

ERCIM NEWS

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

Number 52

January 2003



Special:
**Embedded
Systems**

Contents

JOINT ERCIM ACTIONS

- 4 ERCIM New European Host of W3C**
- 5 About the World Wide Web Consortium**
- 5 W3C Activities**
- 6 Recommendations published by W3C in 2002**
- 7 World Wide Web Consortium launched Hungarian Office at SZTAKI**
by László Kovács and Éva Megyaszi, SZTAKI
- 8 Human Resources Managers Task Force**
by José Koster, CWI and Heather Weaver, CLRC
- 8 Euro-Legal**
by Heather Weaver, CLRC
- 9 Andreas Rauber Winner of the 2002 Cor Baayen Award**

Cover:

The European Ariane 5 launcher crashing after takeoff on 4 June 1996.

The explosion was the result of a software error. Embedded systems play an important role in many safety critical systems.

Photo: courtesy by AFP.

Next issue:

April 2003

Special theme:

Cognitive Systems

SPECIAL THEME

- 10 Embedded Systems - Introduction**
by Erwin Schoitsch, ARC Seibersdorf Research (AARIT)
- 12 An Approach to Dependable Architectures and Components for Railways**
by Andrea Bondavalli, Felicita Di Giandomenico and Luca Simoncini, ISTI-CNR
- 13 Monitoring and Controlling Scientific Experiments: Anywhere - Anytime**
by Christian Wattering, Peter Fornaro, Martin Guggisberg, Tibor Gyalog, and Helmar Burkhart, University of Basel (SARIT)
- 15 Emerging Frontiers in Miniature Computer and Embedded Systems**
by Angelos Bilas and Evangelos Markatos, ICS-FORTH
- 16 Wearable Systems for Everyday Use**
by Spyros Lalis, Anthony Savidis and Constantine Stephanidis, ICS-FORTH
- 18 New Solutions for In-Vehicle Embedded System Development**
by Françoise Simonot-Lion, Laboratoire Lorrain de Recherche en Informatique et ses Applications, France and Yvon Trinquet, Institut de Recherche en Cybernetique de Nantes, France
- 20 Remote Control and Monitoring over the Internet - Wireless Construction Site**
by Leila Rannanjärvi, VTT
- 21 Creating a Network for Those working with Embedded and Real-Time Systems**
by Barrie Ross-Dow, Institute of Electrical Engineers, UK
- 22 DECOS - Dependable Embedded Components and Systems**
by Dieter Donhoffer and Erwin Schoitsch, ARC Seibersdorf Research (AARIT)
- 24 Time Triggered Architecture**
by Hermann Kopetz, Vienna University of Technology, Austria
- 26 Automatic Distribution of Synchronous Programs**
by Alain Girault, INRIA
- 28 synERJY: a High-Level Language for Embedded Controllers**
by Reinhard Budde, Axel Poigné and Karl-Heinz Sylla, Fraunhofer-Gesellschaft

ERCIM has been chosen as European host of the World Wide Web Consortium (W3C). Read about this event and the W3C on pages 4-7.

- 29 STOREQ: Storage Requirement Estimation and Optimisation Tool for Data-Intensive Embedded Applications**
by Per Gunnar Kjeldsberg, NTNU
- 31 Composing Embedded Real-Time Software Components: the PECOS Data-Centric Approach**
by Stéphane Ducasse, Oscar Nierstrasz and Roel Wuyts, University of Berne (SARIT)
- 32 Time Machines and Black Box Recorders for Embedded Systems Software**
by Henrik Thane, Malardalen University, Sweden
- 33 Design of Embedded Systems: New Challenges**
by Gilbert Cabillic and Isabelle Puaut, INRIA
- 35 Transforming Specifications to verify Embedded Systems**
by María del Mar Gallardo, Jesús Martínez, Pedro Merino and Ernesto Pimentel, University of Malaga, Spain
- 36 EUDEMES: Component-Based Development Methods for Small-Size Embedded Systems**
by Nicolas Guelfi, Luxembourg University of Applied Sciences (FNR), Aloyse Schoos, IEE S.A, Luxembourg, Gérard Sébastien and François Terrier, National Institute of Nuclear Sciences and Techniques (CEA), France
- 37 Model Checking of Embedded Systems**
by Stefania Gnesi, ISTI-CNR
- 39 Systems Validation Centre shows Benefits of Formal Methods**
by Henk Eertink, Telematica Instituut, Wan Fokkink, CWI, Holger Hermanns, University of Twente, and Izak van Langevelde, CWI
- 40 REDEST — Sharing Experiences in the Field of Embedded Software Requirement Gathering**
by János Nacsa, SZTAKI
- 42 EmbLab — A Laboratory for Teaching Embedded Systems**
by Amund Skavhaug, Trygve Lunheim and H. Skinnemoen, NTNU
- 43 Master Course Embedded Systems at Eindhoven University of Technology**
by Jan Friso Groote, Eindhoven University of Technology, The Netherlands


R&D AND TECHNOLOGY TRANSFER

- 45 SOMLib - New Approaches for Information Presentation and Handling**
by Andreas Rauber, Vienna University of Technology
- 46 Hybrid Caching of Search Engine Results**
by Tiziano Fagni and Fabrizio Silvestri, ISTI-CNR
- 48 Personal Space Station handles Virtual 3D Objects**
by Jurriaan Mulder, CWI
- 49 Augmented Virtual Table Tennis with Real Racquets**
by Charles Woodward and Petri Honkamaa, VTT
- 51 Scene Understanding using Hierarchical Neural Networks**
by Gabriele Pieri and Ovidio Salvetti, ISTI-CNR
- 52 A Digital Film-Saver for Archives**
by George L. Kovács and Iván Kas, SZTAKI
- 54 Knowledge Technologies in Support of e-Government and e-Democracy**
by Jan Paralic and Tomas Sabol, Technical University of Kosice (SRCIM)
- 55 Rapid Prototyping for Embedded Parallelism**
by David Johnston, Martin Fleury and Andy Downton, University of Essex, UK
- 57 Discretizing Continuous Attributes in AdaBoost for Text Categorization**
by Pio Nardiello, Fabrizio Sebastiani, and Alessandro Sperduti, ISTI-CNR

EVENTS

- 58 The 7th ERCIM Workshop on 'User Interfaces for All'**
by Noëlle Carbonell, INRIA and Constantine Stephanidis, ICS-FORTH
- 59 Announcements**

63 IN BRIEF



Tim Berners Lee,
Director of W3C (left)
and Gerard van
Oortmerssen,
President of ERCIM.

ERCIM New European Host of W3C

The change of W3C European Host from INRIA to ERCIM took place on 1 January 2003. This organisational change aims to strengthen research relationships throughout Europe to better support Web technology development.

As the Web has grown into a more widespread and essential resource for industry, academia, governments, and their citizens, there is a pronounced need for an even more diverse infrastructure and support network. Establishing the W3C's European base at ERCIM will be beneficial to all involved. W3C's move to ERCIM opens new opportunities to integrate Web research and industry in Europe.

"The Web grows stronger and more useful when there are real commitments to engaging international presence and influence in Web technology development," explained Tim Berners-Lee, W3C Director. "For W3C, it means having one headquarters supported by three global partners — MIT, Keio University and now ERCIM. With the move to ERCIM, there is the potential for considerable growth and synergies of Web technologies across Europe."

"As a network of IT research centers, ERCIM encompasses over 10,000 researchers. W3C will benefit from this reservoir of expertise. For ERCIM, this change brings an opportunity to enhance cooperation on a global scale," noted

Gerard van Oortmerssen, President of ERCIM.

Since its inception in the United States at MIT's Laboratory for Computer Science in October 1994, W3C has made sure that the World Wide Web lives up to its name and its promise. At the insistence of its Director, Tim Berners-Lee, the W3C has equal presence in North America, Europe and Asia. W3C identifies not-for profit institutions to serve as regional headquarters and provide physical locations for staffing. These institutions are identified as W3C 'hosts'. Moving the European host to ERCIM allows W3C to expand its base from one country to 16 for improved research and development opportunities throughout Europe, while maintaining its historically strong relationship to INRIA, one of the ERCIM founders.

"INRIA provided the necessary foundations for European involvement in Web infrastructure development, and now we have the opportunity to expand into a new phase," explained Bernard Larrourou, President of INRIA and Manager of ERCIM. "Moving the host to ERCIM is consistent with INRIA's

culture of incubating new initiatives and our commitment to strengthening the IT community in Europe." When INRIA became the first W3C European host in 1995, W3C had 50 Members in four countries. Today, the W3C Membership exceeds 450 organizations, with representation in nearly every country in Europe. INRIA's resources attracted both member organisations and technical staffing to W3C. INRIA also provided crucial links to European research communities. During this time, W3C received financial support from the European Commission to build a network of Offices throughout Europe. Currently, W3C has 13 Offices worldwide, with eight in Europe. Seven of the European Offices are already based at ERCIM Institutes, including CWI (Benelux), Fraunhofer IMK (Germany and Austria); FORTH (Greece), SZTAKI (Hungary), CNR (Italy), SICS (Sweden) and CCLRC (UK and Ireland). Offices work with W3C hosts promoting local languages, broadening W3C's geographical base, and encouraging international participation in W3C Activities.

About the World Wide Web Consortium

The World Wide Web Consortium was created in October 1994 to lead the World Wide Web to its full potential by developing common protocols that promote its evolution and ensure its interoperability. W3C has nearly 450 member organizations from all over the world and has earned international recognition for its contributions to the growth of the Web.

By promoting interoperability and encouraging an open forum for discussion, W3C commits to leading the technical evolution of the Web. In just over seven years, W3C has developed more than 40 technical specifications for the Web's infrastructure. However, the Web is still young and there is still a lot of work to do, especially as computers, telecommunications, and multimedia technologies converge.

W3C's Goals

W3C's long term goals for the Web are:

- *Universal Access*: To make the Web accessible to all by promoting technologies that take into account the vast differences in culture, languages, education, ability, material resources, and physical limitations of users on all continents
- *Semantic Web*: To develop a software environment that permits each user to make the best use of the resources available on the Web
- *Web of Trust*: To guide the Web's development with careful consideration for the novel legal, commercial, and social issues raised by this technology.

W3C's Role

As with many other information technologies, in particular those that owe their success to the rise of the Internet, the Web must evolve at a pace unrivaled in other industries. Almost no time is required to turn a bright idea into a new product or service and make it available on the Web to the entire world; for many applications, development and distribution have become virtually indistinguishable. At the same time, easy customer

W3C Activities

W3C Activities are organized into groups: Working Groups (for technical developments), Interest Groups (for more general work), and Coordination Groups (for communication among related groups). These groups produce the bulk of W3C's results: technical reports, open source software, and services (eg, validation services). There are currently over thirty W3C Working Groups. To facilitate management, the Team organizes W3C Activities and other work into four domains:

Architecture Domain

The Architecture Domain develops the underlying technologies of the Web. Activities include:

- The Document Object Model
- Internationalization, to make sure that Web technology meets the needs of the global community
- Jigsaw, W3C's Web server platform
- Uniform Resource Identifier (URI), the Web's naming and addressing technology
- Web Services, programmatic interfaces for application to application communication
- Extensible Markup Language, XML.

Interaction Domain

The Interaction Domain seeks to improve user interaction with the Web, and to facilitate single Web authoring to benefit users and content providers alike. Activities include:

- Amaya, W3C's own editor/browser
- Device Independence
- Graphics
- Hypertext Markup Language (HTML)
- Math (primary focus is the MathML language)

- Multimodal Interaction
- Style
- Synchronized Multimedia
- Voice Browser.

Technology and Society Domain

The W3C Technology and Society Domain seeks to develop Web infrastructure to address social, legal, and public policy concerns. Activities include:

- Privacy
- Semantic Web
- XML Encryption
- XML Key Management
- XML Signature.

Web Accessibility Initiative

W3C's commitment includes promoting a high degree of usability for people with disabilities. The activities are:

- WAI International Program Office
- WAI Technical Activity.

In addition, the Quality Assurance Activity and Patent Policy apply to all domains.

Further information:

<http://www.w3.org/Consortium/Activities/>

feedback has made it possible for designers to fine tune their products almost continually. With an audience of millions applying W3C specifications and providing feedback, W3C concentrates its efforts on three principle tasks:

- *Vision*: W3C promotes and develops its vision of the future of the World Wide Web. Contributions from several hundred dedicated researchers and engineers working for Member organizations, from the W3C Team (led by

Tim Berners-Lee, the Web's inventor), and from the entire Web community enable W3C to identify the technical requirements that must be satisfied if the Web is to be a truly universal information space.

- *Design*: W3C designs Web technologies to realize this vision, taking into account existing technologies as well as those of the future.
- *Standardization*: W3C contributes to efforts to standardize Web technologies by producing specifications

(called ‘Recommendations’) that describe the building blocks of the Web. W3C makes these Recommendations (and other technical reports) freely available to all.

Design Principles of the Web

The Web is an application built on top of the Internet and, as such, has inherited its fundamental design principles:

- *Interoperability*: Specifications for the Web’s languages and protocols must be compatible with one another and allow (any) hardware and software used to access the Web to work together.
- *Evolution*: The Web must be able to accommodate future technologies. Design principles such as simplicity, modularity, and extensibility will increase the chances that the Web will work with emerging technologies such as mobile Web devices and digital television, as well as others to come.
- *Decentralization*: Decentralization is without a doubt the newest principle and most difficult to apply. To allow the Web to ‘scale’ to worldwide proportions while resisting errors and breakdowns, the architecture (like the Internet) must limit or eliminate dependencies on central registries.

These principles guide the work carried out within W3C Activities.

Activities

W3C organizes the work necessary for the development or evolution of a Web technology into activities. Each activity has its own structure, but an activity typically consists of one or more working group, interest group, and coordination group. Within the framework of an activity, these groups generally produce recommendations and other technical reports as well as sample code. Important to every W3C activity is quality assurance (QA) and patent policy. To manage related activities, the W3C team groups them into four domains: Architecture, Interaction, Technology and Society, and the Web Accessibility Initiative.

Guided by these design principles, W3C has published more than forty Recommendations since its inception. Each Recommendation not only builds

Recommendations published by W3C in 2002

- User Agent Accessibility Guidelines 1.0
17 December 2002, Ian Jacobs, Jon Gunderson, Eric Hansen
- XML Encryption Syntax and Processing
10 December 2002, Donald Eastlake, Joseph Reagle
- Decryption Transform for XML Signature
10 December 2002, Merlin Hughes, Takeshi Imamura, Hiroshi Maruyama
- XML-Signature XPath Filter 2.0
8 November 2002, John Boyer, Merlin Hughes, Joseph Reagle
- Exclusive XML Canonicalization Version 1.0
18 July 2002, John Boyer, Donald E. Eastlake 3rd, Joseph Reagle
- The Platform for Privacy Preferences 1.0 (P3P1.0) Specification
16 April 2002, Massimo Marchiori
- XML-Signature Syntax and Processing
12 February 2002, Donald Eastlake, Joseph Reagle, David Solo

Further information: <http://www.w3.org/TR/#Recommendations>

on the previous, but is designed so that it may be integrated with future specifications as well. W3C is transforming the architecture of the initial Web (essentially HTML, URIs, and HTTP) into the architecture of tomorrow’s Web, built atop the solid foundation provided by XML.

Challenges for Tomorrow

In other specifications, W3C is addressing a number of challenges for the Web of tomorrow:

- Ensure access to the Web by many devices.
- Promote best practices.
- Coordinate with international regulatory bodies
- Account for cultural diversity
- Encourage research.

W3C Organization

To meet its goals (universal access, semantic Web, Web of trust) while exercising its role (vision, design, standardization) and applying its design principles (interoperability, evolution, and decentralization), W3C process is organized according to three principles:

- *Vendor neutrality*: The W3C hosts are vendor and market neutral, as is the Team. W3C promotes neutrality by encouraging public comment on specifications during their entire life cycle.
- *Coordination*: The Web has become phenomenon so important (in scope and investment), that no single organi-

zation can or should have control over its future. W3C coordinates its efforts with other standards bodies and consortia such as the IETF (Internet Engineering Task Force), the Unicode Consortium, the Web3D Consortium, and several ISO committees.

- *Consensus*: Consensus is one of the most important principles by which W3C operates. When resolving issues and making decisions, W3C strives to achieve unanimity of opinion.

W3C Team

The W3C Team includes more than seventy researchers and engineers from around the world who lead the technical activities at W3C and manage the operations of the consortium. Most of the team works physically at the three host institutions. ERCIM currently hosts 19 team members.

More information:
<http://www.w3.org/>

World Wide Web Consortium launched Hungarian Office at SZTAKI

by László Kovács and Éva Megyaszi

W3C strengthens presence in Central Europe through Hungarian outreach by W3C Hungarian Office, located at SZTAKI's Department of Distributed Systems. The Office was launched in Budapest, on September 24, 2002.

The opening ceremony was a public event, with presentations and tutorials from the W3C Team, including Daniel Dardaillier (W3C Deputy Director for Europe), Marie-Claire Forgue (W3C European Communications Officer), Ivan Herman (Head of Offices at W3C), Vincent Quint (W3C Document Formats Domain leader) and Max Froumentin (Math Activity Lead and XSL Working Group Team Contact); and presentations of projects at SZTAKI (CORES, SOTF, PVX). Péter Bakonyi, Deputy State Secretary for Information Society Strategy, Hungary and Péter Inzelt, Director of SZTAKI were also among those attending the opening ceremonies at the Hungarian Academy of Sciences, in Budapest.

The Hungarian Office is the first office in Central Europe, joining an active roster which includes a number of European Offices of W3C as well as the

W3C European host site at ERCIM. Its role is to act as a local point of contact, and to make sure that W3C and its specifications are known in the region. The mission of the Office is to promote adoption of W3C recommendations among developers, application builders, and standards setters, and to encourage inclusion of stakeholder organizations in the creation of future recommendations by joining W3C.

Hungarian economy shows the most dynamic and attractive growth in the region, especially with respect to information technology and telecommunication. SZTAKI has well-established co-operations with leading institutions and firms covering different areas of Information Technology, Computer Science and Control. The Institute maintains more than 100 contracts yearly with industrial, governmental, and other partners, including joint projects with European Union funding.

W3C looks forward to a long-term partnership with SZTAKI, leading to a stronger Hungarian voice within W3C and a greater awareness in Hungary of W3C's Activities.

Links:

W3C Hungarian Office: <http://www.w3c.hu/>
W3C: <http://www.w3.org/>
SZTAKI, Department of Distributed Systems: <http://dspd.sztaki.hu/>

Please contact:

László Kovács, SZTAKI
Tel: +36 1 279 6212
E-mail: office@w3c.hu

Éva Megyaszi, SZTAKI
Tel: +36 1 279 6204
E-mail: office@w3c.hu

ERCIM Fellowship Programme

ERCIM offers 18-month fellowships in leading European information technology research centres. Fellowships are available in areas covered by the ERCIM Working Groups and other fields. This offer is available for PhD holders from all over the world.

**Deadline for application:
30 April 2003**

Conditions

Applicants must:

- have a PhD degree (or equivalent) or be in the last year of the thesis work
- be fluent in English
- be discharged or get deferment from military service.

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

Topics

Topics include:

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

Detailed information and the online application form is available at:

<http://www.ercim.org/fellowship/>



Opening ceremony.

Human Resources Managers Task Force

by José Koster and Heather Weaver

The ERCIM Board of Directors has set up a Human Resource Management task force with the aim to share information and experience among ERCIM institutes and to co-ordinate ERCIM internal mobility of researchers and other issues related to human resources.

The ERCIM Board of Directors has emphasised the growing importance of human capital in ICT and Mathematics research. The ICT and Mathematics area as a whole is growing rapidly. The European heads of government have clearly expressed an ambition to make Europe the best performing continent in ICT, and therefore ICT is a main item on the agenda of the next Framework Programme. At the same time, ERCIM partners are experiencing a very tight labour market in ICT and Mathematics, and in view of low student numbers a significant shortage of talent is to be expected in the coming years.

Another development of interest is the increasing co-operation across Europe, thus encouraging opportunities for greater mobility of researchers in Europe. Most ERCIM institutes have seen significant increases in co-operation patterns, both nationally and internationally. Consequently, ERCIM is determined to increase its internal mobility programme. This development towards network organisations puts new demands on Human Resources management.

The ERCIM Human Resource Management task force (HR-TF) was established in September 2001 with the mission to attract and retain the best human capital at ERCIM institutes by:

- encouraging talent and enhancing skills in the area of ICT and Mathematics through promoting mobility opportunities among researchers and supporting staff
- encouraging and establishing best practice through sharing information and experience over a wide range of HRM issues.

In addition, the HR-TF has the mission to:

- develop an Internal Mobility programme
- improve fellowship programmes and recruitment outside Europe
- initiate a European portal for job opportunities.

The HR-TF will analyse the the different mobility programs within the ERCIM organisations. As a result, this should provide new ideas how to improve the mobility of researchers, administrative and supporting staff among the ERCIM member institutes.

As a next step, the HR-TF will produce a web-based template available to all ERCIM members outlining practical advice on such issues as salary, tax, national insurance, travel, accommodation, etc, to provide partners wishing to participate in the various types of mobility schemes with more information. Furthermore, the HR-TF is exploring the funding-possibilities for mobility and management development. The upcoming VIth Framwork Programme and the EU Leonardo da Vinci scheme can be useful for this purpose.

The HR-TF invites all ERCIM member institutes to participate on a regular basis.

Please contact:

José Koster, CWI
HR-TF Secretary
E-mail: Jose.Koster@cw.nl

EURO-LEGAL

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

Metatags that cause Trademark Infringement

The UK has seen a recent landmark case that contains the first judicial analysis of trademark infringement where the use of the mark was made via metatags. Website owners rely on metatags to ensure that search engines bring visitors to their sites. The search engines use spider software to create indexes of catalogued websites by reviewing the sites' metatags for keywords. In some cases these metatags could include a trademark and website owners can use this as a method of redirecting internet users to their site. In order to establish trademark infringement, it is necessary to show that the infringing name or mark has been used 'in the course of trade'. This landmark case rules that despite the fact that these metatags are invisible to the website visitor they are visible when translated into a search result on the user's browser. This case provides further protection to registered trademark owners from other websites using their brands to attract visitors to their site.

Internet Sales — New Rules for EU Consumers

The Brussels Regulation, which came into force on 1 March 2002, introduced new rules for determining where claims arising from civil and commercial disputes involving parties from different countries can be heard within the EU. The rules apply to claims made by consumers arising out of sales of goods or services over the Internet.

Court proceedings must be brought within the Member State where the consumer lives. This right of jurisdiction cannot be waived even if the terms of the contract imply something else. Until now the contract must have been concluded while the consumer was in his or her Member State. The new regulation, however, introduces an additional jurisdiction that deals specifically with cross-boarder sales to consumers over the internet, where the trader "directs such activities to that Member State or to several Member States including that Member State". In other words, if a consumer from the UK purchases goods on the internet whilst on holiday in another Member State (eg France) he or she can bring a claim against the trader in France providing the trader "directed his activities" to the home State of the consumer (in this case, the UK).

Although this has not yet been tested in the courts it is likely to impact more heavily on traders who offer their goods or services in a number of languages and accept payment in different currencies, including the Euro. If traders wish to avoid claims in the future, they may have to consider accepting orders only from consumers residing in their own Member State and in the currency of that Member State. This is likely to impact on the value of trading on the internet. It will be interesting to see how case law develops in this area.

by Heather Weaver, CLRC

Tel: +44 1 235 446151
E-mail: H.Weaver@rl.ac.uk

Andreas Rauber Winner of the 2002 Cor Baayen Award

The annual ERCIM Cor Baayen Award was presented to Andreas Rauber during a ceremony in Nice on 7 November 2002. The award is given every year to a most promising young researcher in computer science and applied mathematics having completed the PhD-thesis in one of the 'ERCIM countries'.

Andreas Rauber is a bright, young researcher at the Department of Software Technology, Vienna University of Technology. Andreas Rauber received his PhD in 2001 from Vienna University of Technology. He started his research career in the area of neural networks where he worked on improvements of the self-organizing map architecture. One of his assets, however, is his wide spread interest which made him progress into and integrate different fields of research. His current focus is on digital libraries where he investigates novel paths for semantically classifying information from heterogeneous sources, as for example text and music. His scientific curiosity paired with his sense for technical rigour is responsible for his high-quality research achievements.

Besides his core research and his more than 60 peer-reviewed publications, Andreas Rauber is also an active member of the international scientific community exemplified by his membership in various program committees at international conferences and on the board of the IEEE Technical Committee on Digital Libraries. See also his article on page 45: 'SOMLib – New Approaches for Information Presentation and Handling'. Andreas Rauber is currently an ERCIM Research Fellow at INRIA.

Link:

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

Please contact:

Lubos Brim, CRCIM
Cor Baayen Award coordinator
E-mail: lubos.brim@ercim.org



Andreas Rauber at the award ceremony.

COR BAAYEN AWARD 2003

The Cor Baayen Award, created in 1995 to honour the first ERCIM President, is open to young researchers having completed their PhD thesis in one of the 'ERCIM countries', currently: Austria, Czech Republic, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Luxembourg, Norway, Slovakia, Sweden, Switzerland, The Netherlands and the UK. The award consists of a cheque for 5000 Euro together with an award certificate. The selected young researcher will be invited to the ERCIM meetings in autumn. A short presentation of the winner together with the list of all candidates nominated for each country will be published in ERCIM News and on the ERCIM website.

Rules for Nomination:

Nominations for each country are made by the corresponding ERCIM Executive Committee member (also referred to as 'national contact'). Those who wish that a particular candidate be nominated should therefore contact the ERCIM Executive Committee member for their country (see the ERCIM Executive Committee page at <http://www.ercim.org/contacts/execom/execom.html>).

Nominees must have carried out their work in one of the "ERCIM countries". On the date of the nomination, nominees must have been awarded their PhD (or equivalent), but this PhD shouldn't be older than 2 years. Each ERCIM institute is allowed to nominate up to two persons from its

country. A person can only be nominated once for the Cor Baayen Award. The selection of the Cor Baayen award is the responsibility of the ERCIM Executive Committee.

How to nominate:

For proposing a nomination to your national contact you should fill out the Cor Baayen Award Nomination Form. Please ensure that all the required information is provided. With your proposal please include a copy of the candidate's PhD thesis (preferably provided as a link to an electronic document) and copies of the candidate's best papers (max. 5 and preferably provided as links to electronic documents).

Deadlines:

30 April 2003:

Submission of nominations to the national contacts.

15 May 2003:

National contacts to submit their two selected nominations to the coordinator

Information:

Further information can be obtained from your national contact or from the Cor Baayen Award coordinator Lubos Brim, CRCIM (lubos.brim@ercim.org).

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

10 Embedded Systems - Introduction

by Erwin Schoitsch

APPLICATIONS

12 An Approach to Dependable Architectures and Components for Railways

by Andrea Bondavalli, Felicita Di Giandomenico and Luca Simoncini

13 Monitoring and Controlling Scientific Experiments: Anywhere - Anytime

by Christian Wattinger, Peter Fornaro, Martin Guggisberg, Tibor Gyalog, and Helmar Burkhart

15 Emerging Frontiers in Miniature Computer and Embedded Systems

by Angelos Bilas and Evangelos Markatos

16 Wearable Systems for Everyday Use

by Spyros Lalis, Anthony Savidis and Constantine Stephanidis

18 New Solutions for In-Vehicle Embedded System Development

by Françoise Simonot-Lion and Yvon Trinquet

20 Remote Control and Monitoring over the Internet - Wireless Construction Site

by Leila Rannanjärvi

NETWORKS AND EU PROJECTS

21 Creating a Network for Those working with Embedded and Real-Time Systems

by Barrie Ross-Dow

22 DECOS - Dependable Embedded Components and Systems

by Dieter Donh offer and Erwin Schoitsch

SYSTEM ARCHITECTURES, METHODS, LANGUAGES AND TOOLS

24 Time Triggered Architecture

by Hermann Kopetz

26 Automatic Distribution of Synchronous Programs

by Alain Girault

28 synERJY: a High-Level Language for Embedded Controllers

by Reinhard Budde, Axel Poigné and Karl-Heinz Sylla

29 STOREQ: Storage Requirement Estimation and Optimisation Tool for Data-Intensive Embedded Applications

by Per Gunnar Kjeldsberg

31 Composing Embedded Real-Time Software Components: the PECOS Data-Centric Approach

by Stéphane Ducasse, Oscar Nierstrasz and Roel Wuyts

32 Time Machines and Black Box Recorders for Embedded Systems Software

by Henrik Thane

33 Design of Embedded Systems: New Challenges

by Gilbert Cabillic and Isabelle Puaut

35 Transforming Specifications to verify Embedded Systems

by María del Mar Gallardo, Jesús Martínez, Pedro Merino and Ernesto Pimentel

36 EUEMES: Component-Based Development Methods for Small-Size Embedded Systems

by Nicolas Guelfi, Aloyse Schoos, Gérard Sébastien and François Terrier

37 Model Checking of Embedded Systems

by Stefania Gnesi, ISTI-CNR

39 Systems Validation Centre shows Benefits of Formal Methods

by Henk Eertink, Wan Fokkink, Holger Hermanns and Izak van Langevelde

EDUCATION

40 REDEST — Sharing Experiences in the Field of Embedded Software Requirement Gathering

by János Nacsa, SZTAKI

42 EmLab — A Laboratory for Teaching Embedded Systems

by Arund Skavhaug, Trygve Lunheim and H. Skinnemoen, NTNU

43 Master Course Embedded Systems

at Eindhoven University of Technology
by Jan Friso Groote, Eindhoven University of Technology, The Netherlands

Embedded Systems - Introduction

by Erwin Schoitsch

Each day, our lives become more dependent on 'embedded systems', digital information technology that is embedded in our environment. This includes not only safety-critical applications such as automotive devices and controls, railways, aircraft, aerospace and medical devices, but also communications, 'mobile worlds' and 'e-worlds', the 'smart' home, clothes, factories etc. All of these have wide-ranging impacts on society, including security, privacy and modes of working and living. More than 98% of processors applied today are in embedded systems, and are no longer visible to the customer as 'computers' in the ordinary sense. New processors and methods of processing, sensors, actuators, communications and infrastructures are 'enablers' for this very pervasive computing. They are in a sense ubiquitous, that is, almost invisible to the user and almost omnipresent. As such, they form the basis for a significant economic push.

These applications are 'vision driven', as in the following examples:

- *Automotive*: Accident free driving
- *Aerospace*: A free, safe sky for all
- *Medical Devices*: Robotic surgeon
- *Communications*: Seamless connectivity
- *e-Life*: ubiquitous/pervasive computing

The European Context

The European Commission has recognised the importance of embedded systems by creating a new unit in the IST Directorate. The visions surrounding the AMI-space (embedded systems everywhere, described in the context of human life as 'ambient intelligence') have considerably influenced the 6th Framework Programme of the IST domain. In this issue we focus on hard real-time, dependability/safety and AMI-scenario applications. There also exists a separate strategic objective on embedded systems in the work program 2003-2004, namely, to develop the next generation of technologies and tools for modelling, design, implementation and operation of hardware/software systems embedded in intelligent devices. An end-

to-end systems vision should allow cost-efficient systems to be built with optimal performance, high confidence, reduced time to market and faster deployment. The focus is on the following:

- Middleware and platforms for building networked embedded systems that aim to hide the complexity of underlying computing, communications, sensing and control, while at the same time providing efficient and effective distribution of resources at low cost. The emphasis will be placed on middleware for small wireless devices, eg mobile phones or PDAs, which make the design, programming, verification and maintenance of systems including such devices easier. Effort will also be devoted to scalable and self-organising platforms that offer services for ad-hoc networking of very small devices and for mastering complexity through perception techniques for object and event recognition and advanced computing and control.
- Concepts, methods and tools for system design and development of warrantable software components and implementation of systems, with an emphasis on the correct handling of complex real-time constraints. Work includes unification of computational models and composition methods, holistic design addressing event and time constraints, interface technologies in hard- and software addressing real-world and legacy issues, and techniques and integrated validation tools to ensure ultra-stable, dependable embedded systems.
- Advanced controls for real-time systems with an emphasis on hybrid systems theories including non-linear processes with both constraints and switching modes. Advanced controls for networked embedded systems, focussing on networked autonomous and fault-adaptive control and management, as well as on reasoning, behaviour, global performance and robustness.

This strategic objective is covered in the second call (15 June to 15 October 2003). The recently founded ERCIM Working Group 'Dependable Software-Intensive, Embedded Systems' is involved in a large IP proposal (Integrated Project, one of the new instruments of European funded research) called DECOS, Dependable Embedded Components and Systems (see article on page 22).

New Research Challenges

A new set of research challenges is presented due to the enormous dimensions of deployment and connectivity, which imply new levels of complexity and risk and new failure modes:

- context-awareness (identify, localise and interact with persons and objects, in a non-location-dependent manner)
- intelligence (the digital environment adapts to mobile objects and persons, and learns and interacts independently, thus providing fascinating new services)
- natural interaction (human language, gestures, speech synthesis)
- personalisation (user-centred, dynamic adaptation to changing situations and user profiles/preferences)
- dependability (time dynamics, timely responsiveness, security, safety, availability, maintainability, robustness etc)
- restricted resources (low power management, size, weight, memory, processing power/speed, interfaces etc)
- wireless/mobile, seamless communication
- hard real-time applications (automotive, aerospace, railways, process control, medical devices, communications etc)
- challenges of composability, COTS, (ultra-)high confidence in design and certification, security and predictability.

Important sub-areas are (rigorous) hardware/software co-design, smart new sensors/actuators, continuous connectivity issues and limited resource management. Horizontal issues include dependability, system integration, software technologies, standardisation, certification, and critical infrastructures.

ERCIM Working Group Dependable Software-Intensive Embedded Systems

The dependability aspect of embedded systems is of the utmost importance, and ERCIM members have identified a huge potential with respect to this. Accordingly, the new ERCIM member AARIT, within which ARCS is an important driving force, has proposed and initiated an ERCIM Working Group on 'Dependable Software-Intensive Embedded Systems'. The inaugural meeting took place in Vienna during the ERCIM board meetings, and a joint workshop with the EU projects on embedded systems (real-time) such as AMSD and ARTIST was organised in Grenoble on 3 October 2002. A second meeting was held in Sophia Antipolis on 6 November. Future plans include a workshop on 'Synchronous Languages, Time- and Event-Triggered Systems' around March/April this year, and a meeting in Trondheim on 2 June 2003.

The Working Group currently includes members from Austria (ARCS/AARIT), France (INRIA), Germany (FhG), Greece (FORTH), Italy (CNR, PDCC), Luxembourg (IST), Netherlands (CWI), Norway (NTNU) and Sweden (MRTC/SICS). Others are very welcome.

At the moment, all themes relating to embedded systems are of interest to the Working Group (WG), which focuses on specific topics case by case. The work program includes:

- experience exchange
- joint position papers on relevant issues
- information exchange with organisations, other WGs and committees in which ERCIM members are already active
- decision building processes
- WG meetings adjacent to other relevant meetings/conferences
- discussion board (via ERCIM Website or a member's Website)
- contributions to ERCIM News
- one major event per year (may be combined with some important meeting or conference)
- cooperation within the EU Framework Programs (NoE, IP); currently Dependability (1st Call) and Embedded Systems (2nd Call)
- cooperation in standardisation and awareness-building processes.

About this issue

This issue of ERCIM News covers all aspects of embedded systems, technical as well as broadly socio-economical, and gives an overview of the scientific and engineering challenges, and fascinating possibilities and risks offered by this enabling technology.

The contributions are grouped into:

- Applications (6): Railways, Remote Control of Scientific Experiments, a more general article on 'Emerging Frontiers in Embedded Systems', Wearable Systems for Everyday Use, Embedded Systems in Cars, and Application in Transport Logistics
- Networking and EU-Projects (2)
- System Architectures, Methods, Languages and Tools (11)
- Education (3).

These contributions clearly demonstrate that the scientific community is attracted primarily by the interesting system, design and implementation challenges of embedded systems. It is also becoming evident that our educational systems require new curricula, labs and courses, as well as networks for experiential exchange, in order to keep up with such innovation. The section on applications shows that the driving force behind this comes from critical control applications (automotives, railways etc) and the 'everyday ubiquitous computing (AMI)' community.

Please contact:

Erwin Schoitsch
ARC Seibersdorf research (AARIT)
Tel: +43 50550 3117
E-mail: erwin.schoitsch@arcs.ac.at

An Approach to Dependable Architectures and Components for Railways

by Andrea Bondavalli, Felicita Di Giandomenico and Luca Simoncini

Scientists at ISTI-CNR study methodologies and techniques to improve the definition and development of systems for railway applications. The envisaged solutions must guarantee the dependability and real-time requirements imposed by the application field while at the same time addressing the needs both of the market and of the technology.

Automatic systems for railway applications have traditionally been developed starting from proprietary architectures, in which the processors and memories were assembled using ad hoc components whereas the electronic components have generally been developed independently. A feature of this approach is that the components are designed in function of the overall system structure; this type of approach has positive and negative aspects.

Advantages are that:

- the design and implementation of ad hoc components facilitates validation of the system
- the entire system is controlled by the designer and there are no single parts protected by third party intellectual property rights; again this facilitates validation and procuring, mandatory actions for safety critical systems
- system redesign and updating is not dependent on third parties.

On the other hand:

- components and implementation technologies change and evolve very quickly; this means that some may be obsolete by the time that the design stage is completed and the system is ready for operation
- the operational life of the system tends to be extended which means that component upgrading may be required
- the strict dependence between system and components (through the design) means that it may be difficult to adapt the system to different contexts or to interface it with other systems, ie, re-configuration may be very difficult if not impossible
- systems with even slightly different requirements and specifications cannot reuse existing components; this means that new systems generally require a complete redesign, and experience gained when operating previous systems cannot be exploited
- any new system or major revision needs to be revalidated ex-novo.

Another major issue is the need to reduce both commissioning time and development and operational costs. This makes the use of COTS (components off-the-shelf) an increasingly popular choice.

It is clear from the above scenario that a strategic R&D activity aimed at the definition, prototyping, partial verification and validation of a generic, safety-critical and real-time architecture for railway systems is needed.

Such an activity should be able:

- to reduce design and development costs
- to reduce the number of components used by the subsystems
- to simplify the product evolution process and reduce associated costs
- to simplify product validation (and certification) through an incremental approach based on reuse
- to make the design of safety-critical parts more flexible.

The resulting architecture should have the following characteristics:

- use generic components (possibly COTS) which can be replaced when necessary without the need for system redesign or revalidation
- reliability/availability and safety properties must be associated with the overall architecture rather than only with intrinsic properties of its components, so that techniques for error detection, diagnosis and error recovery are as far as possible independent from specific hardware or software components
- openness of the system, so that it can interface and communicate with other systems through different communication systems (eg GSM-R, radio, ISDN, etc.)



A strategic R&D activity is needed to design a generic, safety-critical and real-time architecture for railway systems.

- adoption of a hierarchical approach for functional and non functional properties, to facilitate validation
- strict conformance with railway standards; in particular with CENELEC EN 50126, 50128, 50129, 50159-1, 50159-2.

For this reason, we are currently investigating architectures/component solutions able to satisfy the dependability and real-time requirements imposed by the high criticality of the application field while also addressing the market and technology demands.

In recent years, the Esprit Project GUARDS 'Generic Upgradable Architecture for Real-time Dependable Systems', which studied the development of methods, techniques, and tools to support the design, implementation and validation of safety-critical real-time

systems, has already investigated a number of these issues. This project involved major industries from the space, nuclear and railways sectors. A careful specification of the needs of industrial end-users was used by the academic partners to identify and define suitable safety-critical mechanisms and architectures; these were then implemented with the support of the expertise and tools of the technology providers.

Within the GUARDS project, a series of R&D activities were performed, spanning from the definition of components/mechanisms for fault tolerance to the set-up of a multi-techniques validation framework and the development of architectural instances appropriate for different application fields. However, more extensive research is needed today, to properly account for current technological competitiveness and other

emerging challenges, such as integration and interoperability. The world market of railway systems has a strong growth trend, both at national and international levels, because of the need to optimise urban and suburban transport, to modernize the existing railway systems, to develop high-speed trains and to render European transport networks interoperable. In the next ten years, the European railway market plans to construct high speed lines for 12,500 km. In addition, 16,500 km of existing lines will have to be upgraded. The expected investment is in the range of 24 billion Euro.

Please contact:

Andrea Bondavalli, ISTI-CNR

Tel: +39 050 315 3068

E-mail: a.bondavalli@cnuce.cnr.it

Monitoring and Controlling Scientific Experiments: Anywhere - Anytime

by Christian Wattinger, Peter Fornaro, Martin Guggisberg, Tibor Gyalog, and Helmar Burkhart

Monitoring and controlling a scientific experiment from your mobile device anywhere and anytime — this is the goal of the MURCI (Mobile Universal Remote Control Interface) project at the Department of Computer Science of the University of Basel.

Mobile devices, such as mobile phones and PDAs, include always-on wireless networking, persistent memory, computational capabilities and enhanced user interfaces, and allow the uploading of small programs that exploit these abilities. This opens the concept of remote, real-time control of experiments.

As mobile devices gain more and better technical features, their role as location-independent networking tools becomes more significant. Only a few months ago the sole use of a mobile phone was voice communication, with cumbersome procedures and slow and expensive connections making data transmission unfeasible. This situation has changed significantly with the introduction of GPRS (General Packet Radio System)

networks. In the meantime, PDAs have also matured, now offering Bluetooth connectivity and the ability to be readily programmed, eg using Java. Despite these positive advancements, today's mobile devices still suffer from slow connection speed, small screen sizes and limited memory and overall performance. Both these new capabilities and restrictions make it advisable to use a 'thin client'-based model for developing mobile applications. A 'thin client' is a lightweight program running on a mobile device, which makes use of networking capabilities as well as the display and user interface. All resource-intensive program parts are run on a server computer which is accessed by the mobile device.

A project of the newly founded Mobile Technologies Group at the Department of Computer Science of the University of Basel is deploying this paradigm in the realisation of a system framework for remote experimentation. The notion of an experiment in this context has a broader meaning, which can include physical experiments, simulations, and long-term calculations in high-performance computer systems.

Our first demonstrator site, 'wireless Nano-World', implements the monitoring and control of an Atomic Force Microscope (AFM) simulator. An AFM is an instrument used in the nanosciences to explore surfaces at the nanometre scale, such that single atoms and molecules become observable and can



Figure 1: Scanned image of an atomic surface (with a change of scan range) and the parameter screen of the thin client.

even be manipulated. The simulator has been developed as part of the Virtual Campus Switzerland project ‘Nano-World’ and is used in teaching, learning, and training exercises. This is necessary because real AFMs are complex and expensive instruments, and very few of them are readily accessible to students. The use of simulators for training is popular with students, and being hands-on, is an effective and motivational learning tool.

The use of the mobile thin-client model to monitor and control a real experiment or simulation offers many advantages to the experimentalist. One advantage of having a tool to remotely observe and steer an ongoing experiment in a laboratory is the ability to control long-time measurement from the office, home or even while travelling, eg on a train. Another big advantage is the ability to easily notify collaborators or an administrator. This notification mechanism can be triggered actively or automatically by the remote-controlled system itself, eg in the case of a malfunction or after a self-test. This message can alert a human supervisor or trigger another process on a receiving machine.

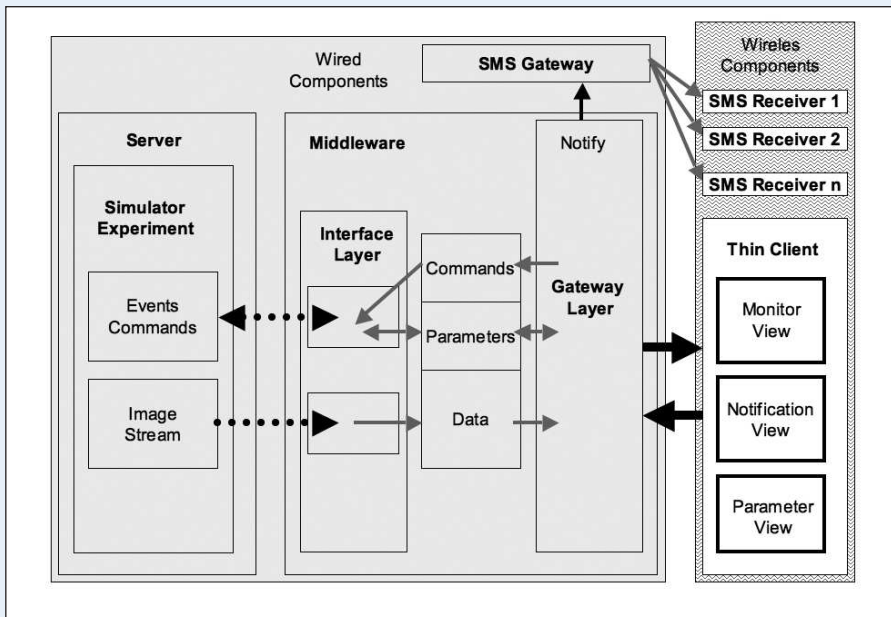


Figure 2: Architecture of the thin client model.

Technically, the model used is a layered architecture. The experiment or simulation defines the input/output ports to which a middleware layer is linked. The middleware processes the data and translates it so as to be usable by the respective end-points. The mobile thin client presents the arriving data and accepts control commands for the experiment.

The thin client itself has three basic components:

1. The monitoring module presents the incoming visualized measurement data. On this screen the user can monitor the ongoing experiment at a glance. Later options for visualization and image capture will also be included (Figure 1 on left).
2. The command module lets the user monitor the current experimental parameters and also allows their values to be changed. Additionally it is possible to send ‘start’, ‘stop’ and ‘pause’ commands to the experiment (Figure 1 on right).




Figure 3: The remote controllable Atomic Force Microscope (a) with rotating sample table (b) and the positioning axes.

3. The third component is the communication module. At present, this is included in the command module, but in future will be decoupled and built as a separate component as befits its importance. It allows the user to inform other mobile users about the experiment. Standard messages for cases of alert, special interest, or help requests are available to be sent to a specified list of receivers. Messages are currently sent as SMSs (Short Message Service) because of this medium's reliability. In the future other protocols like instant messaging (IM, eg Jabber) or Multimedia Message Service (MMS) may be better solutions. Voice and video channels would also be useful, as soon as the transmission speed and cost issues are resolved.

The thin client connects to the real-world experiment or simulation via a middleware layer, which mediates between the experiment and thin client as seen in Figure 2. The middleware in the second

layer must present measurement data to the wireless client in a bandwidth- and transmission-cost-saving and standardised manner, while accepting continuous data streams from an experiment. Commands from the client on the other hand must be converted by the middleware to make them understandable by the experiment/simulator. This middleware layer is advantageously implemented as a Java servlet to easily introduce scalability for many simultaneous clients, as well as optional authentication elements. At the moment the middleware is still Common-Gateway-Interface-(CGI)-based and leaves much to be desired in terms of performance and reliability.

As feedback from the research community has been very positive, we want to rewrite and extend the prototype towards an operational mobile client to be used by researchers of the nanoscience competence centre. Thus the next goal is a real remote-controlled Atomic Force

Microscope with a controllable sample table (Figure 3). This microscope allows standard sample positions to be defined, which can then be remotely reselected with micrometre precision. In this way, the instrument can be shared according to a schedule or notification by a number of researchers, each of whom can independently access his or her samples.

With regular self-tests and corresponding alerts to inform users in case of problems, it will be possible to run such research instruments in an administrated mode, where human interaction is only necessary for the setup of the probe and sample and in case of a failure notification. Valuable measurement time is therefore efficiently shared between research groups.

Link:

<http://www.ifi.unibas.ch/research>

Please contact:

Christian Wattinger

University of Basel, Switzerland

E-mail: christian.wattinger@unibas.ch

Emerging Frontiers in Miniature Computer and Embedded Systems

by Angelos Bilas and Evangelos Markatos

Although over the last several years computer systems research has focussed on building general-purpose PCs, current emerging technologies and societal needs are shifting the focus towards ambient intelligent systems that monitor, adapt, and react to humans and objects.

Over the last twenty years, computer systems and architecture research has addressed the performance of general-purpose microprocessors and desktop computers, with impressive results. Processor performance has been doubling every eighteen months, and today's general-purpose desktop systems are used in most applications to address single-user needs or as components of much larger computing and storage platforms. However, there is now an evident need to turn the focus of systems research from the above targets to new directions, moving from general-purpose to application-specific systems and from performance to reliability,

availability, security, autonomy, scalability, and manageability.

Current research at ICS-FORTH, which is partly funded by the 'Excellence' project initiative of the General Secretariat for Research and Technology of the Greek Ministry of Development, builds on existing expertise and addresses fundamental problems. These include:

- miniaturisation of computer and embedded systems
- interconnection of such systems
- autonomy and coordination of very large numbers of such systems.

It is expected that such autonomous, networked, miniature computer and embedded systems will find application in many new areas, eg diagnosis of epidemic diseases, forest and crop protection, human safety and monitoring of fine-grain parameters in the human environment. Overall, miniature computer and embedded systems are predicted to be the most effective carrier of ambient intelligence. Such systems can both perform intelligent functions and be placed or integrated in locations where traditional systems cannot fit, due to form factor, autonomy, and cost restrictions. Moreover, they will have the ability to communicate with each other, providing a means of transferring,

monitoring and controlling information everywhere in a ubiquitous and transparent manner. These directions in computer systems research at ICS-FORTH are discussed in more detail below.

Going small and going in numbers

In this regard, our research is searching for ways to build very small (say, coin-size) and flexible smart devices and embedded systems. Looking at today's state-of-the-art computing systems, two of the main factors restricting miniaturisation are the requirements on versatile I/O interfaces and rich runtime systems. The first is a result of various standard interfaces that need to be supported, whereas the second is a consequence of requirements for general-purpose functionality. Our goal is to re-think both of these issues. Firstly, we are interested in designing interfaces that can be used to connect components within very small systems. Of principal interest here is connectivity to non-logic-based devices (MEMS). Second, we would like to explore architectural support for novel runtime systems. Today's runtime systems provide flexibility that is not needed in application-specific devices, leading to high requirements in compute cycles. Our goal is to investigate how such devices should be designed to best support novel runtime systems that exhibit the necessary characteristics.

Connectivity is the only way in and the only way out

This focus area examines how such systems should be interconnected among themselves and also to the outside world. Building interconnects for large numbers of miniature devices introduces constraints that are not present when building traditional interconnects. The physical medium used is constrained not only by size and speed but also by power consumption, and may well limit the distance at which such devices can transmit and receive data. Moreover, given the requirements of future miniature embedded systems, how their protocol stack should look is unclear. Traditional protocol stacks are inadequate due to the size, complexity, and the requirements they need to fulfil. In addition, security may be jeopardised since, on the one hand, the network is the only means of gaining access to such systems, while on the other hand, a full access and management API exposed over the network makes these systems vulnerable to attacks.

All for one and one for all

Miniature systems are unique compared to existing computers because they are able to contribute to new application domains simply by being replicated in large numbers and by collaborating on solving a specific problem. However, this aspect of collaboration, which

assumes a high level of autonomy, is new territory that needs to be explored. Previous work in distributed systems has addressed issues in systems that are more closely supervised, more independent in performing various tasks, and rather loosely coupled. The runtime system is a key component that implements most of the functionality relating to achieving collaboration and providing autonomy, but also that relating to tolerating and recovering from faults, providing reconfigurability, security, power management etc. Given that these systems will be less capable than modern computers, to a large extent the runtime system needs to be redesigned.

In summary, computer systems research is shifting its focus to new and exciting directions. We believe that among the most promising new aspects of computer systems research is the miniaturisation of networked computer and embedded systems and the implications of using large numbers of such devices. Large areas of uncharted territory remain in these fields for computer architecture and systems research.

Link:

<http://archvlsi.ics.forth.gr>

Please contact:

Angelos Bilas and Evangelos Markatos,
ICS-FORTH
E-mail: {bilas, markatos}@ics.forth.gr

Wearable Systems for Everyday Use

by Spyros Lalis, Anthony Savidis and Constantine Stephanidis

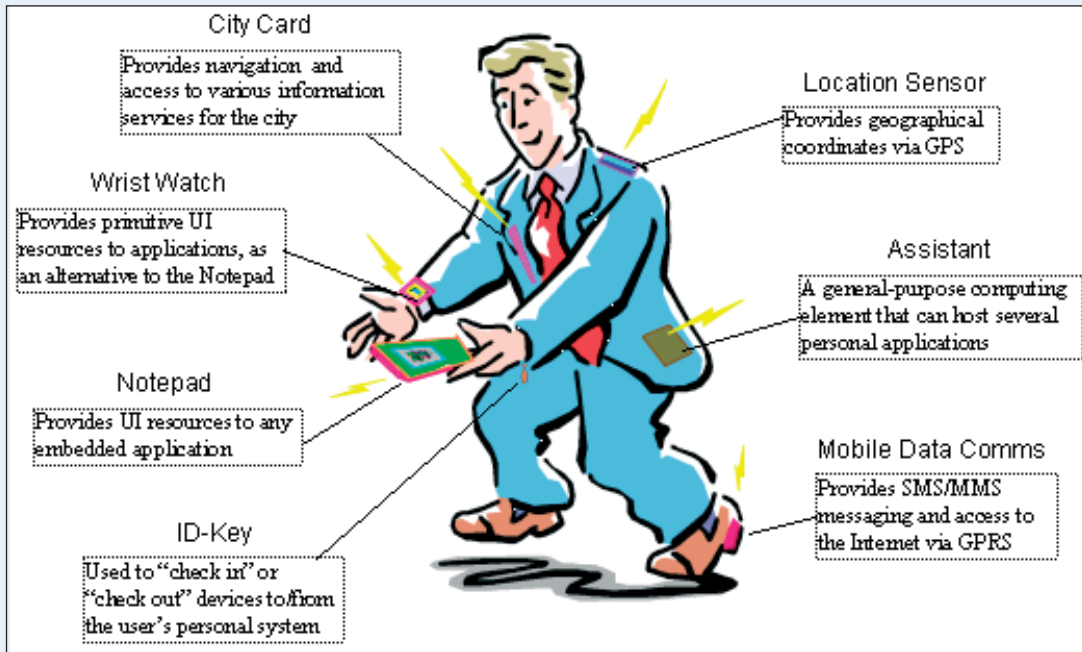
You have just landed in Heraklion. As you leave the airport you are handed a CityPass, a small gadget hosting a map application for the city, which you touch with your ID-Key and put away in your pocket. A few seconds later your wristwatch beeps with a message: "Welcome to Heraklion - the MapGuide can be controlled using your PDA or your wristwatch".

As you stroll downtown, you shoot several pictures, which are tagged with time and location information retrieved from the GPS on your wristwatch. Storage in the camera begins to run out, and the camera informs you that photos are now being stored in the Assistant in your backpack. At the same time, the Assistant silently uploads the photos to

your home repository whenever you pass near a public network access point. You meet a friend at a café and together you browse through the photo collection. You decide to have dinner at an old tavern displayed in one of the pictures. You pull the Notepad out of the backpack and use the CityMap application to display the location of the tavern, using

location information from the photo coordinates and your current position provided by the GPS on your wristwatch.

This brief scenario gives a flavour of the 2WEAR project and the vision shared by FORTH, ETHZ (Switzerland), MA Systems & Control (UK), and



Wearable systems.

NOKIA/NRC (Finland) of the future personal computer put together as an agglomerate of small, wearable and physically distributed devices. 2WEAR is a three-year EC-funded project that started on January 1, 2001 with the objective to build such a system, both in terms of software architecture and devices that communicate using short-range radio technology. The main characteristic of the envisioned system is its dynamic extensibility and adaptation, which would allow users to combine devices in a flexible manner according to their current needs, and to exploit the potential of surrounding infrastructure within buildings, vehicles, etc. The 2WEAR system is exploring the following directions in parallel.

Core Runtime Mechanisms

In this new paradigm, the personal computer, rather than being fixed in a box with a well-known amount of hardware resources and a given number of peripherals attached to it, is actually a system that is dynamically composed by bringing together different devices. To convert this physical versatility into tangible flexibility for the application and user, several problems must be addressed at the system level. Besides being a provider and manager of local resources, the runtime now becomes responsible for detecting and exploiting resources that become available via other devices that are added to the system. It must also take corrective

actions should such resources become unavailable, due to failures or the user removing devices from the system.

The main system aspects that are being researched to achieve the desired functionality are ad-hoc discovery, flexible remote communication in the form of dialogue channels, and adaptive system services. Adaptive system services correspond to high-level resource abstractions provided to the application programmer, while dealing with low-level runtime issues such as resource discovery, remote resource access, resource switching and compensation in the case of resource loss.

Human-Computer Interaction

In dynamic, transient environments, the main architectural objective 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. To that end, the User Interaction Framework specifies: (a) how applications communicate with the user, and (b) the mechanisms allowing for the abstraction of Input Output (I/O) devices as well as for concurrently running applications to share I/O hardware resources. In practice, this means that framework services specific to the human-computer interaction dynamically allocate 'pooled'

system resources to user activities according to user preferences.

These mechanisms operate at a level between applications and the runtime, defining an infrastructure of services that take over all decisions relating to the administration of I/O resources and their allocation to active applications. As a result, application developers are able to produce user interfaces that remotely utilise multiple I/O resources (eg audio devices, different types of displays, input devices) over a wireless communication medium. In cases where loss of connection with some I/O resources or discovery of new I/O resources may occur as a result of mobility, running interfaces dynamically employ such new I/O resources and/or possibly reallocate existing ones to ensure interaction continuity. In addition, applications continuously inform the user of changes in the interactive space, such as the detection of newly available computing units and/or the loss of connection with devices. In this way, the user can engage or disengage computing units on a task-oriented basis.

Links:

Project Web Site:
<http://2wear.ics.forth.gr>

Please contact:

Spyros Lalis, ICS-FORTH
Tel: +30 2810 391746
E-mail: lalis@ics.forth.gr

New Solutions for In-Vehicle Embedded System Development

by Françoise Simonot-Lion and Yvon Trinquet

While automobile production is probably to increase slowly in the coming years, the part of embedded electronics and more precisely embedded software is growing. New solutions for the development of in-vehicle embedded systems was the purpose of the French ‘Embedded Electronic Architecture’ cooperative research and development program ‘AEE’. The results are the identification of embedded component classes, the specification of a generic embedded architecture, as well as the definition of a declarative language.

Today, functions embedded in vehicles include braking assistance, active suspension, steering functionalities, etc. They are subject to stringent timing constraints and more generally to dependability constraints. In the close future, these constraints will be more critical with the generalisation of X-by-Wire technology. Therefore the development of such systems must define an eligible system, ie, satisfying these constraints, and must provide the best one according to cost criteria. Furthermore, the development process of an embedded system is shared between several actors: carmakers and suppliers; the development of reusable components is a main way for the reduction of costs.

In this context, the French ‘AEE Project (EEA in English: ‘Embedded Electronic Architecture’) specified new solutions for the development of in-vehicle embedded system. This program (September 1999 - December 2001) was

granted by the French Under-Ministry for Industry and involved the French carmakers (PSA and Renault), OEM suppliers (SAGEM, SIEMENS, VALEO), the company EADS LV and the research centres INRIA, IRCCyN and LORIA. The European ITEA project (EAST-EEA) – currently in progress - extends and generalises these results.

Electronic Embedded Architecture and its Components

Most of the hard and software embedded in a car are specified and developed separately. Each one is dedicated to a particular feature and designed by a supplier with respect to carmaker requirements. On the one hand, this is a bar to the reusability of solutions in other projects and on the other hand, this leads to oversize the resources (hardware, buffers, etc). To solve this problem, the AEE project characterised formally the basic embedded components and defined the perimeter of the reusable ones. Furthermore it provided a generic archi-

ture for an Electronic Control Unit (ECU), a station connected to one or several network(s) and supporting the embedded application (see Figure 1).

Some components are independent of a particular ECU; this means that they can be implemented on any ECU in a distributed architecture:

- sensors and actuators (hardware components) and software components (Local Device Manager) realising the signal processing for these devices
- software components implementing the specific embedded application (Application Software Components).

On the contrary, the Input/Output drivers, the Software Components implementing the Operating System (OS) or the Communication Services are dependent of a specific ECU.

Finally, in order to ensure an entire independence of Application Software Components, a specific component, named Inter Component Exchange Manager was specified. It plays the role of a middleware, in particular by providing transparent communication services. This component is developed specifically for each ECU with a common Application Interface.

AIL_Transport: a Language for the Design of Embedded System

The project defined a specific development method for embedded systems to reduce costs and optimise the use of hardware elements. At the first step, functionalities are defined and validated independently of their implementation (Functional Architecture and Software

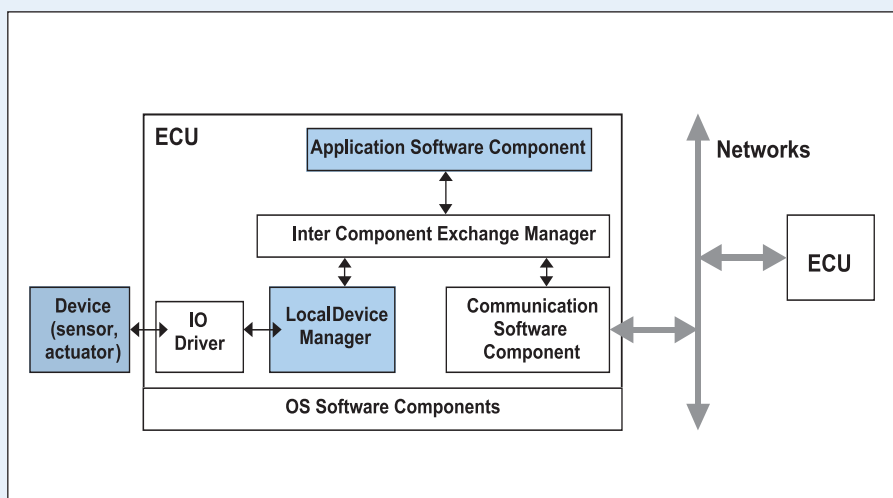
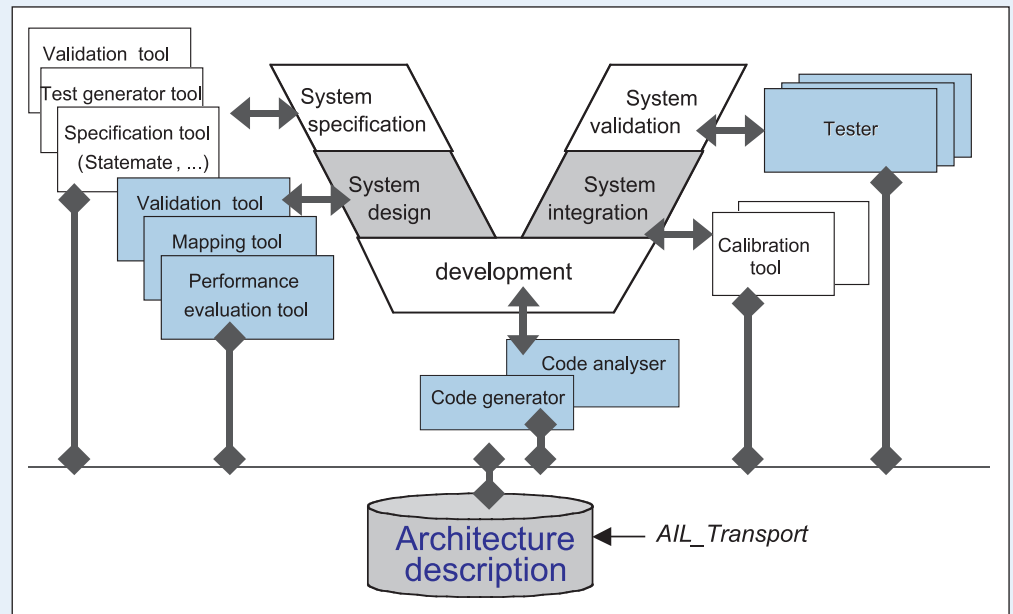


Figure 1: component classes and generic architecture.

Figure 2:
AIL_Transport principles.



Architecture). Then an allocation mechanism maps the specified functions onto the ECUs of the embedded architecture and subsequently, some exchange flows onto communication networks (Hardware Architecture). Finally, local task execution and frame transmissions are optimised (Operational Architecture). With this approach, capitalisation no longer focuses on ECUs, but on the implemented functions through validated hardware and software modules.

Strong co-operation between OEM suppliers and carmakers in the design process implies the development of a specific concurrent engineering approach. In order to specify this process, synchronisation points (rendezvous) across the co-operative development model have to be identified and information exchanged at these points must be characterised. Furthermore, a unique syntax of the exchanged information has to be defined. For this, the AEE project has specified a business model specifically adapted to the architecture development conjointly by carmaker(s) and OEM supplier(s).

Performance of a vehicle embedded system may be evaluated from different points of view, according to the system analysis (wholly or partly) necessary at each development step. Usually carmakers attempt to optimise the number of ECUs, which are used to

implement the vehicle functionalities. Furthermore, system designers attempt to optimise performances of communication networks. Finally, OEM suppliers have to demonstrate that their COTS are compliant with the carmaker requirements, etc. The AEE approach improves these different analyses and optimisations by enabling the use of various industrial and academic software tools. These tools are dedicated to analyse, test, simulate, validate, comment, and generate code of the electronic architecture. For that, each tool extract a specific and coherent model from the architecture description by means of a repository integrating all the pertinent data of the architecture modelling. This repository is the skeleton of the AEE development process, as shown in Figure 2. In order to build this repository, a language for the specification of any electronic architecture has been defined. It is called Architecture Implementation Language (AIL_Transport). The AIL_Transport language integrates the AEE design process, and thus is used by all designers as the backbone of the architecture development. Moreover, AIL_Transport is the source language to define the reusable architecture objects.

Associated to AIL_Transport, a development process has been specified for defining and harmonizing the exchanges of partial architectures between carmakers and OEM suppliers. The main benefit of this study is to allow the design of flexible architectures while

reducing costs and increasing the quality of the development. At present, the results obtained in AEE project are one of the entry points for an ITEA European project (EAST – EEA), gathering the main actors of European automotive industry.

Link:

AEE: <http://aee.inria.fr>.

Please contact:

Françoise Simonot-Lion, LORIA
E-mail: Francoise.Simonot-Lion@loria.fr

Yves Sorel, INRIA
E-mail: Yves.Sorel@inria.fr

Jean-Pierre Elloy, Yvon Trinquet
Institut de Recherche en Communications
et Cybernétique de Nantes (IRCCyN), France
E-mail:
Jean-Pierre.Elloy@ircsyn.ec-nantes.fr,
Yvon.Trinquet@ircsyn.ec-nantes.fr

Remote Control and Monitoring over the Internet – Wireless Construction Site

by Leila Rannanjärvi

At the ‘wireless construction site’ trucks and bucket loaders complete transportation tasks cooperatively. The truck drivers send acknowledgements to the work-site server when a transportation task has been completed. The loaders send messages to the server after each loading event. This system enables the real-time management of trucks, bucket loaders, loads and even material administration. The driving force in this case is in the expected saving, since no material is lost and there is no need for extra kilometres or work hours to be charged.

This case is not a typical Internet application, because it focuses on enhanced production and not on entertainment or context sharing. The communication is more peer-to-peer and the system is a typical nomadic system. Nomadic systems are based on a core of fixed routers, switches and hosts; at the periphery of this fixed network, base stations with wireless communication capabilities control message traffic to and from dynamic configurations of

mobile hosts. Nomadic distributed systems sit in between traditional fixed networks and ad-hoc networks.

Wireless Construction Site

Together with communication technology, modern software development enables the integration of previously totally isolated work machines into the expanding global Internet and mobile communication community.

With this vision in mind, we started a project called ‘Wireless Construction Site’ with Finnish industrial partners. On this first-tier solution the main added value will be found in:

- accumulated information of transportation tasks containing information about the truck, the driver, the load and the driving distance
- accumulated information on loading tasks containing information about the bucket loader, the loader (the operator of the bucket loader) and the load (material and volume)
- quality 3D information about the tasks carried out by the road construction machinery.



Figure 1: Vision of the wireless construction site.

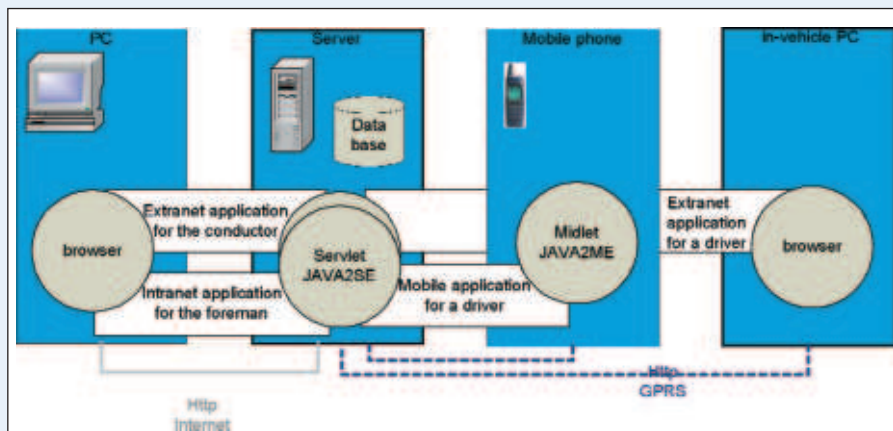


Figure 2: Implementation architecture of the wireless construction site.

With this information the administration of materials, work machines (trucks, bucket loaders, road scrapers) and drivers becomes easier and more punctual! The vision includes various kinds of work machines, localisation by GPS or by GSM cell location and different communication modes depending on the available network and the current work phase (Figure 1). If no operator network is available in the gravel pit, either the machines must save pending messages or a local license-free network has to be established. Messages will be sent to the server using GPRS network and http messages. We want to test the usability of mobile phones at the work site, because the investment is lower than in the case of an in-vehicle PC. In most cases the driver already has a mobile phone, so adaptation to a newer mobile phone is easier than to an in-vehicle PC.

The wireless construction site application has three types of user roles: the

driver, the contractor and the site foreman.

The driver has the mobile application in his mobile phone (or in the in-vehicle PC). He receives the transportation task from the site foreman and starts driving. When his truck enters the depot, he orders the load and the bucket loader commences loading. When the load is on truck, the loader sends the load message to the server. The message contains information about the truck, the load volume, the material loaded, the position of the bucket loader and the loading time. When the driver unloads, the unload message will be sent to the server. It contains information about the load, the unloading position and time. The localisation of truck must be transparent. Both GPS and GSM cell location methods are possible.

The site foreman allocates tasks to drivers, trucks and bucket loaders with an Intranet application running in a browser on a PC. The user identification and password will of course be checked beforehand. The server is able to combine the information for materials

administration from the load and unload messages.

For the contractor an Extranet interface provides the information about trucks and work machines (hours, kilometres and tonnes), drivers and materials. Access is tightly restricted to contractor's own tasks.

There is also the server software, which contains the core wireless construction site database and also the service for database access (Web server, html form pages, WAP server and servlets). Between the Extranet user and the server is of course a firewall.

Figure 2 describes the implementation architecture. The application for the driver is a midlet developed by Java2ME™ Wireless toolkit. The midlet is downloaded from the WAP server and saved into the mobile phone. The mobile phone must support Java, ie contain a Java K virtual machine and CLDC (Connected Limited Device Configuration) allowing external midlets. We are currently using Nokia 6310i phone. The carrier is GPRS and the protocol is http.

The protocol between the server and Extranet or Intranet applications is also http. The Intranet application implies direct hard-wired connection to the server as the Extranet application may be used from an in-vehicle PC. In both cases access is limited. The server allows only http protocols and the contractor or the site foreman has a set of forms that make up the database interface. The database communicates with the users only through these http forms.

The wireless construction site vision has been described and the first experiments with mobile phone and server completed. When the driver uses a midlet in a GPRS phone, he can send messages to the servlet in the server. The application development is ongoing and experimental use with several trucks and bucket loaders will take place in north Finland during spring 2003.

Please contact:

Leila Rannanjärvi, VTT
Tel: +358 8 551 2314
E-mail: Leila.Rannanjarvi@vtt.fi

Creating a Network for Those working with Embedded and Real-Time Systems

by Barrie Ross-Dow

Recognising the growing importance of Embedded Real-Time Systems the Institute of Electrical Engineers (IEE) created a Professional Network to cover this field. This new network offers support to the development of industry by promoting co-ordinated research, knowledge sharing and educational support.

Embedded Real-Time Systems are important components in a wide range of application areas including the Process Industry, Manufacturing Automation, Transportation, Telecoms, Medical, Power Generation and Distribution, and Defence. They act as an enabling technology allowing systems to sense their environments and directly influence them through controlled and timely actions.

The combination of temporal requirements, limited resources, concurrent environmental entities and high dependability requirements, (together with distributed processing) presents the system engineer with unique problems. Embedded Real-Time systems are now recognised as a distinct discipline. It has its own body of knowledge and theoretical foundation. The next generation of Embedded Real-Time Systems are likely to be subject to even greater demands than those that are currently placed on

them. However, effective support for such technology is currently lacking: there needs to be a co-ordinated research, technology transfer and educational programme with the long term goal of providing support for the engineering of intelligent, flexible and dependable Embedded Real-Time Systems.

By creating a Professional Network (PN) in the field of Embedded and Real-Time Systems, the IEE has taken a major step

towards establishing and co-ordinating such a programme. The new Embedded and Real-Time Systems PN organises seminars, lectures, residential courses, contributes to IEE publications and provides a interactive community website with access to relevant knowledge and information. The website also facilitates communication between participants. Registration to the PN is currently free and there is no requirement to be a member of the IEE.

There are certain distinct advantages in having the backing of the IEE behind this initiative. As Europe's largest engineering body, with over 130,000 members, the institution has the resources to support the project with key

contacts in industry, government and academia. Guidance for the running of operations is steered by a panel of professionals who make up the PN's Executive Team. The current team for the Embedded and Real-Time Systems PN is chaired by Dr Iain Bate a member of the Real-Time Research Group at the University of York.

Content for the PN website is supported by the Publications division of the IEE, an established provider of specialist publications and the INSPEC database. They provide detailed editorial work as well as supply content for a regular industry related newsletter, sent to all registrants. In August 2002, the IEE devoted an entire issue of its Computing

& Control Engineering Journal to papers given at an event organised by the Professional Network. This helped highlight many of the important issues which concern the industry.

Future events are planned on: Developing Real-Time Embedded Systems, Real-Time Software in Formula 1; Real-Time UML and Military & Aerospace applications of FPGAs.

Link:

<http://www.iee.org/pn/embeddedrealttime>

Please Contact:

Barrie Ross-Dow
Institute of Electrical Engineers, UK
E-mail: brossdow@iee.org.uk

DECOS - Dependable Embedded Components and Systems

by Dieter Donhoffer and Erwin Schoitsch

A proposal for the 'Integrated Project' DECOS for the 6th Framework Programme of the EU is being prepared by ARC Seibersdorf Research. Research on dependable embedded systems will result in certified generic commercial-off-the-shelf (COTS) components for safety-critical systems.

Embedded systems already play an important role not only in consumer electronics but also in many important and safety-critical systems in applications such as avionics, space, railway and transport, process control and medical systems. There are, for instance, already many embedded systems in cars with critical control functions (eg ABS braking systems, airbags), and these will become much more widely used in the automotive industry once they can be delivered at prices acceptable to the automotive market.

A study by Allied Business Intelligence Inc, 'X-by-Wire: A Strategic Analysis of Time-Triggered Architectures and Next-Generation Automotive Control Systems' (http://www.the-infoshop.com/study/ab8859_x_by_wire.html), expects the global market for automotive x-by-wire systems to grow from approximately \$600 million in 2001 to over \$27 billion in 2010.

The DECOS project, which is being prepared in cooperation with Prof. Hermann Kopetz at the Vienna University of Technology, is intended to develop dependable components and systems that will be applied in various control systems in the automotive and aerospace/avionics industries, as well as in railways and transport, industrial automation and medical instruments and systems. These 'embedded' computer systems, which realise various distributed, safety-relevant functions, will have enormous economic and societal relevance in the future.

The mission of DECOS is to develop an architecture-based design methodology and the associated COTS hardware and software components together with certified development tools and advanced hybrid control technologies. This will significantly reduce the design, deployment and life-cycle cost of dependable embedded applications and will integrate these applications seamlessly into the

global information infrastructure. DECOS will contribute to strengthening Europe's leading position in highly developed control systems in the avionics (Airbus) and automotive industries and in the area of dependability of software-intensive systems.

The industry vision of DECOS was expressed by one of our partners, Mr. J. Rennhack from Airbus Germany, at the Embedded Systems concertation meeting in Brussels: "Aerospace and automotive industry have to work together to achieve Aerospace Safety at Automotive Cost."

The specific integration aspects of DECOS are twofold (see Figure 2):

- Horizontal integration over a number of application sectors: automotive, aerospace, railways/transport, industrial automation/process control, medical devices and systems.
- Vertical integration over the following levels:




Figure 1: Embedded Systems in modern cars.




Figure 3: Mixed Traffic Integration of different networks.

- highly integrated chips (SoC – systems on chip)
- hard real-time core systems (integrating several available time-triggered protocols (TTP) such as TTA, FlexRay, TTCan, TT-Ethernet etc)
- integration of time-triggered and event-triggered systems (see Figure 3)
- integration of hard real-time, soft real-time and high-level environments (IP, other buses, legacy components, RT-CORBA etc).

DECOS will deal with the following research areas:

- application architecture design and validation (design methodology, formal design methods, UML, design tools, composability, partitioning, security, dependability modelling, configuration management, tractability of requirements, interface specification)
- component design and validation (programming models and tools, synchronous languages, worst-case execution time analysis, programmers' work bench for TT systems, operating systems, scheduling, tractability of requirements, component parameterisation, reuse of components)
- mixed traffic integration (ET-TT integration, gigabit time-triggered network with integrated guardian in the star, integration of legacy protocols, dependable real-time networking issues)

- middleware for dependable systems (dual-processor hardware, operating systems, fault-masking algorithms, fault-tolerance layer, standard API, dynamic reconfiguration, security, predictability, power awareness, CORBA)
- smart transducers (hardware, meta-level description by XML, wireless transducer networks, EMI robustness, dynamic reconfiguration, remote diagnosis, location-aware sensors, low-power devices, CORBA)
- distributed real-time control networks (system identification, model building, control algorithms, automatic code generation, hybrid systems, validation strategies)
- distributed simulation (real-time simulation of a design, hardware in the loop, application demonstrators, eg Virtual Car including cockpit and

visualisation, other simulation case studies)

- certification (safety case, formal analysis, failure mode and effect analysis (FMEA), testing strategies, collection of evidence, fault-injection, software quality assurance, certification tools)
- standardisation and training (project management standards, technical standards, quality assurance, version control, development of a training course, embedded system laboratory)
- time-triggered applications (application-specific research projects in the field of dependable embedded systems along the TT paradigm, eg in the automotive industry, railway, aerospace/avionics, process control and medical electronics).

The time-triggered architecture assumes that:

- a large distributed control system is structured into clusters of components
- every component has access to a fault-tolerant sparse global time-base of known precision. Important time-critical actions are triggered by the progression of this global time
- components communicate by the exchange of messages with a priori known latency and minimal jitter across well-specified (in the domains of time and value) interfaces
- a component is a fault-containment unit (FCU).




Figure 2: DECOS Integration.

DECOS builds on the experience and results of more than twenty years of research in the field of dependable distributed real-time computing carried out by Prof. Kopetz at the Vienna University of Technology and, among others, the following EU-sponsored research projects:

- PDCS: Predictably Dependable Computing Systems
- DEVA: Design for Validation
- TTA: Time-Triggered Architectures
- X-by-WIRE: Safety Related Fault-Tolerant Systems in Vehicles

- SETTA: Software Engineering for Time-Triggered Architectures
- PAMELA: Prospective Analysis for Modular Electronic Integration in Airborne Systems
- DSOS: Dependable Systems of Systems
- FIT: Fault-Injection into the Time-Triggered Architecture
- NEXT TTA: High-Confidence Architecture for Distributed Control Applications
- HRTC: Hard Real-Time CORBA.

Over 100 interested partners from research and industry have already declared their interest in participating in the project. Close cooperation with the ERCIM Working Group on Embedded Systems will be maintained and the participation of ERCIM WG members in the DECOS project is encouraged.

Link:
<http://www.decos.at/>

Please contact:
 Dieter Donhoffer
 ARC Seibersdorf Research (AARIT)
 E-mail: dieter.donhoffer@arcs.ac.at

Time Triggered Architecture

by Hermann Kopetz

Time Triggered Architecture meets future requirements for building cost-effective dependable embedded systems from components.

Dependable Embedded Systems are an enabling technology for a huge set of critical technical applications, eg in automotive, aerospace, railways and other transportation systems, industrial automation and process control, medical systems and the like, where hard real-time requirements have to be met in a dependable, predictable manner, because people's life may depend on the services provided by this critical systems, subsystems and components. At the moment, sectoral or proprietary solutions are the basis of highly dependable systems, and systems are designed case-by-case almost from scratch.

The Challenge

The challenge is to facilitate the systematic design of large dependable control systems out of components. The mission is to develop a basic methodology and technology which allows to significantly reduce the design, deployment and life-cycle cost of critical embedded applications. A crucial part of the problem is the interaction of the components which is realized by the exchange of messages across linking interfaces (LIFs) to a real-time communication system.

The driving forces for the composition of a large System of Systems (SOS) out of a

set of components (component systems) are:

- cognitive complexity reduction in order to reduce the design and development effort
- reuse of components: the components may be newly designed according to a given architectural style or may be already existing systems (legacy systems)
- simplified diagnostics and repair.

Silicon Trend and New Failure Modes

The trend for systems capable of cost-efficient mass-deployment leads to 'Systems on a Chip (SOC)', where at the end of the day complete nodes of a distributed system are on a single die. According to semiconductor industry studies, the further shrinkage of silicon building blocks is still progressing. The further shrinkage leads to new failure modes of SOC, such as:

- transient multi-bit failures caused by a single fault event
- intermittent failures of the interconnect that can affect different functions on the die simultaneously.

It is expected that in future the rate for permanent failures will remain unchanged, but that the rate for intermittent and transient failures will increase.

Therefore, the assumption that a fail-silent node can be implemented on a single die that hosts two independent FCUs is not sustainable in future high-dependability applications, which severely influences system architectures and composability of embedded SoS.

What is a Component?

In the abstract, a component is an encapsulated building block that is of use when building a large system. The focus here is on system components. Components are characterized by their interfaces with respect to composability and are described by:

- their data properties, ie, the structure and semantics of the data items crossing the interface; the semantics are expressed by an interface model
- their temporal properties, ie, the temporal conditions that have to be satisfied by the interface: control and temporal data validity.

Event triggered vs. Time Triggered

It implies a reliable global notion of time throughout the system if we want to give warranties on timelessness. Simplified, an event triggered system follows the principle of reaction on demand, where temporal control is enforced from the environment onto the system in an

unpredictable manner (interrupts), with all the undesirable problems of jitter, missing precise temporal specification of interfaces and membership, scheduling etc, but good for sporadic action/data, low – power sleep modes, best-effort soft real-time systems with high utilization of resources. Time-triggered systems derive control of follows the principle global progression of time, such allowing precise temporal specification of interfaces and ‘temporal firewalls’ to protect from unpredictable outside interference, membership identification, interoperability and replica determinism.

Time-Triggered Architectures

A properly designed time-triggered architecture (TTA) can provide generically at the level of the architecture:

- strong composability: independently developed functions can be integrated with minimal integration effort
- effective fault propagation barriers
- fault-tolerance by active replication of components
- strong diagnosability: the loss of consistency of the distributed computing base can be promptly detected and diagnosed
- formal analysis of critical architecture functions.

We do not know how to provide these necessary characteristics if the base architecture is event-triggered.

The TTA assumes that a large distributed control system:

- is structured into clusters of components
- every component has access to a fault-tolerant sparse global time-base of known precision. Important time-critical actions are triggered by the progression of this global time
- components communicate by the exchange of messages with a priori known latency and minimal jitter across well-specified (in the domains of time and value) interfaces
- a component is a fault-containment unit (FCU)
- in a properly configured cluster any one component of a cluster can fail in an arbitrary (Byzantine) failure mode without affecting the proper operation of the components not affected by the fault.

The TTA distinguishes cleanly between fault containment and elimination of error propagation:

- a node is an FCU (fault containment unit) that can fail in an arbitrary failure mode
- control errors are detected by an independent FCU, the guardian
- data error masking is in the responsibility of the application, not on TTA level.

In the past few years, a number of time-triggered communication protocols have appeared or have been announced that

provide the clock synchronization service needed for the TTA at the protocol level:

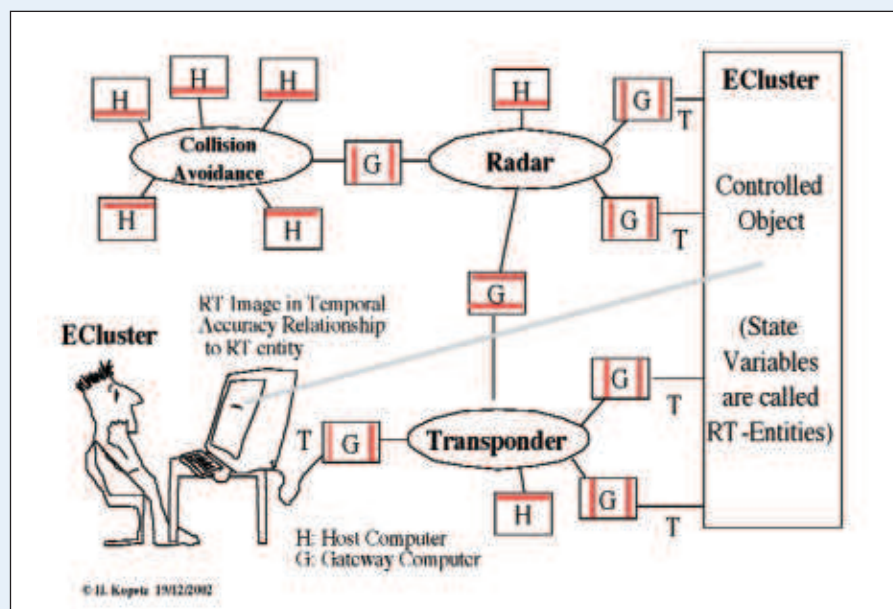
- TTP/C (silicon available since 1998)
- TT-CAN (available since 2002)
- TTP/A (standardized by the OMG in 2002)
- FlexRay (silicon planned end of 2004)
- TT-ETHERNET (in planning phase).

The main difference among these protocols is the attitude towards the inherent design conflict between safety, flexibility and cost. A component-based approach has to take into account deployment of different TTP protocols.

Such an architecture designed for composability must support:

- independent development of components — relates to the architecture
- stability of prior services— relates to the components
- performability of the Communication System — relates to the communication system
- replica determinism — to support transparent implementation of fault tolerance
- diagnostics — it must be possible to identify the sending FCU (Fault Containment Unit) of every message.

The Figure shows such an example for an ATC system. There will be in future always several buses and ‘clusters’ in an airplane, car etc, and there must be support for mixed traffic and integration of or co-existence with legacy systems in less critical parts. This can be achieved by putting the less critical application parts and protocols on top of TTA, or by providing gateways or ‘firewalls’ between the critical and less critical parts by provision of appropriate interfaces.



Composability – a hard real-time Air Traffic Control system of five clusters with maybe different TT protocols.

Link:

<http://www.vmars.tuwien.ac.at/projects/ttp/ttpmain.html>

Please contact:

Hermann Kopetz
 Technical University Vienna
 E-mail: hk@vmars.tuwien.ac.at

Automatic Distribution of Synchronous Programs

by Alain Girault

Synchronous languages have recently emerged as a very elegant, safe, and efficient solution for programming embedded systems. The success with which they have been implemented in safety-critical industries is impressive (eg Airbus planes, nuclear plants etc). Scientists at INRIA Rhone-Alpes are currently performing research in the field of automatic distribution of synchronous programs.

Reactive systems are computer systems that react continuously to their environment at a speed determined by the latter. This class of systems contrasts with transformational systems on the one hand (classical programs whose inputs are available at the beginning of their execution, and which deliver their outputs when terminating; eg compilers), and interactive systems on the other (programs which react continuously to their environment, but at their own speed; eg operating systems). Reactive systems include most industrial embedded real-time systems, such as control, supervision, signal-processing, etc.

The synchronous approach has been proposed to ease the design of reactive systems. It is based on the so-called synchronous abstraction. Without entering into details, this is similar to the abstraction made when designing synchronous circuits at the gate level. Concerning the implementation, synchronous programs are embedded in a periodic execution loop, where inputs are read at the beginning of each cycle, before the next state, and the outputs are computed.

There are numerous languages based upon the synchronous abstraction, including data-flow languages (Lustre, Signal, and Lucid Sychrone), imperative languages (Esterel), and graphical languages (Argos and SyncCharts).

Our work deals with the automatic distribution of Esterel programs. The purpose of automatic distribution is, given a centralised source program and some distribution specifications, to build as many programs as required by the distribution specifications. These programs must be able to communicate harmo-

niously, such that their combined behaviour will be functionally equivalent to the behaviour of the initial centralised source program.

To achieve automatically a distributed implementation of an Esterel program, we first compile it into some intermediate format, and then distribute it according to the user's specifications. The main advantage of this approach is that directly designing a distributed system is always more difficult and error-prone. This explains the recent success of automatic distribution methods. The other advantage of this approach is the ability to debug and

deterministic finite state automaton. This state graph can be cyclic, but in each state there is sequential acyclic code, represented by a directed acyclic graph of actions. With these actions, the program manipulates input, output, and local variables. This internal format is quite general since programs written with a classical imperative programming language can be compiled into it.

In the circuit format, the target program is a sequential circuit, with Boolean gates and registers connected by wires. Some wires are tied to an action in order to manipulate the input, output, and local variables of the program. This control

Format	control structure		
Automation	sequential	explicit	static
Circuit	parallel	implicit	dynamic
Control points	parallel	explicit	static

Table: The three target code formats for Esterel compilers.

formally verify the centralised program before its distribution, which is always easier and faster than debugging a distributed program.

Over the years, a number of compiling methods have been invented for Esterel. These methods vary according to the successive language semantics and the target. More than the compiling method, we are interested here in the format of the target code. Three of these exist, as summarised in the Table above.

In the automaton format, the target program generated by the compiler is a

structure is dual to the automaton structure, the internal state being implicitly stored in the registers. Furthermore, it is parallel in the sense that there are several control paths, and dynamic in the sense that the control depends on the data.

In the control points format, the target program is an ordered list of blocks, each block being a sequence of C instructions linked to one control point. The program maintains two ordered lists of control points, one for the current instant and one for the next instant. At the beginning, a number of control points are true, and the first of those is executed. This

execution can modify the status of other control points, either later in the same instant, or anywhere in the next instant. The next true control point is then executed, and so on until the end of the list. Finally, the list of control points of the next instant is copied into the list of the current instant. This control structure is parallel in the sense that each block contains C code enhanced with parallel constructs.

We have developed three distribution algorithms for the three target formats of Esterel. They all share a common structure that involves the following steps:

1. Assign a unique computing location to each sequential action, thanks to the designer's specifications. These specifications are a partition of the set of inputs and outputs of the program into N subsets, one for each computing location of the final parallel program.
2. Replicate the program on each location.
3. Prune from the program of each location the actions not belonging to the considered location.
4. Insert send actions in order to solve the data dependencies between any two distinct locations. These data-dependencies were created by pruning the actions during the previous step.
5. Insert receive actions in order to match the sending actions.

The algorithms involved in each of these steps vary according to the control structure of the program (automaton, circuit, or control points).

The communications take place through a fully connected network of FIFO queues. Queues allow send actions to be advanced, and receive actions to be postponed, thereby minimising the waiting time induced by the communication network. The send and receive actions perform both the data-transfer and the synchronisation between the source and the destination locations: when the queue is empty, the receive action is blocking (in contrast, the send action is never blocking). The only requirement on the network is that it must preserve the integrity and the ordering of messages.

These three distribution methods have been successfully implemented and tested:



- The ocrep tool automatically distributes Esterel programs compiled into the automaton format (it also works for Lustre). It acts as a post-processor of the standard Esterel compiler. The user must specify the desired distribution in a separate file.
- The screp tool automatically distributes Esterel programs compiled into the circuit format. It acts as a post-processor of the standard Esterel compiler. The user must specify the desired distribution in a separate file.
- The Saxo-RT tool compiles Esterel programs into the control point format, distributes them automatically, and finally generates one C file for each computing location specified in the distribution.

Cockpit of an Airbus plane. Synchronous languages have recently emerged as an elegant, safe, and efficient solution for programming embedded systems and are already successfully applied to safety-critical industries such as Airbus planes.

Links:

Synchronous languages:
<http://www.synalp.org>
 Esterel:
<http://www.esterel-technologies.com>
 ocrep:
<http://www.inrialpes.fr/bip/people/girault/Ocrep>
 screp:
<http://www.inrialpes.fr/bip/people/girault/Screp>
 Saxo-RT:
<http://www.rd.francetelecom.com>

Please contact:

Alain Girault, INRIA
 Tel: +33 476 61 53 51
 E-mail: Alain.Girault@inrialpes.fr

synERJY: a High-Level Language for Embedded Controllers

by Reinhard Budde, Axel Poigné and Karl-Heinz Sylla

A team at Fraunhofer Institute for Autonomous Intelligent Systems has developed and implemented 'synERJY', a Java-like synchronous object-oriented language and programming system for the efficient design of embedded systems.

The design of embedded systems poses challenges reminiscent of those in the early days of computing: namely, to use as few resources in terms of memory and time as possible. New challenges have also been added, such as the distribution of functionality across a network and dependability of systems. For example, a modern car has a number of microprocessors communicating via field busses, and consumers are not amused if let down by a car due to a software error. Migration of functionality from hardware to software poses another challenge, as does sheer size. A mobile phone today carries software of a million lines of code in C, and substantially more at the high end of the market. Economic considerations are also indisputable: 99% of all processors are said to be used in embedded systems. Most of these are small microprocessors, but even 75% of high-end processors are claimed to

reside not in workstations but in embedded systems.

In spite of the scientific challenges and the economic importance of embedded systems, software is often handcrafted as in the early days of computing. Even assembler languages are still used. High-level languages are often considered too costly in terms of memory and time consumption. On the other hand, it is fair to note that coding in assembler or even in C often turns out to be costly, since there is a mismatch between the design and the abstractions supported by the languages. The last fifteen years have seen the evolution of some languages and paradigms to close this gap. For instance, SDL has been adopted in telecommunications for programming asynchronous soft real-time systems, and tools such as STATEMATE have been used for programming synchronous hard real-time systems.

Engineers use tool sets like MATLAB/Simulink for modelling systems and for designing and analysing controllers on a wide scale. The underlying paradigm is that of synchrony. If we consider the difference equation of a filter

$$y(n) = a_0 \cdot x(n) + a_1 \cdot x(n-1) + b_1 \cdot y(n-1)$$

we note that:

- computation proceeds in discrete steps
- that the result depends on the actual input $x(n)$ and the system state $x(n-1)$, $y(n-1)$
- that the result $y(n)$ must present at the same instant n as the input has been read.

The latter implies in practice that, while executing, input and system state should be frozen and that execution should be timely, ie terminate before the next step

is supposed to take place. This is exactly the execution model of synchronous programming. In fact, we believe that synchrony is pervasive as a paradigm in the world of embedded control.

With synERJY we promote a high-level reactive Java-like language combining


- object-oriented modelling for a robust and flexible design
- synchronous execution for a precise modelling of reactive behaviour with a well-defined semantics.

Highlights of synERJY are that it

- provides a smooth embedding of the reactive behaviour into the object-oriented data model
- offers fine-grained integration of synchronous formalisms such as
 - the imperative style of ESTEREL
 - the dataflow-based style of LUSTRE
 - hierarchical state machines in the style of STATECHARTS
- generates efficient code in space and time even for small microcontrollers.

There is little need to discuss the merits of object-oriented design. There are, however a couple of notable benefits to be gained from the combination of object-oriented design with synchronous programming:

- due to the underlying broadcast mechanism of synchronous programming, reactive objects (ie objects that exhibit reactive behaviour) behave like components, in that they may be 'plugged in' without changing their behaviour, regardless of the environment
- extending the causality analysis of other synchronous languages, time races between method calls can be detected enabling deterministic scheduling at compile time including the scheduling of all data actions



A mobile phone today carries software of a million lines of code in C. In spite of the scientific challenges and the economic importance of embedded systems, software is often handcrafted as in the early days of computing because high-level languages are often considered too costly in terms of memory and time consumption.

- the component-based approach supports hiding of hardware details such as irrelevant properties of, for example, specific CAN hardware.

Although it is a flexible high-level language, synERJY generates efficient code in contrast to standard expectations for Java-like languages. The language compiles to a subset of ANSI-C, for cross-compiling to the target architecture. Efficiency is gained by optimisation, eg replacing dynamic by static binding and avoiding object allocation whenever feasible. The basis of the optimisation is the closed world assumption that is reasonable for embedded systems. At present, the project is focusing on further optimisation.

The reactive code is compiled to a set of Boolean equations, which have a dense representation due to Boolean optimisation techniques such as using jump tables similar to BDDs. In fact, although the language targets microcontrollers, hardware formats can be generated, for instance for FPGAs, given that the data structures used are elementary. This is an additional benefit in that implementation decisions may be revised using hardware instead of software, without revising the design.

The language features a particular type of 'time' that allows system time to be accessed in a way consistent with the synchronous execution model. As a consequence, simulation using the synERJY tool set exactly coincides with the execution on a target machine

including the timing, and is reproducible due to the guaranteed deterministic behaviour, an obvious benefit when testing systems.

SynERJY is a relatively new but by now stable language, the development of which has been supported by the ESPRIT Projects SYRF and CRISYS. At present it is mainly used in-house for robotics applications and externally for teaching purposes.

Links:

<http://www.ais.fraunhofer.de/~ap/sE>

Please contact:

Reinhard Budde, Fraunhofer-AIS

Tel: +49 2241 142417

E-mail: Reinhard.Budde@ais.fraunhofer.de

STOREQ: Storage Requirement Estimation and Optimisation Tool for Data-Intensive Embedded Applications

by Per Gunnar Kjeldsberg

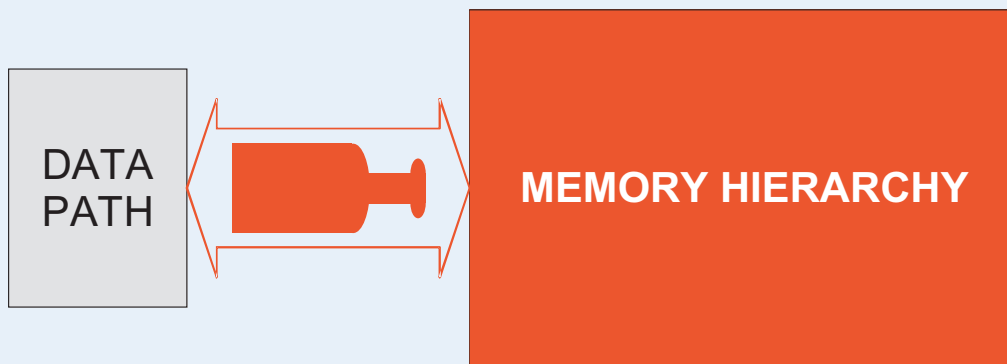
Scientists at NTNU have developed a tool for storage requirement estimation and optimisation of data-intensive applications, called STOREQ. The tool can guide the designer during the early system design steps towards an implementation with low memory usage.

Many embedded systems, especially in the multi-media and telecom domains, are inherently data dominant. For this class of applications, data transfer and storage largely determine cost and performance parameters. This is the case for chip size, since large memories are usually needed, performance, since accessing the memories is often the main bottleneck, and power consumption, since the memories and buses consume large quantities of energy. Even for systems with caches, the overall storage requirement has a vital impact on the performance and power consumption, since it greatly influences the number of slow and power expensive cache misses. Consequently, during the system development process, the designer must concentrate first on exploring the data

transfer and storage to produce a cost-optimised end product. The STOREQ tool assists the designer during these early steps of the system design.

This work has been performed as cooperation between NTNU in Trondheim, Norway, and IMEC in Leuven, Belgium. At IMEC, a Data Transfer and Storage Exploration methodology has been in development for a long time. Since 1999, the Department of Physical Electronics at NTNU has cooperated with IMEC within this framework in the area of storage requirement estimation. The cooperation is ongoing and is now intensifying.

At the system level, no detailed information is available regarding the size of the memories required for storing data in the alternative realisations of an application. To guide the designer and assist in selecting the best solution, estimation techniques for the storage requirements are needed very early in the system design trajectory. The simplest estimates use the maximal size of the intermediate array data as declared in the application code. This is, however, not representative for the effective size required for their storage during the actual execution, since arrays and parts of arrays may not be alive simultaneously. To achieve accurate estimates, the so-called in-place mapping opportunity generated by these non-overlapping lifetimes must be taken into account. For scalars, a relatively



For embedded data-dominated applications, data transfer and storage forms the bottleneck for system size, performance, and energy consumption.

simple lifetime analysis suffices, but for arrays, this is extremely complex due to the huge number of signals and the often very complex interdependencies between them.

For our target classes of data-dominant applications the high-level description is typically characterised by large multi-dimensional nested loops and arrays. Within the loop nests, statements access the arrays using read and write operations. At the beginning of the design process, no information about the execution order of these loops is available, except that which is given from the data dependencies between the statements in the code. As the process progresses, the designer makes decisions that gradually fix the ordering, until the full execution ordering is known. This execution ordering determines the lifetimes of the array elements, and hence the storage requirements of the arrays. To guide the designer it is therefore essential to have storage requirement estimates that can take the available partially fixed execution ordering into account during exploration. Previous work has either not taken execution ordering into account at all, resulting in large overestimates, or has required a fully specified ordering. In the latter case, a full exploration of all alternative orderings of the unfixed loop dimensions is needed, which is unfeasible for fast feedback purposes.

The storage requirement estimation methodology implemented in STOREQ solves these important design problems. The methodology is divided into four steps. In the first step, a data-flow graph is generated that reflects the data dependencies in the application code. The array accesses and the dependencies between them are described using a polyhedral model. The second step

places the polyhedral descriptions of the array accesses and their dependencies in a so-called common iteration space. The third step estimates the upper and lower bounds on the storage requirement of individual data dependencies in the code, taking into account the available execution ordering. As the execution ordering is gradually fixed, the upper and lower bounds on the data dependencies converge. Finally, simultaneously alive data dependencies are detected. Their combined maximal size at any time during execution equals the total storage requirement of the application. An important part of the estimation technique utilises loop-ordering guidance to estimate upper and lower bounds on dependency sizes.

The STOREQ tool has been developed using MATLAB. It takes as input a polyhedral description of the array accessing and dependencies in the application. A prototype interface exists between STOREQ and the ATOMIUM tool being developed at IMEC. This enables automatic generation of the polyhedral input description from the original application code. The focus of the current version of the tool is on the third step of the methodology. Hence the main output consists of the upper and lower bounds on the storage requirement of individual dependencies, given a partially fixed execution ordering specified by the designer. In addition, STOREQ employs the guiding principles of the methodology on a more global basis and suggests an ordering for the still unfixed dimension.

The feasibility and usefulness of the methodology and the tool are confirmed using several representative application demonstrators. It is for instance shown how the designer is guided into reducing

the memory size of the major arrays in the MPEG-4 Motion Estimation Kernel from 262400 to 257 memory locations. Similar results are achieved for a Cavity Detection algorithm. In applying the methodology to an Updating Singular Value Decomposition algorithm, it is also demonstrated how estimation feedback during global loop reorganisation can approximately halve the application's storage requirement.

Links:

<http://www.fysel.ntnu.no/~pgk/STOREQ-Users-Manual.pdf>
<http://www.imec.be/atomium/>

Please contact:

Per Gunnar Kjeldsberg, NTNU
 Tel: +47 7359 4400 / 4405
 E-mail: Per.Gunnar.Kjeldsberg@fysel.ntnu.no

Composing Embedded Real-Time Software Components: the PECOS Data-Centric Approach

by Stéphane Ducasse, Oscar Nierstrasz and Roel Wuyts

Although component-based software development (CBSD) has become mainstream for conventional applications, it has remained elusive for embedded applications due to non-functional constraints. The PECOS project has demonstrated that CBSD can also be applied to severely constrained embedded devices, with timing and memory requirements being taken into account.

Most component models and systems, such as (D)COM, .Net, JavaBeans and CCM, target the development of mainstream desktop applications with abundant hardware resources. None of these existing models explicitly addresses the constraints imposed by small, embedded devices. Although CBSD would bring numerous advantages to the embedded systems world, including shorter development times and the ability to secure investments through reuse of existing components, in practice, standard component models are too heavyweight to be applied to embedded systems development.

The European IST project PECOS (IST-1999-20398), which stands for Pervasive Component Systems, enables CBSD for small embedded systems by providing an environment that supports the specification, composition, configuration checking, and deployment of embedded systems built from software components. The PECOS approach was applied in the context of field devices - small, embedded devices developed by the prime contractor of the project, ABB.

A Data-Centric Model

The PECOS component model is unusual in that PECOS components have very simple interfaces consisting solely of data ports. Components are one of three types: active, passive or event; interaction occurs only where data ports are connected. This design choice makes it particularly easy to work out the synchronisation and timing aspects of components. First of all, it enables timing analysis of component compositions using rate monotonic analysis, a technique for verifying that a set of tasks

can meet their deadlines. Second, it enables schedule verification of component compositions using a constraint logic programming approach. Last but not least, the model minimises the number of concurrent tasks needed to implement the system.

The model is the conceptual part of the toolchain used by developers to compose, deploy and test applications for field devices. Specifying components and compositions is done with CoCo, the component description language. From the CoCo specification, skeleton code is generated that tailors a runtime environment for the device. The skeleton code has hooks where the actual behaviour of the components is manually implemented. The result is then compiled and deployed on the device.

When components are composed, different validation actions are performed. First of all, compositions can be checked structurally for adherence to developer-specified patterns. Second, a timing analysis can be performed to determine whether a composition can meet its deadlines by mapping the composition to sets of tasks that can be checked by rate monotonic analysis. Note that the components need to have certain runtime figures (worst-case execution time and maximum blocking time) associated with them for which a profiling environment was built. Third, a constraint logic programming approach can be used to check whether compositions have a valid schedule (or a schedule can be generated if none was specified).

Although a full validation of the PECOS approach was not possible within the scope of the project, the ideas have been convincingly tested in the context of the field device demonstrator, and plans are underway to apply the model to further case studies.

Participants

PECOS was composed by ABB Corporate Research Centre, Ladenburg (DE), Research Centre for Information Technologies (FZI), Karlsruhe (DE), the Software Composition Group (SCG) at the University of Bern (CH), and Object Technology International (OTI), Amstelveen (NL).

Links:

Results of the project are available at <http://www.pecos-project.org/>
Software Composition Group,
University of Berne:
<http://www.iam.unibe.ch/~scg>

Please contact:

Oscar Nierstrasz, University of Berne
E-mail: nierstrasz@iam.unibe.ch

Time Machines and Black Box Recorders for Embedded Systems Software

by Henrik Thane

Black-box recorders and virtual Time Machines make it possible to deterministically re-execute embedded real-time software. It is possible to jump back and forth in time while debugging the system offline.

The concepts of time machines and time travel have enticed people’s fantasies for hundreds of years. Since the publication of H.G. Wells’ ‘The Time Machine’ in 1895, hundreds of books and movies have followed. Even in theoretical physics, research on the subject has been taken seriously. Esoteric ideas like making use of black holes for time travel (Kip Thorne at Caltech) or consequences of Einstein’s theory on relativity have been considered. However, time travel has never been proven to work in practice or theory. It has been speculated that time travel might even be prohibited by

the laws of the universe. Nonetheless, if we consider man-made constructs such as computers and computer programs, it is possible to make ‘time machines’ that allow us to achieve time travel for a specific purpose: debugging.

The cost for verification and debugging of embedded software typically exceeds half the development budgets. Debugging of embedded systems software is difficult, and for multi-tasking real-time software especially so. Embedded systems have few interfaces for diagnostic observation, precisely

because they are embedded. What makes matters worse is the fact that the actual act of observation may change the behaviour of the system, especially if the observation is performed using some software other than the application code (causing a probe-effect). Another large problem inherent in the concurrency of multitasking real-time software is that it is very difficult to reproduce executions and observations.

In solving these problems, our research (at Malardalen Real-Time Research Centre) has led us to solutions based on black-box recorders (similar to those in aeroplanes) and ‘time machines’. By recording significant events online like task-switches and interrupt hits as well as data from the external process and internal state, we can deterministically re-execute the embedded system software offline, as dictated by the recording. From a user’s point of view, this deterministic replay will behave exactly like a regular sequential program, mimicking the exact execution of the recorded multitasking real-time




Figure 1: A robotic assemblyline with black-box recorders. Retrieval of black-box contents allow remote deterministic post mortem replay debugging using virtual Time Machines.

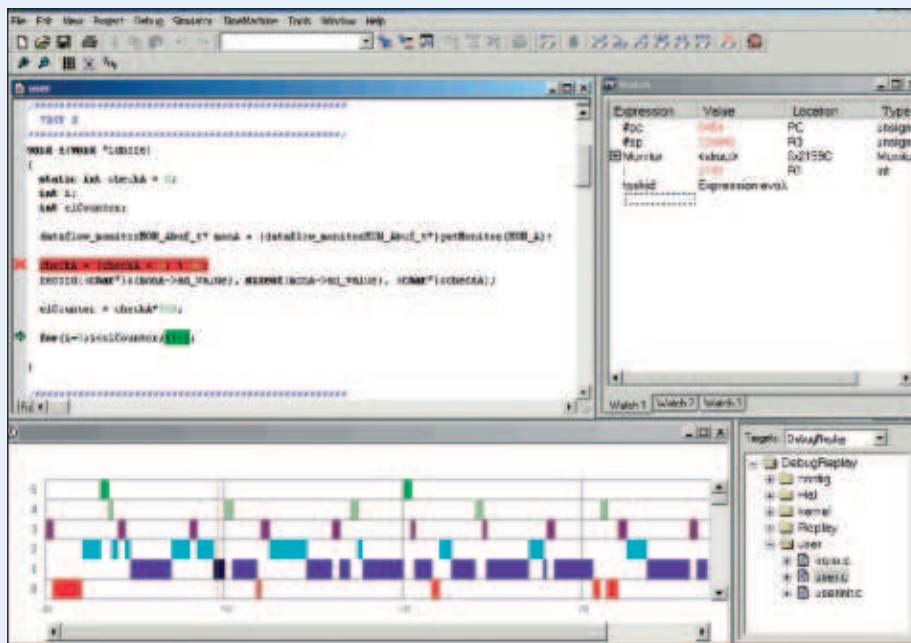


Figure 2: A commercial IDE with an instruction level simulator debugger, into which we have integrated our Time Machine technology (the lower left window). The time line illustrates the recorded control-flow for six tasks; task priorities on the vertical axis. Selecting any instance of a task re-executes the system from an idle point (the red lowest priority task) up to the selection (it is possible to jump back and forth in time). The debugger window shows the current state. From here it is possible to single-step, watch variables, and set new additional breakpoints.

application. We can single-step, insert any number of breakpoints and inspect data without introducing the probe-effect. We can even jump back and forth in time using the debugger (therefore named the Time Machine). Since we have eliminated the dependency on real time and replaced the temporal and functional context of the application with the recording, we can replay the system history repeatedly.

We have applied our method to a number of systems, but the most recent and most complex is an industrial robot control system from the largest industrial robot manufacturer in the world, ABB Robotics. Their system consists of several computing control systems, signal processing systems and I/O units. We applied our Time Machine to the

motion-control part of the system, which consists of approximately 2.5 million lines of C code and is run on the VxWorks real-time operating systems. The motion-control part is a hard real-time system, with about 70 tasks running (the most frequent task is activated every 4ms) and multiple interrupts driving an assortment of device drivers.

The control flow of the system (the task-switches) was captured by a task-switch hook and recorded in a cyclic buffer of programmable length (the black-box). We also transparently instrumented the system calls that could change the system control flow by making use of an existing operating system abstraction layer. The only manual instrumentation that had to be inserted into the source code were calls to data-flow monitors

after blocking system calls, in order to capture messages and the state of the task (represented by specified local and global variables). Worth noting is that we needed only to record the start conditions, since we re-executed the code offline. In total, our black-box recorder introduced an overhead of less than 2% of the processor utilisation and a few hundred kB of data in order to capture the last few hundred events before major failures, which could subsequently be replayed in the Time Machine.

Links:

<http://www.mrtc.mdh.se/>
<http://www.real-timemachine.com/>

Please contact:

Henrik Thane
 Malardalen University, Sweden
 E-mail: henrik.thane@mdh.se

Design of Embedded Systems: New Challenges

by Gilbert Cabillic and Isabelle Puaut

Ongoing work at INRIA Rennes concerns embedded system software, and addresses the problems presented in providing an easy and safe framework for estimating software's resource consumption, which is a difficult task when done manually by a benchmarking approach. Further, by enabling dynamic download of Java applications, a discontinuity in the embedded application development life cycle is introduced. This is a new and significant contribution towards opening up the use of embedded systems.

Traditionally, the development of embedded software has had to take into account several constraints stemming from either the software or the embedded architecture. Avionic or robotic embedded software has inherent real-time constraints that need to be supported by the underlying operating system. Regarding the architecture, limited and customised resource features (eg shared memory multiprocessor with multiple frequencies, a small capacity of several models of memory, limited battery capacity) increase the complexity of building an embedded operating system.

To facilitate the development of embedded systems, we address two different challenges not met in existing

embedded systems. Our first focus is on a resource consumption evaluation of embedded software (CPU, energy, or memory). Our second challenge aims to enable the use of Java for small embedded devices like cell-phones or Personal Digital Assistants. This will support the dynamic download of applications and will radically change the embedded application development life cycle.

Predicting Resource Consumption of Embedded Software: Static Worst-Case Execution Time Analysis and Beyond

Predicting the amount of resources required by embedded software is of prime importance for verifying that the system will fulfill its real-time and

resource constraints. A particularly important point in the framework of hard real-time embedded systems is to predict the Worst-Case Execution Times (WCETs) of tasks, so that it can be proven that task deadlines will be met.

Our ongoing research concerns methods for automatically obtaining the upper bounds of application execution times on a given hardware. Static analysis of the source code of applications is used to identify worst-case execution scenarios; static analysis methods have been preferred to testing because the latter class of methods requires the exploration of all possible inputs for a piece of software to identify its longest execution path. Hardware models of processors are used to obtain WCET instruction

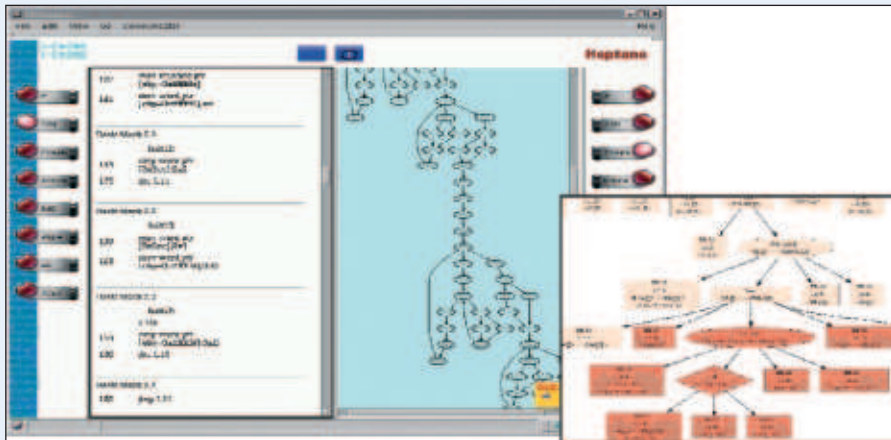


Figure 1: Heptane static WCET analysis tool.

sequences. The use of hardware models instead of the actual hardware allows static WCET analysis methods to be used earlier in the application development life cycle. In addition to allowing the verification of deadlines, static WCET analysis methods can help in selecting or dimensioning the hardware used in hard real-time embedded systems, and can serve in the comparison of different application implementation strategies. A result of our research on static WCET analysis is the open-source analyser Heptane (see Figure 1) aimed at obtaining WCETs on processors with in-order execution, equipped with caches and pipelined execution. The modularity of Heptane allows it to be ported to different target processors and programming languages.

Our belief is that an early estimation of the resource consumption of embedded software allows its suitability to the constraints of embedded systems (limited time, memory or energy) to be proven early in the application development life cycle. In the future, our ongoing work on static time estimation will be extended to the other scarce resources of embedded systems, such as energy consumption. This is a new and challenging issue for hard real-time embedded systems.

Bringing Java to Small Embedded Devices

Java offers several serious advantages that could broaden the use of the Wireless Personal Digital Assistants (WPDA) for the user. Firstly, Java is portable and hence is independent of the

hardware platform. This is very important for reducing the cost of application development in the pervasive embedded market. As Java can be run anywhere, the application development can be done on a desktop without the need for the target hardware platform. This is a strong discontinuity on the application development life cycle for embedded architectures. Secondly, Java supports dynamic loading of applications and can significantly contribute to extending the use of WPDA (see Figure 2).

For these reasons, Java is of great interest for these embedded environments. Nevertheless, even if Java shows much potential, one of its main drawbacks is the need for resources in order to run an application. These include memory volume, execution time and energy consumption, which are the typical resources examined for embedded system tradeoff conception. It is therefore clear that the success of Java is conditioned by the availability of a Java execution environment that will manage these resources efficiently.

Over the last four years, our ongoing research has been done in cooperation with Texas Instruments, and aims to provide a Java execution environment for WPDA architectures that permits a good tradeoff between execution time, energy and memory consumption. During this research, we first identified Java opcodes that will significantly influence the energy consumption. We then designed a new approach to construct a Java execution environment based on a modular decomposition.



Figure 2: Bringing Java on embedded devices.

Using modularity, it is possible to specialise some parts of a Java Virtual Machine (JVM) for one specific processor (eg to exploit low power features of Digital Signal Processing processors in order to minimise energy consumption). Lastly, we implemented our JVM from scratch and validated it on a WPDA architecture based on an Omap Texas Instruments platform (a shared memory heterogeneous multiprocessor). Our current work is focusing on how to provide a high-performance Java embedded platform, based on strong hardware and software cooperation.

Nowadays, an embedded system is closed to the outside world and its development is difficult. Many systems are implemented using assembler or C language. Thanks to our research on Java, and in particular its performance, Java could become a very valuable tool for designers building embedded systems. For users, it will bring the openness of an embedded system. These features could also make it easier to deploy and design ubiquitous applications. For designers, it enables the use of a strong typed object language instead of basic languages like C, or assembler.

Please contact:
 Gilbert Cabillic, INRIA
 Tel: +33 2 99 84 72 92
 E-mail: gilbert.cabillic@irisa.fr

Isabelle Puaut, INSA de Rennes, France
 Tel: +33 2 99 84 73 10
 E-mail: isabelle.puaut@irisa.fr

Transforming Specifications to verify Embedded Systems

by María del Mar Gallardo, Jesús Martínez, Pedro Merino and Ernesto Pimentel

A team from the GISUM research group at the University of Málaga is working on extending current tools to analyse the functional correctness of software for embedded systems.

Interest in these tools, which are based on formal languages, is now moving from academia to industry. In this new context, the systems to be analysed are more realistic and the problems of complexity appear more frequently, especially for automatic verification. One promising method for maintaining the interest in verification is the use of reduction techniques to obtain smaller specifications. We are using a robust syntactic transformation method that reduces the specification and is compatible with existing tools.

The use of formal specifications to analyse the correctness of software for embedded systems has been a hot research topic in the last 25 to 30 years. Many design errors have been detected by using the specification as a way of manually revising the design, or more interestingly, as a way of performing an exhaustive inspection of the potential executions of the system. This last auto-

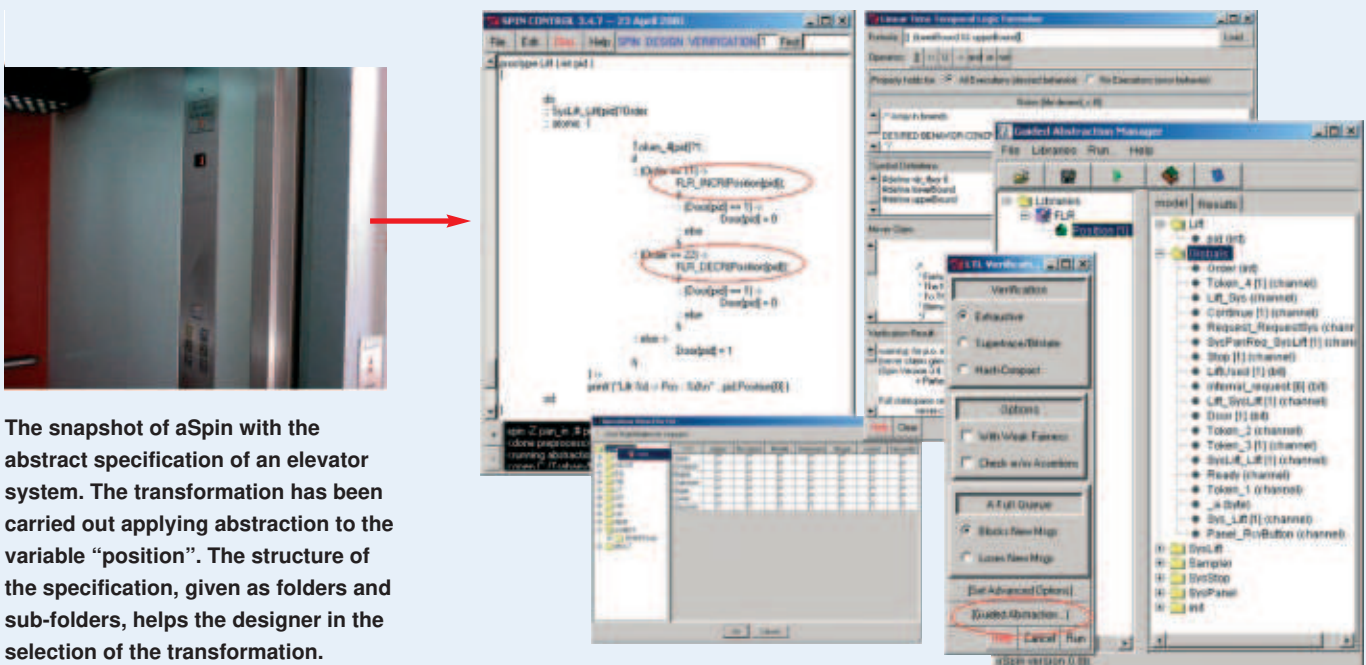
matic method, usually called model checking, has been mainly implemented for academic use. More interest now exists in industry, however, in having robust tools supporting verification facilities for the development of this kind of critical system.

Although the available hardware and software technology has considerably increased the size of the specifications that can be exhaustively inspected, the state explosion still presents a major problem. The algorithms usually implemented produce a low quality analysis when the tool runs out of memory.

The automatic verification team in GISUM is working on areas of model checking, static analysis, abstract interpretation, and verification of software for critical systems, including communication protocols and embedded systems. A current topic of our research is the combination of model checking and

static analysis techniques to improve the quality of verification with the abstract model-checking approach. This method consists in replacing the specification of the system being analysed with an abstract one that produces a smaller state space. When the abstraction is sound, then the verification of the abstract version provides information about the correctness of the original specification for given kinds of properties. Our aim is to develop both theoretical frameworks and tools to support abstraction by automatically transforming the specification.

One of the main current results is the development of a new approximation method to verify properties represented with temporal logic formulas using abstract models. The abstraction of the model is guided by the change in the type of some selected variables. This change reduces the domain of values for the variables, and must be followed by updating the instructions that manipulate



The snapshot of aSpin with the abstract specification of an elevator system. The transformation has been carried out applying abstraction to the variable "position". The structure of the specification, given as folders and sub-folders, helps the designer in the selection of the transformation.

the variables. Both changes can be implemented by program transformation. Similar ideas about approximation and code transformation have also been proposed for more commercial specification languages like SDL/MSD and UML StateCharts/Sequence Diagrams.

Regarding tools, our purpose is to develop reusable modules to incorporate abstraction in existing model-checking toolsets. One main topic is to define standard representations for the specification of systems and properties. XML is now being used for this purpose. The availability of APIs to process XML makes the transformation work easier.

The current visible result is aSpin that extends the model checker Spin to support the theoretical proposals developed in the group (see 7th ERCIM FMICS Workshop).

We now plan to extend other tools. The verification of Messages Sequence Charts (MSC) against SDL (presented at 5th ERCIM FMICS Workshop) can be implemented on top of Telelogic Tau. The method for verifying StateCharts against Sequence Charts (see paper in JOT) can be implemented in VisualState, Rapshody or Rational Rose.

A further project is the development of transformation schemes for standard programming languages, in order to take advantage of future model-checking tools for these languages.

Links:

GISUM research group.
<http://www.lcc.uma.es/~gisum/>
 Verification tools and related papers.
<http://www.lcc.uma.es/~gisum/fmse/tools>

Please contact:

María del Mar Gallardo, Jesús Martínez,
 Pedro Merino, Ernesto Pimentel,
 University of Málaga, Spain
 E-mail: gallardo@lcc.uma.es,
jmcruz@lcc.uma.es, pedro@lcc.uma.es,
pimentel@lcc.uma.es

EUDEMES: Component-Based Development Methods for Small-Size Embedded Systems

by Nicolas Guelfi, Aloyse Schoos, Gérard Sébastien and François Terrier

EUDEMES is a recent joint France-Luxembourg project in the field of embedded systems engineering. The first targeted systems are small embedded systems developed by the Luxembourg company IEE S.A. for the control of airbags through sensors or infrared cameras.

The main objective of this project is to study how component-based development methods can be applied to small-size embedded systems engineering. These approaches are most often applied to large and complex embedded systems, but their principles should be applicable to small systems in order to provide a high degree of maintainability and to foster product family development. These approaches are of real interest when dealing with software processes in which the specification phase is often subject to iterations. An important issue of the project is also the development of a component model that could be integrated into an existing system with a minimum of effort for interfacing and validation. The formalism retained for this study is the Unified Modelling Language (UML).

The first phase of EUDEMES will be to transfer the ACCORD/UML development method (mainly developed by the project AIT-WOODDES, <http://wooddes.intranet.gr>)

to the project partner IEE. This first phase should:

- provide an impact study of CBSE (Component-Based Software Engineering) on IEE development processes
- state a list of research subjects to be addressed and mastered in order to provide new solutions in the field of CBSE as applied to small-size embedded systems.

In the initial stage of the project, the concept of components, and particularly the description of their dynamic behaviour, will be identified and formalised in UML. This provides the basis for defining techniques and constraints, which will ensure the compatibility of the operation properties of the components.

The concept of 'component' clearly seems necessary for the development of software systems that are less expensive, more reliable and more evolutionary. In

this context, research is directed towards the improvement of object-oriented technologies, in particular through the use of the standard formalism UML.

A sticking point identified by the research community is the management of the dynamics of the components and their interaction with other system elements. In this context, it is essential to provide tools which make it possible to model the components' dynamic behaviour in a synthetic yet precise way. For that, the principle is to enhance and formalise the concept of 'quality of service', which is already well established in telecommunications research. Its adaptation to the constraints of other embedded systems fields, such as the automotive industry, is an essential first step. A further point is the capacity to validate the assembly of components starting from properties attached to each component, the assembling topology, and dynamic properties imposed on the overall system.

EUMEDES is a joint project of LIASIT (Luxembourg International Advanced Studies in Information Technologies), IEE Luxembourg (International Electronics & Engineering S.A), CEA-LIST (Laboratory for System and Technologies Integration of the French Atomic Energy Agency) and the French NINST (National Institute of Nuclear Sciences and Techniques). The project started in January 2003 and is of four years' duration. The project is within the scope of work carried out at the LLSP (Laboratory on Software for Process Safety) at the Laboratory for System and Technologies Integration of the French Atomic Energy Agency (CEA-LIST) on the reliability of embedded systems, with the main concern being the integration of new development techniques by components. The work will be based on the assets of two CEA research projects: modelling of the dynamic behaviour of real-time systems (ACCORD) and formal specifications analysis based on automata (AGATHA).

The results will make it possible to develop the work that has been carried out by CEA over a number of years in object-oriented modelling, UML, real-time systems such as those relating to formal models analysis, and the automatic generation of tests starting from industrial formalisms (such as SDL, statecharts, UML etc). This work will enhance the current proposals of the ECA-LLSP, which are in progress at OMG, the international organisation dealing with the standardisation of object-oriented technologies, which is in charge of CORBA and UML.

The developments will be carried out such that they can be integrated into the software engineering tools that support the ACCORD platform. This approach will facilitate the evaluation within an industrial framework of the suggested solutions and, in the long term, their industrial transfer.

This project will cooperate with international partners, in particular with:

- DECOS – Dependable Embedded Components and Systems (<http://www.decos.at>)
- ARTIST – Advanced Real Time Systems (<http://www.systemes-critiques.org/ARTIST/>)
- ERCIM Working Group 'Dependable Software-Intensive (Embedded) Systems' (Erwin.Schoitsch@arcs.ac.at)

Please contact:

Nicolas Guelfi, Luxembourg University of Applied Sciences (FNR)
Tel: +352 420 101 251
E-mail : Nicolas.Guelfi@ist.lu

Aloyse Schoos, International Electronics & Engineering S.A, Luxembourg
E-mail: aloyse.schoos@iee.lu

G rard S bastien, Fran ois Terrier
National Institute of Nuclear Sciences and Techniques (CEA), France
E-mail: Sebastien.gerard@cea.fr,
Francois.Terrier@cea.fr

Model Checking of Embedded Systems

by Stefania Gnesi

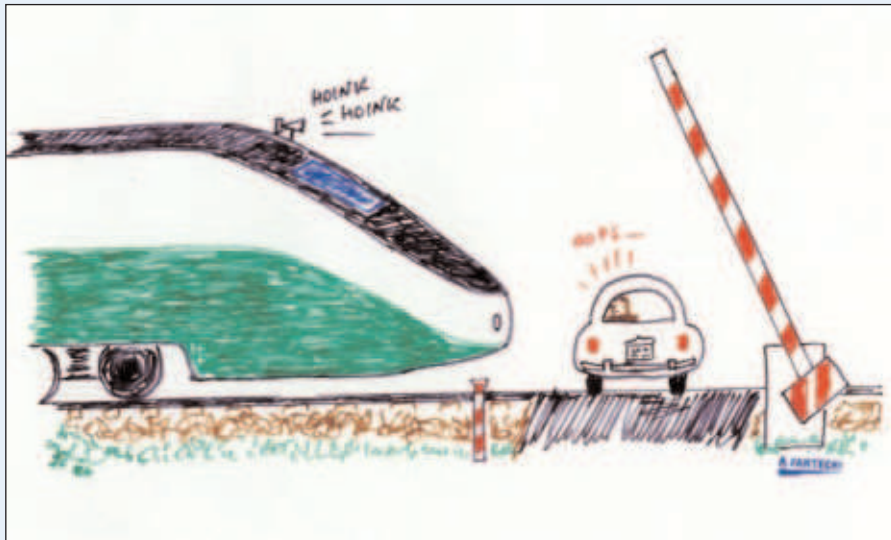
The integration of different dependability techniques is an open research question. We address problems that arise when attempting to combine fault tolerance mechanisms with formal methods and formal verification tools in the development of an embedded system.

In recent years, the wide spread deployment of embedded systems on which human activities depend has raised many concerns about safety issues. A combination of fault prevention, fault tolerance, fault removal and fault forecasting techniques are commonly used in order to achieve a high degree of dependability. However, there is no common agreement on a standard method to combine and integrate individual techniques. For example, industries with different backgrounds and application fields tend to adopt their own particular development trajectories when applying techniques aimed at enhancing dependability.

The application of formal methods in the rigorous definition and analysis of the functionality and the behaviour of a system means that the system is designed according to a set of predefined abstract properties that guarantee its 'correct' behaviour. It is thus astonishing to see how seldom formal methods are actually used by the safety-critical system industry, despite the fact that their adoption is increasingly required by the international standards and guidelines for the development of such systems. The truth is that industrial acceptance of formal methods is strictly related to the investment needed to introduce them, to the maturity of the tool support available, and to the ease of use. For these reasons, the current industrial trend is to adopt

formal verification techniques to validate system design and integration within the existing development process. Industries prefer to use formal verification techniques assessing the quality attributes of their products, obtained by a traditional life cycle, rather than adopting a fully formal life cycle development, simply because it is cheaper to do so.

Several approaches to the application of formal methods in the development process have been proposed; they mainly differ with respect to the degree of involvement of the method. Starting from rigorous specifications, formal methods can be used for the derivation of test cases, as a validation technique



Embedded computer-controlled systems often include fault tolerance techniques. These are, for example, applied to a Railway Interlocking System.

aimed at proving that the specification satisfies the requirements, or just as an auxiliary technique in the automated generation of code.

Formal verification methods based on model checking are applied on a finite state representation of system behaviour. Verification is usually carried out by using model checking algorithms to demonstrate the satisfiability of certain properties formalized as logical formulae over the model of the system. For example, safety and liveness requirements can be expressed as temporal logic formulae and can be checked on the model of the system. Unfortunately, this approach suffers from the so-called ‘State Space Explosion’ problem that can arise when a system is composed of several subsystems. In this case, a finite state model with a number of states, which is exponential to the number of the component subsystems, can be generated. Systems that are highly dependent on data values share the same problem, producing a number of states exponential to the number of data variables. Hence, traditional model checking techniques have proved to be insufficiently powerful for many ‘real’ systems, when their models are larger than 100000 states.

Recent advances in model checking techniques, however, have managed to deal with very large state spaces by using symbolic manipulation algorithms inside model checkers. Such tools have been

successfully applied to very large state spaces in the realm of hardware verification.

Embedded computer-controlled systems often include fault tolerance techniques. Fault tolerance is the property of a system to provide, through redundancy, a service that complies with the specification despite the occurrence of faults. The rigorous definition and verification of this class of systems is extremely important since it makes it possible to demonstrate that a system is correct even in the presence of faults and failures.

We have applied model checking verification techniques to embedded systems illustrating how certain characteristics of embedded systems, such as the use of redundancy, can help to reduce the state space explosion problem. In this work, we have considered several case studies. Two interesting examples were the verification of the safety requirements of a Railway Interlocking System developed by Ansaldo Trasporti and the verification of some fault tolerant mechanisms defined inside the EU project GUARDS (Generic Upgradable Architecture for real-time Dependable Systems).

Both studies have shown that:

- the application of model checking formal verification methodology is feasible and well accepted in the industrial context of embedded fault tolerant systems

- the formalization process strictly depends on the application domain. Some standard rules for the passage from the semi-formal description of the system to its formal specification can be successfully applied to the field of embedded fault tolerant systems. This passage is generally recognized as one of the critical points in the introduction of formal methods in the software development cycle
- the reduction in the state space due to the phased structure of redundant systems makes the model checking approach viable in this application domain
- the use of finite state machines as the specification language has the advantage of ensuring the adherence of the formal specification to the original semi-formal one.

Link:
<http://matrix.iei.pi.cnr.it/FMT/>

Please contact:
 Stefania Gnesi, ISTI-CNR
 ERCIM Formal Methods for Industrial
 Critical Systems Working Group Chair
 Tel: +39 050 3152918
 E-mail: gnesi@iei.pi.cnr.it

Systems Validation Centre shows Benefits of Formal Methods

by Henk Eertink, Wan Fokkink, Holger Hermanns and Izak van Langevelde

Embedded systems and network protocols have been formally verified by the Dutch Systems Validation Centre (SVC), using state-of-the-art techniques developed at international research institutes. This has led to substantial improvements in their design.

SVC was founded September 1998 as a cooperative effort of the Telematica Instituut, the Embedded Systems Group at CWI, and the Formal Methods and Tools Group at the University of Twente. Industrial support was provided by companies like CMG, IBM, KPN and Lucent. The central aim of SVC is to apply the fruits of fundamental research and tool development concerning system verification at CWI and the University of Twente to industrial cases handled by the Telematica Instituut. Thus, SVC contributes to building bridges between theory and practice, tearing down the walls of misconception about the applicability of formal methods.

This unique project builds on a strong foundation of theoretical research. On the one hand, μ CRL, developed at CWI, is a language for specifying and verifying distributed embedded systems in an algebraic fashion. It is based on the

classic theories of process algebra and abstract data types. On the other hand, the theory of interactive Markov chains (IMC), a conservative extension of stochastic models and, again, process algebra, forms the basis of a clean language to specify performance related properties of distributed embedded systems.

Directly resting on these theories is the development of tools. The μ CRL tool set supports the analysis and manipulation of μ CRL specifications. These specifications are automatically converted into a symbolic format, where a range of tools is available for optimisation, theorem proving, state space generation and model checking. Some of these tools belong to the CADP toolset developed at an ERCIM partner site at INRIA Rhone-Alpes. The same format is also supported by a novel method to specify stochastic models in terms of IMC.

Dedicated analysers for performance and dependability have been developed.

In the context of SVC, the use of these tools and techniques in the development of complex embedded systems and state-of-the-art telematics standards was assessed in a number of real-world case studies. The results of these were reflected back into foundational research and tool development, by posing new theoretical questions and further challenging the developers of automated tools.

One successful case is the verification of the IEEE P1394.1 Draft Standard for High-Performance Serial Bus Bridges, which aims at the standardisation of hot-pluggable 'bus bridges' as a flexible means to connect FireWire serial buses. The net update fragment of this standard, which ensures the unhampered continuation of network functionality throughout addition or removal of bridges, has grown towards a protocol which is too complex to be understood by mere human wit. This research project, in cooperation with researchers from Eindhoven University of Technology, already led to the detection of and solutions to many flaws in this standard. This project continues, side by side with the IEEE standardisation body.

Another case is the verification of an Erlang optimisation of the Transaction Capabilities Procedures of the Signalling System No. 7, a protocol for intelligent network services, in close cooperation with Ericsson. Both the original and the optimised version of the protocol were formally specified, in order to verify that the latter is a correct optimisation of the former. This effort revealed a number of bugs in the optimisation, the most noticeable of which resulted in a



SVC puzzles over the IEEE P1394.1 standard for high-performance serial bus bridges.

Illustration: I.A. van Langevelde, CWI.

memory leak which would have been hard to localise by more conventional means.

Together with Lucent Technologies, the stability and control of an SDH data communication network was analysed. The performance of an experimental network at Lucent was measured. Moreover, this network was formally modelled and analysed using both simulation and numerical methods. Although the main focus was on performance issues, the measurements revealed system traces that did not adhere to the specification of the intended network behaviour.

The experience gained within SVC demonstrates how formal methods have grown mature. The theoretical foundation has been proven a solid base for the construction of tool sets, the applicability of which is wide enough to span the gap between the theory of formal specification and model checking and the practice of modern embedded systems and telematics standards. SVC continues until January 2003, and has in the meanwhile stimulated various dedicated successor projects, including projects on formal testing, on performance and dependability analysis, and on improving the quality of embedded

systems and of real-time distributed shared data spaces.

Links:

<http://fmt.cs.utwente.nl/svc/>
<http://www.cwi.nl/sen2>
<http://www.inrialpes.fr/vasy>
<http://svc.telin.nl/>

Please contact:

Wan Fokkink, CWI
 Tel: +31 20 592 4104
 E-mail: wan@cwi.nl

Henk Eertink, Telematica Instituut,
 The Netherlands
 E-mail: Henk.Eertink@telin.nl

Holger Hermanns, University of Twente,
 The Netherlands
 E-mail: hermanns@cs.utwente.nl

REDEST — Sharing Experiences in the Field of Embedded Software Requirement Gathering

by János Nacsa

Requirement gathering is an initial but critical phase of any software development. In embedded systems it is even more difficult, since many of the requirements have effects that cannot be fulfilled with software elements. The aim of the REDEST project is to improve the requirement-gathering phase of the embedded software development process by introducing methodologies to help companies augment the quality of their embedded software.

The REDEST consortium consists of fourteen small and medium enterprises (SMEs) from Spain, Germany and Hungary that develop embedded software for their products, three coordinating institutions, and a technology expert company, to harmonise the different activities and provide the SMEs with continuous support.

Although all the SMEs in REDEST develop embedded software, there are significant differences between them:

- the company size ranges from ten to more than one hundred employees
- the software part of the product varies greatly from company to company
- some companies sell their products with minimal software modification for a certain order, while others develop one-of-a-kind software solutions
- while some companies have previously used requirement-gathering (RG) technologies or forms, others

have no previous systematic method for handling these

- the cultural aspects of the given application areas and the targeted markets of the companies are very different.

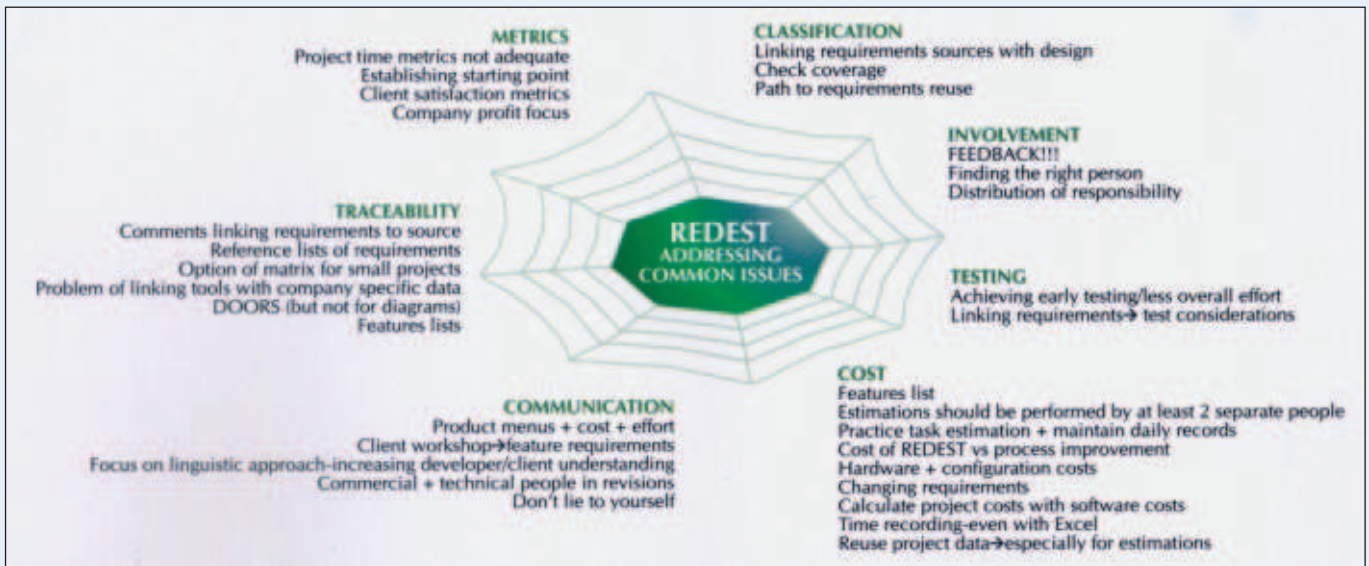
The SME partners in REDEST perform experimental tasks consisting of two parts: a baseline project on a development common to the given company, and — in conjunction with this — an innovative experiment in RG, with clear objectives and measurement facilities. Each experiment has its own independent strategic, business and technical objectives.

The coordinating/expert partners supervise the experiments and provide the local partners with advice. They also evaluate the results and elaborate the Best Practice cases for identifying and exploiting synergies within the experimental results. The independent method produced by each experiment is then

used to measure the resulting business and technical benefits, as explained in their descriptions.

Based on both the fourteen SMEs' experiment definition documents and the information gained from the technical meetings, many common problems were identified. To overcome these difficulties, the coordinators suggested the following actions as a minimum set for all REDEST members:

- courses on the new RG technique to be introduced in the company, and training materials to be prepared by the company (or the coordinator upon request)
- a template/form to be worked out according to the special needs of the given company to support its requirement gathering
- strategic indicators to be defined to allow the effects of the new RG techniques to be measured, and



Common requirement gathering problems of the project partners.

- simple tracking methods to be elaborated for the requirements within the whole development (eg in design documents, source codes and test patterns).

It is expected that a clear connection will be identified between the previous problems and the actions proposed.

The Hungarian Experiment

Although REDEST is about requirement gathering, it was necessary to look at the companies' complete development cycles. For some SMEs, the whole process was found to be rather unstructured, and many operational problems were related to this fact. An immediate positive effect of REDEST was that the management staff and the key developers of the three Hungarian SMEs had to rethink their conventional design and development methods. As a consequence, they realised that advanced software engineering methods (eg RUP) can be more effective than the ad hoc ones they had used previously. It was also recognised that the review itself of their usual tasks and the development of a structure for these have generated good new ideas.

The first Hungarian SME is a research and development company. They have high-tech products in the field of medicine, which use embedded software technology. Because of the nature of the

company, the requirement specification is not formal and they have no adequate solution for handling changes in the requirements. In the frame of REDEST, they work on the formalisation of their whole development cycle, from the requirement gathering to the final tests.

The second Hungarian SME has a relatively large market share in a special field of traffic control. As their products are mainly hardware-oriented, they employ only a few software specialists and thus depend greatly on the expertise of these people. In the future they intend to introduce a more standardised way of working, including requirements gathering, design and tests. They have problems with tracing the effects of the different requirements on the product life cycle, including the software version control during maintenance.

The third Hungarian SME produces special high-capacity, high-cost visualisation equipment in relatively small quantities. As they have very remote and culturally different markets, the RG must include communication with their agents, which needs to be more formal and precise. Requirement forms in this case means the ability to select a given product feature from among various options.

Examining the different problems of the SMEs in REDEST, the following conclusions were arrived at:

- since the problems vary between companies, it is not possible to suggest unique solutions for requirement gathering
- the companies are keen to learn from the experiences of others and within the framework of meetings and workshops are ready to share their views (including problems and possible solutions etc)
- it is worth managing and comparing the requirement-gathering problems of companies as user cases, meaning there is a need for user-specific help from the expert companies
- efforts should be concentrated on learning (to help to apply the new RG techniques), formalisation (to clarify and standardise the RG) and measurement (to prove to each stakeholder the benefits of RG).

REDEST is designed to be a growing modular structure, to which other sets of local supporting organisations and experiments, or even independent experiments, may be added. The project is sponsored by EC, under the contract no. IST 2000-29425.

Link:

<http://www.redest.net>
<http://www.sztaki.hu/~nacsa>

Please contact:

János Nacsa, SZTAKI
 Tel: +36 1 279 6224
 E-mail: nacsa@sztaki.hu

EmbLab — A Laboratory for Teaching Embedded Systems

by Amund Skavhaug, Trygve Lunheim and H. Skinnemoen

EmbLab is a new large-scale modern teaching laboratory designed to increase Norwegian IT-industry to compete in the future.

NTNU was quite recently created from two very different university systems in the same city, one traditional run-of-the-mill university for the 'soft' sciences, and the other being the technological university for the whole country, NTH. Much of the Norwegian industry has been fathered by students graduating from NTH, and NTH has always been the institution from where Norwegian industries have recruited their most qualified engineers. Due to this strong relationship with the industry, one can argue that at NTNU, there has traditionally always been a healthy activity related to the study and training for embedded and industrial computer systems.

In the design of embedded systems it is necessary to make a number of technologies play together. For a number of reasons, faculties are often hedging 'their' areas, extending them into overlapping areas, and as a result, they are often duplicating others' efforts and achieving less. NTNU is no exception here. To overcome such departmental boundaries we have tried to create a laboratory facility with enough capacity and with such a variety of SW/HW resources that it can be of interest for several departments, thus lessening the amount of redundant work, and easing the growth of this cross-boundary discipline. As everyone working with embedded systems knows: We need real-time and operating system theory, low-level software design and development, know-how for specialized electronics, design methodologies and so on.

Based initially on work in the Department of Engineering Cybernetics, a large effort has been made to collect and integrate all the necessary resources needed to be able to present a complete and up-to-date environment. Today we have a laboratory where 40-50 students can work concurrently with state

of the art tools and equipment for embedded and real-time systems.

The most important commercial players in the embedded- and real-time software world have made significant contributions. We see it as important that the graduate students are allowed to run the complete, 'professional' versions of any software, without limitations often associated with the 'student' or 'academic' versions. As an example, the QNX RTOS (version 6.2) is freely available for non-commercial use, but only in a limited version lacking support for several hardware platforms. Through the qnx-in-education program our embedded lab has been equipped with the latest version of the professional distribution. And just to name a few others:

Rational Software has contributed licenses and support for Rational Rose Real-Time. And as rapid development will become increasingly important also for the smallest embedded systems in the future, VisualState from IAR Systems, an UML based graphical code generator for small microcontrollers has been a welcome complement.

Wind River has contributed licenses for their RTOS VxWorks/AE and development environment Tornado, in almost any configuration possible regarding host/target systems. And for the Windows world: The most recent addition of sponsors is Microsoft, which have given all relevant tools and operating systems for practically unlimited use at the lab without any cost at all, while VenturCom has provided licenses for RTX, a real-time extension for Windows NT.

On the hardware side, Atmel Norway has made available a really generous amount of equipment for their AVR parts.

A fast growing Scandinavian company, DataRespons, who is a definite marked leader in commercial consulting, sales and development of embedded/real-time systems in Norway, has recently agreed to further supply new equipment for larger systems, where we today has been using PC/104 (or just old pc's).

It is important to have hands on experience with assignments that demonstrate the theory taught for example in real-time programming courses. The software-only simulation often used might be accurate, but not really seen as motivational by the students. Since, by definition, an embedded system is linked to a physical system, there really should be an external system connected to any computer system used for teaching topics in this domain.

With the aid of assigned persons from both an electrical- and a mechanical lab, we have designed an elevator model, of course duplicated on each working station, complete with 4 floors, buttons, sensors and lights mimicking a real elevator from both the inside and the outside. This has been a huge success when it comes to motivate student groups for practicing rapid development with state diagrams and code generation.

To run this software, we use both ordinary PCs as targets, but also micro-PCs, depending on the choice of OS. Due to design of interface logic that can be switched from 5v to 24v 'high' level, these models are also used with another very different type of embedded computers, namely the PLC. Common practice for industrial control is, of course, to use PLCs. Rugged, relatively low-cost and actually also easy to program with modern tools. We were therefore very pleased that Siemens donated a modern PLC with necessary IO and software for every seat in the lab.

Logic control may not be the most fashionable topic around, but invaluable for most industries.

A recent addition to our inventory of physical systems is a ping-pong game. To achieve the goal of controlling the game with a joystick, the students have to build a control system consisting of two microcontrollers connected together via CANbus. The complete system also features A/D and D/A conversion, Human-Computer Interaction through an LCD-display, 7-segment LEDs and a matrix-keyboard, optical sensors and

various inter-IC buses (memory-mapping, SPI and I2C). A third microcontroller-based node with a PC/104 interface is planned, the goal being to demonstrate how to integrate microcontroller-based embedded systems with (micro-)PC systems.

Most educational institutions have some sort of laboratories where aspects of real-time and/or embedded systems development are studied. In my opinion, it seems that these are quite often limited in both functionality and size. We believe that the possibility of handling more than 200

students a week, each receiving a full day of work, together with the extent of available modern SW and HW is quite unusual. The comments from Norwegian industry has been very encouraging, and signals further growth in quality, size and importance.

Link:

<http://www.itk.ntnu.no/lab/embedded>

Please contact

Amund Skavhaug, NTNU

Tel: +47 7359 4396

E-mail: Amund.Skavhaug@itk.ntnu.no

Master Course Embedded Systems at Eindhoven University of Technology

by Jan Friso Groote

The area around Eindhoven in the Netherlands houses one of the largest concentrations of embedded systems industry in Europe. Therefore, Eindhoven University of Technology decided to offer an international master programme especially focused on the design and construction of embedded systems. The master course will start in September 2003.

Higher education in the Netherlands is quickly transformed according to the guidelines of the Bologna convention. All studies used to take five years. Now these are split into a bachelor phase of three years and a master phase of two years. Especially with the master phase this allows for more focused education. Therefore, instead of having a general master in computer science, the division of computer science offers four master programmes: Computer Science and Engineering (CSE), Embedded Systems (ES), Business Information Systems (BIS), and Information Security Technology (IST).

Embedded Systems will be offered by the division of Computer Science and the faculty of Electrical Engineering. Computer Science in Eindhoven is strong in program design and system architecting, semantics and analysis. Electronics has its strengths in system theory, optotronics, and IC-design. These are rather disjoint fields. And it was observed that students from one of the faculties had difficulties communi-

cating to students in electrical engineering.

Clearly, as was also clearly indicated by several industries in and around Eindhoven, this was not a desirable situation. The embedded systems industry is in need of integrators. Engineers that can translate high level requirements into a high-level architectural system design such that it can be realised with appropriate electronic technology. This is exactly the purpose of the master programme Embedded Systems.

The study takes two years and has three major parts of equal size, each consisting of 60 ects-credits. One part consists of mandatory courses. These comprise courses on software architecting, requirement analysis, software testing, performance modelling, multiprocessor systems, platforms, IC-design and the societal consequence of embedded system technology. The second part of the master programme consists of elective courses. These include an internship and if needed, elementary courses in case a student did enrol with almost

sufficient basic knowledge. The third part of the master programme consists of a master thesis that can be performed at an embedded industry or internally at the university.

The whole idea of the Bologna convention is that students throughout Europe can select those places where they will receive the most appropriate education. It is explicitly intended that people with appropriate bachelor degrees will be able to attend our master programme. Therefore, the courses are offered in English. Admission to this master programme requires a bachelor degree in Electronics, Computer Science and Engineering, or a comparable degree.

Please contact:

Desiree Meijers

Tel: +31 40 247 3902

E-mail: d.w.t.meijers@tue.nl

Jan Friso Groote

Eindhoven University of Technology,
Director of education in Embedded Systems

Tel: +31 40 247 5010

E-mail: J.F.Groote@tue.nl

<http://www.win.tue.nl/~jfg>



Articles in this Section

- 45 SOMLib - New Approaches for Information Presentation and Handling**
by Andreas Rauber, Vienna University of Technology
- 46 Hybrid Caching of Search Engine Results**
by Tiziano Fagni and Fabrizio Silvestri, ISTI-CNR
- 48 Personal Space Station handles Virtual 3D Objects**
by Jurriaan Mulder, CWI
- 49 Augmented Virtual Table Tennis with Real Racquets**
by Charles Woodward and Petri Honkamaa, VTT
- 51 Scene Understanding using Hierarchical Neural Networks**
by Gabriele Pieri and Ovidio Salvetti, ISTI-CNR
- 52 A Digital Film-Saver for Archives**
by George L. Kovács and Iván Kas, SZTAKI
- 54 Knowledge Technologies in Support of e-Government and e-Democracy**
by Jan Paralic and Tomas Sabol, Technical University of Kosice (SRCIM)
- 55 Rapid Prototyping for Embedded Parallelism**
by David Johnston, Martin Fleury and Andy Downton, University of Essex, UK
- 57 Discretizing Continuous Attributes in AdaBoost for Text Categorization**
by Pio Nardiello, Fabrizio Sebastiani, and Alessandro Sperduti, ISTI-CNR

SOMLib - New Approaches for Information Presentation and Handling

by Andreas Rauber

While tools exist that allow us to search through vast amounts of text within seconds, most systems fail to assist the user in getting an overview of the information available or maintaining orientation within an information space, and fail also to convey meta-information in an intuitively graspable way. SOMLib is a digital library system addressing these issues by providing automatic content-based organisation and metaphor-graphics-based visualisation facilitating exploration and understanding of information spaces.

With the increasing availability of information in electronic form, be it online magazines, legal or medical document databases, project archives, or documentation on a company-internal intranet, advanced digital library systems that support users in interacting with large information repositories are gaining in importance. Yet, while databases and search engines help us in retrieving snippets of information, current tools fail to provide us with a feeling of ‘where’ information is available, and how different facts relate to each other. The ability to keep an overview of factors such as the information available, the topics covered by a given site and the amount of information available on a given topic is only poorly supported. In addition to the powerful search methods offered by modern information systems, it seems difficult to provide equally powerful means of organising and structuring the information.

What we would like are ways of information organisation and representation that allow us to make use of the concepts that we are using constantly, unconsciously, when handling and navigating

real-world information spaces. Libraries, bookstores, project documentation in binders, working material and paper collections are all conventionally organised (also) by thematic criteria. This allows us to immediately get an overview of which kind of information is available in which section of an archive, how many reports have been filed on a specific topic in a binder, and so on. Due to the spatial location it is also easier to find a paper, report etc for the second time, as it is easier to recall roughly where a given document was located than to remember sufficiently precise search criteria, or its relative position within a listing.

With the SOMLib digital library, we created a system providing content-based organisation of document repositories, facilitating intuitive browsing and exploration of the information space. It builds on and incorporates works in the fields of information retrieval, neural networks, information visualisation, and usability analysis.

Low-level features based on word frequencies are extracted from the text to

provide a domain- and language-independent content representation of text documents in a high-dimensional vector space. The ‘self-organising map’ (SOM), a popular unsupervised neural network model, is further used to cluster the document feature vectors, performing a topology-preserving mapping of the documents from the high-dimensional vector space onto a two-dimensional map space. Documents are thus grouped according to their mutual similarity, having documents on similar topics mapped onto neighbouring locations on the map. Using the ‘growing hierarchical self-organising map’ (GHSOM), a novel extension of the SOM, we can further detect subject hierarchies in a document collection, with the neural network adapting its size and structure automatically during its unsupervised training process to reflect the topical hierarchy. Individual SOMs can further be integrated to form a network of referencing maps.

By mining the weight vector structure of the trained maps using the ‘LabelSOM’ technique, the system automatically extracts keywords describing the various

The SOMLib system: from text collections via content-based organisation, to metaphor-graphics-based representation of document repositories facilitating intuitive browsing and exploration.

File	Last Modified	Size
file1.txt	24-Oct-2000 24:52	-
file2.txt	24-Oct-2000 24:53	10
file3.txt	24-Oct-2000 24:53	10
file4.txt	24-Oct-2000 24:52	10
file5.txt	24-Oct-2000 24:53	10
file6.txt	24-Oct-2000 24:53	10
file7.txt	24-Oct-2000 24:53	10
file8.txt	24-Oct-2000 24:53	10
file9.txt	24-Oct-2000 24:53	10
file10.txt	24-Oct-2000 24:52	10
file11.txt	24-Oct-2000 24:52	10
file12.txt	24-Oct-2000 24:52	10
file13.txt	24-Oct-2000 24:52	10
file14.txt	24-Oct-2000 24:52	10
file15.txt	24-Oct-2000 24:52	10
file16.txt	24-Oct-2000 24:52	10
file17.txt	24-Oct-2000 24:52	10
file18.txt	24-Oct-2000 24:52	10
file19.txt	24-Oct-2000 24:52	10
file20.txt	24-Oct-2000 24:52	10



topical clusters. This is based on the analysis of the feature distributions within each cluster. It helps users in identifying which topics are present in a given document collection and where they are located on the map.

Finally, the 'libViewer' provides an intuitive representation of the documents in a repository by using real-world metaphors such as different document types, spine widths, dust etc to convey metadata in an intuitively graspable way.

Using the 'SOMLib' system, users can browse a document collection in the form of bookshelves and find clusters of documents on similar topics located in neighbouring boxes, with the topic of

each box being described by a set of automatically extracted keywords, and metadata being depicted in the form of different document representations. In combination with conventional approaches for searching and dynamically sorting text archives we thus have a powerful tool at our disposal, which allows us to obtain and maintain an overview of the amount and type of information available, and to detect relationships between different documents. This means we can better handle, interact with, and use the available information.

The 'SOMLib' system has been applied in numerous different domains in a variety of languages, such as the organisation of legal databases, newspaper

archives, scientific document collections and Web search results. Recently, the principles of this work have been expanded for use in digital music archives as part of the 'SOM-enhanced Jukebox' (SOMeJB) system. By analysing frequency spectra of audio files and transforming them to time-invariant representations while incorporating psycho-acoustic models, organisation and exploration following musical genres is facilitated.

Andreas Rauber is currently an ERCIM Research Fellow at INRIA

Please contact:

Andreas Rauber
Vienna University of Technology
E-mail: rauber@ifs.tuwien.ac.at
<http://www.ifs.tuwien.ac.at/~andi>

Hybrid Caching of Search Engine Results

by Tiziano Fagni and Fabrizio Silvestri

The High Performance Computing Lab of ISTI-CNR in Pisa, has developed an efficient caching system aimed at exploiting the locality present in the queries submitted to a Web search engine. The cache adopts a novel hybrid strategy, according to which the results of the most frequently submitted queries are maintained in a static cache of fixed size, and only the queries that cannot be satisfied by the static cache compete for the use of a dynamic cache.

Caching is a very effective technique to make a service that distributes data/information to a multitude of clients scalable. As suggested by many researchers, caching can be used to improve the efficiency of a Web Search Engine (WSE). This is motivated by the high locality present in the stream of queries processed by a WSE, and by the relatively infrequent updates of WSE indexes that allow us to think of them as mostly read-only data.

The architecture of a modern WSE may be very complex and includes several machines, since a query may require various sub-tasks to be executed. For example, a large and scalable WSE, which is typically placed behind an http server, is usually composed of several searcher modules, each of which preferably runs on a distinct machine and is responsible for searching the index relative to one specific sub-collection of

documents. Of course, in front of these searcher machines there is a mediator/broker, which collects and reorders the results of the various Searchers, and produces a ranked vector of the most relevant document identifiers (DocIDs), eg a vector usually composed by 10 DocIDs. These DocIDs are then used to get the associated URLs and page snippets included in the html page returned to the user through the http server (see Figure 1).

On the basis of an accurate analysis of the locality present in three large query logs, we designed and implemented a novel hybrid caching strategy where the results of the most frequently accessed queries are maintained in a static cache of fixed size, which is completely rebuilt at fixed time intervals. Only the queries that cannot be satisfied by the static cache compete for the use of a dynamic cache.

The superiority of our hybrid strategy over purely static or dynamic caching policies was demonstrated by several experimental tests, conducted to measure the cache hit-rate as a function of the size of the cache, the percentage of static cache entries (f_static), and the replacement policy (LRU, LRU2, LRU2S, 2Q, FBR) used for managing dynamic cache entries. Figure 2 shows the cache hit rate obtained on the largest of our query logs by using different replacement policies as a function of the ratio between static and dynamic cache entries. As it can be seen, our hybrid caching strategy always outperforms purely static and dynamic policies.

Moreover, we showed that WSE query logs exhibit not only temporal locality, but also a limited spatial locality, due to the presence of requests for subsequent pages of results. While most user searches are satisfied by the first page of

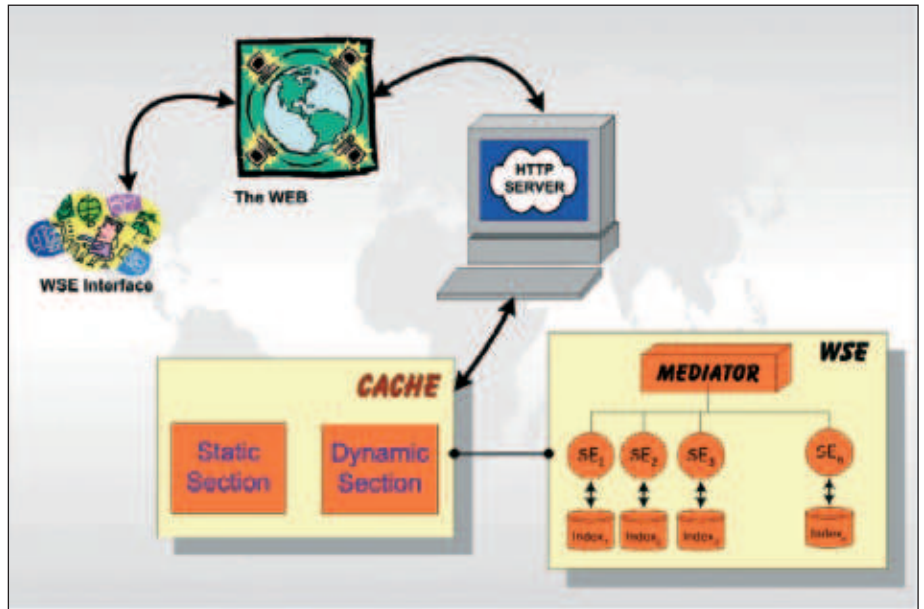


Figure 1:
The caching system architecture and its positioning within a typical distributed WSE environment.

results returned by the WSE, about 10% of searches are less focused and require additional pages. Our caching system exploits this user behavior by anticipating the requests for the subsequent pages by means of an adaptive prefetching heuristic that improves the hit rate achieved and, at the same time, limits the additional loads on the core query service of the WSE. Figure 3 shows the improvement on the cache hit rate for different prefetching factors as a function of the ratio between static and dynamic cache entries. While caching reduces the load over the core query service of the WSE and improves its throughput, prefetching aims to increase the cache hit rate and thus the responsiveness of the WSE.

Finally, differently from other work in this field, we evaluated cost and scalability of our cache implementation when executed in a multi-threaded environment. Our hybrid cache implementation resulted very efficient due to the accurate software design, and to the presence of the read-only static cache that reduces the synchronizations between the multiple threads concurrently accessing the cache (see Figure 4).

Please contact:
Fabrizio Silvestri, ISTI-CNR
Tel: +39 050 315 3011
E-mail: fabrizio.silvestri@guest.cnuce.cnr.it

Figure 2:
Hit rate obtained on a large query log by using different replacement policies as a function of the ratio between static and dynamic cache entries.




Figure 3:
Hit rate for different prefetching factors as a function of the ratio between static and dynamic cache entries with adaptive prefetching.


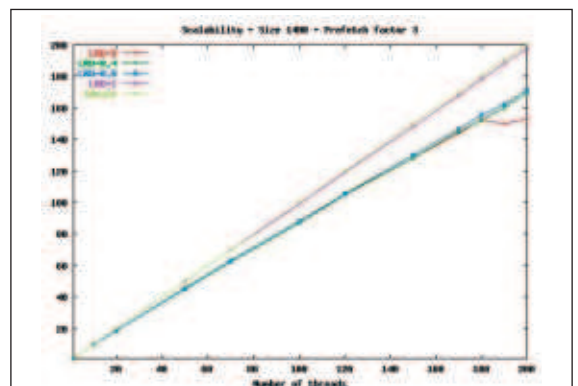


Figure 4:
Scalability of the caching system as a function of the number of concurrent threads used.



Personal Space Station handles Virtual 3D Objects

by Jurriaan Mulder

CWI developed the Personal Space Station (PSS), an affordable construction for 3D interaction and visualization. Key questions were: how can we realize 3D-interaction and navigation in a virtual reality environment and how can we solve visual perception problems?

Virtual reality offers a user the possibility to look at and manipulate a virtual three-dimensional world, generated by a computer. Besides obvious applications in the amusement industry, virtual reality can also contribute considerably to (scientific) research or product development in various areas of application, like biomedical modelling.

An example of a virtual reality environment is the 'Personal Space Station' (PSS) that is developed at CWI. This environment enables the user to interact with the virtual world in a direct and intuitive way. The environments used so far were often indirect, hard to operate and they lack precision which results in lesser performance, discomfort and weariness for the user. During the devel-

opment of the Personal Space Station special attention was paid to aspects like image quality, ergonomic use, natural interaction, multi-user cooperation and small initial expense.

The Personal Space Station is built of standard components, entailing minimal costs. The user looks at the monitor via a mirror. This enables the user to bring his hands into the same environment as the virtual 3D-objects without interrupting the visual image. Thus, interaction with virtual objects can take place in a direct, natural and intuitive way. Hand-eye coordination and 'proprioception' (the subconscious perception of the position of muscles and joints) are used to the full, which has a favourable effect on user convenience and task performance.

The user is 'attracted' by the objects he sees; his natural reaction is usually to grab and manipulate them.

To facilitate wireless tracking of the interaction a custom optical tracking system was developed. A user has interaction devices in his hands, like a marked pen, cube or thimble. Using infrared lighting, two cameras with infrared-pass filters, and retro-reflective markers these interaction devices can easily be constructed and applied. Device recognition is based on projection-invariant pattern characteristics. Reconstruction is realized through stereo correspondence and so-called epipolar geometry. Once it is determined which point in the left camera image corresponds with which point in the right camera image, and the internal and external camera parameters are known, the position of the matching marker in the 3D environment can easily be determined.




Figure 1: Scientific visualization with the Personal Space Station, developed at CWI.




Figure 2: Depth of field simulation.

A unique aspect of the Personal Space Station is that it enables more people to look at and manipulate the same virtual world simultaneously. By cascading several Personal Space Stations a virtual reality environment is created in which multiple users share the same physical and virtual workspace. This is of the utmost importance for applications where several persons want to cooperate, communicate and interact. For the communication and synchronization of the connected stations, a software model was developed based on a 'publish and subscribe' paradigm. Through a central 'Data Manager' the important simulation, tracker, and synchronization parameters are communicated.

Personal Space Stations will be placed at various universities and research institutes, like the Swammerdam Institute for Life Sciences (SILS) and the departments for Technology Management (former IPO), Biomedical Technology (BMT) and Mathematics and Computer Science of the Technische Universiteit Eindhoven (Eindhoven Technical University). These Personal Space Stations will be applied in different branches of scientific research, like cell biology, biomedical image processing, man-machine interaction and scientific visualization.

Apart from the development of the Personal Space Station, research is done

on 3D computer graphics and 3D interaction techniques for virtual reality. In the field of 3D computer graphics, for instance, new techniques were developed to draw transparent objects (by way of 'screen-door' transparency based on generated pixel masks), to simulate depth of field (based on perception models and Gaussian pyramids) and to improve depth perception (through elimination of faulty depth cues for objects on the screen edges).

Link:

<http://www.cwi.nl/ins3/>

Please contact:

Jurriaan Mulder, CWI

Tel: +31 (20) 592 4317

E-mail: Jurriaan.Mulder@cwi.nl

Augmented Virtual Table Tennis with Real Racquets

by Charles Woodward and Petri Honkamaa

A group of researchers at VTT have developed a computer system for natural interaction in an augmented virtual environment, enabling people to play table tennis over Internet/LAN with real racquets. Apart from the web cameras, no special hardware is required. The players can see each other in the camera image, which is streamed in real time compressed over the network. The position of the racquets is computed by marker detection from the image. The multicast implementation also enables a network audience to view the game.

Among the various fields of computer graphics, real time marker detection from a camera-produced video stream is most closely related to Augmented Reality (AR). In AR applications, virtual objects are superimposed on the user's view of the real world, which is displayed, for example, through a computer monitor or HMD devices.

A popular application of AR technology is collaborative gaming. For example, AR table tennis for two people has been previously implemented using magnetic trackers and sharing the same computer processing unit. In another case, a virtual reality (VR) table tennis game has been implemented over a computer network, but it does not involve video image

augmenting and is based on the traditional mouse interface.

Our solution combines the virtues of these different approaches for an augmented virtual table tennis game in a computer network system. Natural user interaction is accomplished by real-time marker detection from the camera image. Thus two people at remote locations can play virtual table tennis against each other using real racquets, with no special hardware required other than the web cameras.




Figure 1: Computer system and virtual objects for augmented virtual table tennis.

Figure 1 shows a schematic drawing of the computer system for the game. The hardware for each player consists of a PC workstation and a web camera. The PCs are connected by Internet or LAN network. The cameras are aimed towards the players, who hold in their hands real




Figure 2: Player's view of the game.




Figure 3: Audience view.

table tennis racquets. Each racquet has a marker, which enables the racquet's position (distance and angle) to be computed from the video image.

A computer program is synchronised between the players, defining a shared virtual ball, table and net. When the racquet's location meets the virtual ball within given distance, a hit is registered, the new ball trajectory is computed, and the ball parameters (time of impact, new trajectory) are sent via multicast address to the other player. Additionally, hits and bounces are accented with sound effects.

Marker Detection

The marker is a coloured rectangle of given size (10x10 cm), glued to both sides of the racquet. Green was chosen as the marker colour, as it occurs less frequently than other basic colours in room interiors or on clothing. Depending on lighting conditions, the image brightness and saturation can be adjusted to preserve bright green colours.

Marker corner points are located from the video image by fitting two diagonals to the detected marker image, and refined by fitting four lines to the marker edges. The marker position in the virtual object's coordinate system is then calculated using laws of weak perspective projection. A simple camera calibration routine is performed when installing the system.

Player's View

Figure 2 shows a screenshot of the player's view of the game. Each player

sees on the computer monitor a virtual racquet corresponding to the position of his or her real racquet. The video image of each player is compressed in real time, streamed continuously to the other player, who sees their opponent at the other side of the table. VTT's proprietary MVQ (Motion Vector Quantization) software is used for efficient video compression and streaming.

Solo Game

A simple solo game mode is provided for playing against a 'wall'. In the more interesting version, a player can actually play against him/herself, as the video image of the player is shown at the other end of the table (see Figure 2 again). But which side should then hit the ball, the 'virtual me' (the racquet), or the 'video me' (the image at the other end of the table)? Our solution is to have both sides hit the ball at the same time. Consequently, the game contains two balls, which the user has to hit twice as frequently as normally.

Network Audience

A network audience option can facilitate, for example, the viewing of game tournaments. Thus the video stream and ball parameters from both players are multicast to people in the network who have joined to watch the game. The audience members see on their screens the same virtual/video game content as the players, with the ball trajectory as an added component. The audience may also rotate the view to different sides and angles (see Figure 3).

Simplifications

We have not attempted to implement a full correspondence with the real game, and we do not account for air resistance, friction etc. Adding more physical realism would be possible, but may not be worth the effort. For one thing, the inevitable imprecision in marker pose detection sets some limitations in reaching full realism. In our experience, the game can actually be more fun with some simplifications. Thus, to keep the ball 'live' longer than normal, the ball is directed slightly towards the table, and hits are registered at a larger diameter than that of the actual racquet. Slow motion balls can also be useful for practice, and they provide a nice exercise tool for elderly and handicapped people.

Future Work

The game is available for free download from our Web pages <http://www.vtt.fi/multimedia>. Ongoing and future work includes detecting a serve by gesture, a foursome game using different marker types, an audio component for chat between players, and implementing the game for video HMD glasses with stereo virtual objects. The table tennis implementation described here can also be generalised to a wide variety of other augmented virtual game concepts.

Link:

<http://www.vtt.fi/multimedia/>

Please contact:

Charles Woodward,
VTT Information Technology, Finland
Tel: +358 9 4565629
E-mail: charles.woodward@vtt.fi

Scene Understanding using Hierarchical Neural Networks

by Gabriele Pieri and Ovidio Salvetti

Scientists at ISTI-CNR are using Hierarchical Neural Networks in order to develop a methodology for the automatic identification of characteristic patterns — representative of particular phenomena — in a given scene. This approach can be adopted in a wide range of applications.

When changes to a dynamic or 3-D scene represented by sequences of two-dimensional images occur with some frequency, a system which learns and adapts to these changes can be developed using neural networks (NN). Each image of the sequence is represented by a set of morphological structures with textural properties, and every image is analysed in order to compute a set of characteristic features, that can then be used to understand the scene. In order to do this, we are adopting a strategy that uses Hierarchical Neural Networks (HNNs), which are global networks composed of a hierarchy of single neural networks.

The hierarchical approach guarantees both specialisation and adaptability at the same time, since the single levels of the network can be finely tuned on the specific characteristics of the problem to be solved and the entire architecture can be easily modified if the problem description changes, by simply training only those levels involved in the modification.

The advantages of using an NN-based hierarchical architecture can be summarised in two main points: the modular organisation facilitates analysis of the networks and the hierarchy enables the network to finely tune itself towards recognising the most promising directions to look for relevant patterns.

Important requirements to take into account when using this approach are the possibility to exploit differences between extracted features to improve the classification capability and to change easily the number of features. The hierarchical NN architecture has been studied to meet these requirements.

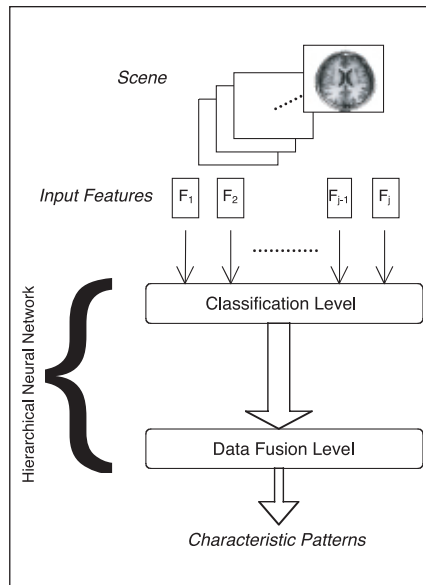


Figure 1: Scene understanding using a hierarchical neural network architecture.




Figure 2: Example of 3D neuro-image density analysis for cerebral anomalies diagnosis. Eight density classes are classified: bone, white matter, grey matter, liquor, air, blood, hypo- and hyper-density.

Scene understanding can be achieved by combining K parallel neural networks, each trained to extract a specific property class from the source images. The use of a set of parallel-specialised networks instead of a single complex net is appropriate to implement efficiency and flexibility.

Furthermore, this model also makes it possible to optimise the computational complexity of the single levels independently. In particular, the addition of a new feature has only a partial effect on a single part of the HNN architecture.

In this context, scene understanding mainly consists of two phases (see Figure 1): classification of the single features extracted from each image used to define a scene (first and lower level); data fusion for a more complete understanding of the characteristic patterns (second and higher level).

The global network does have to be homogeneous, since its sub-nets can differ for typology and topology in order to face complex problems with high flexibility.

A large number of different application fields can be modelled using an approach of this type. In our study, we have applied HNN architectures to a number of problems ranging from real-time monitoring of the oscillation states in the flame front of gas combustors in power plants to the three-dimensional classification of brain tissue densities for the diagnosis and follow-up of cerebral anomalies (see Figure 2).

Please contact :
Ovidio Salvetti, ISTI-CNR
Tel: +39 050 3153124
E-mail: o.salvetti@iei.pi.cnr.it

A Digital Film-Saver for Archives

by George L. Kovács and Iván Kas

Pressure is growing fast for film archives to restore and save their priceless collections before the original material has decayed beyond repair. In the framework of the Hungarian National Research and Development Programs, SZTAKI leads a project consortium that will develop a complete and affordable digital film-to-film restoration system for archives that is not currently available on the market.

According to UNESCO, there are 2.2 billion metres of 35mm movie film in archives all over the world. In spite of the huge efforts of special, air-conditioned storage spaces, most of these films will soon be completely useless due to their extreme sensitivity to various physical and chemical effects and their current level of degradation. The amount of film to be saved in Hungary is about 60 million metres. Ninety percent of these are sound films, two-thirds are positive, one-third are negative and about fifty percent were taken in Hungary. These films are part of the national heritage and as such must be restored and saved.

Similar efforts exist worldwide, with some results having been achieved in the USA, Japan and Europe, including Kodak's CINEON system and the FRAME system developed by a European project consortium. However, complete systems are not available on the market, and parts such as scanners and film-writers are extremely expensive.

In Hungary, limited results were achieved by physical and chemical cleaning, rewinding and making one-to-one film copies, but in general these procedures finally 'kill' the 'wounded' originals. The reasons for commencing such an R&D effort in Hungary included the fact that earlier we did not have the appropriate mathematical/software knowledge for complete restoration, and the necessary high speed/high memory capacity computational tools have become available only recently. Furthermore, the success of the project can be guaranteed by our strong, well-balanced consortium (Hungarian National Film Archives, our academic research institute, the Veszprém

University and the Cortex Ltd. – an SME), which has experience and knowledge in films, project management, laser technique, hardware and software making and in picture-, sound- and colour-processing.

Goals, Expectations and Preliminary Results

The project started in mid-2001 and will last for three years. Its goal is to produce an intelligent and powerful workstation with dedicated peripherals to restore and save - in real time - archive films damaged in their physical and/or informational condition. We began to create

results however, all media will be used in our work.

To demonstrate the difficulties inherent in this project, we list some problems and the planned solutions.

One of these is the issue of film resolution. Recent TV screen resolution (about 1K) is insufficient, and while 2K or 4K are accepted for electronic storage as there is no better solution, filmmakers would prefer a resolution of 16K, which is not yet a realistic goal. Our basic aim is 6K resolution, as we are convinced that this is good enough to preserve all

Figure 1:
A typical worn-out film-frame with damaged perforation.



programs with mathematical methods – image- and sound-processing algorithms that will also be useful in other fields (eg medicine, biology, chemistry etc).

Movie picture experts do not accept electronic storage such as video, CD, DVD, HDVD etc, as final media for saved or restored films. As such, our solution should be based on film material. One of the reasons for this is that electronic media may have a limited lifetime, and it is obviously desirable to avoid compatibility problems wherever possible. As intermediate and/or mass production

the picture/colour parameters of an original film. To provide this resolution, the storage of one hour of film material would require between 10.5 and 74.5 TBytes capacity (see Table). It is also evident that in the case of a resolution of 6K with the given colour densities, one metre of film needs 6.4 Gbytes minimal background memory, which leads to approximately 2 Tbytes for 300 metres. The minimal background memory capacity is defined by the fact that film professionals do not accept anything less than 300m of film as a basic unit. This value – 2 Tbytes – is still a minimal

RESOLUTION	MASS STORAGE SIZE REQUEST				
	1 frame [MB]	1 sec. [MB]	1 meter [GB]	1 min. [GB]	1 hour [TB]
2K	9	216	0,5	13	0,8
4K	36	864	1,9	52	3,1
6K	122	3 000	6,4	176	10,5
16K	864	21 000	40,8	1 242	74,5

Table: Some resolution/mass storage capacity request data.



Figure 2:
The experimental set-up of the film-saver with scanner and sound correction (in the middle) and film-writer (completely covered on the left-hand side).



Figure 3:
Experimental result with colour and contour corrections.

value, and is more or less the maximum available today.

Another critical issue is the speed of the film restoration process. The processing time of a frame means the input, output and total processing time of a frame, plus the time required for data transfers within the system. Our plan is to bring this down to 60 seconds or less. The best result so far is 72 seconds, achieved with the FRAME system developed by a European project consortium led by Joanneum Research. Using this speed, one system is able to restore ten movies per year.

We expect to compensate for missing or damaged frames and to correct scratches and patches on the films. We will be able to compensate for shakes or vibrations due to camera movements, which affect the viewing quality and the value of the film. We expect to produce a film that is almost as good as the original could have been in every sense, including richness of detail, colour, contrast etc. We plan to correct and compensate for all dirt, patches and perforations due to bad handling and chemical erosion of the

film. Earlier, these could only be partially corrected by dangerous chemical processes, which have several negative side effects. Emulsion problems, copying errors, wet materials effects etc, will also be corrected and eliminated with the digital processing.

The basic architecture of the system is rather simple. An input unit (scanner: high resolution line-camera) reads the film to be improved, a processing unit (high-capacity workstation with appropriate software to meet the above expectations) performs the restoration job and an output unit (film-writer based on a high-speed rotating mirror with modulated RGB laser beams) produces the new film. A terabyte-sized background memory (discs) stores data and is used as the basis for electronic read/write procedures, and a high-speed data network makes communication possible. Since we deal with sound restoration as well, sound management is symmetrical to picture management, but it works completely differently.

Since the project has been running for over a year, several results have already

been achieved which prove that we are on the right track. All basic error types can be managed with the experimental software packages, and the software system plan is ready. All basic hardware units have been designed, several experimental parts are working already, and other parts are currently being constructed.

The consortium is open for a broader cooperation in the future.

Please contact:

George L. Kovács, SZTAKI
Tel: +36 1 279 6140
E-mail: gyorgy.kovacs@sztaki.hu

Iván Kas, SZTAKI
Tel: +36 1 279 6270
E-mail: kas@luna.ikk.sztaki.hu

Knowledge Technologies in Support of e-Government and e-Democracy

by Jan Paralic and Tomas Sabol

'Webocrat' is a Web-based system developed in the framework of the EC-funded R&D project Webocracy. This system is at the heart of the project that aims to establish efficient systems that provide effective and secure user-friendly tools, working methods and support mechanisms to ensure the efficient exchange of information between citizens and public administration institutions.

The Webocrat system comprises modules and functions for content management, computer-mediated discussion, discussion forums, organisational memory, information retrieval, data mining, and knowledge modelling. The system will support communication and discussion, publication of documents on the Internet, browsing and navigation, opinion polling on questions of public interest, intelligent retrieval, user alerts, and convenient access to

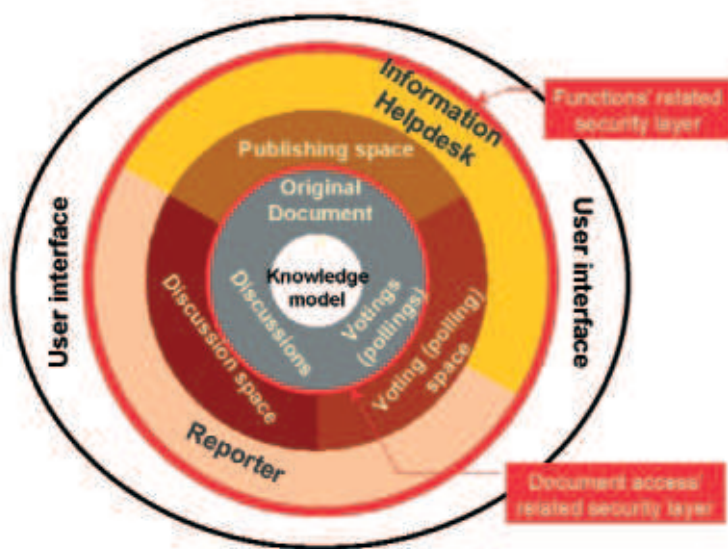
ties (user partners) involved in the project. The project objectives are refined in order to reflect user partners' needs.

Based on the results of user needs analysis, subsequent analysis, design, and implementation of the Webocrat system have been done. The Webocrat system consists of the following modules (see Figure):

including all documents and guidelines for competitive tendering

- Opinion Polling Room, enabling electronic polling on issues or questions of public interest published and discussed within the Webocrat system
- Citizens' Information Helpdesk, enabling users to query the system through a user-friendly interface; the system will retrieve those documents that are relevant to the user query
- Reporter Module, providing the means for generating different statistics and user alerts, and summarising information contained in published documents, submitted contributions and recorded opinions
- Knowledge Module (the core of it is an ontology-based knowledge model), supporting management of the whole system, intelligent retrieval and personalisation.

The project's objectives are being assessed and further refined within the framework of pilot applications by the deployment and evaluation of the developed technologies and methodologies in naturalistic settings. This should ensure that the tools and methodologies delivered by the project will correspond to citizens' interests and needs, and will have the potential for further development.



Webocrat functional scheme.

information based on individual needs. The development of the system will be complemented by the development of a methodological framework for implementing Webocrat-like systems.

The approach of the project partners to implementation of the project objectives is 'user-driven' and 'user-centred'. All work within the project is based on the user requirements of three local authori-

- Discussion Forum, supporting intelligent communication of citizens, elected representatives and public servants in several areas
- Web Content Management Module, for publication of various document types of documents that are of interest to citizens and different interest groups; additionally, there will be Electronic Public Procurement, supporting the publication of tenders

Technology provided by the Webocrat system is tested within three pilot applications. Three public administration institutions – the metropolitan council from the UK (WCC) and two local authorities from Kosice, Slovakia (LATA and LAFU) – play the role of user partners and testers of the Webocrat system.

Three modules of the Webocrat system were available for testing during the first

Trial (May-July 2002). These were the Discussion Forum (tested by WCC) and its specific part, the Communication Module (tested by LAFU), the Web Content Management Module (tested by LATA) and the Opinion Polling Module (tested by all three user partners).

Briefly, it can be concluded that the functionality provided by the above-mentioned modules proved interesting and quite attractive for citizens. It is worth emphasising that this is only a minor part of the functionality that will

offered by the whole Webocrat system, which will integrate all the above-mentioned modules.

Hence, even the small number of available and tested Webocrat modules have the potential to increase citizen participation in local democracy and political processes in Wolverhampton and Kosice. On the other hand, issues with design, functionality, administration, support structures, targeting and reliability will need to be addressed before moving forward and beginning to deliver

significant benefits in the future as part of a meaningful, valued and integrated consultation and participation process.

Link:

<http://esprit.ekf.tuke.sk/webocracy/>

Please contact:

Jan Paralic
 Technical University of Kosice, Slovakia
 Tel: +421 55 602 4128
 E-mail: Jan.Paralic@tuke.sk

Tomas Sabol (Project co-ordinator)
 Technical University of Kosice, Slovakia
 Tel: +421 55 602 3259
 E-mail: Tomas.Sabol@tuke.sk

Rapid Prototyping for Embedded Parallelism

by David Johnston, Martin Fleury and Andy Downton

While an embedded system has a more specialised role to play than a PC or workstation, the diversity of embedded applications is no less wide and their requirements are no less demanding. The RaPPID project (Rapid Parallel Prototyping in the Image/Multimedia Domain) is developing a methodology to rapidly prototype high-bandwidth, compute-intensive embedded applications which are amenable to parallelisation. Example application areas are radar and image processing.

The approach is based on software components, whereby the software overhead of handling parallelism; of interfacing to proprietary parallel hardware; and of performance monitoring and tuning are abstracted into application-oriented (cf. system-oriented) 'classes'. Parallelisation has traditionally been a dead-end route for application software, as parallel hardware becomes redundant and particular parallel languages became less favoured or unsupported.

Instead the aim of the RaPPID project is to support the parallel execution of applications, that are expressed in a traditional sequential form, such as C++. The choice of the RaPPID classes selected and the (composable) manner in which they are used, gives clues to a 'harness' (the underlying software framework) as to the parallelism that may be exploited.

The resulting application code can run on any conventional serial machine, or on any parallel machine for which a suitable harness has been written, because the form is uncommitted to any type of

parallel architecture. The harness dynamically monitors performance and optimises execution parameters such as granularity.

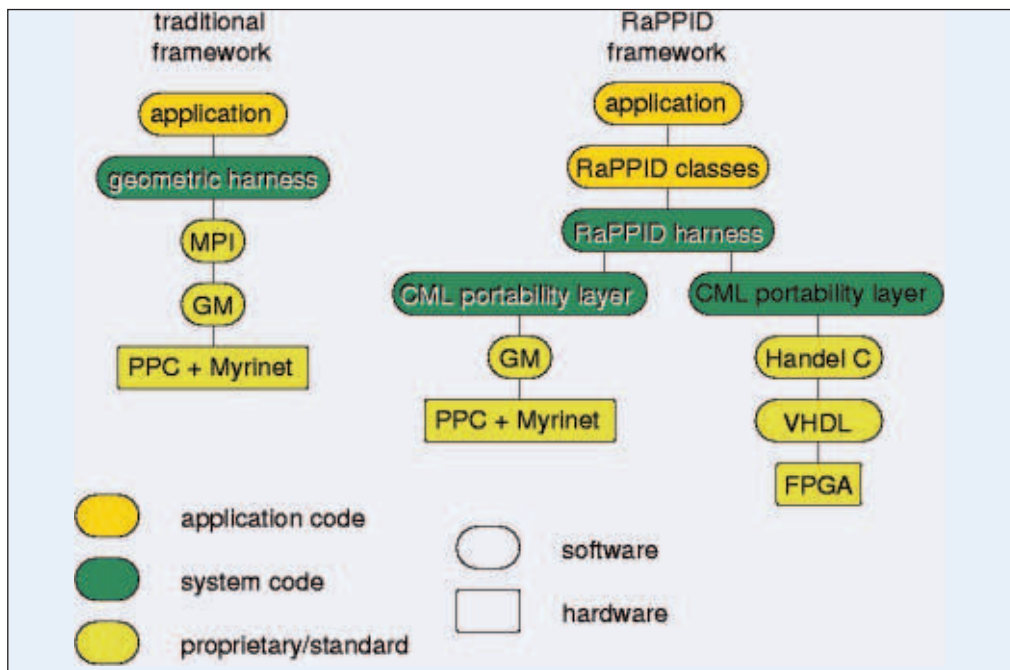
A variety of harnesses have been developed in CML (Concurrent Meta-Language) as the power of a functional language (code or 'functions' can be manipulated with the facility of data) provides an efficient mechanism for system prototyping. The very simplest case of data decomposition provides an illustration. If a user writes a software module to match the following CML type signature, then the user's software can be executed in parallel:

```
signature DATA_PARALLEL =
sig
  type input;
  type output;
  val process : input -> output;
  val combine : output * output ->
output;
  val split : input -> input * input;
end;
```

Here, the harness manages the parallel execution of the task, splitting the input recursively until the required granularity is achieved or until further splitting would become meaningless. Note, the user's code is purely sequential, and is in C++ rather than CML. Similar separations of concerns are possible for both geometric and pipeline parallelism.

Compare the protocol stack of a traditional parallel execution harness with that proposed by the RaPPID project (the left and right of Figure 1 respectively). The objective is to move towards a more uncommitted form of parallelism in order to provide unified support for a wider range of parallel and indeed hybrid parallel architectures with such components as Digital Signal Processors (DSPs), conventional Central Processing Units (CPUs) and Field Programmable Gate Arrays (FPGAs).

The traditional protocol stack model was used to port a geometrically parallel Augmented Reality (AR) application to a message-passing architecture. The



Protocol Stacks.



Target hardware

real-time performance and parallel efficiency of 90% achieved constitute a baseline system for future comparison. The hardware, shown in Figure 2, consists of four PowerPC (PPC) processors connected by a high speed Myrinet network.

The harness was written on top of the MPI (Message Passing Interconnect) standard which in turn uses the Myrinet proprietary GM (Glenn’s Messages) communication library.

In contrast, the portability layer of the RaPPID harness being developed supports the small and elegant CML model which is based on CSP (Communicating Sequential Processes). Adopting a clean CSP model allows the harness to be supported on a variety of architectures, and Figure 1 shows how this may unify the programming of FPGAs and message-passing architectures. A suitable mechanism for bringing FPGAs into the fold is Handel C. This is a concurrent C-based and CSP inspired programming language which compiles down to an EDIF (Electronic Design Interchange Format) net list, suitable for programming an FPGA.

In conclusion, the path to embedded parallelism described in this paper has proven encouraging on two fronts:

- good parallel efficiency of tricky-to-parallelise applications written to the proposed framework
- positive software engineering experience of functional language use in the concurrent domain.

RaPPID is a 3-year project (EPSRC Contract: GR/N20980) that started in October 2000 at Essex University in association with QinetiQ, Malvern.

Link:
http://www.essex.ac.uk/ese/research/mma_lab/rapid/index.html

Please contact:
 David John Johnston
 University of Essex, UK
 Tel: +44 1206 872912
 E-mail: djjohn@essex.ac.uk

Discretizing Continuous Attributes in AdaBoost for Text Categorization

by Pio Nardiello, Fabrizio Sebastiani, and Alessandro Sperduti

Researchers from the University of Padova and from ISTI-CNR, Pisa, are undertaking a collaborative effort aimed at producing better best text classification strategies through the design of methods for the discretization of continuous attributes. This will make it possible to exploit the rich information contained in the non-binary weights produced by standard statistical or probabilistic term weighting techniques, in the context of high performance learners (such as AdaBoost and its variants) requiring binary input.

In the last ten years an impressive array of learning techniques have been used in text categorization (TC) research. Among these, the two classes of methods that most seem to have caught the attention of TC researchers are boosting (a subclass of the classifier committees class) and support vector machines. The reasons for this attention are twofold, in the sense that both classes exhibit strong justifications in terms of computational learning theory and superior effectiveness once tested on TC benchmarks of realistic size and difficulty. It is on the former class of methods that our work focuses.

Classifier committees (aka ensembles) are based on the idea that, given a task that requires expert knowledge to perform, several experts may be better than one if their individual judgments are appropriately combined. In TC, this means applying several different classifiers to the same task of deciding whether a given document belongs or not to a given category, and then combining their outcome appropriately. Boosting is a method for generating a highly accurate classifier by combining a set of moderately accurate classifiers ('weak hypotheses'). In this work we make use of two algorithms, called AdaBoost.MH and AdaBoost.MH(KR), which are based on the notion of "adaptive boosting", a version of boosting in which members of the committee can be sequentially generated after learning from the classification mistakes of previously generated members of the same committee. AdaBoost.MH is a realization of the well-known AdaBoost algorithm, which is specifically aimed at multi-label TC

(ie the TC task in which any number of categories may be assigned to each document), and which uses 'decision stumps' (ie decisions trees composed of a root and two leaves only) as weak hypotheses. AdaBoost.MH(KR) is a generalization of AdaBoost.MH previously designed and implemented by these authors (see ERCIM News 44, p. 55) and based on the idea of learning a committee of classifier sub-committees; in other words, the weak hypotheses of AdaBoost.MH(KR) are themselves committees of decision stumps. So far, both algorithms have been among the best performers in text categorization experiments run on standard benchmarks.

A problem in the use of both algorithms is that they require documents to be represented by binary vectors, indicating presence or absence of the terms in the document. As a consequence, these algorithms cannot take full advantage of the 'weighted' representations, consisting of vectors of continuous (ie non-binary) attributes that are customary in information retrieval tasks, and that provide a much more significant rendition of the document's content than binary representations. In this work we address the problem of exploiting the potential of weighted representations in the context of AdaBoost-like algorithms by discretizing the continuous attributes through the application of entropy-based discretization methods. These algorithms attempt to optimally split the interval on which these attributes range into a sequence of disjoint subintervals. This split engenders a new vector (binary) representation for documents, in which a binary term indicates that the

original non-binary weight belongs or does not belong to a given sub-interval.

Although the discretization methods we present can also be used in connection with learners not belonging to the boosting family, we focus our experiments on AdaBoost.MH and AdaBoost.MH(KR). Our experimental results on the Reuters-21578 text categorization collection (the standard benchmark of TC research) show that for both algorithms the version with discretized continuous attributes outperforms the version with traditional binary representations. This improvement is especially significant since AdaBoost.MH and AdaBoost.MH(KR) are nowadays in the restricted lot of the peak text categorization performers, a lot where the margins for performance improvement are slimmer and slimmer.

Link:

<http://faure.iei.pi.cnr.it/~fabrizio/Publications/ContAtt.pdf>

Please contact:

Fabrizio Sebastiani, ISTI-CNR
Tel: +39 050 315 2892
E-mail: fabrizio@iei.pi.cnr.it

The 7th ERCIM Workshop on 'User Interfaces for All'

by Noëlle Carbonell and Constantine Stephanidis

The 7th Workshop of the ERCIM Working Group 'User Interfaces for All' took place in Paris (Chantilly), France, on 23-25 October 2002, building upon the results of the six previous Workshops held in Heraklion, Crete, Greece, 30-31 October 1995; Prague, Czech Republic, 7-8 November 1996; Obernai, France, 3-4 November 1997; Stockholm, Sweden, 19-21 October 1998; Dagstuhl, Germany, 28 November -1 December 1999; and Florence, Italy, 25-26 October 2000.

The vision of User Interfaces for All (UI4ALL) advocates the proactive realisation of the 'design for all' principle in the field of Human-Computer Interaction (HCI) and involves the development of user interfaces to interactive applications and telematic services, which provide universal access and usability to potentially all users. In the tradition of its predecessors, this year's workshop aimed to consolidate recent work and to stimulate further discussion on the state of the art in User Interfaces for All and its increasing range of applications in the emerging Information Society.

The emphasis of this year's event was on 'Universal Access'. The requirement for Universal Access stems from the growing impact of the fusion of the emerging technologies, as well as from the different dimensions of diversity intrinsic to the Information Society. These dimensions become evident when considering the broad range of user characteristics, the changing nature of human activities, the variety of contexts of use, the increasing availability and diversification of information, knowledge sources and services, the proliferation of technological platforms, etc. In this context, Universal Access refers to the accessibility, usability and, ultimately, acceptability of Information Society Technologies by anyone, anywhere, anytime, thus enabling equitable access and active participation of potentially all citizens in existing and emerging computer-mediated human activities. The user population includes people with different cultural, educational, training and employment background, novice and experienced users, the very young and the elderly, as well as people

with different types of disabilities, in various contexts of use. As people experience technology through their interaction with the user interfaces of applications and services, the field of HCI plays a critical and catalytic role towards a universally accessible, usable and acceptable Information Society.

This year's workshop has attracted the strongest ever interest worldwide, with a total of nearly 100 registered participants. The Draft Workshop Proceedings appeared as an INRIA research report and ERCIM Workshop Adjunct Proceedings. The official Workshop Proceedings will be published by Springer as part of LNCS (Lecture Notes in Computer Science) series, in the subline 'State-of-the-Art Surveys', and will be embedded in the LNCS digital library. Keynote speakers in this year's workshop were Alfred Kobsa, University of California, Irvine, USA, with a speech entitled 'Universal Access and Privacy: Making AVANTI Legal', and Steven Pemberton, CWI with a speech entitled 'The Kiss of the Spiderbot'.

The Workshop's technical programme comprised 80 submissions accepted for presentation as long/short/position papers or posters, covering a wide range of topics that include novel interaction paradigms and contexts of use, universal access to multimedia applications, software design and architecture for User Interfaces for All, design issues in User Interfaces for All, new modalities and dialogue styles, accessibility issues in novel interaction paradigms, design and usability in mobile computing and privacy issues in the context of Universal Access.

Some of these topics constitute core elements in a joint programme of research for a proposed Network of Excellence (NoE) submitted by the ERCIM Working Group 'User Interfaces for All' in the context of the European Commission's 6th Framework Programme. The focus of the associated expression of interest is on short-, medium- and long-term activities towards a new foundation for designing and developing user interface software and technologies to empower citizens' interaction with distributed and context-aware environments of use in the context of an inclusive Information Society.

From the year 2001 onwards, the UI4ALL Working Group Workshop has become a bi-annual event, in alternation with the newly established 'Universal Access in Human-Computer Interaction' (UAHCI) Conference, which is also held every two years. The 2nd International Conference on Universal Access in Human-Computer Interaction will take place on 22-27 June 2003, in Crete, Greece (<http://hci2003.ics.forth.gr/>).

Link:

ERCIM Working Group UI4ALL Web Page:
<http://ui4all.ics.forth.gr/>

Please contact:

Noëlle Carbonell, LORIA, UMR CNRS & INRIA, Workshop Programme Chair and Local Organiser
Tel: +33 3 83 59 20 32
E-mail: Noelle.Carbonell@loria.fr

Constantine Stephanidis, ICS-FORTH Working Group and Workshop Chair
Tel: +30 2810 391741
E-mail: cs@ics.forth.gr

CALL FOR PARTICIPATION

HCI International 2003

Crete, Greece, 22-27 June 2003

HCI International 2003, the 10th International Conference on Human-Computer Interaction, is jointly held under one management and one registration with the Symposium on Human Interface (Japan) 2003, the 5th International Conference on Engineering Psychology and Cognitive Ergonomics and the 2nd International Conference on Universal Access in Human-Computer Interaction.

The Conference aims to provide an international forum for the dissemination and exchange of scientific information on theoretical, generic and applied areas of HCI. This will be accomplished through the following modes of communication: plenary presentations, parallel sessions, poster sessions, tutorials, workshops, and other meetings of special interest groups.

Five major thematic areas are in the focus of the program:

- Human-Computer Interaction

- Ergonomics and Health Aspects of Work with Computers
- Human Interface and the Management of Information
- Universal Access in Human-Computer Interaction
- Engineering Psychology and Cognitive Ergonomics.

This is a truly international in scope event, representing work from the world's foremost leaders in the field. It will be held under the auspices of a distinguished international board of 114 members from 24 countries and we expect over 1500 participants from over 60 countries.

Keynote Speakers

- Ben Shneiderman, University of Maryland, USA
"Leonardo's Laptop: Human Needs and the New Computing Technologies"
- Jennifer Preece, University of Maryland Baltimore County, USA
"Designing Sociable, Universally Usable Online Communities"

Cooperating Societies

Chinese Academy of Sciences, Japan Management Association, Japan

Ergonomics Society, Human Interface Society (Japan), Swedish Interdisciplinary Interest group for Human Computer Interaction – STIMDI, Asociación Interacción Persona Ordenador - AIPO (Spain), Gesellschaft fuer Informatik e.V - GI (Germany), and ERCIM.

Exhibition

A major exhibition, open from Tuesday, June 24 through Friday, June 27, 2003, will provide an excellent opportunity to interested organizations to promote state-of-the-art products and services and enhance their visibility to professionals and researchers in the field of HCI from all over the world.

Sponsorship opportunities

Sponsorship in the context of HCI International 2003 is an ideal opportunity for interested organizations to raise their profile to an international targeted audience. Sponsorship is not limited to financial support, but can also take the form of provision of material and services.

Further information:

<http://www.hcii2003.gr/>

CALL FOR PARTICIPATION

Indexing and Searching
Audiovisual Material
at the D-Lib Competence Center

Pisa, Italy, 20-21 January 2003

The aim of this course is to provide students with a theoretical and experimental background on the techniques and methodologies for the organization, creation, and management of an Audio/Video Digital Library.

The ECHO system (<http://pc-erato2.iei.pi.cnr.it/echo/>) will be used as a testbed during the course. ECHO provides a digital library service for historical films. It provides tools to index and retrieve audio/video material using speech transcripts, automatically extracted video features, and metadata

manually associated by the user. The metadata are described using an audio-video metadata model based on the IFLA-FRBR standard.

The course will address the following topics:

- Audio/Video Digital Libraries: an introduction
- How to design and build an A/V Digital Library
- Examples of applications of an A/V Digital Library
- Manual indexing of A/V documents
- Training on the indexing of A/V documents
- Training on the usage of the DL System documentation will be provided.

Teaching Staff:

Giuseppe Amato, Claudio Gennaro, Pasquale Savino (contact person: savino@iei.pi.cnr.it). The D-Lib

Laboratory is available for use during the course.

The course is organized by the Digital Library Competence Center (IST-2001-32587), an FP5 project of the European Community. Courses at the Competence Center are designed for librarians, archivists, scholars and technicians, and offer access to advanced DL testbeds, services, expertise and knowledge.

Future courses also include:

- Open Access to Digital Archives
- Digital Library based Collaboration.

Courses are held free-of-charge and the maximum number of participants that can be accepted is 10. Registration should be made on-line via the website.

More information:

<http://www.cu.lu/hcp2003/>

CALL FOR PARTICIPATION

SLAP'03 – Second International Workshop on Synchronous Languages, Applications, and Programming

Porto, Portugal, 1 July 2003

Synchronous languages have been introduced in the 80s to program reactive systems. Such systems are characterised by their continuous reaction to their environment, at a speed determined by the latter. Synchronous languages have recently seen a tremendous interest from leading companies developing automatic control software for critical applications. For instance, Schneider Electric uses a Lustre-based tool, named Scade, to develop the control software for nuclear plants. Aerospatiale also uses this tool to develop the flight control of the new Airbus planes. Dassault Aviation uses Esterel Studio to program the flight control software of the Rafale fighter. Snecma uses Sildex, a Signal-based tool to develop airplane engines. ST Microelectronics, Texas Instrument, Motorola, Intel, are also interested in the Esterel technology for chip design. The key advantage pointed by these compa-

nies is that the synchronous approach has a rigorous mathematical semantics which allows the programmers to develop critical software faster and better.

Indeed, the semantics of the languages is used as formal model upon which all the programming environments are defined. The compilation involves the construction of these formal models, and their analysis for static properties, their optimisation, the synthesis of executable sequential implementations, the automated distribution of programs. It can also build a model of the dynamical behaviours, in the form of a transition system, upon which are based the analysis of dynamical properties, e.g., through model-checking based verification, or discrete controller synthesis. Hence, synchronous programming is at the cross-roads of many approaches in compilation, formal analysis and verification techniques, and software or hardware implementations generation. The approach is related to formal methods for reactive systems like Statecharts, StateFlow, UML StateCharts.

The workshop topics are covering all these issues:

- synchronous model of computation

- synchronous languages and programming formalisms
- compiling techniques
- formal verification
- test and validation of programs
- case-studies.

This year's (non-exclusive) focus will be on implementation aspects, especially:

- code generation and execution schemes,
- synchronous models of Real-Time Operating Systems and languages,
- implementation-related analyses, eg, performance evaluation (WCET).

SLAP'03 will be a satellite event of the Fifteenth Euromicro Conference on Real-Time Systems (ECRTS-03).

Important Dates

- Paper submission: February 1st 2003
- Notification of acceptance: April 1st 2003
- Final paper due: May 1st 2003

The proceedings will be published electronically by Elsevier in the Electronic Notes in Theoretical Computer Science series (ENTCS).

More information:

<http://www.inrialpes.fr/bip/people/girault/Slap03>

CALL FOR PARTICIPATION

Twelfth International World Wide Web Conference

Budapest, 20-24 May 2003

The conference will start with a day of tutorials and workshops and followed by a three-day technical programme. The fifth day will be a 'Developers Day.' The tutorials and workshops will provide in-depth looks at specific areas of current interest. The technical programme will include refereed paper presentations, alternate track presentations, plenary sessions, panels and poster sessions. Developers Day will be devoted to in-depth technical sessions designed specifically for Web developers.

Further information:
<http://www2003.org/>

CALL FOR PARTICIPATION

14th Mini-EURO Conference HCP'2003: Human Centred Processes: Distributed Decision Making and Man-Machine Cooperation

Luxembourg, 5-7 May 2003

Topics include:

Tools and Methods: Distributed decision making, group negotiation and decision making, collaborative work, Intelligent assistance for decision, knowledge extraction, representation and modelling, Intelligent management of multimedia documents, perception, recognition and interpretation, Human expertise centred

decision support, intelligent operator or user guidance and assistance systems.

Application Areas: practice and integration, quality control, management of industrial and administrative processes, risk management, scheduling and planning, industrial production, supervision and control, e-government, e-administration, transportation, health care, e-banking, insurance, telecommunications, management of the environment, NTIC.

A special publication of selected papers is foreseen in a feature edition of EJOR and the In Cognito Journal.

More information:

<http://www.cu.lu/hcp2003/>

CALL FOR PARTICIPATION

CLEF 2003



The Cross-Language Evaluation Forum (CLEF) provides an infrastructure for the testing and tuning of mono- and cross-language information retrieval systems working with European languages. The main focus of the activity is to stimulate the development of multilingual information retrieval systems, ie systems capable of matching a query in one language against document collections in many languages.

The three main tasks in this year's agenda will assess text retrieval system performance on a multilingual comparable collection of news documents in eight languages.

Multilingual Information Retrieval

There will be two distinct tasks:

- 'Small-multilingual' (Multilingual-4)
- 'Large-multilingual' (Multilingual-8)

The corpus for Multilingual-4 will contain English, French, German and

Spanish documents; Multilingual-8 will add Dutch, Finnish, Italian, and Swedish.

Bilingual Information Retrieval

The 2003 bilingual track will encourage the tuning of systems running on challenging language pairs. Runs will be accepted for the following source -> target languages:

- Italian -> Spanish
- German -> Italian
- French -> Dutch
- Finnish -> German

Newcomers only (ie groups that have not previously participated in a CLEF cross-language task) can choose to search the English document collection using any of the topic languages.

Monolingual (non-English) Information Retrieval

The CLEF experience has demonstrated the importance of monolingual system performance for multiple languages as a first step towards cross-language work. CLEF 2003 will offer tasks for Dutch, Finnish, French, German, Italian, Spanish and Swedish.

Other tracks have been designed to test system performance on different tasks and data types:

- Mono- and Cross-Language IR for Scientific Collections
- Interactive Cross-language Information Retrieval — iCLEF
- Multiple Language Question Answering (QA at CLEF)
- Cross-Language Retrieval in Image Collections (Image CLEF)
- Cross-Language Spoken Document Retrieval (CL-SDR).

Important Dates

- Registration Opens: 15 January 2003
- Data Release: 30 January 2003
- Topic Release: 1 March 2003
- Submission of runs by participants: 15 May 2003
- Release of individual results and relevance assessments: 1 July 2003
- Submission of paper for Working Notes: 20 July 2003
- Workshop: 21-22 August 2003

More information:

<http://www.clef-campaign.org/>

CALL FOR PAPERS

INTERACT 2003

Zurich, 1-5 September 2003

The ninth IFIP TC 13 international conference on Human-Computer Interaction will be organized at the Swiss Federal Institute of Technology (ETH) in Zurich (Switzerland), from September 1st to 5th 2003.

INTERACT 2003 is the next in a series of international IFIP HCI (Human-Computer Interaction) conferences. The conference will act as a platform for HCI research and practice which seeks to embrace the demands of the 21st century, while continuing to include state-of-the-art approaches to traditional HCI issues.

The program will include high quality papers, posters, industrial presentations and system demonstrations, as well as workshops and tutorials. INTERACT 2003 is intended as a conference for everyone involved in HCI, whether designing software or systems, managing the development of new systems or doing research or teaching.

More information:

<http://www.interact2003.org>

SPONSORED BY ERCIM

ECOOP 2003

Darmstadt, 21-25 July 2003

The European Conference for Object-Oriented Programming 2003 is the premium forum in Europe for bringing together practitioners, researchers, and students to share their ideas and experiences in a broad range of disciplines related to object technology. Events include outstanding invited speakers, carefully refereed technical papers, real world experiences in the form of practitioner reports, panels, workshops, demonstrations, and an interactive posters session.

Further information:

<http://www.ecoop.tu-darmstadt.de/>

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.

Copyright Notice

All authors, as identified in each article, retain copyright of their work.

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: Wolfgang Hawlik
w.hawlik@ocg.at
CLRC: Martin Prime
M.J.Prime@rl.ac.uk
CRCIM: Michal Haindl
haindl@utia.cas.cz
CWI: Henk Nieland
Henk.Nieland@cw.nl
CNR: Carol Peters
carol@iei.pi.cnr.it
FORTH: Yannis Pachoulakis
yannisp@ics.forth.gr
FhG: Michael Krapp
michael.krapp@iuk.fhg.de
FNR: Patrik Hitzelberger
hitzelbe@crpg.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 Csuha-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/>

EVENTS

CALL FOR PAERS

IEEE International Conference on Software Maintenance 2003 - ICSM-2003

Amsterdam, 22-26 September 2003

The International Conference on Software Maintenance (ICSM) is the world's major international conference for software and systems maintenance, evolution, and management. The focus of this year's conference will be 'the architecture of existing systems'.

Keynote Speakers

Alexander Rinnooy Kan (ING Group, The Netherlands) will share his insights in managing multi-billion dollar IT budgets. Chuck Howell (MITRE, USA) will present his work on assurance arguments: formal documentation of why we should trust a piece of software. Henk Obbink (Philips Research, The Netherlands) will discuss work on software architecture and its connection to existing systems.

Topics

Topics of interest include but are not restricted to the following aspects of maintenance and evolution:

- maintaining reusable components
- agile methods for maintenance
- maintenance of agile-produced software
- maintaining requirements
- financial / economical aspects of software maintenance
- software architecture recovery
- the architecture of existing systems
- system understanding
- automated software maintenance.

Important Dates

- Submission deadline: 1 March 2003
- Acceptance notification: 16 May 2003
- Final version due: 5 July 2003.

More information:

<http://www.cs.vu.nl/icsm2003>

CALL FOR PAPERS

3rd ESA Workshop on Millimetre Wave Technology and Applications: Circuits, Systems, and Measurement Techniques

Espoo, Finland, 21-23 May 2003

The third European Space Agency (ESA) workshop on the technology and applications of millimetre waves will be organised by MilliLab jointly with ESA/ESTEC, VTT and Helsinki University of Technology. The aim of this international workshop is to bring together people involved in research and industrial development of millimetre wave components and systems, and to explore common areas and synergies in the development of millimetre wave techniques for commercial and scientific ground-based and space-borne applications.

Submission deadline: 7 February 2003

More information:

<http://www.vtt.fi/millilab/pages/SecondCall.htm>

CALL FOR PAPERS

Eighth International Workshop on Formal Methods for Industrial Critical Systems (FMICS 03)

Trondheim, Norway
5-7 June 2003

The aim of this workshops of the ERCIM working group on Formal Methods for Industrial Critical Systems is to provide a forum for researchers who are interested in the development and application of formal methods in industry. In particular, these workshops are intended to bring together scientists who are active in the area of formal methods and interested in exchanging their experiences in the industrial usage of these methods. These workshops also strive to promote research and development for the improvement of formal methods and tools for industrial applications.

Invited Speakers


- Werner Damm, Oldenburg University, Germany
- Reiner Hähnle Chalmers, Sweden


Topics include:

- Tools for the design and development of formal descriptions
- Verification and validation of complex, distributed, real-time systems and embedded systems
- Verification and validation methods that aim at circumventing shortcomings of existing methods in respect to their industrial applicability
- Formal methods based conformance, interoperability and performance testing
- Case studies and project reports on formal methods related projects with industrial participation (eg safety critical systems, mobile systems, object-based distributed systems)
- Application of formal methods in standardization and industrial forums.

More information:

<http://www.inrialpes.fr/vasy/fmics/workshop-8/>

 **INRIA — INRIA has published a report summarising the institute's activities in the fields of Grids.** In France, an important part of computing grid research is done at INRIA. At least five of INRIA's research teams are deeply involved, in collaboration with various academic or industry partners. In addition to this, many other INRIA teams are doing work that is more or less closely related to the grid topic, either in the framework of Institute projects, in that of the incitative concerted initiative GRID (Globalization of computer resources and data) launched in 2001 by the Ministry of Research, in that of the National Network of Research in Telecommunications (RNRT, set up in 1997) or in that of the National Network for Research and Innovation in Software Technology (RNITL, set up in 1999). A similar report about the institute's activities in Telecommunications will also be available shortly. The report is available at <http://www.inria.fr/presse/dossier/gridcomputing/index.en.html>

 **Fraunhofer-Gesellschaft - A German-Chinese research institute** for Information and Communication Technology will be founded in order to intensify the bilateral co-operation in this field. This was decided at the occasion of the visit of the Chinese Minister for Science and Technology, Mr. Xu Guanhua in Germany, in June 2002. Two joint labs are expected to be created in Berlin and Beijing. In Berlin the lab will concentrate on mobile communication, in Beijing the focus will be on software technology, in particular on middleware and embedded systems. On the German side Fraunhofer is in charge of this activity, on the Chinese side it is the High Technology Research and Development Center of the Chinese Ministry of Science and Technology (MoST HTRDC).



Over 200 Job Opportunities at INRIA in 2003

INRIA offers a large variety of positions in 2003, among them are :


- visiting positions from six months to two years duration for French and foreign specialists, with academic or industrial background ;
- tenured positions for junior and senior research scientists.

In addition to these job opportunities, INRIA offers more than 200 hosting possibilities within one of INRIA's research teams, for preparing a PhD thesis or for postdoctoral visits.

A competitive selection for recruiting four experienced research scientists (chargés de recherche 1ère classe) has been opened recently. The deadline for applications is January 17.

A competitive selection to recruit some 30 junior research scientists (chargés de recherche 2ème classe) or senior research scientists (directeurs de recherche 2ème classe) on tenured position will start in February 2003.

For more information, see: <http://www.inria.fr/>



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.



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/>



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/>



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/>



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/>



Consiglio Nazionale delle Ricerche, IISTI-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.pi.cnr.it/>



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/>



Czech Research Consortium
for Informatics and Mathematics
FI MU, Botanická 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>



Swiss Association for Research in Information Technology
Dept. Informatik, ETH-Zentrum, CH-8092 Zürich,
Switzerland
Tel: +41 1 632 72 41, Fax: +41 1 632 11 72
<http://www.sarit.ch/>



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/>



Slovak Research Consortium for Informatics and
Mathematics, Comenius University, Dept. of Computer
Science, Mlynska Dolina M, SK-84248 Bratislava,
Slovakia, Tel: +421 7 726635, Fax: +421 7 727041
<http://www.srcim.sk>



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>



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/>



Foundation for Research and Technology – Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
Tel: +30 81 39 16 00, Fax: +30 81 39 16 01
<http://www.ics.forth.gr/>



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>



Fraunhofer-Gesellschaft,
Information and Communication Technology Alliance
Schloß Birlinghoven, D-53754 Sankt Augustin, Germany
Tel: +49 2241 14 0, Fax: +49 2241 14 2889
<http://www.fraunhofer.de/>



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

Name:

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.

Organisation/Company:

Address:

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

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/