio Paternò, ISTI-CNR
Tel: +39 050 315 3066
E-mail: fabio.paterno@isti.cnr.it

Carmen Santoro, ISTI-CNR
Tel: +39 050 315 3053
E-mail: carmen.santoro@isti.cnr.it

Interoperating with Heterogeneous Mobile Services

by Paul Grace and Gordon S. Blair

Mobile applications are now developed upon a wide range of service development platforms, commonly referred to as middleware. However, the diversity of those available presents a problem for mobile client development. How can a single client implementation interoperate with heterogeneous service implementations?

The emergence of mobile computing has created new classes of applications dependent upon the user's location, context and interaction with their current environment. However, given the constraints of the wireless environment (ie weak connection, poor network quality of service and mobile devices with limited resources) developing distributed applications within this domain is a complex task. Therefore, new middleware has emerged to mask such problems from the user. These encompass synchronous (eg remote method invocation) and asynchronous (eg publish-subscribe and tuple spaces) communication paradigms. However, it is the heterogeneous nature of the solutions that has created a new problem. Only applications and mobile services developed upon the same middleware style can interoperate with one another. For example, a tourist guide client application interacts with services available at the user's current location that provide tourist information. If the client is implemented upon a particular middleware platform (for example, SOAP) it will only function in locations that offer tourist services across SOAP.

Similarly, services are advertised using one of the contrasting service discovery protocols. At present, four main service discovery protocols exist: Jini, Service Location Protocol (SLP), Universal Plug and Play (UPnP) and Salutation. In addition, new protocols are emerging to better support the discovery of services in mobile environments and across ad-hoc wireless networks (eg Service Discovery Protocol in Bluetooth and

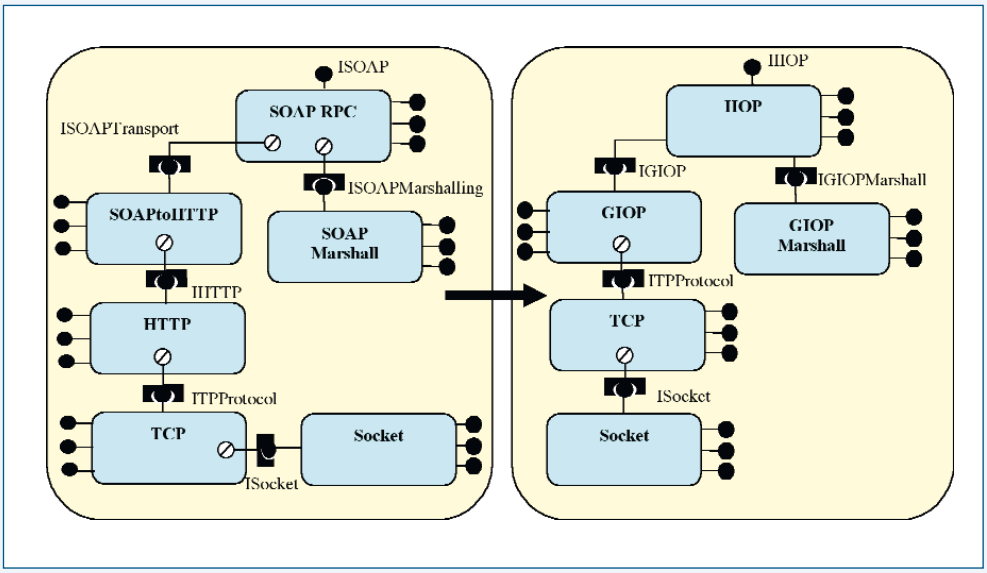

Dynamic change from SOAP client configuration to IIOP client configuration of OpenCOM components.

Salutation Lite). Mobile clients using a single discovery protocol will miss services advertised by alternative protocols.

Hence, we argue that adaptive middleware is required to support the interoperation of mobile clients with heterogeneous services. Using this approach, middleware dynamically alters its behaviour to: i) find the required mobile services irrespective of the service discovery protocol, and ii) interoperates with services implemented by different types of middleware.

The ReMMoC (Reflective Middleware for Mobile Computing) project, which began in October 2000 and concludes in September 2003, is being carried out at Lancaster University in collaboration with Lucent Technologies UK. The aim of the project is to design and implement such an adaptive middleware, which allows mobile clients to be developed independently from the underlying middleware implementations that may be encountered at different locations.

For applications to dynamically operate across different middleware they must be programmed independently from them. Hence, abstract definitions of mobile services are required. The client application, which requests this service, can then be developed using this "interface" in the style of IDL programming. A request of the abstract service is mapped at run-time to the corresponding concrete request of the middleware implementation. The emerging Web Services Architecture includes a Web Services Description Language (WSDL) that provides such an approach based on abstract and concrete service definitions. We employ WSDL as the basis of service description and programming model.

To build the adaptive middleware, we utilise three key techniques: components, component frameworks and reflection. We use OpenCOM (developed at Lancaster University) as our component model to develop components for Windows CE platforms. OpenCOM is a lightweight, efficient and reflective component model, built atop a subset of Microsoft's COM. The motivation behind component frameworks is then to constrain the scope for evolution, ensuring all dynamic changes between component configurations are legitimate. Finally, reflection is a principled method that supports introspection of the underlying middleware and provides techniques to dynamically adapt component architectures.

The ReMMoC platform consists of two key component frameworks for binding and service discovery, whose behaviour can be dynamically altered using reflection. The binding framework re-configures between different interaction protocols (we have implemented IIOP, SOAP and publish-subscribe bindings); the figure illustrates a dynamic change of personalities within this framework. The service discovery framework changes between one or more of the different service discovery technologies (we have implemented the SLP and UPnP protocols).

ReMMoC has been fully developed and tested using simple applications eg chat, news and stock quote clients. Ongoing work includes an evaluation of this method on larger, complex applications (eg ubiquitous computing and intelligent home environments) and across a range of further middleware bindings including data sharing and tuple spaces. An evaluation of memory use has concluded that each device cannot store every possible middleware component. Therefore, methods for dynamically downloading components when needed are required. Furthermore, techniques to ensure components are available to start-up before they are used, eg predictive caching based upon context information, are an interesting option.

Link:
OpenCOM and ReMMoC Web Page - http://www.comp.lancs.ac.uk/computing/research/mpg/projects/opencom/

Please contact:
Paul Grace, Lancaster University, UK
Tel: +44 1524 59 3801
E-mail: p.grace@lancaster.ac.uk

Gordon Blair (Project Co-ordinator)
Lancaster University, UK
Tel: +44 1524 59 3801
E-mail: gordon@comp.lancs.ac.uk

NAC: An Architecture for Multimedia Content Adaptation for Mobile Devices

by Nabil Layaïda and Tayeb Lemlouma

In the current Web infrastructure, access to content is achieved from a variety of terminals such as desktop PCs, smart phones, set-top boxes and personal digital assistants (PDAs). The different capabilities of these terminals has led to the development of a variety of parallel Webs, each accommodating a particular device feature. This situation poses a serious problem to Web designers and maintainers, since every piece of information needs to be authored for every type of terminal with the associated protocols. Research in the WAM project at INRIA attempts to unify Web access to information with a more flexible and negotiated approach.

Providing tailored content and presentation for different clients in heterogeneous environments is becoming increasingly important. Devices such as small palm computers, smart phones and pocket PCs are already common components of the computing infrastructure. Providing adapted content that takes into account the complete range of devices and user agents available is challenging, since the capabilities of these terminals are evolving at a rapid pace. On the other hand, Web content is mostly designed for the capabilities (screen resolution, bandwidth etc) of desktop PCs and to a lesser extent terminals with poorer capabilities such as first generation WAP-based systems. Designing an infrastructure flexible enough to cope with this situation requires the modification of several components of the current Web. At the lower layers, one of the basic requirements is the provision of a minimal knowledge about the different entities that compose the content delivery chain, ie client, server, document and media resources, etc. Starting from such knowledge, adaptation mechanisms and content-negotiation strategies could be applied to deliver to the target client content that takes advantage of its capabilities and satisfies its constraints.

The negotiation and adaptation core, called NAC, is an architecture developed in order to provide a solution for the delivery of multimedia content in heterogeneous environments. For content generation, NAC uses dynamic and static adaptation techniques. The adaptation is controlled using an adaptation and negotiation module (called ANM), and an optional module, called UCM (user context module), that runs on embedded devices. UCM allows the ANM knowledge about the client description to be enriched through the use of profiles. NAC default organisation is proxy-based (Figure 1), but the proxy entity can be omitted by using (installing) ANM at the server side.

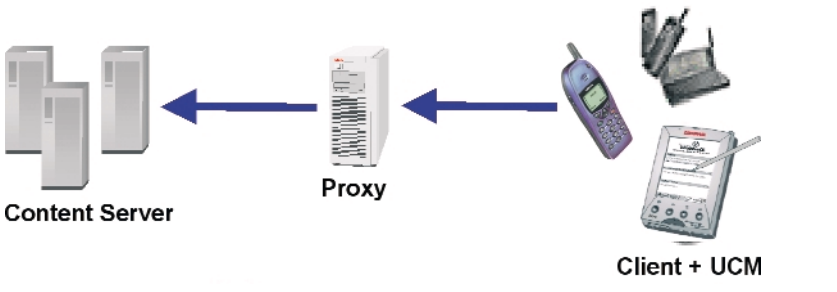

Figure 1: The NAC architecture.

The proxy is a third entity that stands between servers and clients. It allows better handling of the heterogeneity of clients and the variety of content representations on servers. Indeed, proxies allow content adaptation without affecting the existing components in the system. For the content generation by adaptation, the proxy (Figure 2a) is the entity responsible for retrieving client requests and profiles and performing possible adaptation on the content received from the original server. The adapted content is then sent to the target client with respect to its characteristics. The proxy can transform existing multimedia content, meaning content does not need to be authored in multiple versions. In such situations and in order to avoid the adaptation cost, a negotiation strategy is used to select the best variant to be delivered. All the proxy tasks are designed to behave transparently to clients and servers.

In NAC, the description of the environment (device capabilities and preferences, network limitation, content characteristics etc) is described as a set of constraints that the content provider should satisfy. In our approach, the constraints resolution strategy is achieved by adding constraints progressively to the original content.

The universal profiling schemata (UPS) was defined in the NAC framework to have a central role in the generation of adapted content by modelling the different profiles of the environments. UPS extends CC/PP to describe not only the client capabilities and preferences but also the capabilities of servers and proxies, the characteristics of the content

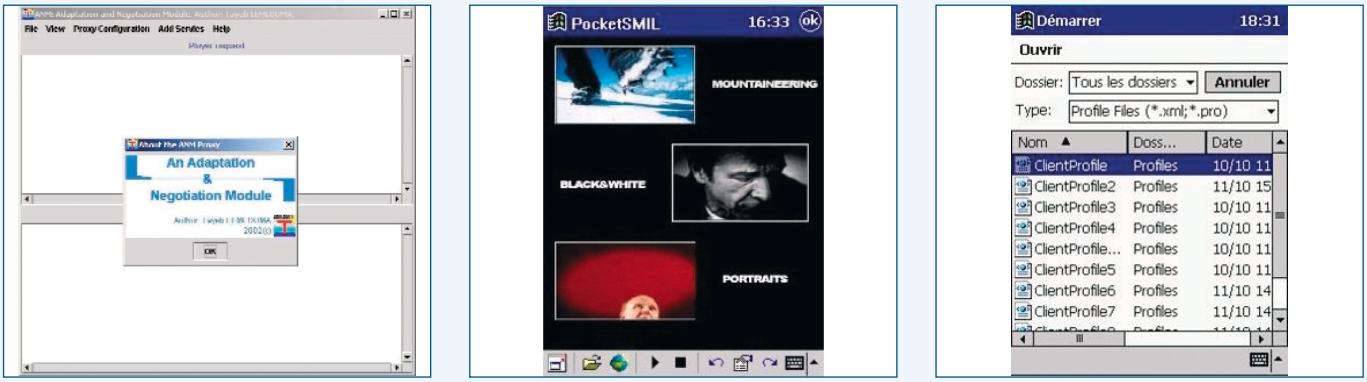

Figure 2: NAC entities: the ANM proxy, Pocket SMIL and the UCM module. From left to right: the proxy; Pocket SMIL, an embedded multimedia player (centre); the User Context Module (right).

and the adaptation methods that may exist at the server or the proxy level. UPS identifies three main categories of contexts: the client category, the server category and the network category. Profiles in the first category contain the description of the client characteristics (software, hardware and user agent) and the constraints related to the use of different media resources (supported media formats, size, resolution etc). The server category includes the document instance profile that describes the document characteristics and functionalities. It includes also the resource profile that describes media resources used in documents and the adaptation method profile that describes available adaptation methods that exist in the server or the proxy side. Finally, the network profile describes the network characteristics. ANM includes a set of adaptation techniques that ranges from real-time (re)encoding of media items such as video, to structural transformations such as XHTML to WML.

Links:
http://wam.inrialpes.fr
http://opera.inrialpes.fr/people/Tayeb.Lemlouma/NAC.htm
http://www.w3.org/Mobile/CCPP/
http://www.w3.org/2001/di/

Please contact:
Nabil Layaïda, INRIA
Tel: +33 476 61 53 84
E-mail: Nabil.Layaida@inrialpes.fr

Integrating GSM Networks and Internet: New Unstructured Services

by Rita Arrighi, Maurizio A. Bonuccelli, Francesca Lonetti and Francesca Martelli

A project at PisaTel, a laboratory located at ISTI-CNR in a joint collaboration between ISTI, Ericsson Lab Italy, Pisa University and the "Scuola S. Anna", aims at integrating wireless, wireline and Internet networks.

Three communication networks, namely the emerging 'mobile' wireless, the classical 'fixed' wireline, and the alternative 'computer-centric' internet, are evolving in a relatively independent way. However, their interactions increase day by day. It is of paramount importance to investigate how these three systems could be integrated in order to form a common communication space, in which the users of any one of the networks can easily and 'transparently' access the other two.

The main purpose of a project currently under way at PisaTel is to investigate the feasibility of such an integration. The work is being carried out within the java JAIN MAP API international project.

JAIN MAP

The aim of the JAIN (Java APIs for Integrated Networks) initiative [1] is to integrate wireline (PSTN), wireless (PLMN) and packet based (IP and ATM) networks, in order to modify the actual proprietary market into an open one. The novelty of JAIN lies substantially in two fundamental aspects: the service portability obtained by inserting Java interfaces over the protocols, and the network convergence ensuring the technology independently of the services.

We are interested in the JAIN MAP API. MAP (Mobile Application Part) is a protocol in the GSM stack, for mobility management and other services. JAIN MAP API provides an abstraction level over the complicated MAP interface. In this way, a MAP service developer does not have to be aware of the specific MAP implementation features. At present, four MAP capabilities are specified in JAIN MAP: transaction (corresponding to the SMS service), session (USSD service), position (MAP location service) and information. We have focused on session capability, realizing the USSD (Unstructured Supplementary Service Data) service, which allows information

exchange between a mobile station and a GSM network application.

The USSD service is very similar to the SMS service. The main difference consists in the session-oriented nature of the USSD service which makes it necessary to establish a session each time a customer (a mobile station) or a network application approaches a USSD service. As the USSD service is not a store-and-forward one, it is faster than SMS. A USSD session can be network initiated or mobile initiated: in the first case, the session is started by a network application; in the second one, the session is started by a mobile station.

Applications

We concentrated on integrating GSM networks and Internet. In particular, we designed and tested two new mobile initiated USSD services for GSM customers allowing them to easily interact with a server connected to Internet. The assumption on the customer side is minimal: customers only need a simple mobile station (eg cellular telephone handset) able to communicate on GSM cellular networks, and the network operator must be able to support the new services. This second assumption is not cumbersome, since the software for these two new services is not large and is quite fast to run.

The first service is e-mail related. The GSM customer first invokes the new service by inserting an appropriate code, followed by an internet ID, the password, and the address of a mail server. These data are delivered to an entity inside the GSM network, called GSMscf, which is also connected to Internet. The GSMscf opens a session with the mail server and, after being authorized by entering the user ID and password, requests the number of unread emails that are in the target mailbox. The answer received by GSMscf is then propagated to the GSM customer that asked for the service, and is shown on the display of the cellular phone.

The second service is designed for more sophisticated customers. It enables the activation of a secure Telnet session with a specified computer connected to Internet. The data exchanged during this session must be reproducible on the display of the available handset. The data exchange is performed similarly to the previous service, with more rounds between the handset and the remote computer, and with a particular emphasis on the security aspects in the Internet part of the data path (security in the GSM part should be guaranteed by the network operator).

In this research we showed that GSM customers can effectively interact with Internet-connected devices by easily developing new simple software applications, without changes in the system architecture. We shall continue in the production of new services for larger integration of Internet and GSM systems.

Links:
JAIN:
http://java.sun.com/products/jain/api_specs.html
USSD:
http://www.mobileussd.com/whatis.asp.

Please contact:
Francesca Martelli, ISTI-CNR
Tel: +39 050 315 3468
E-mail: f.martelli@isti.cnr.it

From Ad Hoc Networks to Ad Hoc Applications

by Benoît Garbinato and Philippe Rupp

A new framework for ad hoc applications is under development at the University of Lausanne in Switzerland. This thin layer, based on existing standards, will help developers build applications that include three features: mobility, peer-to-peer operation and collocation. As the framework is independent of the underlying network technology, the proposed solution will be adaptable to the emergence of the next-generation ad hoc network.

The goal of several recent research projects has been to study the new generation of mobile communication and information services, based on self-organisation (AODV, Terminodes, CarNet). Such systems have lately become very topical, with the advent of the peer-to-peer communication paradigm and the emergence of ad hoc network technologies. Many fundamental technical questions remain open however, and real business applications still need to be rolled out. One specific domain where the peer-to-peer phenomenon has had a major impact is data and media interchange on the Internet. Mobile devices will presumably be the next target, as the minimal requirements for multimedia applications are almost guaranteed.

An ad hoc network is “a transitory association of mobile nodes which do not depend upon any fixed support infrastructure. [...] Connection and disconnection is controlled by the distance among nodes and by willingness to collaborate in the formation of cohesive, albeit transitory community.” [Murphy et al]. But what are ad hoc applications? One could simply answer 'pieces of software that will run on ad hoc networks'. We believe that defining a class of applications (ad hoc applications) by coupling it to only one possible implementation technology (ad hoc networks) reduces the generality and effectiveness of the definition. Therefore we define an ad hoc application as a self-organising application composed of mobile and autonomous devices interacting as peers, whose relationships are made possible because of

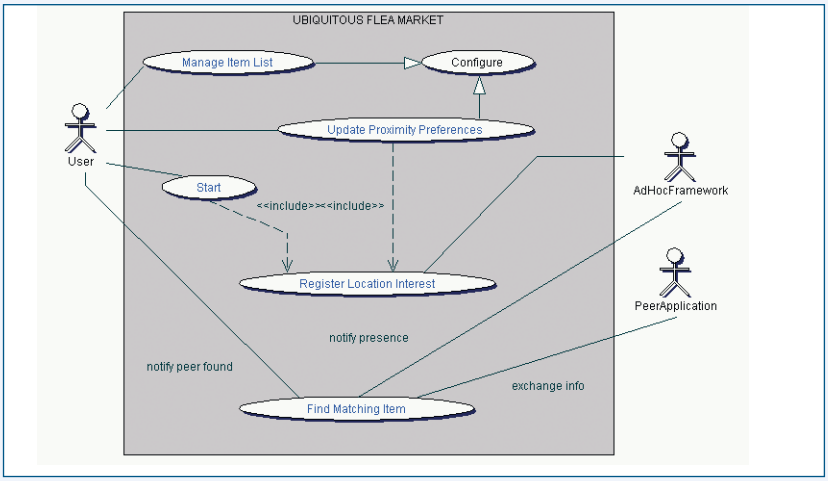

Figure 1:
The ubiquitous flea market use case.

relatively close physical distances (collocation). In addition, this dynamic community, defined by geographical proximity, needs to have a common (application-level) interest. More formally, three features must be present in an application for it to deserve the ad hoc label:

- Mobility: in order to be able to use the application everywhere, the user should not be limited by range. The range limit is set by the business logic of the application.
- Peer-to-Peer: direct communication between peers is mandatory. This means that the client/server relationship is defined in an ad hoc manner by the application logic (direct interaction between the pieces of software).
- Collocation: all logical interactions between applications must result in a physical interaction between users. This means that in order to be called an ad hoc application, the service has to be location-based.

The above definition allows us to abstract the network completely, and focus on the application aspects. In this way, there is a clear decoupling between the application and the network, and any type of network — GSM, WLAN or ad hoc — can be used. This decoupling is fundamental in understanding the key issues underlying ad hoc applications and in providing reusable solutions to solve these issues. This eases the development of ad hoc applications. By abstracting the network level, we are able to build ad hoc applications in the absence of any ad hoc network, as long as the underlying infrastructure can provide support for the three basic aspects described above.

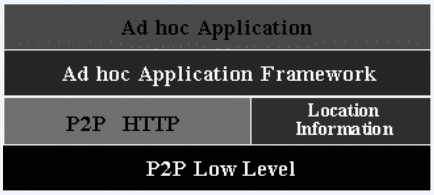

Figure 2: The different layers of an ad hoc application.

Here is an example of a possible ad hoc application called the ubiquitous flea market. This application is available wherever you are and at all times. It is available on many mobile devices and matches buyers and sellers present within a certain range, the latter being previously defined by the user. As the user is walking, this ad hoc application scans its surroundings for possible peer sellers or buyers. It has to be noted that any user can be buyer and/or seller. When the application finds another mobile device that runs the same piece of software, it scans the shared items in order to find a match. If there is a match, the user is alerted and can then ask the peer to get in touch and make the physical transaction.

The ubiquitous flea market (see Figure 1) is based on a framework to ease its deployment. Our so-called ad hoc application framework is based on existing standards, which makes it lightweight, and facilitates its adoption by developers. Some frameworks close to ours actually redefine, for instance, a transport protocol (PeerWare and Proem). This is not necessary as with the advent of Java-enabled mobile devices, most of these basic services are already available and standardised. Figure 2 presents the different layers for an ad hoc application when based on our framework.

The work is still in progress and the first version of the framework will be available in summer 2003. It will be demonstrated using a first implementation of an ad hoc application, the ubiquitous flea market described above.

Links:
Inforge - Institute for Business Information Systems Lausanne:
http://inforge.unil.ch

[Murphy and al.]:
http://citeseer.nj.nec.com/murphy98exercise.html

Terminodes project:
http://www.terminodes.org

AODV:
http://www.cs.ucsb.edu/research/trcs/abstracts/2002-18.shtml

CarNet:
http://www.pdos.lcs.mit.edu/papers/grid:sigops-euro9

PeerWare:
http://www.elet.polimi.it/Users/DEI/Sections/CompEng/GianPietro.Picco/papers/esec01.pdf

Proem:
http://www.cs.uoregon.edu/research/wearables/Papers/p2p2001.pdf

Full version of the paper:
http://isi.unil.ch/radixa/usr/prupp

Please contact:
Benoît Garbinato and Philippe Rupp,
Inforge, University of Lausanne, Switzerland
E-mail: benoit.garbinato@hec.unil.ch,
philippe.rupp@hec.unil.ch

Light Support for Dynamic and Pervasive Services on P2P Networks

by Stéphane Frénot, Anis Krichen and Stéphane Ubéda

As everything becomes dynamic and pluggable, we must envisage applications created through a combination of ambient elementary services without the help of an infrastructure. The WaveCar project investigates the notion of elementary services in a highly dynamic environment, in order to create what are now referred to as ambient and hybrid networks.

ARES – the Architecture of Network of Services – has been set up as a new INRIA Project in 2000 and is now hosted by the INSA of Lyon, France. ARES works on the design, modelling and validation of software and hardware elements for deploying services and protocols into networks. Network architectures are becoming increasingly complex and heterogeneous both in terms of the technologies on which they are based and the areas in which they are deployed. The notion of communicating objects and ambient networks, in which interactions between network elements are spontaneous and transparent to the end user, significantly complicates the deployment of these infrastructures. Thus, new models have to be realised, new approaches have to be validated, and new architectures proposed with two principle concepts in mind: mobility and embedment.

ARES is currently developing dynamic software components that would be useful in middleware approaches, taking into account the specific characteristics of hybrid networks,namely:

- in the design of elementary services, mobility must be considered the rule, rather than the exception
- potential loss of availability of some nodes and strong variation in the quality of service is also the rule, not the exception.

Many concepts have to be revisited because of this particular context, especially the service discovery engine and notions of service deployment and administration.

WaveCar

The aim of WaveCar is the dynamic deployment of services in a distributed environment. In this project, mobility is one of our main interests, since deployment of services will take place on a car-embedded platform. The idea is to equip the car with a system which will provide synchronisation with the requisite client services through an Internet connection.

The other important point is that the car is not always connected to the network; in fact it must be near an access point to get a connection through radio waves. This off-line mode shows the importance of synchronisation between the client expression of needs and the effective deployment of services.

Our architecture separates three important elements of our system - the client administration platform, the service point (access point) and the autonomous physical entity:

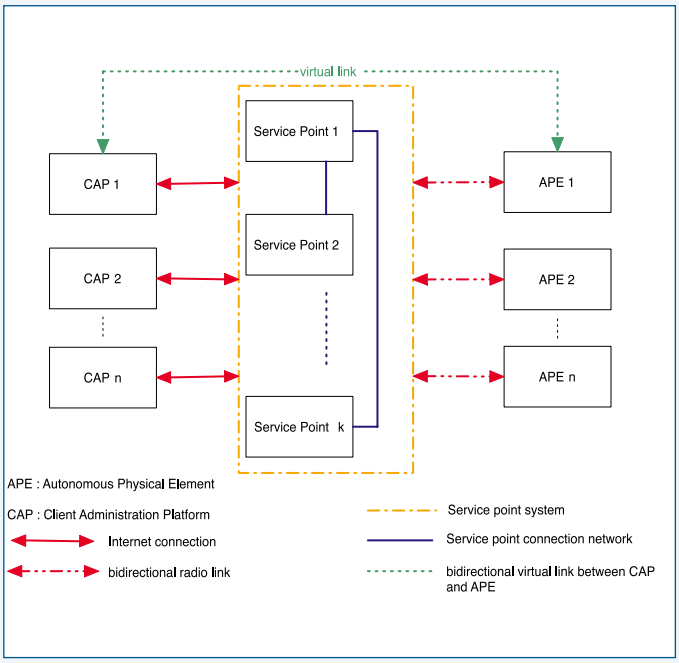

Figure 1: WaveCar functional architecture.

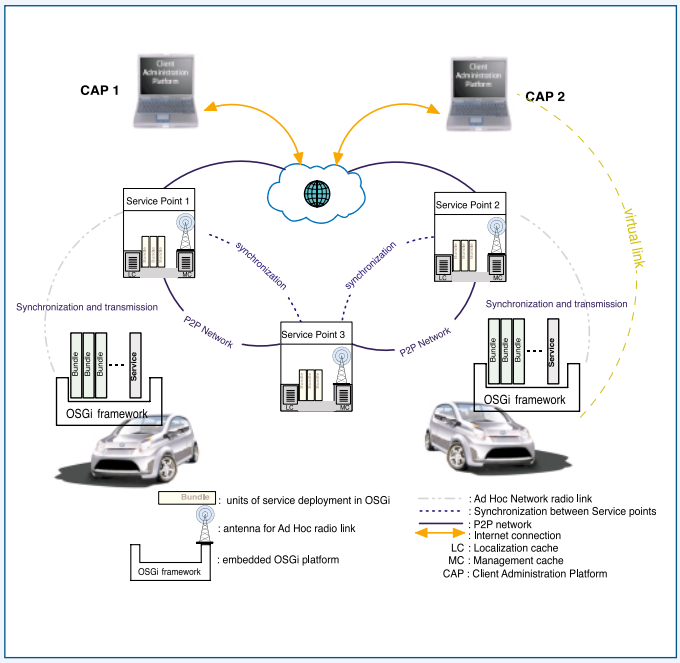

Figure 2: WaveCar general architecture.

- *Client Administration Platform (CAP):* This platform software is installed on the client computer. It has three main functions: to authenticate the client, to localise the client's car and to request management. The client can visualise the actual state of the embedded system and take management actions. These are performed immediately if the car is connected to the network, otherwise synchronisation will be performed next time the car connects.
- *Service Point:* Service points are the core of the system. These are all connected through a P2P network. This is an interface between the CAP and the autonomous physical element. It essentially provides a network connection to the embedded system, a caching system for localisation and requests management. This caching system is essential for synchronisation. All service points must communicate in order to achieve homogeneous management tasks like physical localisation of cars or event-sending for upgrades concerning a special Autonomous Physical Element.
- *Autonomous Physical Element (APE):* This is the car's embedded system. In our case it is a platform for the deployment of services which is in conformity with the OSGi specification. The role of this platform is to dynamically load new services on demand. Services are represented by bundles which can be downloaded from a service point and directly managed through the service platform. These services could be MP3 players, vocal post-it, mail clients etc.

Links:
http://citi.insa-lyon.fr/
http://darts.insa-lyon.fr/
http://www.insa-lyon.fr/

Please contact:
Stéphane Frénot, CITI Laboratory
at INSA-Lyon Scientific and Technical University, France
Tel: +33 04 72 43 64 22
E-mail:stephane.frenot@insa-lyon.fr

Towards a Virtual Home Environment Application Server

by Ali Hafezi and Klaus-Peter Eckert

A Virtual Home Environment (VHE) is a means for the context-dependent provision of telecommunication-related services to nomadic users. The Fraunhofer institute FOKUS has participated in several national and international projects that have defined VHE architectures and implemented appropriate platforms. While the current work focuses on the personalisation and adaptation of services, future work will concentrate on operational aspects of VHE systems including the flexible deployment of services, comparable to the deployment of components in a supporting application server.

The future success of wireless networks like UMTS and WLAN depends on the large availability of innovative services. Such services will drive network usage up for data and voice services and will hopefully make network technologies like UMTS worthwhile over the long term. So far the services provided by a particular operator have been strongly tied to the particular capabilities of the operator's network vendor. This vertical approach to service provisioning undermines innovation; services are excluded from specialised service developers because the services typically have to run in a proprietary execution environment. In the last few years there has been considerable worldwide turmoil in the area of communication service development and deployment, leading to different service architectures, different interfaces and different service environments. On the other hand, the third-generation mobile networks are currently under intensive study by international standardisation bodies such as the ITU and ETSI, as well as the European Commission. Where possible, Fraunhofer FOKUS has tried to merge these resulting architectural and network concepts in order to provide a global service-provisioning architecture, coming from the different solutions and capabilities suggested and provided by various projects.

Starting from TINA-C's service architecture and appropriate inter-domain reference points, Fraunhofer FOKUS has developed a series of service execution environments. These provide well-defined interfaces between services and execution environments on the one hand, and between execution environments operated by different providers on the other. One lesson learned was that it is very hard to agree on a global standard for the interfaces between services and their execution environment. Thus every execution environment defines its own internal interfaces to services. These interfaces become more and more complex when the execution environment provides VHE support. Additional interfaces between the service and adaptation and profile components have to be introduced. Appropriate topics have been investigated in a series of VHE-related projects like Vesper, AlbatrOSS or VHE-UD.

A new category of problems arises in the case of mobile, nomadic users. These

users need personalised access to their services even if they connect from other network or VHE providers. In such situations, either interworking is necessary between service instances running in the home and visited provider's domain, or cross-domain access to profile information is required. These scenarios are investigated in the AlbatrOSS project. Again the lack of globally standardised interfaces between services and execution environments prevents the provision of a comprehensive solution.

One possible solution to these kinds of problems is the introduction of service deployment descriptors that describe the requirements and obligations of services towards their users and execution environments. As in well-known component technologies like EJB or CCM, the services as well as the execution environments describe their required and provided interfaces and the appropriate protocols. Thus it is possible to map both descriptions either manually or automatically and to generate bridging components that adapt the technology and protocols between services and execution environments. First steps in this direction have been made in the VHE-UD project.

This adaptation during the deployment time of a service needs only static information. It can be used to satisfy service-level agreements and to guarantee some agreed quality of service during the start-up phase of a service. Much more complex is the task of controlling and justifying the quality of service during the execution time of the service. Components are required that monitor the status of the network and of the services, the end-to-end evaluation of the provided service quality is necessary and an appropriate end-to-end reconfiguration of the service must be enabled. In a fully operational environment, the monitored status information will also be used for logging and accounting purposes.

In current projects, Fraunhofer FOKUS has developed various building blocks for a comprehensive VHE platform together with a series of component-based, adaptive middleware environments. It will be the challenge of future projects to integrate these ideas, experiences, and software modules into a global VHE architecture built around the concept of VHE application servers.

Link:
http://www.fokus.fraunhofer.de/platin/projects

Please contact:
Ali Hafezi, Fraunhofer FOKUS
Tel: +49 30 3463 7109
E-mail: hafezi@fokus.fraunhofer.de

M-commerce API Package for Mobile Phones

by Ioannis G. Askoxylakis, Diomedes D. Kastanis and Apostolos P. Traganitis

M-commerce, the new buzzword in today's mobile industry, opens new horizons for mobile phones and services. Originally used to place voice calls, today's mobile phone uses data, multimedia, entertainment, and mobile commerce services, which are expected to overtake voice in revenue generation for operators.

The market for mobile phones, handheld computers and wireless PDAs is increasingly driven by multimedia-based Internet applications. New demands relevant to mobile commerce services include the electronic purchase of tickets, goods, or audio/visual content. The Telecommunications and Networks Laboratory (Netlab) of ICS–FORTH is developing and implementing an M-commerce API package for mobile phones, intended to provide mobile commerce services, achieve widespread usage, and offer unique benefits over and above the alternatives. The design of the corresponding API complies with secure hardware modules already installed on mobile phones (eg SIM cards) and upcoming advanced secure memory modules (SecMMCs).

To achieve this, M-commerce has to be based on secure applications and solutions to gain trustworthiness. SecMMC at ICS-FORTH's NetLab places strong security at the centre of this project. The challenge is to implement a security scheme that meets the end-user requirements and which, at the same time, is convenient and simple to use. The potential revenue that multimedia and data services can bring to the industry depends on the end-user perception of security and trust, ie, whether the mobile can become a 'digital wallet'. M-commerce should be viewed in the context of the number of ways that end users will be able to pay for goods and services.

To that end, we are developing secure mobile phone applications, and establishing a framework for secure mobile transactions to support a variety of applications and services. In order to maximise the potential of such services, a number of factors are crucial for success.

Firstly, end users require device ownership: although anyone can pick up the device, only its owner must be allowed to carry out transactions involving personal data. Secondly, end users require secure and reliable networks. Thirdly, all brands and companies supplying services must be able to be trusted. This is true for all components of the value chain, including the financial institution, the retailer, the network oper-

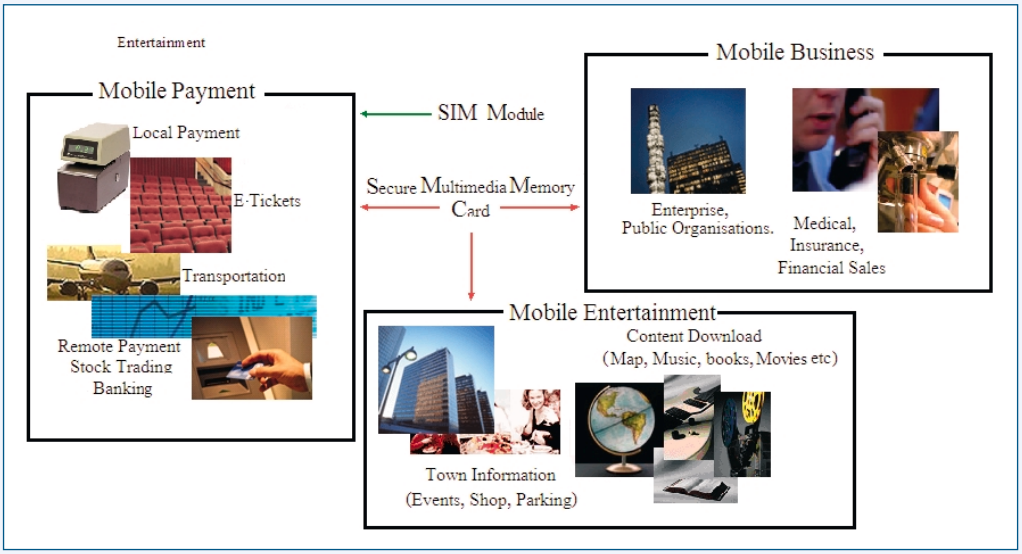

The M-commerce API for mobile phones focuses on three basic areas of application: mobile payment, mobile entertainment and mobile business.

ator and the device manufacturer. Other important factors include allowing end users to record and print transactions, ensuring that services are simple, easy to use, and reasonably priced.

Mobile commerce applications require the existence of a secure module. In our case we consider a variety of security modules, including:

- the broadly used SIM-card module
- the forthcoming Secure Multimedia Card (SecMMC) module
- similar smart memory devices (as specified in the draft standard 'Mobile Commerce Extension Standard-McEX; Core specification' by the 5C group).

(The 5C Group is an industrial standardisation group of the five main Multimedia Card manufacturers: Infineon Technologies Flash Ltd, Hitachi Ltd, Matsushita Electric Industrial Co Ltd (MEI), Toshiba Corporation, and SanDisk Corporation.)

The SecMMC can be viewed as a classical Multimedia Card with an integrated smart-card kernel, which combines the high security of smart cards with the large memory capabilities of classical Multimedia Cards. Obviously the supported services depend on the security module being used (SIM module or SecMMC). For example, whereas both module types can provide an electronic payment system, memory-demanding multimedia applications can only be supported by the SecMMC.

The SecMMC is being developed by a consortium consisting of Guardeonic Solutions AG Germany (project coordinator), Infineon Technologies Flash Ltd Israel, IAIK University of Graz, Technische Universität München, Fraunhofer-Gesellschaft-IGD, FORTH, Infineon Technologies, Flash GmbH&Co KG and Mühlbauer AG.

Our development of the M-commerce API for mobile phones focuses on three basic areas of application:

- mobile payment
- mobile entertainment (or infotainment)
- mobile business.

Mobile Payment applications allow personal payment transactions like e-ticket purchase, local merchant payments, stock trading, and eventually telematic and transportation usage. This type of application can be supported by both SIM and SecMMC modules. Key characteristics of this application are:

- *Anonymity:* during an economic transaction nobody can extract any information concerning the corresponding transaction pair.
- *Off-line operation:* two parties can complete a transaction without the intervention of a third centralised intermediary entity.
- *Transferability:* each entity of the system should have a hybrid role, both as a payer and as a payee.
- *User mobility:* the payment process is independent of the location of the transaction pair.

In addition, Mobile Entertainment applications involve copyrighted content (music, video and eventually navigation content). Consumers want to personalise their mobile phones with ringtones, graphics and picture messages from content providers. Games, downloadable applications, music and video feeds are sure to follow, as secure modules like the SecMMC become available for integration into mobile phones. In some regions of the world this type of system is already in use to a limited degree.

Mobile Business applications securely connect mobile users to both internal company and external public organisations. This type of application is an essential part of comprehensive sales force automation implementation.

All these emerging mobile services involve large volumes of multimedia data and require high levels of security to protect the transmission and storage not only of financial accounts and PIN codes but also of the multimedia data that is downloaded and stored.

Please contact:
Apostolos P. Traganitis
Tel: +30 2810 391724
E-mail: tragani@ics.forth.gr

Ioannis G. Askoxylakis
Tel: +30 2810 391723
E-mail: asko@ics.forth.gr

Diomedes D. Kastanis
Tel: +30 2810 391723
E-mail: diomedes@ics.forth.gr

Wireless Ad Hoc Network for Dublin: A Large-Scale Ad Hoc Network Test-Bed

by Stefan Weber, Vinny Cahill, Siobhan Clarke and Mads Haahr

The migration of wireless networks from hotspots to multihop ad hoc networks is the first step towards self-organized global routing. The WAND project is deploying a large-scale test-bed for ad hoc network protocols in the centre of Dublin to facilitate research in this emerging discipline.

Wireless networks are rapidly becoming common place; upcoming technologies support wireless communication on chip-sets; new laptops have built-in WiFi cards; hotspots are being installed in airports, hotels and coffee shops; and offices are converting their existing local area networks to wireless to allow their workers take advantage of mobility. Soon numerous embedded devices - that form the heart of the ubiquitous computing revolution - will join this domain and increase the number of participants exponentially. Existing network architectures that use dedicated central nodes to coordinate the communication between participants will become more and more infeasible as the number of participants increases by orders of magnitude and where devices will more likely converse with one another in an ad hoc fashion.

WAND Route from Trinity College Dublin to Media Lab Europe.

Ad hoc networks present the next great challenge for distributed systems research. In an ad hoc network, the participants are used to route communication traffic from senders to receivers. Every participating node in these networks executes a routing algorithm that allows messages to be directed towards the next node along a route to the receiver. Protocols that support communication in ad hoc networks have to take into account the mobility of the participants and the variation in the connectivity between associated parties. New applications in these ad hoc communication scenarios can take advantage of the peer-to-peer characteristics of the network and offer location-based services that address specific communities of users.

To support research in this area the Distributed Systems Group in the Department of Computer Science at Trinity College Dublin, in collaboration with Media Lab Europe (MLE), is deploying WAND, the Wireless Ad hoc Network for Dublin, as a large-scale test-bed for ad hoc network protocols and applications. The network will cover the centre of Dublin along a 2km route from Trinity College to MLE. This area will be seeded with a number of custom-build wireless-enabled embedded PCs. These PCs are 3x3x6 inch containers that accommodate a stack of PC/104 boards, WiFi PCMCIA cards, and a set of antennae. The embedded PCs will be hosted in apartments and shops, on traffic lights, and in phone kiosks along the route to provide a minimum level of connectivity.

The embedded PCs form a sparse population of wireless network nodes. This sparse coverage is constantly available and the embedded PCs can be configured to create a variety of network models. The test-bed can be further populated through the introduction of mobile nodes such as laptops, PDAs, and other mobile devices with wireless connectivity. This ability to be configured in various ways and to be populated sparsely or densely enables researchers to develop and investigate protocols and applications that are specific to the area of ad hoc networking in various scenarios.An example of such an investigation is the examination of routing protocols for ad hoc networks. The routing of messages in an ad hoc network represents a difficult problem that is addressed by several mature or recently proposed routing protocols. These protocols have distinct advantages and disadvantages and it is important to determine the exact circumstances under which each of these protocols represents a good choice. The WAND test-bed offers a unique opportunity to explore the behaviour and performance of routing protocols in a real-life environment and to investigate requirements that are necessary to provide services such as location-based information or quality-of-service provision.

Another area of research that is investigated with the help of the WAND test-bed concerns the adaptation of applications to wireless ad hoc networks. Existing applications are rarely designed to take advantage of additional informa-

tion and services that are offered by ad hoc networks. The applications range from peer-to-peer file sharing mechanisms that facilitate the exchange of data in ad hoc networks to location-based services that provide data to users depending on their physical position in the network.

An example of the applications that will be explored with the help of the WAND test-bed is communication between vehicles. This application will help mobile network nodes like cars, buses, and trucks communicate with one another and with static network nodes such as traffic lights, bus shelters, and traffic cameras. Mobile participants may inform one another about obstacles on the roads, actions that are going to be taken to avoid obstacles, or the routes of emergency services vehicles. Static nodes may provide information about the state of the traffic in an area and exploit information from mobile nodes such as the number of passengers or the estimated arrival time of buses.

The WAND test-bed offers researchers a unique opportunity to investigate low-level mechanisms such as routing protocols and low level services, as well as user-level applications in a real-life environment that reflects all the randomness and unpredictability that is extremely difficult to be reproduced with simulation.

Links:
http://wand.dsg.cs.tcd.ie
http://www.dsg.cs.tcd.ie
http://www.mle.ie

Please contact:
Stefan Weber,
Trinity College Dublin, Ireland
Tel: +353 1 608 1543
E-mail: Stefan.Weber@cs.tcd.ie

Lightning Active Node Engine - An Active Network User Service Platform

by Sami Lehtonen

There is a strong trend towards wireless networking and smaller terminal devices. The advent of various wireless terminals such as PDAs, mobile phones, or similar handheld equipment, will result in greater and greater volumes of Internet traffic. The key issues are support for seamless roaming, security and content adaptation.

Lightning Active Node Engine (LANE) is an active network platform developed at VTT Technical Research Centre of Finland. It is based on a decoupled active node architecture which separates the active applications from non-active IP-related routing. This decoupling was used in former active network projects Caspian and MAO.

The LANE architecture is depicted in the Figure. Decoupling is implemented by dividing the Active Node into an Active Router and an Active Server.

The Internet lacks a number of services relevant to mobile users such as automatic conversion proxies, dynamic firewalls etc. Active Networking or programmable networks can answer these service requirements. However, this raises some fundamental implementation issues on Active Nodes.

These issues have been addressed in the system design of LANE, which is based on the Application Layer Active Networking (ALAN) concept. The work is divided between the two parts, the Active Router (AR) and the Active Server (AS). The routing in the Active Router is dynamically configurable, and all time- or processor-consuming activities (or Active Services) are located and executed in the Active Server.

The Active Server handles the creation, execution and management of active services. All active services are on-the-fly deployable Java binaries. For security reasons, users can only reference services by name. Active services are designed and implemented for the users rather than for the servers. As a proof of

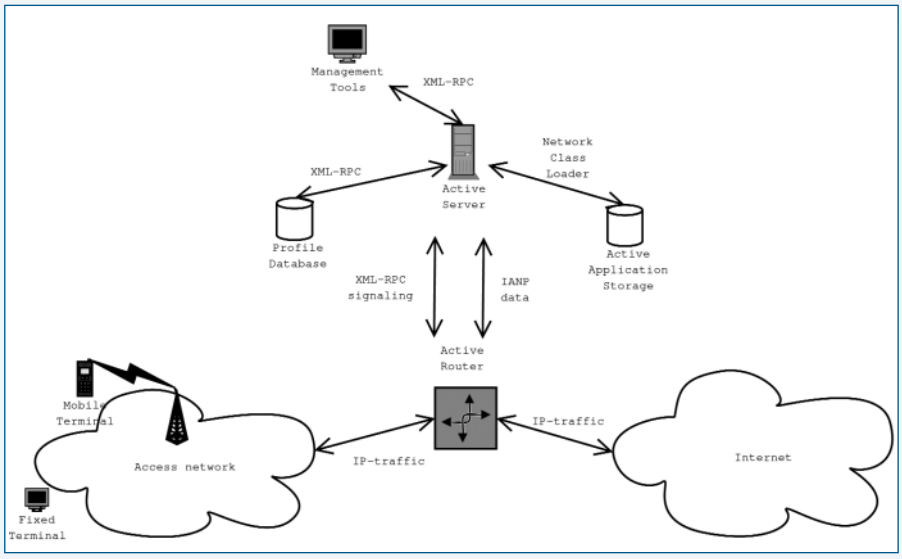

The LANE architecture

concept we have implemented dynamic Mobile IP support for our Active Server framework.

Services needed by mobile users can be located anywhere in the network, follow the user, and balance server loads. In our system, we designed an entity called the User Agent. This entity is created by the Active Server according to the user's profile, which is located in a user profile database. There are User Agents with different privilege levels for different purposes (or users). The security of the design is ensured with separately implemented User Agent classes. These classes implement only the interfaces that are granted for a particular user.

The User Agent represents the user on a particular Active Server. All services for the user at that Active Server are managed by the User Agent, and the user might have his/her User Agent instantiated in several Active Servers around the network. This makes it possible to move services from one Active Server to another, balancing the load between Active Servers. If the user then moves away, abandoning the User Agent, the removal of the User Agent removes all services for the user at that Active Server.

In our system, the services are mainly implemented as atomic, meaning that a single service instance serves only one user. We have followed a strict design principle, namely, there are multiple services for a user and multiple users at a server.

The research group included Kimmo Ahola, Mikko Alutoin, Markus Aroharju, Arto Juhola, Titta Koskinen, Sami Lehtonen, Marko Lyijynen, Sami Pönkänen, and Pertti Raatikainen.

Future work will consist of implementing load-balancing between User Agents located in different Active Servers. We will implement a method to transfer the state of a User Agent and its services to another Active Server (when a Mobile Node moves to another foreign network). Furthermore, implementing other useful services will prove the design of our concept.

Please contact:
Sami Lehtonen, VTT
Tel: +358 9456 7240
E-mail: sami.lehtonen@vtt.fi

Middleware Support for Ubiquitous Computing

by Luc Hogie and Christian Hutter

Mobile devices such as Personal Digital Assistants (PDAs) are currently used mostly for personal information management. However, these devices are becoming powerful enough to run much more sophisticated applications not only in an isolated manner, but even, due to their wireless communication capabilities, in a cooperative way. We are looking at middleware support in the domain known as ubiquitous computing.

In order to highlight some of the challenges inherent to ubiquitous computing, we use the example of ubiquitous gaming. Based upon elements of traditional multiplayer games like 'The Settlers' or 'SimCity', we introduce various ideas of how to adapt them for use in a ubiquitous environment.

For instance, while travelling by train, passengers could play a game together. At train stations, some players might leave while others join the game, possibly taking over existing roles. People could also collaborate, joining forces as well as dividing work. In addition to that, the environment itself could have an influence on the game. For instance, train stations might represent marketplaces in which a large variety of goods are offered and collaboration with other players waiting for their trains is possible. Aside from that, several sophisticated services, hosted for example on a powerful stationary node, might be available only at such special locations (see Figure).

Obviously, middleware support for ubiquitous applications is a very big field. Currently we are focusing on two specific areas within this complex domain, which will be discussed throughout the subsequent sections.

Information Brokering

One problem area to be tackled is that of sharing and distributing information between network nodes, ie, not only between all participants in a single application, but also across different applications, where some information (eg, where other devices are likely to be found) might be of general interest. For that purpose, we introduce the concept of information brokers. An information broker is a service running locally on each device, collecting and disseminating data in the background while within the coverage area of other devices.

From a logical perspective, a single set of information is available via the local broker. Physically, however, data is fully distributed over all nodes, with a local cache holding copies of the most interesting information. Queries are issued against the local broker, which can not only return data stored locally, but also

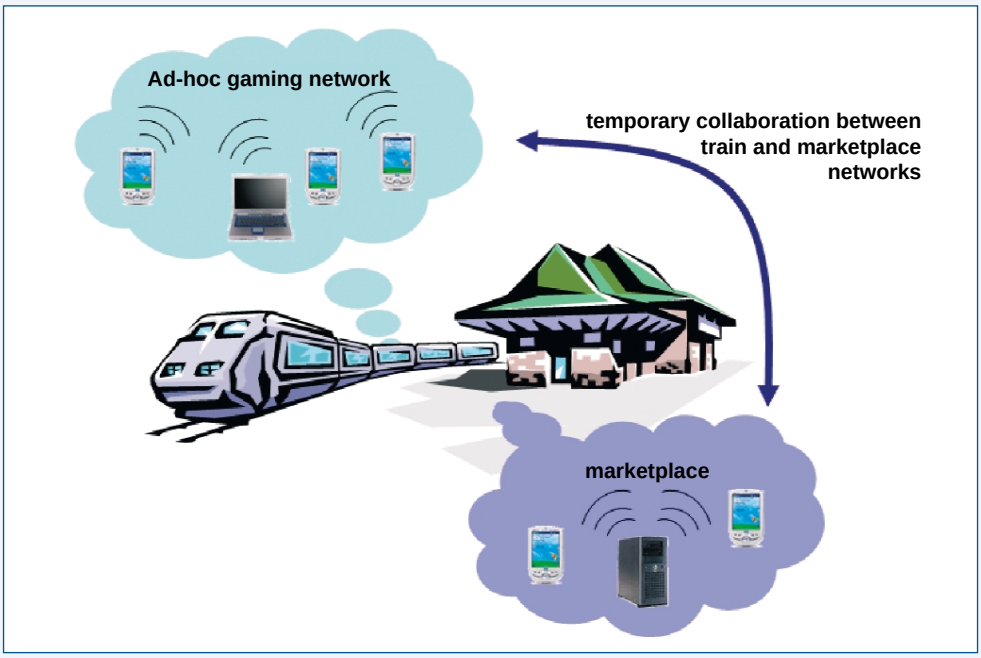

Illustration of ubiquitous gaming scenario.

send agents into the network. These agents collect information from devices in the local surroundings, thereby supplementing the query results in an incremental way. For non-time-critical queries that are allowed to run for some time, it is even possible to collect information from nodes that were not connected to the network at the time that the query was initiated. In the game example, a player might want to look for the best place where to buy or sell items of a specific type, dependent upon the current price and other criteria. As this kind of information might become outdated after a period of time, appropriate strategies, eg for limiting information lifetime, need to be provided. In particular, the trade-off between replica consistency and communication overhead is a complex research topic.

In the context of this task, we are collaborating closely with the University of Trier, Germany, where a macroscopic simulation environment for multihop ad hoc networks has been developed. Information brokers represent a service integrated with that environment. By using the simulation, we can test our algorithms with a large number of devices. The Trier team in turn employs our service in their own applications, thereby giving us valuable feedback.

Service Architecture and Mobile Grid

Aside from sharing information between the nodes in the network, it is also sensible to share resources, eg in terms of computing power and services. Up to now, scientific and business applications have benefited from networking mostly by using distributed resources for storage and retrieval of information or distributing the computational effort on different nodes.

Most research on distributed and parallel computing is based on an existing and fixed network infrastructure and therefore leads to fairly static solutions such as PVM, MPI, CORBA, RMI, etc.

With the emerging ubiquitous computing paradigm, certain significant limits relevant to current middleware must be addressed. In particular, middleware is:

- usually tied up with a central naming mechanism
- built for use of workstations or servers, meaning most current middleware is heavy and therefore unsuitable for use on lightweight devices.

Our aim is to define elegant models of a service architecture for ubiquitous applications. We are interested in the concept of a large-scale operating system using the resources available not only on the local machine but everywhere in the network. This might be considered as an extension of current well-known peer-to-peer networks like Kazaa or Gnutella. These open networks focus on file sharing. The idea is to allow nodes in such a network to act as service providers. Assuming a device can make available services of any type, the network is able to represent a computational grid, a data-sharing system, or even a distributed database. We are currently working on the model (using the metaphor of human society and communication) and its implementation. Once completed, it will be the lowest layer of (and a first step towards) a middleware and runtime environment dedicated to applications for the mobile grid.

This research is done in collaboration with the Computer Science Laboratory at the University of Le Havre (France), where work is currently being done on the specific issue of load balancing for ad hoc networks.

Please contact:
Luc Hogie, Christian Hutter
Luxembourg University of Applied Sciences
E-mail: Luc.Hogie@ist.lu,
Christian.Hutter@ist.lu

Sentient Objects: Towards Middleware for Mobile Context-Aware Applications

by Gregory Biegel and Vinny Cahill

Despite rapid advances in sensor technologies making a wealth of environmental data available to applications, programmer support for the development of context-aware applications remains poor. The majority of such applications are developed in an ad-hoc, application-specific manner and there is an urgent need for middleware and services to support the development of such applications. Mobile, intelligent software components known as sentient objects provide one middleware abstraction that can ease the development of mobile context-aware applications.

The recent proliferation of cheap, small, and increasingly accurate sensor technologies is creating a new information revolution where applications that interact with the physical environment are becoming widespread. This awareness of the physical environment, and its use by applications in fulfilment of their goals, is known as context-awareness and a number of promising applications have appeared which make use of context information. Context-awareness is of particular importance in mobile environments where the operating environment is constantly changing due to the mobility of devices and the characteristics of wireless communication technologies. In such environments, context-awareness can enable applications to respond intelligently to variable bandwidth, unreliable connections and the economy of different connections. One of the greatest challenges in context-aware computing, and one that has not yet been adequately addressed, is the provision of middleware and services to support the application developer. The major problem lies in providing generic support for the acquisition and use of multiple fragments of information gleaned from (potentially unreliable) multi-modal sensors in a mobile environment.

The use of context information by applications in a mobile environment poses a number of challenges arising from the distributed and dynamic nature of sensors, the accuracy and resolution of sensors, and the fusion of output of multiple sensors in order to determine context. In addition, the mobile environment poses further challenges with regard to the dependability, predictability, and timeliness of communication. Middleware is required that provides abstractions for the fusion of sensor information to determine context, representation of context, and intelligent inference. Essential services that provide support for operation in a mobile environment, such as supporting the reliability of communication, are also required.

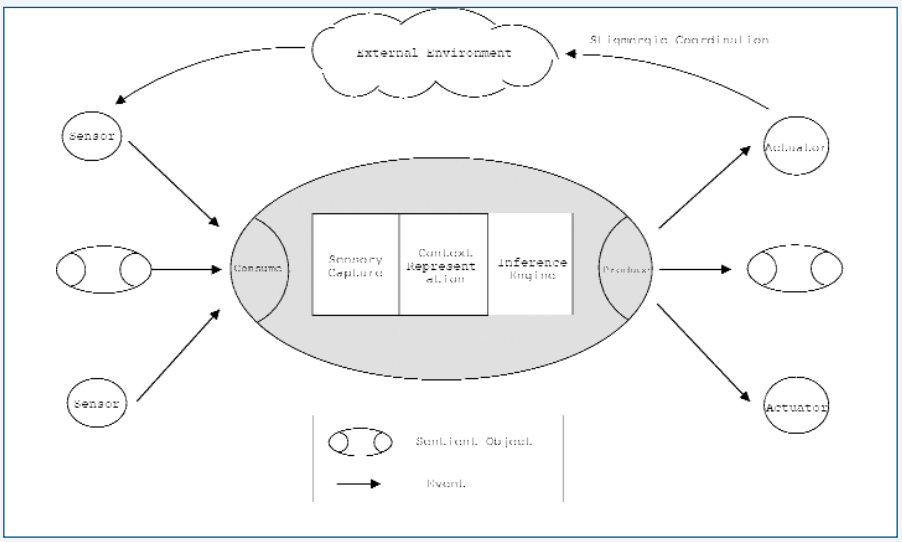

The sentient object model.

The sentient object model defines software abstractions that ease the use of sensor information, and associated actuation, by context-aware applications. At the heart of the model is an event based communication model that permits loose coupling between objects and consequently supports mobility and application evolution. The event-based communication model includes mechanisms for the specification of constraints on the propagation and delivery of events. A sensor is a software abstraction that encapsulates a hardware sensor device and produces software events in response to real-world stimuli. An actuator consumes software events and encapsulates a physical device capable of real-world actuation. A sentient object is an entity that can both consume and produce software events, and lies in some control path between at least one sensor and one actuator.

Internally, a sentient object consists of three major functional components:

- the sensory capture component is responsible fusing the outputs of multiple sensors, and uses probabilistic models, including Bayesian networks, to deal with inherent sensor uncertainties
- the context representation component maintains a hierarchy of potential contexts in which an object can exist, and the current active context

- the inference engine component is a production rule based inference engine and supporting knowledge base, giving objects the ability to intelligently control actuation based on their context.

Sentient objects are cooperative and in addition to traditional forms of communication, the sentient object model uses stigmergic coordination, or coordination via the environment. Stigmergic coordination does not rely on direct communication between objects, say via TCP/IP, but rather depends on objects being able to sense and make changes to their physical environment. Simple rules, embedded in the inference engine, then govern actuation and consequently behaviour, according to what has been sensed from the environment. This type of coordination is extremely robust and is particularly suited to mobile environments where traditional communication channels may not always be available or economical to use.

We are developing middleware and associated services, based on the sentient object model, to support context-aware applications. Such middleware and services will ease the task of developing applications based on context perception in a number of key areas. Applications will be insulated from both the complexities of physical sensors, actuators, and associated protocols, as well as the fusion of sensor data to infer context and reduce uncertainty. The middleware will also provide a high-level rule specification language through which intelligence may be added to objects without requiring knowledge of the intricacies of a specific production rule system. Furthermore, an intuitive visual programming tool is under development that will permit the rapid development of context-aware applications.

The work described in this article was partly supported by the Future and Emerging Technologies programme of the Commission of the European Union under research contract IST-20000-26031 (CORTEX – CO-operating Real-time senTient objects: architecture and EXperimental evaluation).

Links:
Ubiquitous computing at Trinity College Dublin: http://www.dsg.cs.tcd.ie/index.php?category_id=228

Cortex Project: http://cortex.di.fc.ul.pt/

Please contact:
Gregory Biegel, Trinity College Dublin
Tel: +353 1 6081531
E-mail: Greg.Biegel@cs.tcd.ie

Vinny Cahill, Trinity College Dublin
Tel: +353 1 6081795
E-mail: Vinny.Cahill@cs.tcd.ie

Database Components on Chip

by Nicolas Anciaux, Luc Bouganim and Philippe Pucheral

Ubiquitous computing and ambient intelligence introduce the need for embedding data and evaluating database queries in a growing variety of ultra-light computing devices. The SMIS project at INRIA focuses on the design of database components embedded in a chip. Applications range from mobile computing to data confidentiality protection.

Pervasive computing is now a reality, and intelligent devices flood many aspects of our everyday lives. As stated by the Semiconductor Industry Association, semiconductors integrated in traditional computers currently represent less than 50% of a market of $US204 billion. As new applications appear, the need for database techniques embedded in various forms of lightweight computing devices arises.

Sensor networks gathering weather, pollution or traffic information have motivated several recent works (eg the COUGAR project at Cornell). They have brought out the need to execute local computation on the data, like aggregation, sort and top-n queries, either to save communication bandwidth in push-based systems or to participate in distributed pull-based queries. Other forms of autonomous mobile hosts include hand-held devices (eg personal databases, diaries, tourist information, etc). These can be used to execute on-board queries on data downloaded before a disconnection. Personal folders on a chip constitute another motivation to execute on-board queries. Typically, smart cards are used in various mobile applications involving personal data (such as healthcare, insurance, phone books, etc). In this context, queries can be fairly complex and their execution must be confined to the chip to prevent any disclosure of confidential data. Thus, saving communication costs, allowing disconnected activities and preserving data confidentiality are three different concerns that motivate the storage of data and the execution of on-board queries on lightweight computing devices.

The resource asymmetry of lightweight devices entails a thorough re-thinking of database techniques. Our research objective is therefore to design embedded database components that can match highly constrained hardware resources.

Preliminary studies led us to design and validate a full-fledged DBMS, called PicoDBMS, embedded in an advanced smart card platform (powerful CPU, tiny RAM, fast read-time and dramatically slow write-time in EEPROM stable storage). PicoDBMS is based on a compact and efficient ad-hoc storage and

indexation model. The PicoDBMS query evaluator takes advantage of the CPU power and the fast read-time in the EEPROM to process any type of SQL query without using RAM, whatever the volume of the queried data.

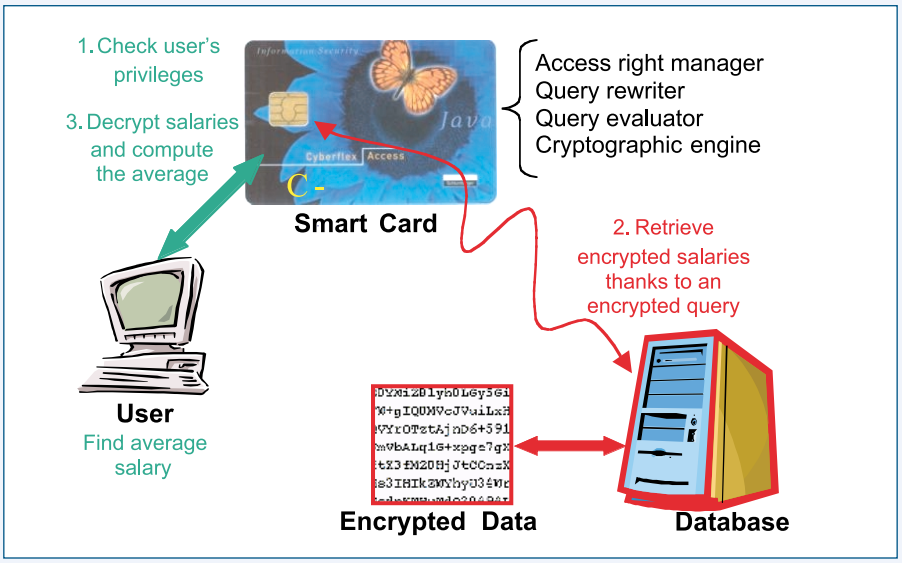

C-SDA functional architecture.

Capitalising on this work, new and important research efforts need to be undertaken:

- to better capture the impact of each device hardware constraint on database techniques
- to propose new storage, indexation and query techniques allowing ad-hoc embedded database components to be built
- to set up co-design rules helping to calibrate the hardware resources of future devices to match the requirements of specific applications.

As a first step in this research agenda, we are conducting a comprehensive analysis of the RAM consumption problem. We try to answer three important questions. First, does a memory lower bound exist whatever the volume of the queried data? Second, how can a query be optimised without hurting this lower bound? Third, how does an incremental growth of RAM impact the techniques devised in a lower bound context? This study provides helpful guidelines to calibrate the RAM resource of a hardware platform according to a given application's requirements, as well as to adapt an application to an existing hardware platform. In the mid-term, we plan to focus on energy consumption and on the usage of future stable-memory technologies.

Besides mobile computing and ambient intelligence, database components embedded in a chip may serve other important purposes in a wide range of applications. Typically, tamper-resistant chips (eg smart cards or tokens) can be used to make information systems safer. Data security is particularly emphasised in mobile environments. For instance, the increasing connection of travelling employees to corporate databases and the need to resort to Web-hosting companies and Database Service Providers to make personal data available to mobile users introduce new threats to data confidentiality. Users have no choice other than to trust the privacy policies of Web companies. However, according to the FBI, database attacks are more and more frequent (their cost is estimated to be more than $100 billion per year) and 45% of the attacks are conducted by insiders.

While traditional database security principles like user authentication, communication encryption and server-enforced access controls are widely accepted, they remain inoperative against insider attacks. Several attempts have been made to strengthen server-based security approaches thanks to database encryption. However, as Oracle confesses, server encryption has not proven to be the 'armour plating' that was expected, because the Database Administrator (or an intruder usurping his/her identity) has enough privilege to tamper with the encryption mechanism and get the clear-text data. Client-based security approaches, where decryption only occurs at the client-side, have been recently investigated but they cannot support data-sharing among users having different access rights.

The client-based solution we are developing at INRIA, called C-SDA (Chip-Secured Data Access), allows encrypted data to be queried while controlling personal privileges. C-SDA comprises an access right manager, a query rewriter, a query evaluator and a cryptographic engine, all embedded into a smart card to prevent tampering. As pictured in Figure 1, C-SDA acts as an incorruptible mediator between a client (potentially mobile) and the encrypted database. To illustrate the benefits of this approach, C-SDA allows complex permissions to be expressed on encrypted data, for example, 'give user X the privilege to compute the average salary of all employees without giving her access to the elementary data'. This cooperation of hardware and software security allows the orthogonality between access-right management and data encryption to be reestablished.

This work is partly supported by the French ANVAR agency (Agence Nationale pour la VAlorisation de la Recherche) and by Schlumberger, which provides us with advanced smart card platforms. This initial work considers relational data, pull-based queries and an SQL-compliant access-right model. New solutions will be devised to tackle more complex data (eg XML), push-based models (selective information broadcasting) and sophisticated access-right models. Some of these solutions will be studied in the context of a MEDEA+ European project scheduled for the fourth quarter of 2003.

Link:
SMIS project: http://www-smis.inria.fr/

Please contact:
Nicolas Anciaux, Luc Bouganim,
Philippe Pucheral, INRIA
E-mail: {Firstname.Lastname}@inria.fr

Object-driven Application Design: from Mobility towards Ubiquitousness

by Davide Carboni, Gavino Paddeu and Stefano Sanna

Techniques and frameworks that assist designers and programmers in the development of applications that are seamlessly portable from one device to another are still a challenge. The e-mate project at the Center for Advanced Studies, Research and Development in Sardinia (CRS4) tackles this issue proposing a technique based on the object-oriented paradigm with a strong focus on the delivery of services on multiple platforms.

At the present, users must carry one or more portable devices that act as data and program repositories or as connection devices. These devices need to be administered, backed-up, and synchronized; their batteries need to be replaced and recharged. This is not a 'pervasive technology', but rather an 'invasive technology'. The next step towards a new concept of services that support user mobility is known as 'ambient intelligence' where computers are no longer repositories for data and programs, but become windows to the environment: services and applications. Ubiquitous applications adapt themselves to the end user delivery context, becoming accessible from any device and giving the user the perception of moving in an empowered environment. Such requirements emphasize the concept of the 'position': 'ubiquitous computers must know where they are'(Weiser).

User mobility implies that applications should be accessible from a number of different channels: from a web browser, from a WAP phone or from an Interactive Voice Responder (IVR) machine. A given channel can provide multiple interaction modalities. For instance, a Windows-Icons-Mouse-Pointing (WIMP) desktop application can make it possible to dictate texts by means of a voice recognizer. Applications should be designed according to the principle of: design once, display anywhere.

The e-mate Project

Integrated development environments are aimed at building applications whose delivery context is known a priori, limiting their reuse in different environments. In the e-mate project, we have designed and implemented an architecture to support the development of applications which are seamlessly deliverable in multiple contexts. Developers can design and implement the software, abstracting from the technology details and focusing only on application concepts and relations between them. For instance, in the development of a 'Personal Agenda', developers only deal with conceptual details such as appointment, person, address, calendar, etc. and conceptual relations such as 'which addresses are relevant to a given person'; 'which persons are involved in a given appointment' and so on. Any other concern about technology constraints such as network protocols, transaction management, service discovery, data persistency, adaptation of components to user terminals, and user interface presentation, should be of no importance as they are already addressed by the e-mate infrastructure.

The rationale behind e-mate is influenced by the object oriented paradigm. Objects are self-contained, behavioural complete pieces of data with variables and procedures. They are reactive entities which are activated whenever someone invokes one of their procedures. The reaction of an activated object is to produce another object (the result) and/or a side effect changing its internal state.

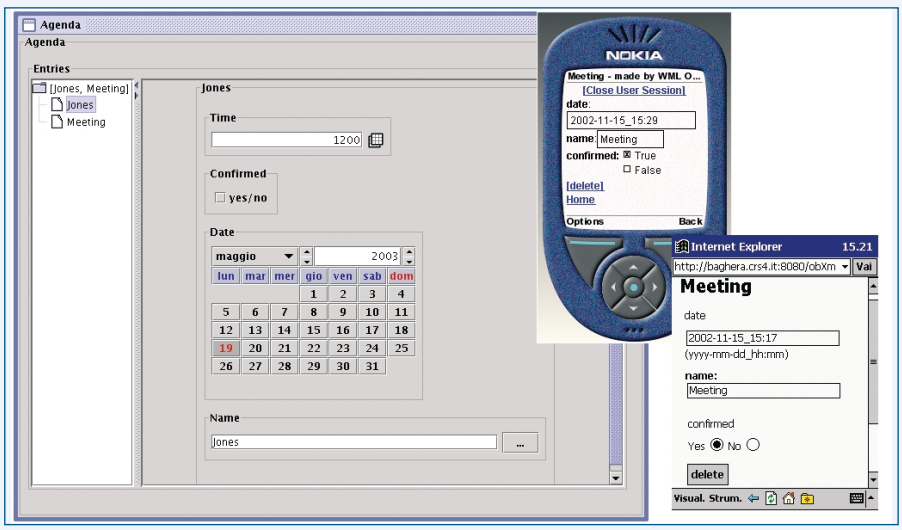

The Agenda application objects running on a desktop PC, a smartphone and a Personal Digital Assistant.

Objects can be considered as building blocks, able to react and therefore to interact. Our system provides a viewing/controlling mechanism to automatically fill the gap between the end-users and the application objects allowing direct interaction between them. This 'bridge' is implemented for any type of terminal used by the user. Thus, we can put into practice a design based solely on the definition of business related objects and their relations.

The issue is to take business objects and generate a user interface aimed at interacting visually with them. To address this critical point we have developed MORE (Multi-platform Object REnderer). The rendering process uses reflective capabilities of the Java language to extract the 'displayable' information directly from the code. Such information consists of the set of public procedures and public variables that are incapsulated in the model object.

Complex models are viewed as composite graphical objects and the rendering process is applied recursively until atomic interactors are obtained. The advantage of this approach is that the same model can be accessed from any device for which a rendering engine has been implemented (see the Figure).

The rendering process takes into account the end-user device. Since displaying and communication capabilities may vary completely on different devices, the engine adapts the model description (in terms of data format and size) to take advantage of the target user interface and available communication channel. The current implementation uses XML as a protocol language to convey the interaction between the end-user and the application business objects. Depending on the processing capabilities of the client platform, the user can perform a simple interaction or eventually access an enhanced set of native graphical components that fully exploit the capabilities of the device.

Mobile devices rely on wireless (and often expensive) connections, therefore smart strategies have been implemented to reduce bandwidth occupation. The rendering engine sends only portions of interface that are truly renderable by the client; on the other side, the client tries to perform most of the computing locally, using the communication channel only when required. Finally, since most mobile devices have small displays and one-hand-operations user interfaces, the rendering engine sorts interface elements by usage, giving faster access to frequently used functions.

Some specific applications require both interface and data adaptation to support mobile client delivery. For example, Location-Based Services require mobile cartography functionalities and present critical issues when deployed on mobile terminals over wireless connections. To address such issues, we have developed a lightweight version of the Geographic Markup Language (called 'compact GML'), which is used by the rendering engine to encode maps according to client displaying capabilities.

The e-mate project was successfully completed in early 2003 and some of the results will be used in future activities. For example, we plan to release the Multi-platform Object REnderer as Open Source software and improve it in terms of usability, and the cGML will be evaluated and used in an EU Project (EurEauWeb IST-34182).

Links:
CRS4 - Networked Distributed Applications: http://www.crs4.it/nda

E-mate project: http://www.crs4.it/nda/e-mate

Compact GML: http://www.crs4.it/nda/cgml

Please contact:
Gavino Paddeu, CRS4, Italy
Tel: +39 070 925 0284
E-mail: gavino@crs4.it

Development Tools for Dynamic Environment-Adapted Mobile Interfaces

by Anthony Savidis and Constantine Stephanidis

One of the main challenges in mobile computing is to maintain continuous high-quality interaction between users and interactive applications. The Voyager toolkit, developed at the Human-Computer Interaction Lab of ICS-FORTH, provides all the necessary implementation components to build interactive applications that are capable of utilizing different types of User Interface input-output resources, as they are dynamically discovered or lost during mobile interaction.

In mobile computing, users are faced with a dynamically changing environment, and are offered a large and varying range of remote wireless devices. In this situation, users may carry small 'pocket style' processing units, emphasising mainly wireless networking (such as BlueTooth), an embedded OS, and some locally running applications, while all interaction may evolve dynamically by utilising such remotely available devices (ie the small machine offers no UI equipment). Dynamic configuration and continuity of interaction are key ingredients in this world of mobile applications, supporting distributed user interfaces where the input and output (I/O) devices (such as buttons, text lines, raster displays, loudspeakers etc) are discovered 'on the fly' during user mobility.

The dynamism and continuity of interaction inherent in these scenarios are at the root of the 2WEAR project, for which FORTH (coordinator), ETHZ (Switzerland), MA Systems & Control (UK), and NOKIA/NRC (Finland) have joined forces to explore mobile computing as an agglomerate of small, wearable and physically distributed devices. Funded by the EC, 2WEAR took off on January 1, 2001 with the objective to build such a system, in terms of both software architecture and devices that communicate using short-range

radio technology, and deliver it by December 2003.

In dynamic, transient environments, the main architectural challenge is to maintain high-quality interaction between users and applications. This principle of continuity emphasises the uninterrupted sequence of dialogue activities and ensures that the human-computer interaction dialogues transform gracefully from the user's perspective. Developed at the Human-Computer Interaction Lab of ICS-FORTH, the Voyager Toolkit is a User Interface (UI) development framework (delivered as a C++ software library) which provides interface programmers with all the necessary implementation components to build interactive applications that utilise different types of UI I/O resources, as they are discovered or lost during interaction. As shown in Figure 1, a 2WEAR UI is implemented as a typical client UI, while the UI I/O resource distribution is handled underneath, through the 2WEAR communication technology.

The Voyager toolkit places abstract interactive objects (the Selector and TextEntry) at the service of the programmer, hiding run-time coordination and control details, and revealing only semantic characteristics at the API level. These objects are internally linked to the lower-level management of remote UI resources to cater for I/O control through various levels of specialisation that are completely hidden to the application programmer. Programmers thus have an effective instrument to easily craft mobile dialogues with diverse dynamic UI I/O resources. Furthermore, the single-focus policy for objects and applications has been enforced, meaning that a single application can have the focus in a mobile system and that only the focus object will actually allocate resources for interaction purposes. This modality, which is concurrently supported in Windowing systems, turned out to be a necessary decision in mobile interactions, since the physical display 'breadth' of the devices which are dynamically engaged does not allow concurrent dialogues to be practically supported.

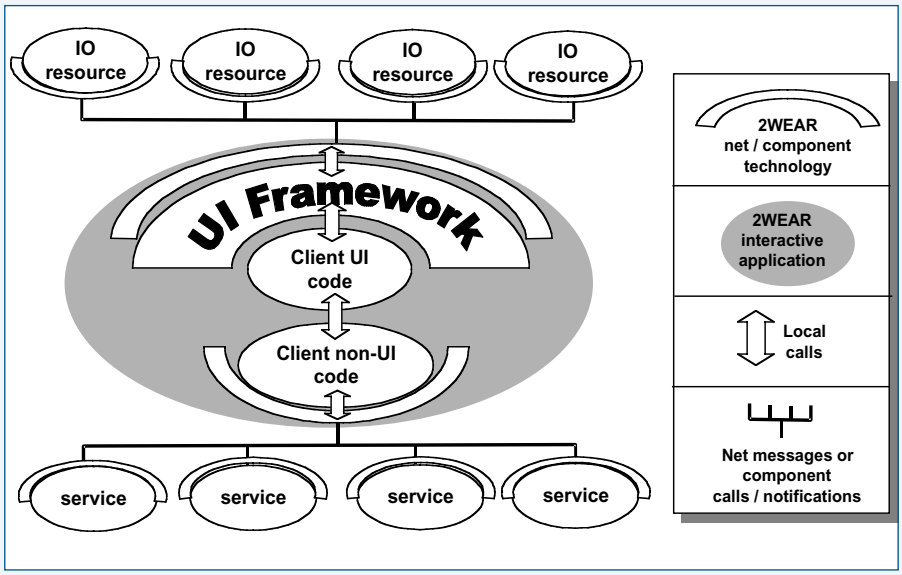

The development role of the Voyager User Interface development framework in developing mobile distributed applications in the 2WEAR project.

When new resources appear, if dialogue is 'suspended', no action is taken, since the application does not own the focus. Otherwise, the Voyager UI framework substitutes the instantiation of the current focus abstract object with the best possible I/O instantiation without affecting the state of the focus abstract object. When new UI resources are discovered, the UI framework finds the best possible I/O instantiations by recycling the resources currently used by the focus, plus those just discovered, as well as those that are not granted and are thus available. If the best instantiations found are the ones currently in use, no actions are taken. Otherwise, the current I/O instantiations are released (assuming dialogue was not 'stalled', in which case those instantiations are null), while the current best-found input and output instantiations are activated (when output instantiations are activated, or when they are associated with a different abstract object, they call their display method automatically). In the case where the dialogue was previously in a 'stalled' state, and is now in a 'working' state, dialogue revival is automatically signalled.

When UI resources are lost (either due to network failure, or to lack of responsiveness, or because the services are gracefully removed), if the dialogue is already 'stalled', no further action is taken. Otherwise, the UI framework asks the current focus object if it happens to use any of the lost resources. If not, no action is taken, but if this is the case, the current input and output instantiations are released, and a configuration round is entered, seeking for the best possible input and output instantiations. If no such instantiations can be found with the currently available resources, the dialogue is signalled as 'stalled'. Otherwise, the input and output instantiations that have been found are activated.

To conclude, the challenges in implementing mobile interactive applications capable of dynamically utilising the UI resources discovered in the surrounding environment are: (a) handling the dynamic resource coordination and control requirements at runtime, and (b) offering a compact and invariant programming model which is intact from the variety of dynamic UI resources. The Voyager UI development framework has been expressly designed and implemented to address these two challenges.

Please contact:
Constantine Stephanidis, ICS-FORTH
Tel: +30 2810 39 1741
E-mail: cs@ics.forth.gr

Introducing New Mobile Services Faster

by Timber Haaker and Oscar Rietkerk

The Business Blueprint gives a broader look at service development, integrating the technical perspective with the end-user, organizational and financial perspective. This method was developed within the Freeband B4U project that aims at accelerating the introduction of new mobile services. The project is a cooperative effort of the Telematica Instituut, TNO Telecom and Delft University of Technology, supported by the Dutch Ministry of Economic Affairs.

Jenny works at 'Care4U', a homecare organisation with about 1100 employees. In order to control costs and boost productivity Care4U has provided every employee with a Personal Digital Assistant (PDA) that has a wireless connection (GPRS= General Packet Radio Service) to the central server. Jenny registers most of her activities on the PDA. The data is used for hour logging as well as updating the medical records of the patients. Jenny trusts the automated system better than the paper predecessors. It saves her time and reduces mistakes as compared to the previous paper forms.

Figure 1: Jenny registers her activities on her PDA.

Such a scenario sounds easy enough but it is far from that. Before it can be realised, a number of technical and organisational problems have to be solved. Medical data is privacy sensitive and needs to be highly accurate. In the Care4U scenario, data is entered and stored on the PDA. Then it is sent to the central database of Care4U in real time to avoid synchronisation conflicts. A back office application processes the data to fit the administrative system and update the medical records of Care4U. The medical records in the ownership of the patients themselves are updated as well. These records are stored and managed by an intermediary organisation acting on behalf of the patients

Data integrity starts at the source, so the PDA application should be carefully designed to avoid mistakes (remember the presidential elections in the US of 2000). The PDA application should check if the data is correct (is the number within the logical range of this data type?). Security starts at the source as well, so in order to be able to register, the user should be identified. The connection between PDA and the central database of Care4U, and between Care4U and the intermediary organisation handling the patients' medical record, should be secure. Care4U should be authorised to update the personal medical records, and the data types and information structure of both databases should match.

Solving these mostly technical problems within the Care4U organisation is relatively easy. In most cases, standard business applications can be adapted to the specific situation and habits of Care4U. The problem becomes much more difficult if integrated healthcare has to be supported using this kind of applications across several institutes.

Without underestimating the technical problems to get these systems to work, this is only part of the problem if you want to bring new mobile services to the market. The Freeband - B4U project has the ambition to accelerate the market introduction of new mobile services and therefore has designed a method to develop business blueprints for such services. This method approaches a business idea from four different perspectives.

To accelerate market introduction of services and service development you need to look not only from a technical perspective. You also need a clear view on:

- *The end user perspective:* The value proposition describes the service offered to specific customers in a particular market segment. Most important is the assumed customer value of the service concept, ie, what is the added value of the service as compared to existing alternatives? To what customer needs does the service appeal? In the homecare example both the care worker and the care institute have to be taken into account when defining the value proposition. What you functionally want to offer determines what you technically need.
- *The organizational perspective:* The value network is the configuration of actors that possess complementary resources and capabilities, which together perform activities to realise the service. Due to financial considerations, increased specialisation in the wireless industry and increased involvement of other sectors like media and IT, it is no longer possible for a single company to develop a mobile service all alone. For example in the Care 4U scenario you need at least a mobile network operator, a device manufacturer and a system integrator. Each actor in the network has some resources, among which technical ones, which need to be compatible and connected. Responsibilities have to be clearly

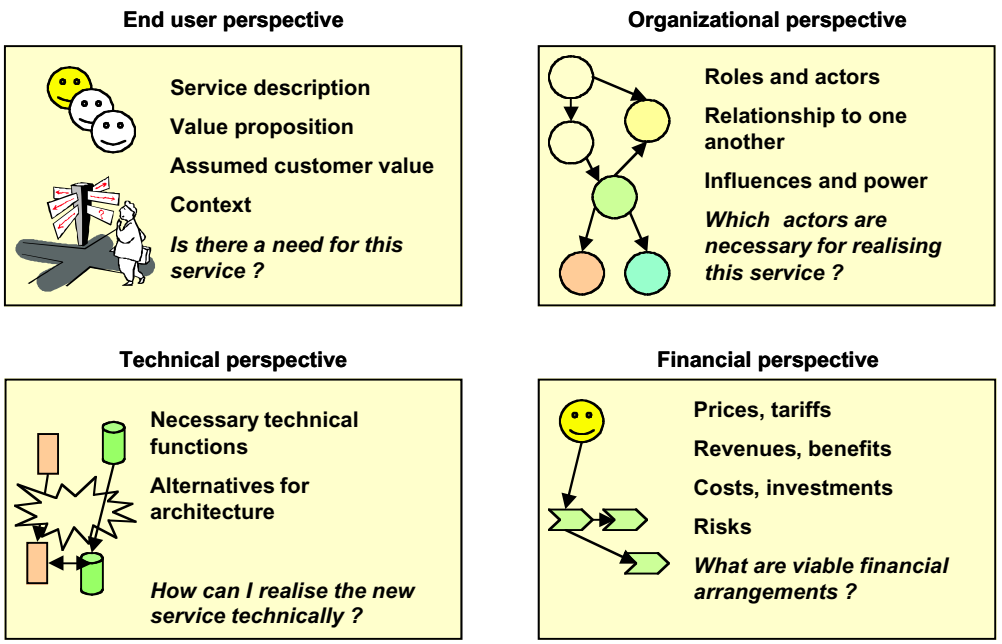

Figure 2: The four perspectives for fast and successful development of new mobile services.

defined. In the Care4U case the situation is complicated due to the existing legacy systems at the care institute. Integration is hard to realise.

- *The financial perspective:* The revenues generated from the service will depend on the price and the tariff structure. Whether customers like Care4U are willing to pay the price depends on the value they perceive. For Care4U that means the cost reduction and productivity boost they may realise plus the advantages they perceive for Jenny. For a sustainable cooperation in the value network the financial arrangements are critical. Only if a fair distribution of benefits, investments, costs and risks can be arranged a sustainable value network can be realised. Costs strongly depend on the chosen technical implementation. Some of the risks are directly related to the technical choices: proven technology, future proof concept and scalability.

Finally, the necessary co-operation between parties and the interdependencies between the perspectives is a complex process to direct through all the stages of service development. That is why the Business Blueprint method was developed to manage this process more easily.

It is a structural approach to simultaneously design the four perspectives and develop an integral business design. The blueprint may be applied:

- *explorative:* to assess a business opportunity of an organization within the value network
- *creative:* to develop a business opportunity with partners
- *evaluative:* to evaluate the strong and weak points of a particular business design.

The Business Blueprint method has been successfully applied to several business ventures, other Freeband projects and prototypes developed within B4U. The Business Blueprint comes in a variety of forms: consultancy, workshops, checklists and a business game. It gives a broader look at service development, integrating the technical perspective with the end-user, organizational and financial perspective. It speeds up the process and increases the success rate of service development while maintaining grip on complexity.

Link:
TNO Telecom: http://www.tno.nl

B4U project:
http://www.freeband.nl/projecten/b4u/

Please contact:
Timber Haaker, Oscar Rietkerk,
TNO Telecom, The Netherlands
E-mail: t.i.haaker@telecom.tno.nl,
o.d.rietkerk@telecom.tno.nl

Uluru: Mobile Interactive Multimedia Experimental Service Environment

by Johan de Heer, Andrew Tokmakoff, Henk Eertink and John Anijs

The ULURU project, initiated in 2001 by the Dutch Telematica Instituut and Mobiview, has its focus on aspects that play a primary role in the value chain between content providers, mobile network operators and end-users. Research has been performed in personalization, payments, digital rights management, content management and content distribution.

From a mobile network operator's perspective, the success of broadband mobile services is highly dependent on the ready availability and acceptance of attractive applications and content, which the end-user is willing to pay for. We have produced an intermediary facilitating platform that is positioned for use in a business-to-business (B2B) setting, between mobile operators on the one hand, and content providers on the other (see Figure 1).

Uluru is a platform that hosts 3rd party application logic and makes common services available to these applications. These services can potentially be quite complex and may also require a certain 'trust relationship' between the specific service and the application using it (eg payment through a mobile operator's billing system). Such trust generally requires contracts/agreements (SLAs) between the mobile operator and parties that wish to charge using the operator's billing infrastructure, which can be difficult and expensive to arrange.

Uluru provides a set of coarse-grained functions to the applications that may use these through well-defined APIs (see Figure 2). These functions fall into the following broad categories:

- content management: handling content and its meta-data
- financial exploitation: accounting, charging, payment and billing
- personalisation (content recommendation): an engine that can be used to optimise application usability
- authentication: used for controlling access of users and applications
- logging: system monitoring
- event messaging: allowing asynchronous messaging between applications and system components

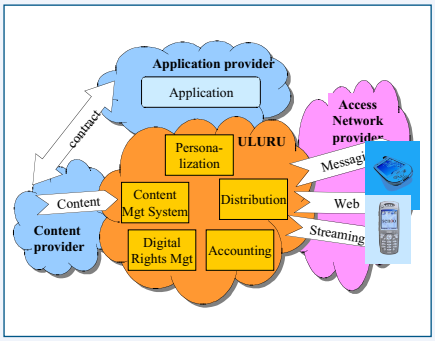

Figure 1: Uluru: Roles and Services.

- subscription: managing user details
- MMS and SMS office: services that can send and receive text-messages from mobile phones.

The Content Management Service Provisioning caters for the intake of the content feeds and performs encoding and storage of these feeds. After transcoding (eg into MPEG-4) and enrichment of the content (eg making it fit for dedicated applications), the content is again stored and ready to deliver to the mobile network.

End-users of Uluru's hosted applications can be charged for the services they consume using existing mobile operator billing services through a Parlay (http://www.parlay.org) interface. This leverages the pre-existing customer billing relationship between the end-user and the mobile operator.

For personalisation we have developed the Duine Toolkit, which is an SDK that allows developers to integrate prediction engines into their own applications. In contrast to most current available personalized information systems that focus on the use of a single prediction technique or a fixed combination of two or three techniques, Duine combines multiple techniques in a dynamic and intelligent way, thereby providing more accurate and stable predictions.

Each of the above-mentioned platform functions are exposed to the hosted applications through the Uluru platform API, and are implemented by 'backend' components using an adapter approach. This adapter approach for decoupling the COTS (Commercial Off The Shelf)

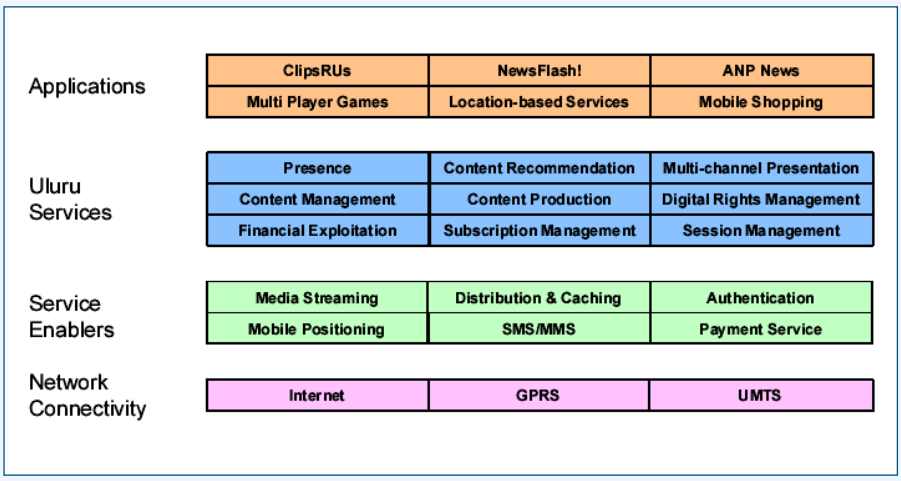

Figure 2: Uluru's functions.

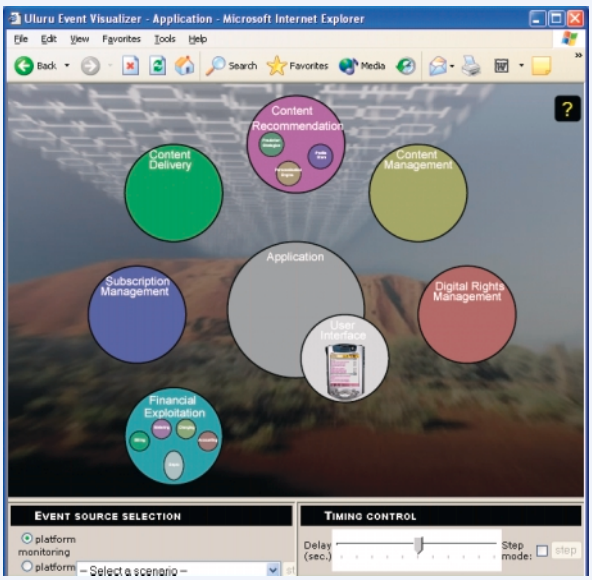

Figure 3: Visualization of System Functions.

components from the platform avoids vendor lock-in and insulates the platform from the underlying technologies. Furthermore, Uluru is communication technology independent; its only assumption is that IP network connectivity is available. This means that the platform can operate using network technologies including 802.11x, GPRS, and UMTS. The implementation of the Uluru platform is heavily based upon distributed Java technologies such as J2EE and the platform itself is distributed over multiple locations. In one example, a NewsFlash application was located at the Telematica Instituut, whilst the Uluru platform itself was located at Hilversum, the MediaPark. In this situation, the NewsFlash application utilized CORBA to access the Platform services.

In order to provide good insight into the underlying workings of the Platform, we have developed a visualization toolkit, which allows system events to be depicted on a graphical representation of the system. The visualiser (shown in Figure 3), depicts the interaction between platform components (such as profiling, content recommendation, content delivery and financial exploitation), based on input of real-time triggers (events) that occur during the process of delivering interactive multi-media services.

All that is needed of the developer is to include some “metering points” within the application and to make use of the package’s event reporting functionality. The implementation of the EventServer interface is based on the Pushlets framework, in which a servlet or JSP streams lines of JavaScript code into a hidden browser frame. The browser executes the lines and performs visualization actions accordingly. In effect, this mechanism provides a callback from Java in the server to JavaScript in the client browser. The client-side graphics have been implemented using Scalable Vector Graphics.

The Uluru platform provides the infrastructure to deliver a wide range of multi-media content and interactive services to mobile end-users (subscribers). The NewsFlash application demonstrates the capabilities of the ULURU platform, including the possibility to offer personalized content and to apply flexible charging for various kinds of content, which are regarded essential for successful business (see Figure 4). NewsFlash offers the latest news in text, image and streaming video format. The ULURU platform dynamically adapts the news content offered to the user based on preset preferences, previously requested news articles and a calculated user interest rating for news articles.

Another major programme goal, the implementation of an experimental platform for mobile multi-media services, has now been achieved. The ULURU platform is now available for user trials, which offers external parties, (such as content providers and mobile network operators) the ability to perform tests on applications providing next generation mobile communication services. The ULURU platform will be available for academic and pre-competitive user trials at the Telematica Instituut site before the end of summer 2003. If you would like to obtain access to Uluru’s platform features, please contact us.

Figure 4: Personalised Teletext news for i-mode (left), and Sony/Ericsson P800 (right), and NewsFlash with ANP content for IPAQ (middle).

Links:

Telematica Instituut: http://www.telin.nl

Mobiview: http://www.mobiview.nl

Uluru project: http://uluru.telin.nl/

Toolkit by Telematica Instituut: http://uluru.telin.nl

Event Visualiser by Telematica Instituut: http://uluru.telin.nl

Please contact:
Johan de Heer, Andrew Tokmakoff, Henk Eertink, John Anijs, Telematica Instituut, The Netherlands
E-mail: {Johan.deHeer;Andrew.Tokmakoff; Henk.Eertink}@telin.nl; John.Anijs@mobiview.nl

Next Generation Public Travel Information Systems

by Karl Rehrl, Harald Rieser and Siegfried Reich

The Sun Technology Research Excellence Center of Salzburg Research and the Ascom Center of Competence are jointly working on a next generation middleware that will allow the integration of mobile devices in Public Travel Information Systems.

Public Travel Information Systems (PTIS) are widely used in train, tramway and subway stations. However, the integration of mobile devices (such as PDAs for personalised traveller information or onboard displays in trains) is difficult due to the heterogeneity of these devices and the fact that the status of connectivity changes rapidly (from high availability of bandwidth to off-line mode). Moreover, configuration and management of these devices and the distribution and delivery of content - increasingly of a multimedia nature - are difficult issues.

During our work we have found a number of requirements for middleware systems dealing with stationary and mobile devices in public information systems. The most important of these are:

- *Heterogeneity of devices:* PTIS must cope with a multitude of devices, including personal computers, public information terminals, videowalls, LED displays and also mobile devices, such as PDAs (Personal Digital Assistants) and mobile phones.
- *Variability in connection modes - the off-line mode:* The use of mobile devices implies that not all devices will be available all the time; in mobile environments, devices are regularly off-line, and as a result, the off-line mode cannot be considered to be an exception. Consequently, the middleware infrastructure must be capable of caching data in the off-line mode and redistributing it upon reconnection through the use of sophisticated data synchronisation mechanisms.
- *Centralised configuration and 'Plug & Work' capability:* Another issue involves the accessibility of mobile devices in large systems for configuration and management purposes. For example, a huge number of information displays in train stations cannot be physically accessed. Consequently, the configuration data for these devices should be minimised, and in general they should be configured, maintained and monitored through a central management tool. Using a central management tool ('Control Center') with automatic data distribution allows the easy maintenance of a huge number of devices with similar configurations. This also allows the so-called 'Plug & Work', ie the dynamic configuration of devices during run-time.
- *Quality of information:* The quality as perceived by the customer, ie traveller, is measured by the quality of the information delivered. This requires information to be synchronised and updated regularly.

© Salzburg Research

Public Travel Information Systems in Graz.

Under consideration of the requirements and aims mentioned above, we have adopted the following technical solutions. The middleware infrastructure is based on an overlay network structure with a hierarchical, two-tier network scheme. This means that, despite the underlying physical network, mobile devices are dynamically assigned to logical sub-domains within one global domain, enabling efficient device management, message routing and allocation of the necessary infrastructure. In order to provide a simple and efficient programming model for developers, software infrastructure on devices had to be structured appropriately. The concept of a service-oriented architecture (SOA), defining services as coarse-grained functional units with well-known interfaces, was therefore adopted. Following the requirements for mobile devices, services are allowed to use different communication modes like synchronous RPC and asynchronous 'Store and Forward' messaging. Reliable data distributed to mostly off-line devices required the development of a data-distribution mechanism called Virtual Database (VDb). With Virtual Database it is possible to provide services with a virtual database connection without establishing a real connection. VDb

hides the necessary communication details from the developer and improves the efficiency of data distribution. Services ask the VDb for database fragments and get these either from the local cache or upon network connection directly from the database. VDb is used for the distribution of configuration data and management data, and is also well suited for multimedia content.

Another of our technical aims is the support of a broad variety of heterogeneous devices. This was achieved through the use of widely adopted standards such as Web Services, and through building upon platforms which are also available for embedded devices. While our first prototype was based on the Java 2 platform, requirement changes demanded that the current prototype be developed using the .net Framework. In the future, we are planning to support the .net Compact Framework on embedded Windows.CE devices as well.

In conclusion, we have argued for requirements for next-generation transport information systems and have outlined how we are addressing these issues with respect to the middleware layer. The concept of Service-Oriented Architecture was carefully chosen in order to overcome platform borders and to simplify the process of porting code to other platforms.

Link:
http://www.suntrec.at

Please contact:
Siegfried Reich, Salzburg Research
Forschungsgesellschaft m.b.H./AARIT
Tel: +43 662 2288 200
E-mai: sreich@suntrec.at

Mobile Communication in Newspaper Distribution and Transport

by Olli Kuusisto and Janne Pajukanta

Newspaper distribution and transport are time-critical operations in which a large number of people are required to work in a coordinated manner. VTT has developed a pilot system based on mobile communication to solve production and communication problems in the newspaper distribution and transport process.

Even though this process is, in principle, repeated every night, varying disturbances and problems can occur each time. By utilising mobile communication, people working in different locations can be reached, and problems solved. It also enables reporting of up-to-the-minute information on the entire production process, and communication between supervision and employees as well aswithin a team. With the help of real-time control and tracking, the entire delivery process can be optimised and the quality of the process improved. Pilot systems and solutions were developed to enable work-related communication in newspaper distribution organisations, and experience was gained in practical field tests. The tests made clear what kind ofbenefits can be offered by the use of these new tools, as well as showing how currently available technology is able to meet the challenges of professional usage in a real working environment.

At present, the flow of electronic information in newspaper production ends at the post-press department. Tailoring of newspapers, new subscription alternatives, new products and predicted shortage of deliverers will increase the challenges for accurate and reliable distribution. In practice, distribution will evolve, and is predicted to function more and more through teamwork, where teams will be autonomous and the need for supervision will decrease. These challenges create a demand for new systems of communication and transmission of data.

The objective of the mobile communication pilots was to plan and implement mobile control and tracking in order to obtain empirical information on the suitability of the system and the portable wireless devices used. The pilot system included following functions:

- on-time and deviation reporting to the distribution organisation
- instant messaging to deliverers
- communication between and among the foreman and the deliverers.

The preliminary requirements specification was made based on interviews with deliverers, their foremen, van drivers and management from several distribution companies. In particular, opinions concerning portable devices and features of the system were gathered. Typical working methods and the need for bilateral communication were also examined. At the pilot stage, the operation of the pilot system was first presented to the pilot companies and deliverers, and feedback was gathered accordingly.

Web-application architecture was chosen since individual applications to different terminals, as well as the control of the state of the application, would have been troublesome, the latter resulting from at the time not so reliable GPRS-connection. The application can be run on any general Web browser. Typical layers in such architecture are client application, www-server and database. Web browsers served as the client application at the different terminals. HTML pages were created dynami-

cally with help of Java Servlet technology based on the database.

The pilots took place in 2002-2003. The pilot system made communication and the flow of information easier to manage between deliverers and their foremen, and within the team of deliverers. It is expected that the schedules of the entire delivery process could be optimised and the quality of the process improved with the aid of real-time control and tracking. The users found the pilot system easy to use and navigate. Information from the production stage at the printing house was useful and helped to plan the deliverers' work. The communication system was found to be useful, as there already exists a high frequency of communications between deliverers. At the moment this is handled by cell phones, but it could be more easily managed with PDAs.

Pilot-like real-time control and tracking gives accurate data on sources of disturbance, which enables a fast reaction and increases quality. Reporting of the start and stop times of delivery gives specific data to foremen and to customer service. If the deliverer has not registered the start time as scheduled, he/she can be contacted without delay. Traditionally it is necessary to either physically check the starting points of delivery, or to wait until customers complain. Fast and accurate feedback can also be given to previous working phases. Registering of the extra copies needed by deliverers enables better and more accurate control of the edition and makes it easier to tackle the causes for ordering extra copies. Also, an electronic control system enables the optimisation of the daily schedules of distribution, which brings savings in decreased waiting time costs.

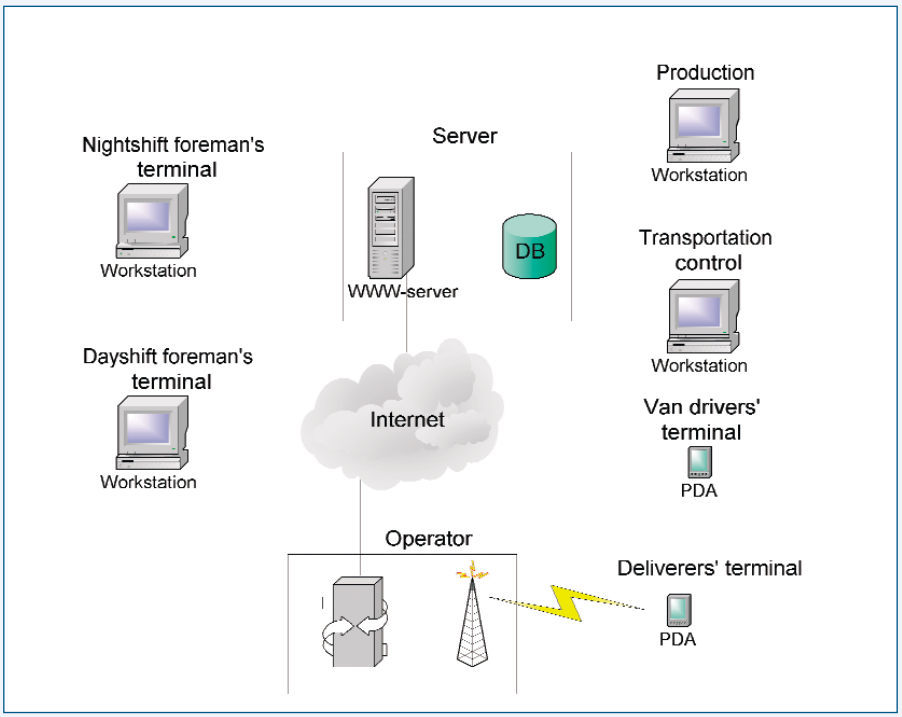

Mobile communication pilot system overview.

Link:
http://www.vtt.fi/tte/projects/JALAN/

Please contact:
Olli Kuusisto
VTT Information Technology
Tel: +358 9 456 4760
E-mail: Olli.Kuusisto@vtt.fi

Values of Mobile Applications to End-Users

by Keng Siau, Fiona Nah and Hong Sheng

What do end-users expect to have in mobile applications? What are the values of mobile applications that are important to end-users? How these values can be achieved? The Mobile Commerce Research team at the University of Nebraska-Lincoln examined the values of mobile applications using the 'Value-Focused Thinking' approach. This research developed means-ends objective networks depicting how mobile applications would be able to achieve values expected by end-users.

Mobile applications are undoubtedly the next wave in the evolution of e-business. Possessing features and functions that are unique to mobile devices, such as mobility, personality, and flexibility, mobile applications are able to provide end-users' added values, including anytime, anywhere access, ability to pinpoint users' locations, and flexibility in arranging tasks. It is predicted that the mobile users will increase dramatically in the near future and the rate for mobile services will drop significantly. However, the promising future of mobile applications has been inhibited by the infant stage of mobile applications, the drawbacks of mobile devices, and the limitations of mobile services. Nevertheless, mobile applications are gaining attention both from academics and practitioners. Understanding the values of mobile applications has become particularly important. It would be helpful to explore the values of mobile applications from the end-users' perspectives (customers and company's employees) and determine how the

values could be achieved, especially for companies embarking on implementing mobile applications or customers who are embracing mobile applications.

To understand end-users' perspectives about mobile applications, the best way is to ask them. We conducted two research projects to examine the values of mobile applications. One studies the values of mobile commerce to customers, whilst another one examines the values of mobile applications in a large public utility company in the US. The same Value-Focused Thinking approach was used in both studies. Value-Focused Thinking approach allows us to explore the end-users' values and define relationships among the values — ie, how one objective can be achieved by other objectives. Open-ended interviews were conducted to collect values of mobile applications. The approach consists of four steps:

1. Identify user wishes, concerns, problems, and values in mobile applications.
2. Convert user input into objectives. An objective is something one wants to strive toward.
3. Distinguish between fundamental objectives and means objectives. If one objective is important because it will influence another objective, it is a means objective. Otherwise, it is a fundamental objective.
4. Build a means-ends objective network depicting specific relationships between means objectives and fundamental objectives.

The results of the two studies presented values of mobile applications from end-users' perspectives and determine the relationships among the values. Some of the interesting findings include:

- Convenience and efficiency are the main advantages of mobile applications both for customers and company's employees. They are the two primary fundamental objectives that end-users expect mobile applications to provide.
- Mobility is the basis for the values of mobile applications. The ability to be used at anywhere, and at any place was highlighted by the interviewees as the main reason to adopt mobile applications. It is also the main advantage mobile applications have over traditional personal computers. In a way, it is a 'necessary condition' for mobile applications.

The researches also suggested some areas that need to be enhanced to ensure the wide diffusion of mobile applications. They are:

- Limitations of mobile devices. Being much smaller, mobile devices have small screen and complicated input mechanism, and are considered to be more difficult to use than personal computers. The low display resolution and small display screen have inhibited information to be displayed completely and clearly. The limited battery life has restricted the mobility of end-users.
- Mobile service quality. The technical restrictions have posed a challenge for the wide diffusion of mobile applications. Low bandwidth, unstable connection, and limited coverage area are the main drawbacks of current mobile technology and services. In our interviews, these are listed as the main obstacles to overcome when using mobile applications. Although we believe that with the advancement of mobile technology, these problems would be alleviated or resolved in the near future, more research needs to be conducted at the current stage.
- Security options. Security is a major concern for mobile applications. Wireless transmission, in a way, biases end-users to perceive mobile applications to be more vulnerable and unsecured. Thus, more security enhancements and options must be provided to ensure the security of data and business transactions.

Our research not only shows the causal relationships among the values of mobile applications, but it also suggests areas for further research and development. This research steam is ongoing and we are currently investigating the human-computer interaction aspects of mobile devices and mobile applications.

Further Reading

- **Siau, K., Lim, E., Shen, Z., "Mobile Commerce – Promises, Challenges, and Research Agenda," Journal of Database Management, Vol. 12, No. 3, Jul-Sep, 2001, pp. 4-13.**
- **Siau, K., Shen, Z., "Mobile Commerce Applications in Supply Chain Management," Journal of Internet Commerce, Vol. 1, No. 3, 2002, pp. 3-14.**
- **Siau, K., Shen, Z., "Building Customer Trust in Mobile Commerce," Communications of the ACM, Vol. 46, No. 4, 2003, pp. 91-94.**
- **Siau, K., Shen, Z., "Mobile Communications and Mobile Services," International Journal of Mobile Communications, Vol. 1, Nos. 1/2, 2003, pp. 3-14.**
- **Krogstie, J., Lyytinen, K., Opdahl, A., Pernici, B., Siau, K., Smolander, K., "Research Areas and Challenges for Mobile Information Systems," International Journal of Mobile Communications (forthcoming).**
- **Siau, K., Sheng, H., Nah, F., Davis, S., "Trust in Mobile Business: An Empirical Investigation," International Journal of Electronic Business (forthcoming).**

Please contact:
Keng Siau
Department of Management
University of Nebraska-Lincoln, USA
Tel: +1 402 472 3078
E-mail: ksiau@unl.edu

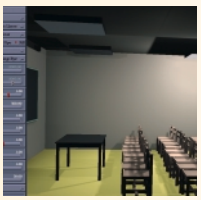

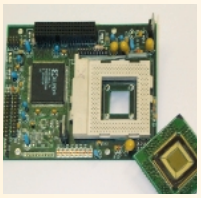

Articles in this Section

Algorithmic Clustering of Music

by Rudi Cilibrasi, Paul Vitányi and Ronald de Wolf

A fully automatic method for music classification has been developed at CWI. It is based only on compression of strings that represent the music pieces. The method uses no background knowledge about music whatsoever: it is completely general and can, without change, be used in different areas like linguistic classification and genomics. It is based on an ideal theory of the information content in individual objects (Kolmogorov complexity), information distance, and a similarity metric.

Recognition of similarities between music pieces is crucial for the design of efficient music information retrieval systems. The ever increasing wealth of digitized music on the Internet calls for an automated organisation of the material based on similarity, which enables users to navigate and which gives them advice. Current methods of automated classification utilize specific musical features related to pitch, rhythm, harmony, etc., just as a human expert categorizes the material manually. One can extract such features using for instance Fourier transforms or wavelet transforms, and classification is done by existing software, based on various standard statistical pattern recognition classifiers, Bayesian classifiers, hidden Markov models, ensembles of nearest-neighbour classifiers, or neural networks. However, this approach requires specific and detailed knowledge of the problem area, since one needs to know what features to look for.

Our approach does not require specific musical knowledge, as it is based on a general mathematical theory of similarity. It uses a metric derived from the notion of 'information distance'. Roughly speaking, two objects (represented as strings of bits) are close if we can significantly compress one given the information in the other. Compression is based on the ideal notion of Kolmogorov complexity, which unfortunately is not effectively computable. Hence we replace this ideal version by standard compression techniques (the program bzip2). We lose theoretical optimality in some cases, but gain an efficiently computable similarity metric. This new metric has been shown to work well in very different applications, including the fully automatic construction of the phylogeny tree based on whole mitochondrial genomes and of a language tree for over 50 Euro-Asian languages, as well as the detection of plagiarism in student programming assignments, and the phylogeny of chain letters.

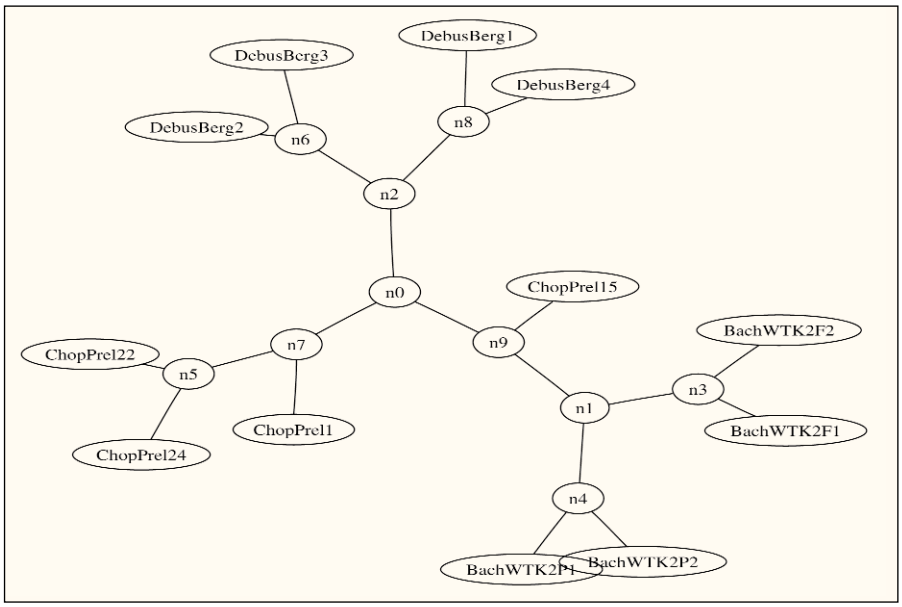

The result for a set of 12 classical piano pieces.

Now we have applied our compression-based method to the classification of pieces of music, and performed various experiments on sets of mostly classical pieces given as MIDI files. We computed the distances between all pairs of pieces, and then built a tree (using the so-called quartet method) containing those pieces in a way that is consistent with those distances. After the successful completion of three controlled experiments we turned to the real world. The method distinguishes quite well between various musical genres (classical, jazz, rock), and in the classical genre even between composers — a remarkable fact indeed, given the method's complete ignorance of music. The Figure shows the result for a set of 12 classical piano pieces: the 4 movements from Debussy's Suite Bergamasque, 4 parts from Bach's Wohltemperierte Klavier II, and 4 preludes from Chopin's op.28.

Future applications may include the use of the program as a data mining machine to discover hitherto unknown similarities between music pieces of different composers or genres, and the selection of a plausible composer for a newly discovered piece of music of which the composer is not known.

Please contact:
Rudi Cilibrasi, Paul Vitányi, or
Ronald de Wolf, CWI
Tel: +31 20 592 9333.
E-mail: {Rudi.Cilibrasi, Paul.Vitanyi,
Ronald.de.Wolf} @cwi.nl

Extension of Electronic Documents to enable Server-Independent Architectures

by Paola Carrara, Giuseppe Fresta, Daniela Poggioli

An eXtended Electronic Document (XED) is an XML document which aggregates data, templates, transformation rules and manipulation tools. It can be used to design complex systems on the Internet in a server-independent way, modifying the traditional asymmetry between the logical units of Web architecture.

The two dialoguing units of traditional Web architecture are a data delivery unit, which is usually in charge of computing activities, and a data display unit, which requests, receives and browses information. Java™ opened the way to a new concept of the Web, by sharing/distributing computation on both sides of the Web architecture, but it delivers binary code which is not extensible and can be only executed.

XML (eXtended Markup Language) makes it possible to exchange open and textual information: It defines structured text files which can be manipulated by programming languages acting at any component of the Internet architecture. XML tags can be freely defined to reflect the semantics of any application universe and the structure of the documents to be dealt with. One example is the Geographic Markup Language (GML), defined by the OpenGIS Consortium to provide a common model and common format for the description and expression of geographic entities.

Recent literature addresses the relationship between semantically-rich Markup Languages, such as GML or GraphML, and SVG (Scalable Vector Graphics), the W3C recommendation for vector graphics. Architectures have been presented in which information is described and managed at the server side by the language(s) suitable for the data model(s) under consideration, and is then translated into a graphic-oriented language at the browser side for visualisation and user interaction purposes.

These approaches maintain the traditional asymmetry between the units of the Internet architectures, and make it impossible to conceive application environments that can manage, manipulate or even exchange information indifferently at either of the two logical units of the network. To overcome this situation, we propose to evolve the concept of electronic document to an eXtended Electronic Document (XED), which is an information aggregation gathering together data items, templates for their representation, and tools for their description/manipulation.

An XED document is an XML-document consisting of:

- a set of fragments describing data to be exchanged (for example GML fragments)
- a set of fragments describing transformation rules
- a set of SVG fragments (or other presentation/interaction languages such as: XHTML, VRML, etc.) specifying the representation templates
- a set of tools (typically ECMAscript code) for the user.

Original data/template fragments can be stored in different repositories, and be compliant with various data models, formats, languages, etc. They can be enriched by meta-information to account for different aspects such as semantics, quality, validity, languages, access rules, etc. The XED can be exchanged between different Internet components, transformed and manipulated at any network node; the entire XED, or its components, can be searched for retrieval or filtering purposes, and aggregation can thus be dynamic. An XED should contain all the instructions necessary for its rendering, manipulation and/or storage, by the preferred Internet data displayer such as a Web browser, a mobile or a sound/voice rendering device.

Our proposal strengthens the possibility of reusing data in different situations and purposes as we do not conceive XEDs as 'static' entities, created by the designer. For example, one 'primitive' XED can be used both by a well-defined group of end users (such as Italian students using a Web browser) and intermediate users who have to enrich it with more information (for instance, translators who must add the English version of the data or programmers who must unite transfor-

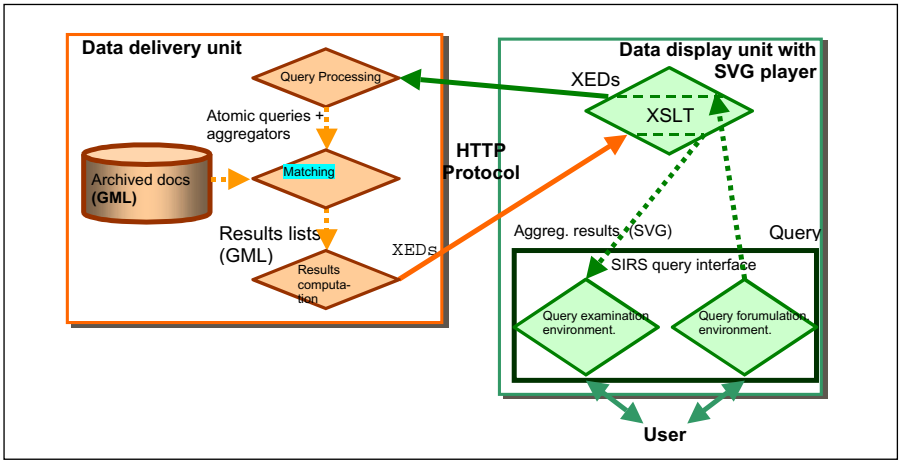

Web-based GIS system architecture with browser activities and delivery of eXtended Electronic Documents.

mations and rules to allow visualisation on a PDA).

In this way, it is possible to design complex systems on the Internet as chains of computing activities on specific XEDs: each activity transforms the document (the data, the templates and/or the tools) and prepares it for a successive activity to be performed by other users at any Internet component. The designer can statically assign the different activities on the basis of criteria such as performance, data dimensions, available hardware platforms, etc.; a dynamic assignment can be also made on the basis of the document status, user profiles, etc.

We tested these ideas, taking as case study, an Information Retrieval system for geographic information to be implemented on the Web. The system design is performed through distinct steps, ie:

a) analysis of the computational activities to be performed, and the data and interaction they require;
b) identification of the criteria guiding their assignment to the two logical units of the Web architecture;
c) identification of the users (not only the end users but also other users who participate, with different roles, in the system creation, tuning, maintenance, etc.), and of the devices they prefer, in order to identify the necessary XED components and/or the chain of actions to create evolved XEDs from primitive ones.

Figure 1 illustrates the result of steps a) and b), ie the activities we identified (only macro-activities are shown), and the assignment to the units where users can formulate queries and look for results. The criterion adopted was the level of interaction with the users: some activities, such as query processing, matching between queries and documents, and the first aggregation of result list(s) are delegated to the data delivery unit, as they do not require direct user manipulation. At this step, data are described by GML. The data display and interaction activities (mainly query formulation and results examination) are included in the display unit and are executed directly on the transformed SVG version of the documents. Therefore the architecture also includes environments to perform the language transformations. The XEDs, delivered between the units by a plain HTTP protocol, are not limited to the GML data description but also contain templates and code to guide presentation, and the necessary interaction tools.

Please contact:
Paola Carrara or Daniela Poggioli,
IREA-CNR
E-mail: {carrara.p, poggioli.d}@irea.cnr.it

Giuseppe Fresta, ISTI-CNR
E-mail: giuseppe.fresta@isti.cnr.it

COSMOS: Safe and Optimal Operation in Complex Industrial Systems

by Péter Inzelt

COSMOS, a three-year project within the framework of the National Research and Development Programs, commenced in the middle of 2001 with the aim of developing a novel, knowledge-intensive technology for safe and optimal operation in large and complex industrial systems. The project, led by SZTAKI, is carried out by a consortium consisting of both academia and industry.

Safety and optimality of operation are decisive factors for large and complex industrial systems due to their complexity and the uncertainty associated with their functioning. The project COSMOS aims to develop a novel information technology applicable in complex industrial systems for process parameter monitoring, early detection and diagnosis of equipment and technology faults, and fast detection of malfunction roots.

The project partners include SZTAKI, the Atomic Energy Research Institute at the Hungarian Academy of Sciences (MTA KFKI), the Department of Control Engineering and Information Technology and the Department of Measurement and Information Systems at the Budapest University of Technology and Economics, Paks Nuclear Power Plant Ltd, and MOL Hungarian Oil and Gas Plc.

The System

The technology to be developed will offer operational support for process devices through advice given to the operator. It will present user-friendly and efficient interactive tools for storing technology-related knowledge, on-line modeling of process phenomena, visualisation of diagnostics information, and the on-line display of advisory documentation to be used during fault-corrective actions.

The kernel of the software system is represented by a central database, with a static part consisting of the following three main elements:

- the 'Object Structure Definition (OSD)', which describes the features and attributes of the different objects handled by the system
- the 'Abstract Object and Knowledge Repository (AOR)', which acts as the repository for the objects described in the OSD
- the external databases, handled in their original form, but using standard interfaces.

The running programs perform their operations over an instant (fast) database (IDB), which constitutes the dynamic part of the central database. An IDB can be formed only from the objects stored in the abstract repository and it always contains as many objects as are required by the processing modules running at the time.

courtesy by NPP

The Nuclear Power Plant in Paks, one of the two application sites.

The modules are used for establishing a structured storage for the engineering-technical information concerning the components of the complex industrial system, and for the preliminary analysis of the safety, operational and ergonomic consequences of the planned reconstructions and changes in the applied technology.

The 'Configuration and Document Handling Module' is used for storing basic, technically relevant information concerning the different objects (buildings, technological systems, equipment etc). For documents directly connected to the operation of the plant, the system ensures access to the on-line database of the process information system, in order to retrieve and display the process parameters (measured or calculated signals) which are referenced in the document.

The 'Virtual Reality Module' is connected to VRML models converted from the 3D CAD drawings of the rooms and the equipment, and is used for performing the following basic tasks:

- *the VR Training Component* models potentially dangerous activities, trains work operations, and plans the optimal work
- *the VR Design and Validation Component* is used for planning and modeling plant modifications and reconstructions in an optimal way, and for checking their conformance to the standards and requirements
- *the VR Decommissioning Component* is used for planning and modeling the dismantling of hazardous facilities and equipment.

Monitoring the technological process, the 'Early Fault-Detection Module' identifies deviations from the normal state using the principle of analytic redundancy, ie, process measurements and on-line calculations based on the component's model.

The 'Fault Diagnostic Module' identifies the root causes of equipment faults or system malfunctions. Having the fault trees, which are defined by an interactive fault-tree editor, the reliability of the minimal cuts can be determined. This information can be used for identifying the most 'fault-sensitive' components, and for determining the relative probabilities of root causes creating specific fault symptoms. For large subsystems, the root causes are determined by an on-line fault-tree evaluation procedure.

The actual operational conditions for technological subsystems and equipment efficiencies are determined by the on-line 'Process Optimisation Module', based on on-line process measurements and background calculations. Using this module, the experts can assess the position of the actual working points with respect to the optimal states.

In its final form, the developed technology will constitute a general-purpose, open-architecture software shell providing the operators of various industrial systems with efficient support. It will realise the following principle: an industrial system should be able to operate safely while simultaneously working close to its optimum capacity.

Results and Applications

The results will be implemented at the Nuclear Power Plant in Paks, firstly for experimental use at the central supervisory computer, and later at the supervisory systems of the nuclear block. Another implementation will be used by MOL Ltd, in the supervisory system of the network of high-pressure natural-gas pipelines in Hungary. According to the schedule of the project, the development of the prototype database, prototype interfaces, prototype functional modules and final interfaces has been successfully finished.

Link:
http://www.sztaki.hu/sztaki/projects/html/cosmos.jhtml

Please contact:
Péter Inzelt, SZTAKI
Tel: +36 1 279 6184
E-mail: inzelt@sztaki.hu

Leading-Edge Knowledge Management for Public Employees

by Simon Lambert

The English word 'pellucid' means 'transparent or translucent; extremely clear in style and meaning' (Collins English Dictionary). That describes the vision of the Pellucid project in aiming to create a software platform for knowledge management for public employees.

Public administrations carry out a huge variety of tasks in the life of their country, but despite this variety they often have certain things in common in their working practices. A frequent phenomenon is 'organisational mobility', in which employees move from one department or unit to another, taking some of their acquired knowledge with them, but at the same time needing to acquire more knowledge about the new working environment. This is what sets the experienced member of staff apart from the newly arrived one, and is a classic knowledge management problem.

The Pellucid project has chosen to use autonomous software agents as the basis for its customisable platform. This approach is extremely well suited to a complex problem such as knowledge management. Different classes of agents have distinct responsibilities in seeking and presenting timely and appropriate information to the user. The Role Agent is responsible for assessing the user's need based on the context in the work process; the Task Agent and Information Search Agent are responsible for searching for information to satisfy that need; while the Personal Assistant Agent deals with presenting it in a personalised way to the user. A Capitalisation Agent has the job of building new knowledge from the organisational memory: for example, using data mining to detect patterns of usage of documents in processes.

The initial focus of the knowledge management offered by Pellucid is on document management, contact management and critical timing management. Knowledge in these areas can be key to effective and efficient working in many routine and non-routine tasks: for example, where to find documents that can be used as models for a new case; who to contact in external organisations, or internally someone who has done a similar task before; how long to expect certain interactions to take. The three pilot applications within the project, though very different in character, all exhibit these three needs for knowledge management. The pilot applications are:

- Comune di Genova (Italy), Traffic and Mobility Directorate: planning and installation of traffic lights in the city;
- Mancomunidad de Municipios del Bajo Guadalquivir (an association of local governments in the south of Spain): management of projects and services;
- Consejería de la Presidencia, Junta de Andalucía (regional government in the south of Spain): management and resolution of problems with fixed telephony systems.

Underlying the operation of the Pellucid agents is a collection of ontologies, representing the structure and relationships of aspects of organisations and processes. These ontologies allow the agents to communicate together and to perform reasoning. A generic ontology plus sub-ontologies for documents and contacts have been developed. A key line of research for the remainder of the project will be the use of domain-specific ontologies and reasoning to allow retrieval and suggestions to the user based on domain-specific knowledge—for example, knowledge about similarities of traffic light installations based on their properties (type of road junction, traffic density, …). The project aims to develop a generic, customisable platform, and the separation between generic and domain-specific parts will be of great importance.

Implementation is being conducted using the Java-based agent toolkit JADE. The possibility of using JESS (Java Expert Systems Shell) is being examined. Java Data Objects, an object-oriented interface to data storage, has been chosen as the basis of the organisational memory. Interfaces are being provided to workflow management systems that are compliant with the Workflow Management Coalition standards. In addition, the concept of a 'workflow tracking system' is being introduced, to allow a lightweight provision of context information to the Pellucid system in cases where a full workflow management system is not needed or desired.

The project is funded under the European Commission's IST programme, and began in March 2002. There are eight partners from across Europe, including three local or regional government bodies as end-users, and two members of ERCIM: CCLRC and the Slovak Academy of Sciences Institute of Informatics (part of SRCIM). The project is currently (May 2003) preparing for the preliminary user evaluations at two of the pilot sites.

Link:
http://www.sadiel.es/Europa/pellucid/

Please contact:
Simon Lambert, CCLRC
Tel: +44 1235 44 5716
E-mail: S.C.Lambert@rl.ac.uk

Radiosity for Realistic Real-Time Rendering

by François Cuny

What influences have different surfaces and materials on the impression of the observer? How do different lighting situations influence the perception of architecture or a design? In which areas is additional artificial illumination required?

Radiosity and ray tracing are two methods to simulate global lighting effects. Although diffuse inter reflections are crucial to produce realistic images, radiosity-like methods are rarely used in production rendering: slow performance, robustness problems, and control difficulty. VSP Technology succeeded in resolving all the industrial constraints.

Founded in 2001, based on research conducted by INRIA, VSP Technology develops and markets a portfolio of lighting simulation software products. With their technologies, it is now possible to physically simulate and render the lighting effect on a 3D scene in real time.

VSP Technology uses wavelet functions in its algorithm in order to resolve the global illumination problem. The use of wavelet functions ensures to have the most accurate approximation of the unknown radiosity function B(x). These recent hierarchical methods combined with a virtual mesh approach have greatly speeded up the computation of radiosity. By carefully designing a parallel algorithm that minimizes synchronization and memory access, VSP Technology was able to obtain speedups of 54 on a 64-processor machine.

Also, an important problem is to provide an efficient solution to display the computed radiosity approximation. Both a meshing and a texturing technique give a solution of this problem.

VSP Technology uses a texturing technique to store the photometric information. This technique consists in separating lighting information from the geometry by storing the illumination information as a texture map. This is commonly called light map. At the rendering step, this set of textures is blended with the initial geometry to render as accurately as possible the lighting information. Since graphic boards can render up to eight textures in only one pass, this technique is actually opportune and allows 3D interactive manipulation.

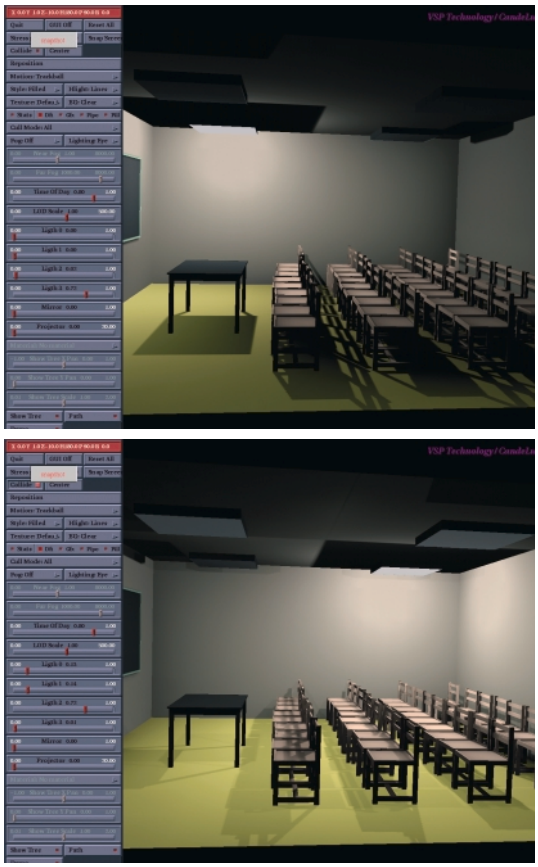

Multiple light maps generation on a same geometry.

To ensure that the lighting information is displayed correctly, the mapping process must avoid any overlap off the 3D model. We are unfolding the model using a segmentation method to decompose the model into charts with natural shapes. This method is based on a new optimization-based parameterization method: the least squares conform maps (LSCM).

The choice of the LSCM parameterization method is motivated by the properties proved by its author (Levy et al.):

Their conformal criterion minimizes angle deformations, which allows often finding a valid parameterization. The parameterization is fast enough to be applied recursively until a valid and low distorted map is found. The natural border extrapolation avoids the distortions implied by the fixed border of other methods. And the result is independent of the resolution of the mesh. This property reduces texture swimming when applied on a multi-resolution structure such as progressive mesh.

In order to avoid overlapping triangles in parameterization space, an extra segmentation is computed. Combined with a validation step, our segmentation method ensures to find a valid solution with a controlled distortion.

The result of our wavelet-based radiosity computation is an illumination function from a 2D space (the parameterization space, also called virtual mesh) to a 3D space (colour space) defined on a local normalized mapping of each mesh element. Two main steps drive to the solution:

- the first step of our algorithm approximates the radiosity function on the initial geometry or its triangulation in the parametric domain with colours stored at the vertices
- the second step is the light map generation to render the illumination function on each mesh element. Then the set of radiosity textures can be mapped

on the initial geometry like other existing textures.

VSP Technology has applied the generation of radiosity textures to industrial models such as architectural and manufacturer design industries. Our radiosity light maps can easily be integrated into existing rendering software and dramatically increase the realism of a virtual scene without slowing down the rendering speed.

Furthermore, we can group light sources together to generate a set of light maps: one light map for each group of light sources. Then, the light maps can be blended together during the rendering step, allowing driving each group of light sources independently. This allows to interact in real-time on the lighting design and to visualize the impact of each light source group in the scene, by controlling independently the intensity of each light group.

The implication of this technology is huge: significant reduction of project planning, enormous cost-saving for the visualization of projects, physically accurate simulation results in combination with breathtaking image quality. This aspect of physical accuracy cannot be overemphasized in the domain of design: looks-as-if-real counterfeits obtained by any means - and approach used in video and movies - are hardly acceptable for serious design jobs.

Links:
http://www.vsp-technology.fr/

Please contact:
François Cuny, VSP Technology, France
Tel: +33 3 83 58 17 51
E-mail : cuny@vsp-tech.com

Ultra-High-Speed Image Processing: Over 10 000 Frames Per Second

by Ákos Zarándy and Róbert Fekete

Cellular Visual Microprocessor (CVM) technology allows extremely high processing speeds to be reached in industrial applications of machine vision systems. These applications are being developed as part of a larger series of development projects carried out by AnaLogic Computers Ltd, a spin-off company of SZTAKI's Analogical and Neural Computing Systems Laboratory.

Machine vision systems are widely used in industry for quality assurance and process control, among other applications. Such systems usually consist of one or more cameras, a framegrabber card, a PC, and the image-processing software. However, compact 'smart' cameras are also appearing on the market, which integrate the above components into a single, compact unit that can autonomously perform certain tasks such as fault detection. In these novel devices, the image-processing algorithms are most often performed by a digital signal processor (DSP), or as in our case, by a Cellular Visual Microprocessor (CVM).

Until now, the automation of multiple processes, production steps, or product inspections was hindered by two main obstacles:

- the algorithms to be utilised were computationally too intensive compared to the capacity of computers;
- often the inspection requirement of the processes required a much higher rate of sampling than the sensors and cameras were capable of handling.

These barriers have now been overcome with the introduction of the CVM technology. CVMs have the ability to capture and process several thousand images per second in real time, while common digital vision systems usually provide only one to two hundred processed images per second.

CVMs are essentially CMOS sensors with an integrated image processor. The image processor is a cellular neural network-based image-processing array, and consists of several identical processor elements. When performing an operation on an image, an everyday processor (like that of a PC) performs the operation on the pixels one by one. In the processor array of the CVM, every pixel of the picture is handled by a different processor element, meaning the operation can be performed simultaneously for every pixel. The individual processor elements are also directly interconnected with their nearest neighbours, enabling high-speed local operations.

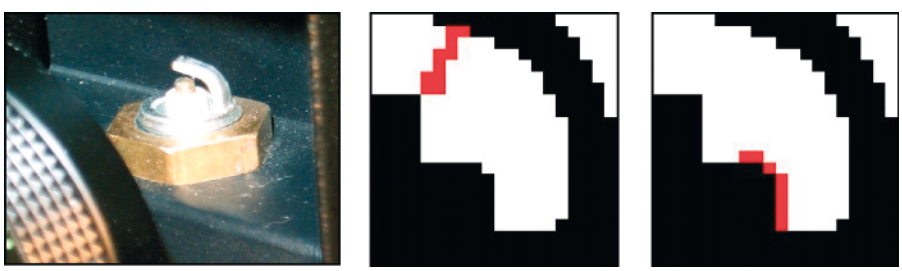

Figure 1: Spark plug and the detected discharge (gas and surface discharge).

Owing to the parallel architecture and the local connectivity, the CVM can perform operations on the full image in a few microseconds. Its computing capacity is the equivalent of 1012 digital operations per second. The CVM can store programs like common computers, and can be programmed from common development environments, using either native or the well-known C and C++ languages.

As an example, a CVM's extremely high speed is capable of detecting and classifying the millisecond-long discharges in the spark plugs of cars (Fig.1). The high-quality time resolution also makes it possible to inspect the build-up of the discharges. The images of the figure were captured by the first prototype of the CVM, a binary system with 20x22 pixel resolution. The testing measurements of a new 128x128 pixel CVM, which is capable of capturing and processing greyscale images, are currently in their final phase (Fig.2).

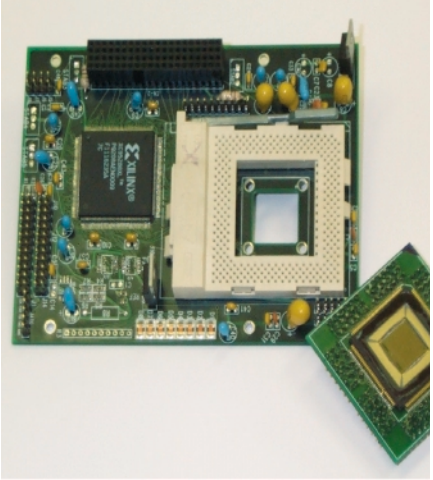

Figure 2:
The 128x128 pixel Cellular Visual Microprocessor and its testing board.

Both of these CVM chips were designed in IMSE-CNM Seville, Spain.

Based on this novel technology, a wide range of applications in a number of areas is being developed. CVMs can be utilised in any context requiring high frame-rate and fast processing. Some examples include:

- monitoring of passenger position and airbag control for the automotive industry, eg, determining if the airbag can be opened without harming the passenger
- quality control in the textile industry, eg, detecting faults in yarn and weaving, as well as knots, dirt etc
- terrain recognition and classification for autonomous navigation of unmanned vehicles
- position detection for robot control.

Link:
http://www.analogic-computers.com

Please contact:
Ákos Zarándy and Róbert Fekete
AnaLogic Computers Ltd., Budapest
Tel: +36 1 279 6131
E-mail: info@analogic-computers.com

Internet-Based Learning and Collaboration for Dispersed Co-Workers

by Dian Tan and Peter Mambrey

The internet-based collaboration platform NetKM, developed at Fraunhofer FIT, supports self-organised learning and knowledge-sharing. It was established in an insurance company to assist a geographically distributed team in organising its daily work practice.

The concepts of learning and information sharing in organisations are rapidly changing: from the idea of learning in advance and the false belief that knowledge can be stored in repositories, to self-organised information sharing within groups and the construction of organisational 'knowledge' bases. We conducted a longitudinal case study to investigate how a group of dispersed co-workers acquires a new technical platform for the purposes of sharing information and self-reliable organisational learning. The daily use of technology and its options evolved over time based on the user's voluntary activity, without direct intervention. We explicitly focused on the self-organisation of these processes.

The users were trainers within an insurance company. They were concerned with internal vocational training, and the appendant back office - 22 people in total. The back office was located at the company's headquarters, while the trainers were dispersed all over the country. Since the back office was responsible for the trainers' assignments, a lively exchange between back office and trainers is implied.

As a technical basis, we chose the BSCW system (Basic Support for Cooperative Work), developed by the Fraunhofer Institute FIT. This groupware system facilitates the coordination and accomplishment of group work in shared workspaces. Since it is web-based it works independently of a specific operating system and can be accessed from any computer with an internet connection. The user can create folders and various types of objects (charts, URLs etc), and download and maintain them (version control, conversion). Object-specific information (eg who has read or changed a document) can be recalled anytime by the members of the workspace by a simple mouse click on a symbol next to that object. They can receive a daily activity report that lists all activities which occurred the day before,

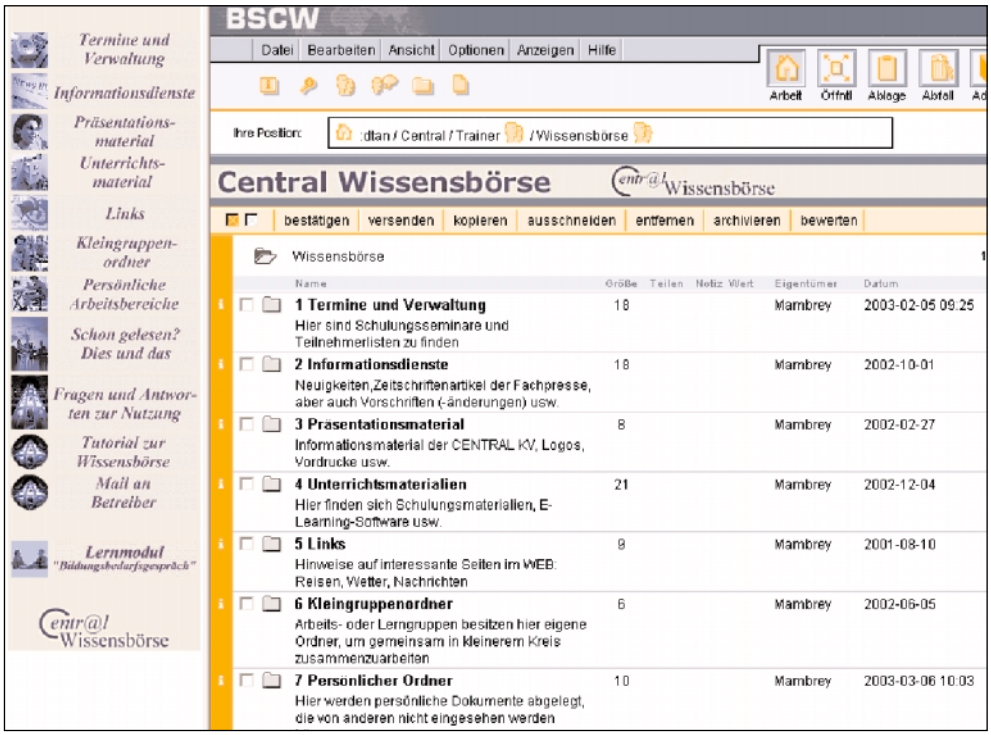

The main structrure of the 'Networked Knowledge Market'.

and the system also offers flexible allocation of rights. Plenty of other functionalities are available. The BSCW system was modified to meet the requirements of the insurance company, and was introduced using the metaphor 'Networked Knowledge Market' (NetKM). Only invited and registered members could enter this portal. Advised by the company, we developed a folder structure consisting of nine main folders, eg organisation, presentations, teaching materials, FAQ or personal folders. We provided them with some material and all users could add their own material. Our aim was to support self-organisation, rather than regulating the content and ways of using the system. The customary methods of communication persisted, and the NetKM complemented but did not replace them.

Evaluation

The period of use started in June 2001 with an introductory workshop and training to explain the handling and functionalities of the system. Our findings refer to the one-year period from the start to July 2002, but the system is still in use by the company. The exploratory study was evaluated using qualitative and quantitative methods. All available data on the use of the NetKM (eg daily activity reports) were documented and analysed. Additionally, all users were asked to answer some questions about the NetKM. We also conducted interviews with seven carefully chosen people.

Users were encouraged to establish as many objects as they wanted within the main folders. In the first year these totalled 337 folders on different levels within the main folders; the folders contained 1043 documents. The number of new objects varies greatly over the period. Back-office employees, trainers and researchers all contributed material to the NetKM. Interestingly, the data suggest that trainers who left the department stored their 'knowledge' in the NetKM shortly before leaving. In this way, important information was transmitted to colleagues and remained available to the company.

The number of downloads (reading activities) per month from the NetKM is even more irregular. The data show that the number of downloads is obviously not dependent on the number of new documents. For example, there was one month (March 2002) when only two new documents were uploaded, but 101 downloads were registered. This may be explained by the fact that this analysis accounted only for the upload of new documents, not the update of existing documents. Documents that were often updated contained current information and were often downloaded by users.

The written and oral questioning brought up some interesting findings. The users mentioned plenty of things they appreciated about the NetKM and also some aspects that could be improved. They also provided us with useful hints for the further development of mobile self-organised methods of knowledge-sharing in dispersed groups.

Links:
http://www.fit.fraunhofer.de
http://www.olvio.de

Please contact:
Dian Tan, Fraunhofer FIT
E-mail: dian.tan@fit.fraunhofer.de

Internet and Isolated Villages: An Opportunity for Development

by Diego Puppin

A recent study proposes a reliable and inexpensive solution for Internet access in isolated villages in Bhutan. The results of the study are to be proposed to the United Nation Development Programme and to the Government of Bhutan for implementation.

Information and Telecommunication Technologies (ICT) are tools that can boost development in emerging nations. In a developing country, ICTs can give access to information resources that are difficult or expensive to reach otherwise. Besides the relevant governments, international agencies such as the United Nation Development Programme (UNDP) and International Telecommunication Union (ITU) are working to bring forth the best technologies to give developing countries a boost in their development and to narrow the digital divide. Nonetheless, the lack of cheap Internet connection, or of electric power, can prevent the implementation of digital services in rural and isolated communities. Simple technological solutions can be effectively deployed to face these limitations: the transfer of solutions from research laboratories to the villages can bring opportunities for growth.

Here, we want to highlight some of the contributions of a recently published study on the use of ICT solutions to offer a reliable and inexpensive Internet connection to isolated villages in Bhutan. The analysis was completed by a group of four MIT students, and is to be proposed to the UNDP and to the Government of Bhutan for implementation.

courtesy of Per K. Sorensen

When the village is online, a number of initiatives can be planned.

In our analysis, we quickly realized that access to the Internet in a developing country can be extremely costly. Various strategies to reduce the required bandwidth or the on-line time have been developed. Offline Internet is a very promising approach, explored, among others, by the TeK (Time equal to Knowledge) group at MIT. The idea is simple: when browsing the Internet, a user will only have access to pages locally stored in a big cache in the local disk. If a page is not locally stored, a request for the page will be queued, and completed the first time the machine goes online.

But a second way is offered by TeK. When users are looking for some data on the Net, they will fill a form with the necessary keywords describing their interest, the same way they would do with a common search engine. This information will be sent to a TeK server when the machine goes online. Later on, when the machine goes online a second time, the compressed results will be retrieved from the TeK server, and the users can navigate them, as they would do from the result page of any search engine. All the results of previous queries, and queries from other users, are kept stored, and they are available to all the users. Over time, TeK will build a local library with knowledge of interest to that community.

This system offers a series of clear advantages, including reduced communication costs, increased reliability, increased opportunities for policy enforcement and control. It can be used as the main tool to offer a reliable connection in the presence of unreliable power supply or Internet connectivity.

When the village is online, a number of initiatives can be planned:

- *Tele-education*: The wealth of knowledge available on the Internet can be an important element to support education in small communities. High-quality syllabi and class materials are cheaply available on the Internet. Educational partnerships can be set up:

big schools in developed areas can offer educational support to smaller schools and communities, by offering remote assistance for grading and teaching. On the other hand, real problems of the village can be brought to classes for solutions.
- *Tele-jobs:* In the globalized world economy, many opportunities for tele-jobs are present. The availability of low-cost, English-speaking, skilled labor, can be exploited by international companies to reduce costs. Customers' data can be sent digitally from abroad to the workers in the community, be processed, and sent back. A small isolated village with some computing capabilities could boost its economy this way.
- *Assisted health care*: Data about an isolated population could be analyzed by a remote staff of doctors in a large hospital. The presence of epidemic problems or rare diseases could be observed by expert doctors, who can warn the local health facilities. Good results have been observed in India for a similar project, focused on eye-care.
- *E-post*: The opportunities for cheap and reliable communication are important even for small communities, where a piece of mail can save a day otherwise spent traveling. Mail messages can be scanned, sent digitally and printed at the destination, with big savings in time, money and energy.

In the study mentioned above, we describe how the TeK software, along with other solutions (wireless connection, local networks, peer-to-peer data transfer, solar panels) can be used within the context of the five-year Bhutan Education Master Plan to reduce the costs of the introduction of Internet connection to the schools. But more than that, we show how simple IT solutions can be used to connect villages with limited or no electricity, and limited or no Internet connection, to the information available on the Internet.

The opportunities offered by the Internet to developing communities are endless, and the means for implementation are available. ICTs can open the way to a de-materialized and more sustainable path of growth for the world's poor: it is the responsibility of the research community to direct its energy where it can be the most effective to the world population.

Please contact:
Diego Puppin, ISTI-CNR
Tel: +39 050 315 3003
E-mail: diego.puppin@isti.cnr.it

SCILAB Consortium launched

by Claude Gomez

Following the tremendous success of INRIA's Open Source Software for Numerical Computations, called SCILAB, a consortium has been created to guarantee the quality and perpetuation of the software and to promote Scilab as worldwide reference software in academia and industry.

SCILAB is a scientific software package developed by INRIA and ENPC (École Nationale des Ponts et Chaussées). SCILAB is one of the numerical computing tools such as Matlab that totally revolutionized research and development in engineering some twenty years ago. It provides a powerful open computing environment for engineering and scientific applications. Distributed freely via the Internet since 1994, SCILAB is currently being used in educational and industrial environments around the world.

SCILAB includes hundreds of mathematical functions with the possibility to add interactively programs from various languages (C, FORTRAN, C++…). It has sophisticated data structures (including lists, polynomials, rational functions, linear systems…), an interpreter and a high level programming language. SCILAB has been conceived to be an open system where the user can define new data types and operations on these data types by using overloading mechanism.

A number of toolboxes are available with the system:
- 2-D and 3-D graphics, animation
- Linear algebra, sparse matrices
- Polynomials and rational functions
- Simulation: ODE solver (ODEPACK) and DAE solver (DASSL)
- Scicos: block-diagram simulator for hybrid dynamic systems
- Classic and robust control, LMI optimization
- Differentiable and non-differentiable optimization
- Signal processing
- Metanet: graphs and networks
- Parallel SCILAB using PVM
- Statistics
- Interface with Computer Algebra (Maple, MuPAD)
- Interface with TCL/TK
- a great number of contributions for various domains.

SCILAB works on most UNIX systems including Linux and on Windows 9X/NT/2000/XP. It comes with source code, on-line help and English user manuals. Binary versions are available.

Today there are more than 10,000 monthly downloads of SCILAB software from SCILAB web site. This success is mainly due to the fact that, by integrating numerous contributions, SCILAB offers tools pertaining to scientific computing (simulation, graphics, modelling, control, etc) that are not only of interest to teachers and researchers, but also increasingly of interest to persons who work in industrial sectors such as the

automotive industry, aerospace, chemistry, finance, etc. Moreover, SCILAB is used all over the world. Books and articles already exist in French, English and Chinese. Workshops and lessons took place in India, Africa and China. Since 2002 an annual SCILAB contest is organized in China.

SCILAB Consortium

To face this success, INRIA proposed to create a Consortium to group academic and industrial members all over the world. The Consortium will guarantee the quality and perpetuation of the software and turn SCILAB into worldwide reference software in academia and industry. Members of the Consortium can contribute to the evolution of SCILAB based on new needs. A scientific committee and a steering committee composed of elected Consortium members will decide on the directions for new developments and play the role of SCILAB architect. The Consortium is hosted by INRIA.

The main objectives of SCILAB Consortium are:

- to group the community of developers to turn SCILAB into a platform to integrate all scientific advances in scientific computations, simulation and automatic control. So SCILAB will be a privileged link between companies and new scientific trends and will rely on a large community of contributors.
- to group the community of users to turn SCILAB into quality software for operational use as well in industry as in education and research.
- to provide resources for a dedicated team for the development of SCILAB.

The dedicated team has already been created as a 'Development Project-Team' at INRIA-Rocquencourt. It is managed by the chief technical officer of the consortium and will be eventually composed of 14 members. In addition to the development of SCILAB, the team will promote SCILAB and animate and federate the community of contributors. INRIA committed to finance 5 people in the team.

On May 16, the SCILAB Consortium was launched with an initial meeting of its members under the chairmanship of Professor Alain Bensoussan, former chairman of INRIA, CNES and Council of ESA. The Consortium founding members are French companies and academics: AXS INGENIERIE, CEA, CNES, CRIL TECHNOLOGY, DASSAULT-AVIATION, EDF, ENPC, ESTEREL TECHNOLOGIES, INRIA, PSA Peugeot Citroën, RENAULT, and THALES. They firmly intend to increase the European and international visibility by attracting new members abroad.

Links:
http://www.scilab.org
Newsgroup: comp.soft-sys.math.scilab

Please contact:
Claude Gomez, Chief Technical Officer
of Scilab Consortium, INRIA
E-mail: claude.gomez@inria.fr

ERCIM-Sponsored Events

ERCIM continues sponsoring up to ten conferences, workshops and summer schools per year. The funding for all types of events is in the order of 2000 Euro.

Conferences

ERCIM invites sponsorship proposals from established conferences with an international reputation, where substantive overlap can be shown 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 an established international science publisher.

Workshops and Summer Schools

ERCIM sponsors workshops or summer schools (co-) organised by an ERCIM institute. The additional funding provided by ERCIM should be used to enhance the workshop by, for example, increasing the number of external speakers supported.

Forthcoming events Sponsored by ERCIM:

- ECOOP – European Conference on Object-Oriented Programming,
  Darmstadt, Germany, 21-25 July 2003
- ECDL2003 – 7th European Conference on Research and Advanced Technology for Digital Libraries,
  Trondheim, Norway, 17-22 August 2003
- FM 2003 – 12th International Formal Methods Europe Symposium,
  Pisa, Italy, 8-14 September 2003
- CP 2003 – 9th International Conference on Principles and Practice of Constraint Programming,
  Kinsale, Ireland, 29 September-3 October 2003
- WISE 2003 – 4th International Conference On Web Information Systems Engineering,
  Rome 10-13 December 2003
- SOFSEM 2004 – The Conference on Current Trends in Theory and Practice of Informatics,
  Merin, Czech Republic, 24-30 January 2004

For detailed information about ERCIM's event sponsorship programme, see:
http://www.ercim.org/activity/sponsored.html

WWW2003 — The Twelfth International World Wide Web Conference in Hungary

by László Kovács

The Twelfth International World Wide Web Conference took place in Budapest, Hungary, 20-24 May 2003, organized by the International World Wide Web Conference Committee (IW3C2) and SZTAKI.

Beginning with the first International WWW Conference in 1994, this prestigious series of the International World Wide Web Conference Committee also provides a public forum for the WWW Consortium (W3C) through the annual W3C track.

The conference with more than 800 attendees from 57 countries started with a day of tutorials and workshops and followed by a three-day technical programme. The fifth day was a 'Developers Day.' The tutorials and workshops provided in-depth looks at specific areas of current interest. The technical programme included refereed paper presentations, alternate track presentations, plenary sessions, panels and poster sessions. Developers Day was devoted to in-depth technical sessions designed specifically for Web developers.

The W3C presented three days of presentations on the Web, future Web browsers, W3C architectural principles, the XML family, Web services, the Semantic Web, new devices, and horizontal essentials.

A number of 602 submissions was received to the refereed papers track this year, a 33% increase over the previous record. 77 papers were accepted (compared to 72 in the previous year) for the proceedings and presentation at the conference, an acceptance rate of 12,8% which makes this World Wide Web conference the most competitive ever. All papers were rigorously reviewed by the technical program committee, which was divided into the following twelve areas:

- Applications
- Browsers and User Interfaces
- Electronic Commerce
- Hypermedia
- Languages
- Mobility and Wireless Access
- Multimedia
- Performance and Reliability
- Search and Data Mining
- Security and Privacy
- Semantic Web
- Web Engineering

Word Wide Web conference in Budapest.

Several changes were made in the area lineup this year: Web Engineering as a new area, and the Electronic Commerce and Security were split into two tracks, Search and Data Mining, and Semantic Web.

One of the main points of the conference was the harmonic symbiosis of the Web Services and Semantic Web in a longer run.

The Conference Co-Chairs were Gusztáv Hencsey (SZTAKI) and Bebo White (Stanford Linear Accelerator Center), and the Programme Committee Co-Chairs were Yih-Farn Robin Chen (AT&T Labs - Research), László Kovács (SZTAKI) and Steve Lawrence (Google, Inc.). The main sponsoring institutions of the event were Ministry of Informatics and Communications, Hungarian National Tourist Office, IBM Research, Alvarion, AT&T Labs, ERCIM, W3C, NIIF and VOA News, and the conference partners were International Federation for Information Processing and W3C.

Links:
WWW2003: http://www2003.org/
IW3C2: http://www.iw3c2.org/

Please contact:
László Kovács, SZTAKI
Tel: +36 1 279 6212
E-mail: laszlo.kovacs@sztaki.hu

ERCIM NEWS

ERCIM News is the 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. This issue has a circulation of over 9000 copies.

ERCIM News online edition is available at http://www.ercim.org/publication/ERCIM_News/

ERCIM News is published by ERCIM EEIG, BP 93, F-06902 Sophia-Antipolis Cedex
Tel: +33 4 9238 5010, E-mail: office@ercim.org
ISSN 0926-4981

Director: Bernard Larrouturou
Central Editor:
Peter Kunz
peter.kunz@ercim.org
Local Editors:
AARIT: Karin Hraby
hraby@ocg.at
CCLRC: Martin Prime
M.J.Prime@rl.ac.uk
CRCIM: Michal Haindl
haindl@utia.cas.cz
CWI: Annette Kik
Annette.Kik@cwi.nl
CNR: Carol Peters
carol.peters@isti.cnr.it
FORTH: Yannis Pachoulakis
yannisp@ics.forth.gr
FhG: Michael Krapp
michael.krapp@scai.fraunhofer.de
FNR: Patrik Hitzelberger
hitzelbe@crpgl.lu
INRIA: Bernard Hidoine
bernard.hidoine@inria.fr
NTNU: Truls Gjestland
truls.gjestland@ime.ntnu.no
SARIT: Harry Rudin
hrudin@smile.ch
SICS: Kersti Hedman
kersti@sics.se
SRCIM: Gabriela Andrejkova
andrejk@kosice.upjs.sk
SZTAKI: Erzsébet Csuhaj-Varjú
csuhaj@sztaki.hu
TCD: Ann McNamara
Ann.McNamara@cs.tcd.ie
VTT: Pia-Maria Linden-Linna
pia-maria.linden-linna@vtt.fi

Free subscription

You can subscribe to ERCIM News free of charge by:

- sending e-mail to your local editor
- contacting the ERCIM office (see address above)
- filling out the form at the ERCIM website at http://www.ercim.org/

CALL FOR PAPERS

IZS — 2004 International Zurich Seminar on Communications

Zurich, Switzerland, 18-20 February 2004

Original technical contributions are solicited on topics including:
Access:

- Wide area mobile radio
- Wireless local area networks
- Wireless local loops
- xDSL and cable modem technologies

Transmission:

- Channel modeling
- Applied information theory
- Coding and modulation
- Joint detection, decoding, and ... (you name it)
- MIMO wireless and space-time coding
- Ultra wide band
- Antenna design
- Fundamental hardware issues

Networking:

- Algorithms, protocols, and systems for routing and switching
- Ad-hoc and peer-to-peer networking of all scales
- High-mobility architectures and protocols
- Next-generation internet
- Optical networks
- Quality of service
- Fault-tolerance and reliability
- Security

As a novelty in 2004, roughly half the contributions will be invited.

Due to its traditional format with at most two parallel sessions, the IZS is an ideal opportunity to learn from, and to communicate to, leading experts in areas beyond one's own speciality. All papers should therefore be presented with this wide audience in mind. Authors of accepted papers will be asked to produce a manuscript not exceeding 4 pages that will be published in the Proceedings. All accepted papers will be allowed twenty minutes for presentation.

Deadline for submissions:
15 September 2003.

More information:
http://www.izs2004.ethz.ch

CALL FOR PARTICIPATION

FM 2003: 12th International FME Symposium

Pisa, Italy, 8-14 September 2003

FM 2003 is the twelfth in a series of symposia organised by Formal Methods Europe, an independent association whose purpose is to promote and support the industrial use of formal methods for computer systems development. The FME symposia have been notably successful in bringing together an international and varied community of users, researchers and developers of precise mathematical methods for software development.

The scope of FM 2003 includes all aspects of the use of formal methods for software development, with particular attention to new application areas that are becoming critical in the new generation information society. The symposium will attempt to cover a wide variety of themes, ranging from fundamental theories to practical experiences. In addition to the three-day symposium, FM 2003 will host 7 Workshops, 8 Tutorials and system and tool demo sessions. There will also be a special day dedicated to industrial concerns: the FME/CoLogNet Industry Day.

More information:
http://fme03.isti.cnr.it/

CALL FOR PARTICIPATION

SIGCOMM 2003

Karlsruhe Germany, 25-29 August 2003

This year, SIGCOMM broadens it traditional scope into a data communications festival. The SIGCOMM conference contains full papers on a variety of topics, including peer-to-peer, overlays, and denial of service, as well as position paper and poster sessions. A set of one-day workshops are part of the festival: networking research methods, scope, education, network-I/O convergence, future directions in architecture, and revisiting QoS. Tutorials will be held on network security and self-similar traffic.

More information:
http://www.acm.org/sigcomm/sigcomm2003

Stelios Orphanoudakis (FORTH) will be the new President of ERCIM from January 2004. He succeeds **Gerard van Oortmerssen** (CWI), who was President from 1998. Van Oortmerssen was appointed director of TNO Telecom from May 2003. He was director of CWI for almost twelve years. At CWI he will be succeeded by **Jan Karel Lenstra** who worked previously at CWI (1969-1989), Eindhoven Technical University (1989-2002) and Georgia Institute of Technology in Atlanta, USA. Lenstra comes from a family of very renowned Dutch mathematicians.

CCLRC — **Rutherford Appleton Laboratory** hosted a launch event including a live launch feed from ESA at the occasion of the Mars Express mission.

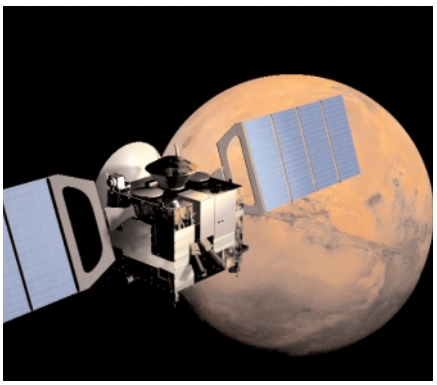

Beagle 2.

Mars Express was successfully launched at 18:45 BST on 2 June 2003 from Baikonur, Kazakstan. Mars Express comprises a number of essential components – the spacecraft and its instruments, the lander, a network of ground and data processing stations, and the launcher itself. RAL has been responsible for the thermal design and analysis of Beagle 2, which transports the lander during its descent from Mars Express.

CNR — **GÉANT, the pan-European multi-gigabit research network, and GARR, the Italian research and academic network are interconnected at 10Gbit/sec** from 6 May 2003. GARR (http://www.garr.it) is the first National Research Network in Europe to be connected at 10Gbit/sec to GÉANT (http://www.geant.net), replicating the event of July 1989, when GARR was the first European network to be connected over an international link at 2Mbit/sec between Bologna and Geneva. GÉANT enables European scientists to compete internationally by providing them with a world-class backbone that offers the bandwidth and the quality of service required for research and development activities at this level. It represents the basis for the introduction of 'virtual laboratories' and 'virtual institutes' in Europe. This will be extremely beneficial for the Italian research community participating in international projects in many domains. The services provided by GÉANT, will also allow the implementation of a powerful network research infrastructure to support GRID applications.

INRIA — '**Internatics**' is the recently launched newsletter of INRIA's Office for European and International Relations. Through short articles and news items, it reports on INRIA's involvement in European and international activites.

Besides the sections dedcated to cooperations in Europe, the Americas, Asia, Africa and Middle East, Internatics regularly features topics related to the European research policy, making it a good source of information for scientists and research managers. Internatics is published quarterly. The printed edition is published in French, the online version is also available in English and offers free subscription to both printed and electronic versions.
http://www-direction.inria.fr/international/internatics/

SARIT — **Professor Ambros P. Speiser** died on May 10, 2003. He was one of the last of the early pioneers of digital computer design and an insightful initiator and staunch supporter of scientific research. Ambros Speiser was the

© Roger Nickel/University Zurich

leading hardware designer for the construction of the legendary ERMETH computer built at the ETH (Swiss Federal Institute of Technology) in Zurich 1950 until 55, ie, during the years when Konrad Zuse's Z4, was used at ETH as the first digital computer at a continental European university. The ERMETH was then used until 1963. The ETH later honored Speiser for this work by awarding him an honorary doctor's degree. Well aware of his talents, Ambros Speiser was chosen by IBM to found its Zurich Research laboratory in 1956; he was its director until 1966. Then BBC (Brown Boveri & Company, now ABB) asked him to found and lead its corporate research laboratory in Baden. This he did until his retirement in 1987. Speiser was also a professor at the ETH in Zurich and often assumed both Swiss and international positions as leader and advisor. He was president of IFIP, the International Federation for Information Processing, from1965 until 1968, and of the Swiss Academy of Engineering Sciences. Ambros Speiser was not only a researcher with the highest aspirations but, in his retirement, added another talent to his repertoire — that of a most articulate and respected scientific journalist.

Carl August Zehnder,
Professor of Computer Science,
ETH Zürich

ERCIM – The European Research Consortium for Informatics and Mathematics is an organisation dedicated to the advancement of European research and development, in information technology and applied mathematics. Its national member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry.

ERCIM is the European Host of the World Wide Web Consortium.

Austrian Association for Research in IT
c/o Österreichische Computer Gesellschaft
Wollzeile 1-3, A-1010 Wien, Austria
Tel: +43 1 512 02 35 0, Fax: +43 1 512 02 35 9
http://www.aarit.at/

Council for the Central Laboratory of the Research Councils, Rutherford Appleton Laboratory
Chilton, Didcot, Oxfordshire OX11 0QX, United Kingdom
Tel: +44 1235 82 1900, Fax: +44 1235 44 5385
http://www.cclrc.ac.uk/

Consiglio Nazionale delle Ricerche, ISTI-CNR
Area della Ricerca CNR di Pisa,
Via G. Moruzzi 1, 56124 Pisa, Italy
Tel: +39 050 315 2878, Fax: +39 050 315 2810
http://www.isti.cnr.it/

Czech Research Consortium
for Informatics and Mathematics
FI MU, Botanicka 68a, CZ-602 00 Brno, Czech Republic
Tel: +420 2 688 4669, Fax: +420 2 688 4903
http://www.utia.cas.cz/CRCIM/home.html

Centrum voor Wiskunde en Informatica
Kruislaan 413, NL-1098 SJ Amsterdam,
The Netherlands
Tel: +31 20 592 9333, Fax: +31 20 592 4199
http://www.cwi.nl/

Fonds National de la Recherche
20, Montee de la Petrusse
L-2912 Luxembourg
Tel. +352 26 1925 31, Fax +352 26 1925 35
http:// www.fnr.lu

Foundation for Research and Technology – Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
Tel: +30 2810 39 16 00, Fax: +30 2810 39 16 01
http://www.ics.forth.gr/

Fraunhofer ICT Group
Friedrichstr. 60
10117 Berlin, Germany
Tel: +49 30 726 15 66 0, Fax: ++49 30 726 15 66 19
http://www.iuk.fraunhofer.de

Institut National de Recherche en Informatique
et en Automatique
B.P. 105, F-78153 Le Chesnay, France
Tel: +33 1 3963 5511, Fax: +33 1 3963 5330
http://www.inria.fr/

Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and
Electrical Engineering, N 7491 Trondheim, Norway
Tel: +47 73 59 80 35, Fax: +47 73 59 36 28
http://www.ntnu.no/

Swedish Institute of Computer Science
Box 1263,
SE-164 29 Kista, Sweden
Tel: +46 8 633 1500, Fax: +46 8 751 72 30
http://www.sics.se/

Swiss Association for Research in Information Technology
c/o Prof. Dr Alfred Strohmeier, EPFL-IC-LGL,
CH-1015 Lausanne, Switzerland
Tel +41 21 693 4231, Fax +41 21 693 5079
http://www.sarit.ch/

Slovak Research Consortium for Informatics and
Mathematics, Comenius University, Dept.of Computer
Science, Mlynska Dolina M, SK-84248 Bratislava, Slovakia
Tel: +421 2 654 266 35, Fax: 421 2 654 270 41
http://www.srcim.sk

Magyar Tudományos Akadémia
Számítástechnikai és Automatizálási Kutató Intézet
P.O. Box 63, H-1518 Budapest, Hungary
Tel: +36 1 279 6000, Fax: + 36 1 466 7503
http://www.sztaki.hu/

Trinity College
Department of Computer Science,
Dublin 2, Ireland
Tel: +353 1 608 1765, Fax: 353 1 677 2204
http://www.cs.tcd.ie/ERCIM

Technical Research Centre of Finland
P.O. Box 1200
FIN-02044 VTT, Finland
Tel:+358 9 456 6041, Fax :+358 9 456 6027
http://www.vtt.fi/tte

Order Form

If you wish to subscribe to ERCIM News ***free of charge*** *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/them to the mailing list.*

Send, fax or email this form to:

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

Name:

Organisation/Company:

Address:

Post Code:

City:

Country

E-mail:

Data from this form will be held on a computer database. By giving your email address, you allow ERCIM to send you email

You can also subscribe to ERCIM News and order back copies by filling out the form at the ERCIM website at http://www.ercim.org/publication/Ercim_News/