ERCIM NEWS

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

Number 63, October 2005

Special:

Security and Trust Management

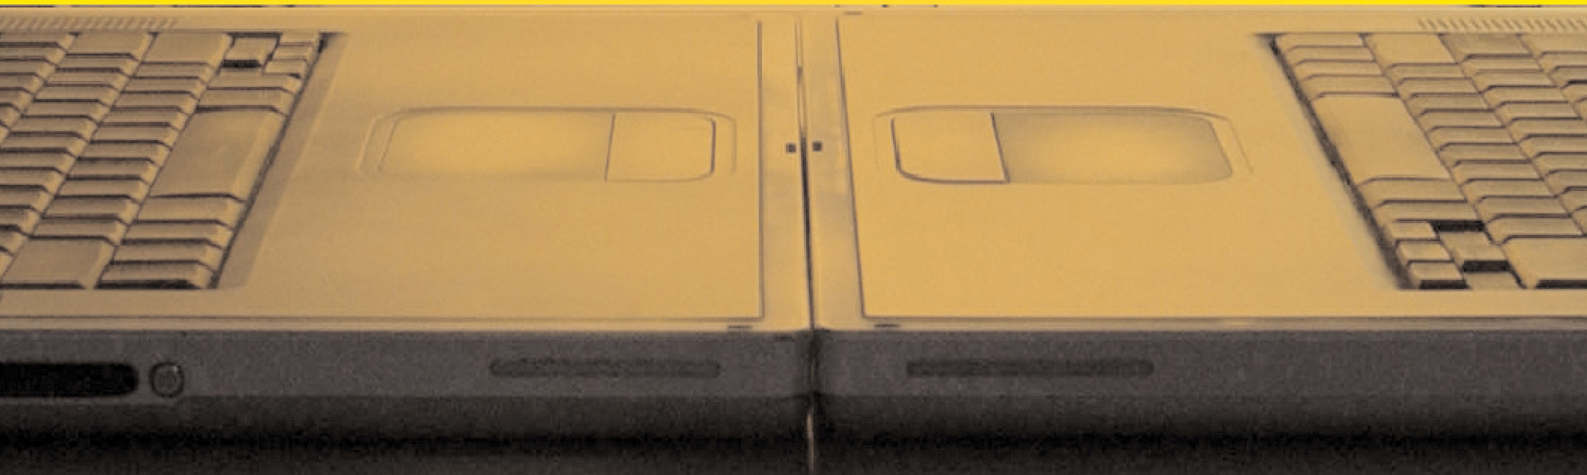

CONTENTS

KEYNOTE

JOINT ERCIM ACTIONS

NEWS FROM W3C

SPECIAL THEME: SECURITY AND TRUST MANAGEMENT

Introduction

Invited article:

Invited article:

Invited article:

R&D AND TECHNOLOGY TRANSFER

EVENTS

Next issue: January 2006 — Special theme: Emergent Computing

Michael Waidner, Head of the IBM Privacy Research Institute, IBM Research Division, Zurich Research Lab

Until a few years ago, talking about information security was a sure way to distinguish between Information and Communication Technology (ICT) and business people. ICT people started to ask questions about bits and bytes, key lengths and signature updates. Business people left the room. Since then the situation has changed dramatically: new regulations and significant security and privacy incidents have pushed the security topic into the board room.

From a business perspective, security is all about managing risk. An ICT system provides the right level of security if it keeps the risk for the business at an acceptable level. What counts for this risk are potential losses due to malicious acts by disgruntled employees, criminal hackers or terrorists. Whether a risk is acceptable or not is a business decision. How to describe this level and how to demonstrate that an ICT system meets that level is one of the fundamental challenges in Computer Science.

While security is getting board room attention, actually securing ICT systems is becoming more difficult than ever: The dependencies between enterprises are rapidly growing: Thanks to the increasing specialization in industry, more and more enterprises need to cooperate to provide a specific service. Moreover, these dependencies are becoming much more dynamic in time. Today most business relations are still based on paper contracts, but the trend clearly goes towards dynamically negotiated and electronically confirmed contracts. As a result, the security boundaries between enterprises are quickly becoming less and less strict. Back-end servers that were formerly carefully protected through multiple protection layers are now directly exposed to the outside, as they offer their services to many enterprises. Applications that used to run on dedicated servers now run on a virtual, shared infrastructure, using physical resources that might be spread all over the world.

But the cause for this amplification of security problems – the trend towards On Demand business – also creates new opportunities to solve these security problems:

Service-oriented architecture (SOA) is the concept that supports the dynamic integration of enterprises into larger, virtual enterprises, based on standardized Web services. This well-defined construction principle allows embedding security and privacy in the fabric of the new infrastructure. Security and privacy play a prime role in forming virtual enterprises – an enterprise can manage its risk only if it knows what risk it takes by interacting with another enterprise, what trust it has to invest in its partners, and how that trust is established and justified. Therefore these risk and trust requirements and guarantees must be negotiated as part of the service-level agreements and the partners must provide service-oriented assurances to one another on which they can base their own risk management.

Virtualization and trusted platforms: Providing services in a cost-efficient and manageable way requires the sharing of technical resources whenever possible. This observation resulted in an astonishing renaissance of virtualization technology. Logical partitions and strong isolation are well-known mainframe security concepts – now they are becoming available on essentially all platforms. Similarly, rooting the security of a platform in a piece of trusted and secure hardware is becoming main stream, thanks to the effort of the Trusted Computing Group. Together these concepts offer the chance to build up distributed computing platforms in a secure way from scratch. In theory this is nothing spectacular, considering that most concepts have been known for decades, but practically this is the first time that there is broad agreement in industry to make this happen.

Business-oriented compliance and risk management: The last part is linking ICT and business. In the same way ICT people are used to express their security requirements in a security policy, business people need to express what they require from their systems in a formalized compliance policy and in the same way business people are used to being informed about their stock price falling or climbing they need to be informed as to whether their security risk level is falling or climbing. Many aspects of this linking are still a matter of research.

'On Demand' amplifies many of today's security and privacy problems, but it also offers the unique chance to build in security and privacy from the beginning, and to make security and privacy first class citizens of the new ICT infrastructure.

Michael Waidner

STM 2005 — First ERCIM Workshop on Security and Trust Management

by Sjouke Mauw

The ERCIM Working Group 'Security and Trust Management' (STM) organized a successful workshop in Milan, Italy on 15 September 2005. The initiative for this workshop was taken at the founding meeting of the STM Working Group in January 2005 in order to fulfil some of the group's goals; in particular to bring researchers together and stimulate scientific discussion.

The organization of the workshop, co-located with ESORICS'05, was in the hands of Fabio Martinelli, Pierangela Samarati (general co-chairs), Valerie Issarny, Sjouke Mauw (Programme Committee co-chairs) and Cas Cremers (vice-chair). Due to the high number of submissions (36) the programme committee selected nine papers of high quality, which were presented at the workshop and will appear in a special issue of ENTCS.

The workshop had a broad scope, ranging from cryptography and formal methods to physical security. The link between the major topics, security and trust, was made by keynote speaker Prof. Dieter Gollmann, who challenged the audience with his presentation entitled 'Why Trust is Bad for Security'. He argued that the notion of trust has many different (often conflicting) interpretations, and there is a need for clarity and precision.

The first session of the workshop covered smart-dust security, a formal approach to multiparty contract signing and the development of security models for mobile agent security. In the second session, the notions of credit and responsibility were formalized, a new scheme for trapdoor hash functions was presented, and access-control mechanisms based on trust were studied. The third and final session was completely dedicated to trust management. It covered the extension of role-based trust management languages with non-monotonicity, assigning trust values to metadata, and a new authorization strategy for distributed environments.

Given the success of this first STM workshop and the interest it generated, we believe that this will be the first of a series of successful workshops.

Links:
http://www-rocq.inria.fr/arles/events/STM2005/
http://www.iit.cnr.it/STM-WG/

Please contact:
Sjouke Mauw, Eindhoven University of Technology and Centre for Mathematics and Computer Science, Amsterdam, The Netherlands
E-mail: sjouke@win.tue.nl

Fabio Martinelli, IIT-CNR, Italy
E-mail: Fabio.Martinelli@iit.cnr.it

Changes in Irish ERCIM Membership

The Irish Universities Association (IUA) signed the scientific agreement with ERCIM becoming the official Irish ERCIM member.

The 55th meeting of the ERCIM Executive Committee took place at Dublin City University (DCU) on 23 September 2005. In conjunction with the meeting, a signing ceremony was necessary as the official Irish ERCIM membership was moved from Trinity College Dublin to IUA and Dublin City University took over the administrative and financial affairs for IUA. Prof. Ferdinand von Prondzynski, president of Dublin City University (for Irish Universities Association) and Eckart Bierdümpel, Fraunhofer Gesellschaft (for ERCIM) completed the signing process.

Left to right: Eugene Kennedy (vice president for research at DCU), Heather Ruskin (DCU and ERCIM board of Directors), Eckart Bierduempel (ERCIM Executive Committee Chair), Mark Roantree (DCU and ERCIM Executive Committee vice-chair), and Ferdinand von Prondzynski (DCU President and IUA President).

The Irish Universities Association represents the seven Irish universities in ERCIM:

- Dublin City University
- Univeristy College Dublin
- Trinity College Dublin
- University of Limerick
- National University of Ireland Cork
- National University of Ireland Galway
- National University of Ireland Maynooth.

Link:
http://ercim.computing.dcu.ie/

Please contact:
Mark Roantree, Dublin City University, Ireland
E-mail: Mark.Roantree@computing.dcu.ie

2005 Cor Baayen Award Winner

Milan Vojnović from Cambridge, UK has been awarded the 2005 Cor Baayen Award for the most promising young researcher in computer science and applied mathematics by ERCIM.

Milan Vojnović, originally from Croatia completed his PhD at Ecole Polytechnique Fédérale de Lausanne (EPFL) in Switzerland and has subsequently worked as an Associate Researcher in the Systems and Networking Group at Microsoft Research in Cambridge.

Milan's work in the area of network modelling impressed the judges with contributions to several topics including congestion control, mobility modelling, queuing performance and filecasting. For a young researcher to work on several topics, come up with novel ideas in each of them, some of which are seen to be valuable contributions convinced the judges of his merit for the award. The work is a clear example of what ERCIM strives to advance in that it is both theoretical in itself, but motivated by practical problems whose solution will have a significant impact. For example, his work on mobility modelling and simulation uses stochastic modelling based on fundamental probability theory, yet has shown that many simulations commonly used in industry are flawed, with potentially serious implications for the results built on them. He has proposed alternative models that can be simulated justifiably, and released code that can be used by practitioners.

The Cor Baayen award is not the first that Milan has received, since he has previously been given an award during his undergraduate studies at the University of Split, a best fellowship award during his graduate studies at EPFL, the ITC-17 best student paper award in 2001, and most recently, the IEEE INFOCOM 2005 best paper award and the ACM SIGMETRICS 2005 best paper award.

About the Cor Baayen Award

The Cor Baayen Award, awarded to a most promising young researcher in computer science and applied mathematics, was created in 1995 to honour the first ERCIM President, and is open to any young researcher having completed their PhD thesis in one of the 'ERCIM countries': Austria, Belgium, Czech Republic, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Luxembourg, Norway, Spain, Sweden, Switzerland, The Netherlands and the United Kingdom.

The award consists of a cheque for 5000 Euro together with an award certificate.

Submitting a Nomination
Nominations should be made by a staff member of the university or research institute where the nominee is undertaking research. Self nominations are not accepted. Nominations must be submitted with an online form, which will be made available when the next Call is launched. Alternatively contact the ERCIM Executive Committee member (the national contact point) for the country in which the nominee is undertaking research.

Further information can be obtained from the national contact or from the Cor Baayen Award coordinator Laszlo Monostori, SZTAKI, (laszlo.monostori@ercim.org).
http://www.ercim.org/activity/cor-baayen.html

ERCIM-Sponsored Events

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

Conferences
ERCIM invites sponsorship proposals from established conferences with an international reputation, where substantive overlap can be shown between the conference topic and ERCIM areas of activity. Typical cases would include annual conferences in computer science with international programme committees, substantial international participation, and proceedings published with an established international science publisher.

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

Next Deadlines for Applications:

- Conferences: 15 January 2006 for conferences later than 15 September 2006
- Workshops and summer schools: 15 January 2006 for workshops and schools later than 15 April 2005

Events sponsored by ERCIM in 2005:

- IEEE Virtual Reality 2005, Bonn, Germany, 12-16 March 2005
- CASSIS'05 - Construction and Analysis of Safe, Secure and Interoperable Smart devices, Sophia-Antipolis, France, 7-11 March 2005
- ECIR-05 - 27th BCS European Annual Conference on Information Retrieval, Santiago de Compostela, Spain, March 21-23, 2005
- IEEE 6th International Workshop on Policies for Distributed Systems and Networks, Stockholm, 6-8 June 2005
- CAiSE 2005 — 17th Conference on Advanced Information Systems Engineering, Porto, Portugal, 15-17 June 2005
- IJCAI-05 - 19th International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, 30 July - 5 August 2005
- INTERACT 2005, Rome, 12-16 September 2005
- ECDL 2005 - 9th European Conference on Research and Advanced Technology for Digital Libraries Vienna, Austria, 18-23 September 2005

For detailed information, see:
http://www.ercim.org/activity/sponsored.html

‘Web Foundations 2005’ Highlights Web Design for All

International experts on accessibility, usability and Web Standards will meet in Spain for the first time at the Web Foundations 2005 Conference, to be held 22-24 November in Gijón, Oviedo, Spain. This conference is the first major event supported by the regional government’s ‘e-Asturias 2007 Plan’, as a strategy for the Information Society development. Organized by the Fundación CTIC and the W3C Spanish Office, this event, open to the public, focuses on conditions to ensure universal access to information regardless of hardware, software, network infrastructure, language, culture, and user capabilities.

The speakers — Shawn Henry (W3C), Jakob Nielsen (NNGroup), Steven Pemberton (W3C), Inmaculada Placencia (European Commission), John Slatin (University of Texas), and Jeffrey Zeldman (Happy Cog) — will discuss ‘Design for All’ as an essential requirement for equitable Internet access. Steven Pemberton also gives an XForms and XHTML tutorial on 24 November in Oviedo.

Link:
Web Foundations 2005: http://www.fundamentosweb.org/

Office in India

W3C World Wide Web consortium®
भारतीय कार्यालय

The Centre for Development of Advanced Computing (C-DAC), is organizing an international conference/workshop on 10-11 November 2005 in New Delhi, coinciding with the launch of the W3C Office in India. Topics such as Internationalization, Mobile Web, Semantic Web, SVG, Voice XML and Multimodal Interaction will be presented by both W3C experts and Indian researchers and industrials.

The objective of the event is to create awareness about the future developments in Web technologies amongst the Indian IT developer’s community, researchers and institutions to leapfrog to software products and services based on international standards. The W3C Office in India is the fifteenth W3C Office in the world, and its launch follows the recent introduction of W3C’s new fee structure for organizations in developing countries.

Links:
W3C India Office: http://www.w3cindia.in/
Launch event: http://www.w3cindia.in/programme.htm
W3C Offices: http://www.w3.org/Consortium/Offices/

W3C Workshop on Internationalizing the Speech Synthesis Markup Language

W3C is organizing a workshop on internationalizing the speech synthesis markup language (SSML) 2-3 November 2005 in Beijing, China. The main objective of this workshop is to identify and prioritize extensions and additions to SSML that will improve its use for rendering non-English languages.

Widely used, SSML is designed to provide a rich, XML-based markup language for assisting the generation of synthetic speech in Web and other applications. The essential role of the markup language is to provide authors of synthesizable content with a standard way to control aspects of speech such as pronunciation, volume, pitch, or rate across different synthesis-capable platforms.

To make SSML more useful in current and emerging markets, W3C’s Voice Browser Working Group is considering enhancements for non-English languages. Suggestions relevant for multiple languages are most welcome, as the scope of the workshop is not limited to Asian languages.

Link:
Workshop agenda:
http://www.w3.org/2005/08/SSML/ssml-workshop-agenda

Upgrades of W3C Quality Tools

W3C Markup Validator

Also known as the ‘HTML validator’, the newest version of the W3C markup validation service includes usability enhancements, improved feedback, support for installation on Windows, and better support for both W3C and non-W3C document types. This popular free service can dramatically help improving and ensuring the quality of Web sites, therefore saving time and money. The Markup validator is an open source/free software, available for download and installation on a variety of systems and platforms.

CSS Validator

The CSS Validation Service has also been updated. Changes include a stable SOAP access and output to the CSS validator, the addition of a profile for the CSS 2.1 specification, and a large number of bug fixes.

Log Validator

A new upgrade of the log validator was released earlier this month. The Log Validator makes it easy to manage the quality of even large Web Sites, step by step, by finding the most popular documents failing Markup or CSS validation, or with broken links.

Links:
W3C Markup Validator: http://validator.w3.org
W3C CSS Validator: http://jigsaw.w3.org/css-validator/
W3C Log Validator: http://www.w3.org/QA/Tools/LogValidator/

W3C's Specification Guidelines to Create Better Specifications

The W3C's Quality Assurance (QA) Working Group concluded its work in August 2005 with the completion of the Specification Guidelines W3C Recommendation, a document which provides clear instruction to writers and editors on creating implementable technical specifications. By identifying both requirements and 'good practices,' the Specification Guidelines help both W3C and other specification authors create and describe technologies in ways that make it easier for developers to implement them as intended.

W3C launched the Quality Assurance (QA) Activity in 2001 with the goal of improving W3C specifications by offering guidelines to W3C groups. Since that time, the QA Working Group produced six useful documents such as the QA Framework Primer, the QA Test FAQ, the Variability in Specifications Working Draft, and the QA Handbook. One of the group's more famous and useful documents is the W3C Quality Assurance Matrix, a list of over 100 W3C specifications, with links to conformance clauses, test suites and validators.

Links:
Quality Assurance at W3C: http://www.w3.org/QA/
Matrix of W3C specifications: http://www.w3.org/QA/TheMatrix
Quality Assurance at W3C: http://www.w3.org/QA/
QA Framework: http://www.w3.org/TR/qaframe-spec/
QA Primer: http://www.w3.org/QA/WG/qaframe-primer
Matrix of W3C specifications: http://www.w3.org/QA/TheMatrix

Mobile Industry Promotes W3C's Mobile Web Initiative

W3C will hold a Mobile Web Initiative (MWI) event at the British Academy of Film and Television Arts in London, UK on 15 November 2005. Mobile access to the Web has been a second class experience for far too long. MWI sponsors aim at promoting Web access from a mobile device to become as simple, easy, and convenient as Web access from a desktop device.

The MWI sponsors are key players in the mobile production chain, including authoring tool vendors, content providers, handset manufacturers, browser vendors and mobile operators: Afilias, Argogroup, Bango.net, Drutt Corporation, Ericsson, France Telecom, HP, Jataayu Software, mTLD Top Level Domain Limited, MobileAware, Nokia, NTT DoCoMo, Opera Software, TIM Italia, RuleSpace, Segala M Test, Sevenval, Vodafone, Volantis.

Link:
W3C's Mobile Web Initiative: http://www.w3.org/Mobile

Web Accessibility Business Case Documents Published

The Web Accessibility Initiative (WAI) Education and Outreach Working Group (EOWG) has published 'Developing a Web Accessibility Business Case for Your Organization.' This document is designed to help organizations develop their own customized business case for Web accessibility. It provides text that can be used as is, as well as guidance on identifying the most relevant factors for a specific organization. The 5-page resource suite describes social, technical, financial, legal and policy aspects of Web accessibility:

- Social Factors addresses the role of Web accessibility in equal opportunity for people with disabilities; the overlap with digital divide issues; and benefits to people without disabilities, including older people, people with low literacy and people not fluent in the language, people with low bandwidth connections to the Internet and older technologies, and new and infrequent Web users.
- Technical Factors addresses interoperability, quality, reducing site development and maintenance time, reducing server load, enabling content on different configurations, and being prepared for advanced Web technologies.
- Financial Factors addresses financial benefits from increased Web site use and direct cost savings.
- Legal and Policy Factors addresses requirements for Web accessibility from governments and other organizations.

Business cases are sometimes accompanied by an implementation plan describing the steps involved in making an organization's Web site accessible. A separate WAI resource suite, Implementation Plan for Web Accessibility, provides information on initial assessment, developing organizational policies, training, selecting authoring tools, and conformance evaluation.

It is essential that the Web be accessible in order to provide equal access and equal opportunity to people with disabilities. An accessible Web also benefits others, including older people with changing abilities due to aging.

Links:
Web Accessibility Business Case Documents:
http://www.w3.org/WAI/bcase/

Implementation Plan for Web Accessibility:
http://www.w3.org/WAI/impl/

Latest W3C Recommendations

- xml:id Version 1.0
  9 September 2005, Jonathan Marsh, Daniel Veillard, Norman Walsh
- QA Framework: Specification Guidelines
  17 August 2005, Karl Dubost, Lynne Rosenthal, Dominique Hazaël-Massieux, Lofton Henderson

A complete list of all W3C Technical Reports:
http://www.w3.org/TR/

Security and Trust Management

Introduction to the Special Theme

by Fabio Martinelli
and Jean-Jacques Quisquater

Modern society is increasingly reliant on the storage, processing and transmission of information. Ensuring the integrity, security and privacy of information is thus paramount, regardless of whether the information is at the level of the citizen or at a national or international level. Moreover, future trends (as outlined in the ISTAG report, for example) in the so-called Ambient Intelligent Space (AmI) will only increase the role of information and our reliance on it. This brings with it great opportunities to enhance our quality of life, but at the same time, presents major challenges in terms of the privacy and integrity of personal information.

There is a common understanding that achieving greater security in information and communications technology (ICT) would increase its development and diffusion, with concomitant benefits in many fields. While this technology is already spreading rapidly, it will only be possible to translate our physical interactions into electronic interactions if sufficient trust and confidence exist in the systems that process our information.

The integrity, security and privacy of information and communication are thus paramount, in everything from personal information transfer to government and critical infrastructures. It is now widely agreed that lack of trust in systems will prevent their widespread adoption. As a consequence, the development and deployment of systems with strong effective security is vital.

In addition, modern ICT systems may consist of several (even thousands or more) computation and communication resources whose number dynamically changes and thus are getting closer to so-called virtual communities. In this new framework, the capability to represent, create, negotiate, monitor and evolve trust relationships in a secure way becomes mandatory.

ERCIM has recently established a WG on security and trust management to foster the research on these issues.

This special issue contains 29 articles on a variety of research topics within the area of security and trust management, authored by ERCIM members and other European research groups. Jacques Bus, head of the 'ICT for Trust and Security' unit within the Directorate-General of Information Society and Media, kindly agreed to provide a contribution on the strategic challenge that trust and security in the information society represents to European research. We also invited two contributions from well-known experts in the USA and Australia, to provide us with a rundown of the research being performed in these countries.

The articles describe research projects and results in the following areas:

- wireless network security in ubiquitous systems
- trust and reputation management in virtual coalitions/organizations
- identity management problems and privacy issues in virtual coalitions; an example of the integration of security technologies with biometric systems is also described
- the area of document security, and the problem of security mechanisms for documents that are platform- and server-independent
- the security of distributed applications, from mobile e-commerce to e-healthcare
- the use of rigorous techniques such as formal languages and automata theory to specify, verify and analyse complex systems
- access-control models, policies and mechanisms for exploiting well-established technology such as X.509 certificates, and new XML-based languages for describing policies and credentials
- research and development activities in cyber-crime protection and detection.

Overall, these articles illustrate current trends not only in ERCIM institutes, but in the whole European research community.

Link:
http://www.iit.cnr.it/STM-WG/

Please contact:
Fabio Martinelli, IIT-CNR, Italy
E-mail: Fabio.Martinelli@iit.cnr.it

Jean-Jacques Quisquater, Université catholique de Louvain, Belgium
E-mail: jjq@dice.ucl.ac.be

ARTICLES IN THIS SECTION

Building Trust and Security in Information Society: A Strategic Challenge for European R&D

by Jacques Bus

Security, dependability and trust plays a key role in building the Information Society. With the move towards the 7th Framework Programme, the European Union will make research more effective and better tailored to the evolving needs and opportunities with which Europe is confronted.

Trust and security are key enablers of the Information Society. For citizens to use and feel comfortable with eGovernment services they must have confidence that their online services are trustworthy and secure. Similarly, for consumers and SMEs to use e-commerce and e-business they need confidence in the security of online transactions. As access to the Internet diversifies, from PCs to digital TVs, mobile phones and wireless devices, people feel increasingly concerned about the protection of their assets and privacy in this networked world. These aspects will become more and more important as we move towards the smart digital environments based on many interacting objects, devices and systems.

In the future, personal area networks and embedded computer chips will be everywhere – in our cars, our homes and even in our clothes. Security in such massively inter-connected environments will require solutions very different to those of today, and its social acceptance will require totally novel approaches to identity and privacy management through user-friendly and trustworthy interfaces, taking into account the privacy needs and data protection regulations in place. Underlying the service and user interface level we must give attention to the information and network security infrastructure. Modern service organisations, such as banking and finance, healthcare, energy, transport and others, rely on ICT for data exchange and control, creating strong mutual dependencies. These critical information infrastructures must be dependable and resilient, protecting against malicious attacks, ensuring tolerance towards and recovery from attacks, and adaptable to the changing security requirements.

The Present: Trust and Security in the 6th Framework Programme

The above describes the focus of the research in ICT for Trust and Security of the Information Society Technologies Priority in the 6th Framework Programme for Community Research (FP6). In this domain we also give particular attention to the promotion of integration of European research and its relation to global activities, given the nature of the challenge which becomes more and more global. In the first part of FP6, we have launched 17 projects (six Integrated Projects, three Networks of Excellence, six Targeted Research Projects and two Coordination Actions) with a total Community funding of about 75 million Euro. These activities cover

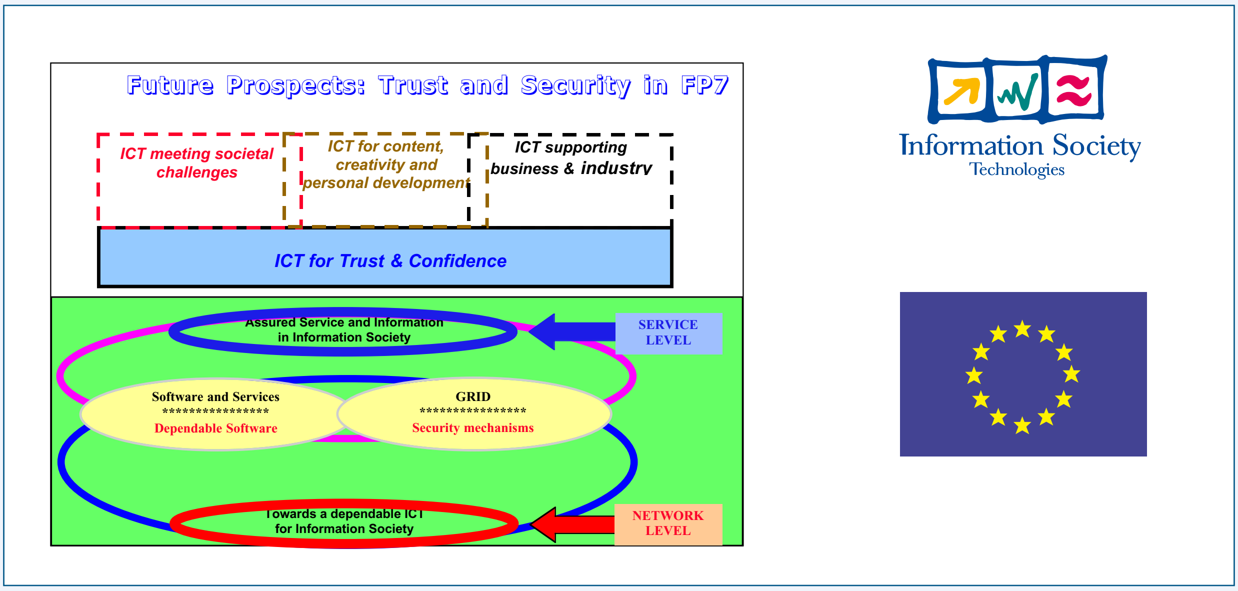
Future Prospects: Trust and Security in FP7
ICT meeting societal challenges
ICT for content, creativity and personal development
ICT supporting business & industry
ICT for Trust & Confidence
Assured Service and Information in Information Society
SERVICE LEVEL
Software and Services
****
Dependable Software
GRID
****
Security mechanisms
Towards a dependable ICT for Information Society
NETWORK LEVEL
Information Society Technologies

Trust and Security in ICT in the the 7th Framework Programme.

advanced and sophisticated research. There are strong interconnections with policy developments in trust and security, ie in multimodal and secure biometrics; identity and privacy management; electronic authentication; secure digital assets management; virtualisation of security resources for advanced and seamless security. It gives also strong attention to the support of standardisation and interoperability. More recently, as a result of the IST Call 4, we have retained and started negotiation of 19 new projects for an expected funding of about 70 million Euro. This set of new projects would strongly extend the technical coverage in this domain. It includes activities on the development of knowledge and technologies to manage and control complex and interdependent networks and systems, so as to enhance security and resilience in the information society infrastructure; provision of interoperable and open trusted computing platforms; advanced mechanisms and models for security, privacy and trust in mobile environments; and sophisticated technologies to fight malware on Internet.

The Future: Security and Dependability - Trust and Confidence the 6th Framework Programme

Whereas the key role of security, dependability and trust in building the Information Society is unquestioned, the move towards FP7 imposes the need to rethink how to make the EU intervention on and funding to research more effective and better tailored to the evolving needs and opportunities with which Europe is confronted. This need, which takes into account also the technological and market trends, brought the European Commission services to identify for FP7, which is planned to start in 2007, new avenues and synergies to renew and strengthen research impetus and momentum in this area. The nature of security and trust relates the subject to many IST domains, infrastructural, as well as application oriented. In order to be effective, dependability and security must be part of the system design, starting at the lowest level. But it should also ensure trust in the applications for the end-user, for example in e-Government, e-Health and consumer services. This rationale is the basis for the proposed structure for research on security, dependability and trust in the FP7 Information and Communication Technologies (ICT) Theme.

The overall structure and research priorities of FP7 were proposed in the Communication of April 2005 'concerning the seventh framework programme of the European Community for research, technological development and demonstration activities'. One of the 'Technology Pillars' (TP) proposed under the ICT theme of the Specific Programme 'Collaboration' covers the research activities on 'Software, Grids, Security and Dependability: dynamic, adaptive, dependable and trusted software and services, and new processing architectures, including their provision as a utility'. In addition, one of the domains of 'Applications Research' (AR) is defined to be 'ICT for trust and confidence: identity management; authentication and authorization; privacy enhancing technologies; rights and asset management; protection against cyber threats'. Although TPs and AR are both part of the Specific Programme 'Collaboration' in the theme ICT, they represent different approaches to the research in Trust and Security in ICT.

The TP on 'Software, Grids, Security and Dependability' deals with the key technological challenges and components that underpin the provision of both 'assured service and information handling' and 'dependable ICT systems'. It addresses the constitutive fabric of Information Society systems and services. The TP supports and enables applications research (eg e-government, e-business, e-health) where the application drive is the engine for future technological progress. The requirements on trust and confidence in these application areas are however often of a generic nature, exploiting 'security and privacy' technologies all across the different application domains. For this reason the choice has been made to include AR on 'ICT for Trust and Confidence', which would build upon the TP described above, but work in close cooperation with the other AR areas to ensure optimal synergy. Figure 1 shows how 'Trust and Security' is covered in the ICT Theme of FP7. It depicts the two important technology levels: network and services, as well as the crucial role of security and trust in the development of software and services and its infrastructures (ie, GRIDs). This technology development forms the fundament for the application in various domains through the domain 'ICT for Trust and Confidence'. Of course, such visualisation is limited. It may however help us in stimulating and managing the discussion with the European stakeholders such as researchers, industries, users, etc. in order to develop a strategic agenda for security and trust in FP7.

Conclusion

We intend to organise in the next few months specific consultation events that would be instrumental, together with the other institutional consultation processes, in the further development of the FP7 workprogramme. We trust that the above proposal can form a sound basis for a fruitful discussion with European researchers, leading to an IST workprogramme that will effectively cover the urgent research to be done for building a secure, dependable and trusted Information Society.

The content of this paper is the sole responsibility of the author and in no way represents the view of the European Commission or its services.

Links:
Specific information about events will be made available at the unit D4 web site:
http://www.cordis.lu/ist/trust-security/index.html

'ICT for Trust and Security' web site:
http://www.cordis.lu/ist/trust-security/index.html

The list of on-going projects:
http://www.cordis.lu/ist/trust-security/projects.htm

Please contact:
Jacques BUS, Head of Unit, European Commission, DG Information Society and Media "ICT for Trust and Security"
E-mail: jacques.bus@cec.eu.int

Flexible Multi-Factor Authentication in an Uncertain World

by Ravi Sandhu

Most organizations require high-grade multi-factor authentication for their high-end users, yet few can justify the high cost for all their users. Flexible multi-factor authentication enables organizations to issue the appropriate grade of credential for each user class from a single reusable infrastructure, and make adjustments as needed.

Weak authentication, and inadequately protected identity data in on-line services are resulting in increasing instances of fraud. Organizations suffer reputational and financial losses and consumers loose confidence in on-line services. Organized criminal activity has just about started to move on-line, but is growing very fast. New attack vectors such as phishing have rapidly gained prominence in a matter of a few months. It is impossible to predict the nature of attack activity in 2006, let alone in 2007, 2008 … Enterprises are faced with uncertainty about the nature of real cyberattacks, yet must make funding decisions today on how to improve their existing authentication infrastructure. Therefore solutions that allow adjustment of credential strength as neededis are gaining prominence.

The TriCipher Armored Credential System (TACS) provides a single platform that can issue and support a flexible range of credentials from a single infrastructure. The only system of its kind, the TACS Vault can be used to issue authentication credentials of many different types and can also serve as a vault for identity data (or encryption keys), providing a comprehensive solution to the problems of weak authentication and inadequately encrypted identity data.

The authentication ladder (see Figure 1) shows TACS issued credentials in increasing order of strength as we go up the ladder. This flexibility is achieved due to the use of 3-key RSA as the underlying cryptographic engine for all forms of authentication. In 3-key RSA there are 2 private keys for each user, one which only the user knows and the other securely stored on the TACS Vault. (In contrast conventional 2-key RSA uses a single private key known only to the user.) Both private keys are used to generate partial signatures which are then combined to authenticate the user. Flexible multi-factor authentication is achieved by generating the user's personal private key from a variety of factors depending upon the credential strength.

In the simplest credential called Armored Passwords, the user's personal private key part is derived from a password. This eliminates the need to store encrypted passwords on the server and thereby eliminates password cracking by theft of these encrypted passwords. In Browser 2 Factor credentials the 2nd factor is stored in an encrypted cookie which comes to the server. Based on information in the cookie, an appropriate personalized "welcome message" is displayed to the user before they enter their password. The password and the cookie are then combined for two factor authentication. In Device 2 Factor credentials the user's PC (more than one can be registered) is used as a low cost and effective 2nd factor. The PC stores a non-exportable 2-key RSA private key which is cryptographically combined with the user's password to generate the user's personal private key. In Portable 2 Factor any removable device, ranging from USB memory sticks to iPods can be used as a convenient second factor. The approach is similar Device 2 Factor in that it combines a 2-key RSA private cryptographically with the user's password to generate the user's personal private key. The information on the portable device is protected both by the fact that the attacker does not have access to the key on the TACS Vault, as well as a special 'rolling key' which cryptographically encrypts the second factor. Armored Tokens add one-time passwords to the credential. The dynamic password travels from the client to the TACS Vault, which performs the verification. The protected channel from the client to the TACS Vault eliminates the man-in-the-middle attack to which one-time passwords are extremely vulnerable. For Smartcard-based credentials the private key stored in the smartcard

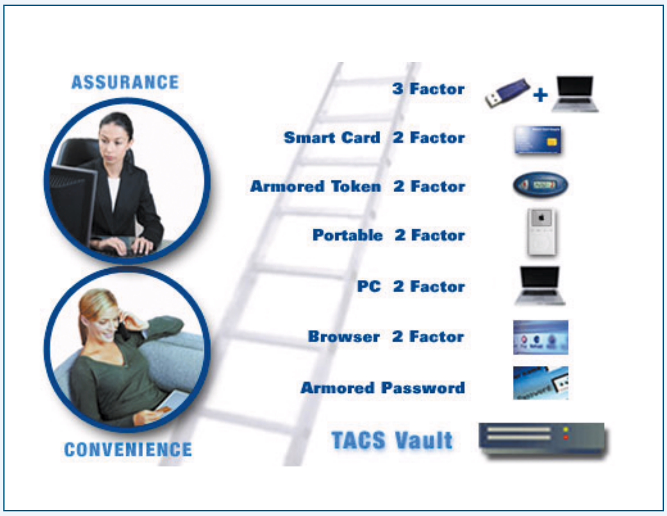

Figure 1: The authentication ladder.

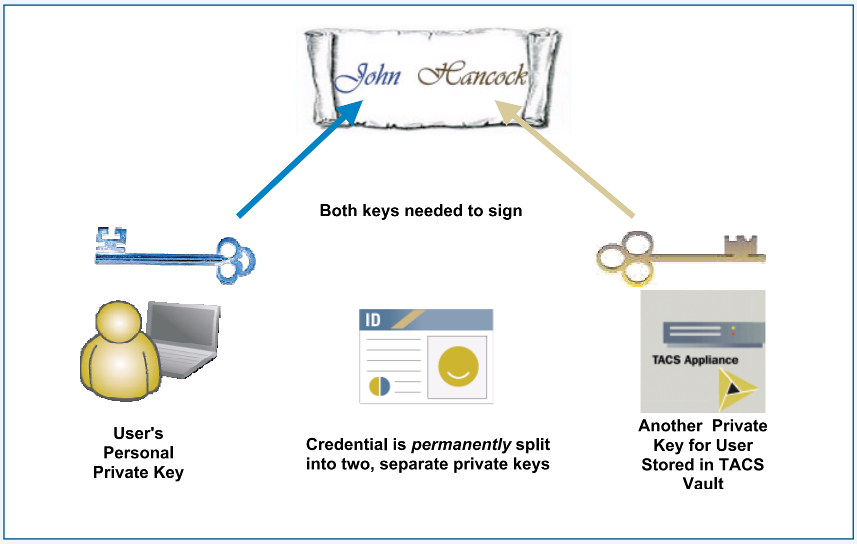
John Hancock
Both keys needed to sign
ID
TACS Appliance
User's Personal Private Key
Credential is *permanently* split into two, separate private keys
Another Private Key for User Stored in TACS Vault

Figure 2: Flexible multi-factor authentication.

(which is a conventional 2-key RSA private key) is used as the second factor, similar to the Device 2 Factor and Portable 2 Factor cases. This key never leaves the smartcard providing hardware protection for it. Most of the above 2 Factor solutions can be easily combined to form a 3 Factor credential. For instance Password + PC + USB Disk or Password + PC + Smartcard.

The TACS Vault itself is specially designed to afford a very high degree of assurance. It is protected using three layers of defenses: (i) it has a locked down, dedicated hardened OS, (ii) all system and user administration is strictly compartmentalized on least privilege, need to know, basis, and (iii) it uses FIPS 140-1 Level 2 rated cryptography. It is also highly scalable and fault tolerant, running as a set of 2 or 3 load-balanced and failover appliances. Finally it is a high assurance platform which can act as a secure storage facility to protect identity data such as credit card numbers. Enterprises can choose to either store identity data directly on the TACS Vault, or else can choose to encrypt data in place, and use the TACS Vault as a key management facility. The data is only available to authorized users after successful strong authentication.

Link:
http://www.tricipher.com

Please contact:
Ravi Sandhu, TriCipher, USA
Tel: +1 703 283 3484
E-mail: sandhu@tricipher.com

Australian Research Network in Security

by Ed Dawson

The Research Network for a Secure Australia (RNSA) is a multi-disciplinary collaboration, funded by the Australian Research Council for a period of five years, established to strengthen Australia's research capacity to enhance the protection of the nation's critical infrastructure from natural, human-caused, or accidental disasters, and terrorist acts.

The RNSA will facilitate a knowledge-sharing network for research organisations, government and the private sector to develop research tools and methods to mitigate emerging safety and security issues relating to critical infrastructure. The network will integrate complementary, yet diverse research areas including physical and information infrastructure security, and surveillance and intelligent systems. The RNSA has identified the majority of Australia's leading researchers, Commonwealth and state officers and industry leaders involved in critical infrastructure protection. This includes more than 300 researchers and professionals from 25 Australian research organisations, 15 government organisations, and an industry consortium comprising over 50 companies. The network has also identified a number of relevant international collaborators.

Objectives

The network will identify and facilitate the integration of research programs and collaboration in the areas of critical infrastructure protection (CIP). The RNSA will encourage and support:

- Open exchange of information and sharing of resources across disciplinary, organisational, institutional and geographical boundaries by organising workshops, focus groups and an annual conference.
- Development and implementation of cohesive and integrated research plans among researchers by bringing them together and encouraging communication opportunities for cross-disciplinary research collaboration.
- Nurturing the careers of young investigators and research students through incentives, such as attending an annual summer retreat, as well as opportunities to participate in international and inter-institutional exchange programs.
- Links with actual and potential end-users, and the broader community through an advisory board composed of recognised key stakeholders in Australian CIP.

Priority Research Area Recommendations

The RNSA's research strengths are indicated by the status of its world leading researchers who have extensive national and international linkages in relevant science, engineering and technology research. These researchers have been successful in attracting research funding and have access to research students and young investigator training programs. The RNSA incorporates a wide research base concentrating complementary research expertise that can address CIP challenges. Expertise pooled by the RNSA is broadly grouped into three research areas, composed of various programs as currently identified by the network participants. The research areas and corresponding research programs (see Figure) all contribute to the common research opportunities.

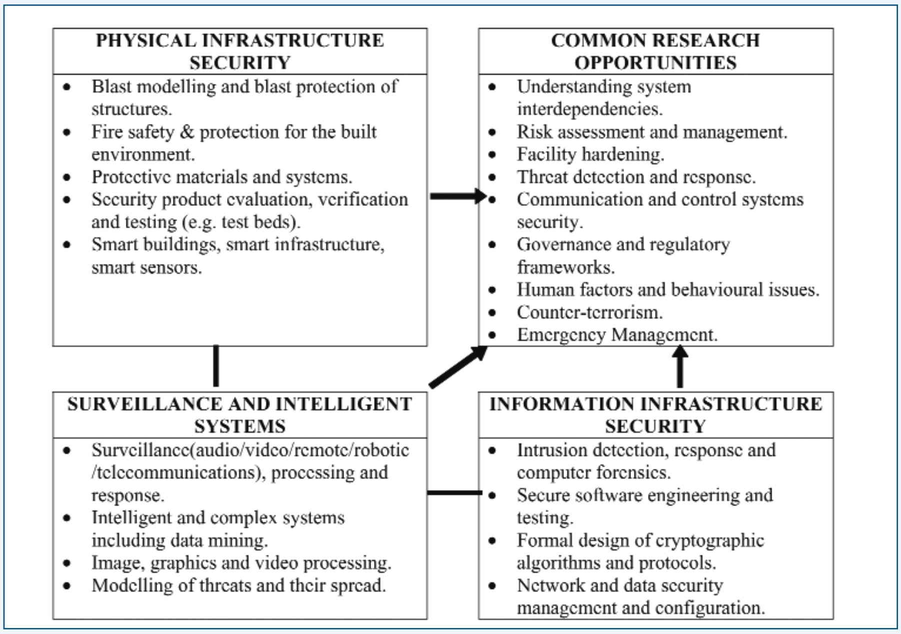

PHYSICAL INFRASTRUCTURE SECURITY
- Blast modelling and blast protection of structures.
- Fire safety & protection for the built environment.
- Protective materials and systems.
- Security product evaluation, verification and testing (e.g. test beds).
- Smart buildings, smart infrastructure, smart sensors.

COMMON RESEARCH OPPORTUNITIES
- Understanding system interdependencies.
- Risk assessment and management.
- Facility hardening.
- Threat detection and response.
- Communication and control systems security.
- Governance and regulatory frameworks.
- Human factors and behavioural issues.
- Counter-terrorism.
- Emergency Management.

SURVEILLANCE AND INTELLIGENT SYSTEMS
- Surveillance(audio/video/remote/robotic /telecommunications), processing and response.
- Intelligent and complex systems including data mining.
- Image, graphics and video processing.
- Modelling of threats and their spread.

INFORMATION INFRASTRUCTURE SECURITY
- Intrusion detection, response and computer forensics.
- Secure software engineering and testing.
- Formal design of cryptographic algorithms and protocols.
- Network and data security management and configuration.

Research areas and corresponding research programs.

National Benefit

The RNSA is a knowledge-sharing network for government, universities and the private sector producing innovative solutions to secure Australia's critical infrastructure from threats that have the potential to cause national security, economic, and/ or social impacts. Facilitating a coordinated approach to CIP will align the efforts of researchers and key stakeholders (ie government organisations and the private sector) in the broad areas of science, engineering and technology. The RNSA has established notable contacts and collaboration with equivalent or similar activities overseas, eg in Europe, Asia, USA, Canada, and NZ. The RNSA activities program and its outreach plan will foster the development of local expertise through the enhancement of postgraduate education and the encouragement of CIP researchers, having particular emphasis on cross-disciplinary approaches. This will ensure that overall security advice in relation to CIP will not need to completely depend on imported or overseas skills. The network serves as a vehicle for the dissemination of best research practices in CIP, as well as a repository of expertise to advise government and industry on CIP matters.

Link:
RNSA web site:
http://www.secureaustralia.org

Please contact:
Ed Dawson, Queensland University of Technology, Australia
Tel: +61 7 3864 9551
E-mail: e.dawson@qut.edu.au

Document Security

by Paul E. Sevinç

For over five years, researchers from around the world have worked on controlling access to XML documents. Typically, the focus has been on server-side mechanisms for record-like documents and the system architecture has assumed a centralized policy administration point. In contrast, we study client-side mechanisms for narrative-like documents with sticky policies.

The objective of this research project is to develop practical and comprehensive technical measures for document security. It commenced in September 2003, and is being conducted by the Zurich Information Security Center (ZISC) – in particular its members Prof. Dr. David Basin (Swiss Federal Institute of Technology Zurich), Dr. Günter Karjoth (IBM Zurich Research Laboratory), Beat Perjés (Credit Suisse), Dr. Gritta Wolf (Credit Suisse) and Paul E. Sevinç (Swiss Federal Institute of Technology Zurich).

Document Security is motivated by the fact that enterprises must secure many of the documents they process for reasons that include protecting a customer's privacy in accordance with the law, and gaining an advantage over competitors by not sharing trade secrets. Currently, these enterprises must resort to organizational measures, since technical ones are impractical, insufficiently comprehensive, or completely lacking.

The (long-term) vision of Document Security is to ensure that information in documents can be protected by mechanisms that enforce a security and privacy policy, and that the mechanisms are not limited to a particular platform or even

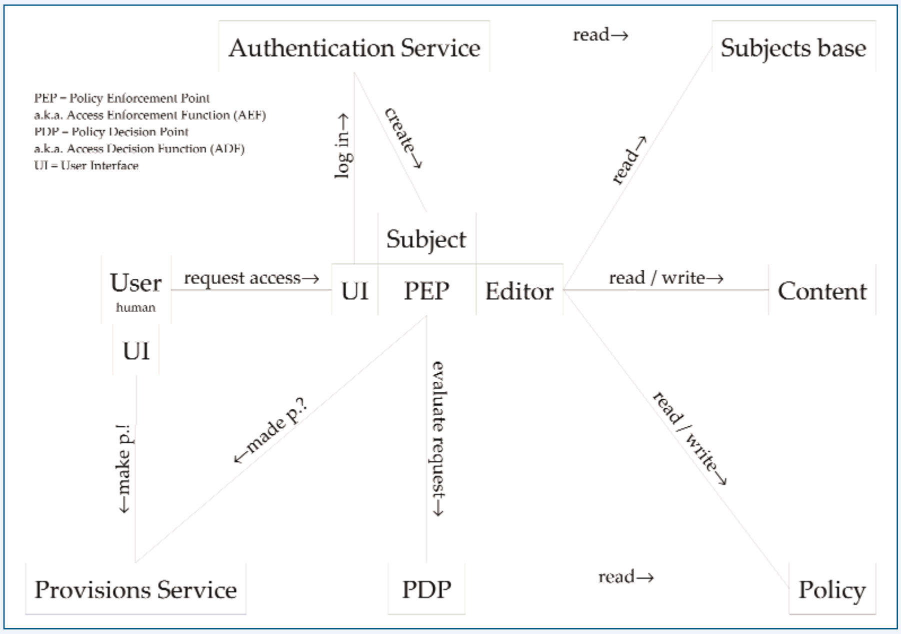

Document Security system architecture.

document processor. The threat model is that a company's stakeholders (employees, consultants, shareholders etc.) who access sensitive documents are not trusted, because:

- they may be careless in their use and distribution of data
- their software might be untrustworthy (eg compromised by a Trojan horse) even when the users are trustworthy or
- some may actually be dishonest.

This lack of trust means strong client-side control is required, which propels Document Security into the realm of rights management and trusted computing.

From the vision and the threat model, we derived the project scope of Document Security: the primary goals are content confidentiality (eg to protect trade secrets) and policy integrity (in order to keep attackers from simply assigning themselves arbitrary permissions). Microsoft coined the term 'enterprise rights management', since the first goal differs from the goal of digital rights management (DRM), which typically involves payment for access to non-confidential data (eg a movie). Privacy is outside the scope of our project.

It is worth mentioning that we derived our requirements not from what we as computer scientists consider challenging or interesting, but from actual business use cases of a major Swiss bank. For instance, users must be able to define different rules (policy) for different parts (content) within one and the same document so that they do not have to create several differently censored versions of a document. Another example is that it must be possible to specify that certain users have the permission to delegate permissions to other users so that they can also delegate certain document-processing tasks.

The documents we consider are a superset of XML documents. We have developed a formal (ie mathematically precise) model of this superset. Finally, we have formally defined the semantics of a policy language that implies the sticky-policy paradigm. That is, the policy sticks with the information and remains enforced when the information is transferred between documents (eg via cut and paste). The subjects are basically sets of properties (ie name-value pairs) that could be used for modelling roles for (hierarchical) role-based access control. The granularity of the objects is that of single attributes and nodes. The set of actions not only includes actions on the content (in particular 'read'), but also on the policy ('add rule', 'delegate permissions'). Furthermore, and in addition to conditions, we also support provisions and thus allow for access decisions to be made dependent on them (eg signing a non-disclosure agreement [NDA] or logging access).

Our next step will be to consider containers other than documents (eg a database). When the information originates from another type of container, the document will take the other container's policy into account, in accordance with the policy-combination or policy-mapping semantics (eg intersection of restrictions).

As for a prototype, we have designed a client system (perceived by the user, for instance, as a word processor) that will be implemented by students under our supervision. Content and policy are the components of a document tuple (see Figure). As the user (red) is not part of the trusted base (green), the integrity of the client system – particularly the confidentiality of secret keys – must be ensured by either operating-system mechanisms (if the computer is administrated by trusted employees of the enterprise) or by hardware-based mechanisms. While the Provisions Service (blue) is part of the trusted base, too, it is merely a local stub. Otherwise, dishonest users could simply trash their computer to destroy access logs or NDAs.

Document Security has resulted in a few spin-off projects such as a Trusted Platform Module (TPM) emulator for Linux and an evaluation of XML-diff algorithms. It also raises interesting questions in the domain of usability and security.

Link:
http://www.zisc.ethz.ch/research/docsec

Please contact:
Paul E. Sevinç, Information Security
GroupD-INFK and ZISC, Swiss Federal
Institute of Technology, Zurich, Switzerland
Tel: +41 44 632 7250
E-mail: paul.sevinc@inf.ethz.ch

Towards a European Malware Containment Infrastructure

by Kostas G. Anagnostakis and Evangelos Markatos

'LOBSTER' and 'NoaH' are two projects designing the necessary infrastructure to support research, development, and experimental deployment of advanced cyber-defence mechanisms.

Over the last few years, we have witnessed increasing levels of innovation among cyber-attackers, which, combined with the increasing penetration of broadband Internet service and the persistent vulnerabilities of host software systems, has led to new classes of rapid and scalable mechanized attacks on information infrastructure. Levelling the playing field requires scalable, automated responses to malicious code that can react as quickly as modern network worms propagate. Traditional approaches have relied on signatures, manual containment and quarantine. However, while tools are improving, progress in the development and deployment of the necessary technology is widely regarded as too slow for a threat that is so clear and imminent.

To address this problem, the Distributed Computing Systems Laboratory at FORTH-ICS has initiated and is currently coordinating two IST-funded projects, LOBSTER and NoaH, whose goal is to roll out the necessary infrastructure to support research, development and experimental deployment of advanced cyber-defence mechanisms.

LOBSTER aims at providing a pilot infrastructure for passive Internet traffic monitoring that will improve current understanding of the Internet, and will contribute towards solving difficult performance and security problems. Based on appropriate abstractions and cooperation among several points of presence, LOBSTER will help to monitor the underlying network, providing early warning of security incidents, as well as accurate and meaningful measurements of performance. The main goal of LOBSTER is to deploy an advanced pilot European Internet Traffic Monitoring Infrastructure, based on passive monitoring sensors at speeds starting from 2.5Gbps and ranging possibly up to 10Gbps. The architecture of the system is heavily influenced by and will make use of the knowledge, software and hardware artefacts obtained by a core group of project partners in SCAMPI, a recently completed IST research project.

NoAH is a three-year project developing infrastructure for security monitoring based on honeypot technology. Honeypots are computer systems that do not provide production services, but are instead intentionally made vulnerable and closely monitored to analyse attacks directed at them. NoAH will use geographically dispersed honeypots as an early-warning system, and will correlate the data received from them to generate automated warnings and possibly trigger appropriate containment measures.

The two projects use complementary approaches towards the same goal, namely, to help ISPs and National Research Networks limit the damage to their networks, allow information security organizations to better assess threats, and provide researchers with a wealth of attack-related data to improve detection techniques. The participants will be able to gather and analyse information about the nature of Internet cyber-attacks by developing an infrastructure to detect and provide early warning of such attacks, so that appropriate countermeasures may be taken to combat them. Both efforts are exploring opportunities for supporting other related initiatives, including Geant and the global Honeynet project.

Links:
http://dcs.ics.forth.gr/
http://www.ist-lobster.org/
http://www.fp6-noah.org/

Please contact:
Evangelos Markatos, ICS-FORTH, Greece
E-mail: markatos@ics.forth.gr

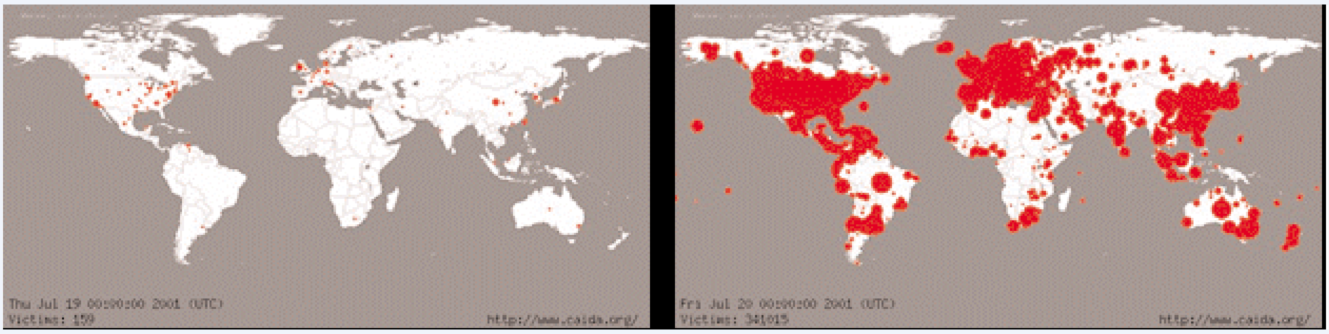

Modern worms have demonstrated that they can infect tens of thousands of computers worldwide in a few hours. Source: http://www.caida.org/ analysis/security/code-red/

iWATCH: Intelligent Watch Based on Networked Smart Sensors and Autonomous Mobile Vehicles

by Haris Baltzakis, Angelos Bilas and Panos Trahanias

The iWATCH activity aims at developing an embedded network of intelligent sensors and autonomous mobile platforms that will enable localization, identification and tracking of mobile assets, goods and people within the context of a real-life monitoring application.

Recent events and changes in society have created an increased demand for security, which in turn has forced governments and organizations to make personal and asset security a priority. For example, London, the most populated city in Europe, has recently instituted some huge security programs, largely based on CCTV surveillance systems. Because of these programs, thousands of CCTV cameras have been installed for surveillance in various areas of interest (eg more than ten thousand cameras in the London metro and Heathrow airport alone). However, cameras are of little use without some means of analysing the data they record. A review in 2002 by the Home Office, the British government department responsible for internal affairs, found that the cameras had a 'small effect' on crime and did not address terrorism at all; other studies have also been inconclusive. London's police admitted that regarding the utility of sensor data, manpower is a major limitation: for the cameras to be useful for spotting terrorists, an army of police would have to be on hand to scrutinize the images.

In addition, the current security boom is to a great extent an export of the United States. The 2004 Athens Olympics was the first summer Olympic Games since the September 11th terror attacks, and the most heavily guarded event in history. The security budget for this single event exceeded one billion euros – more than three times the amount spent on protecting the previous Olympic Games (Sydney 2000). Most security systems for the Athens Olympics were provided by US companies and more specifically by a group of companies under the US-based SAIC (Science Applications International Corporation).

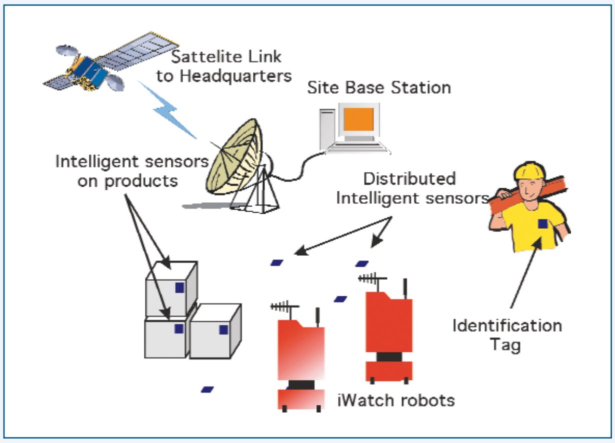

iWATCH scope and system components.

Most existing systems rely on stationary sensors such as cameras and motion detectors for surveillance and tracking. Easy deployment, scalability and customizability are also very important issues for security systems. For this reason, the current trend of taking over the processing, communication and sensing capabilities from workstations and gradually embedding them in the sensors themselves is expected to pick up in the years to come. In the long run, the 'disappearing computer' will hand over its competences to smaller devices embedded virtually everywhere in the environment. Usually such sensors do not possess any processing power and no communication takes place between them. Besides sensors, small passive information tags are commonly used for identification of products and persons as well as for anti-theft control.

iWATCH is an interdisciplinary activity at FORTH-ICS. It aims to use existing state-of-the-art technology and expertise in this important application domain, and conversely, to use the specific application domain to drive further work on the underlying technology. The focus of iWATCH is to investigate how situation awareness and security in large-scale premises and assets of global interest can be improved by integrating various devices under a common framework that involves:

- Low-cost intelligent sensors with RF communication, memory and processing capabilities can by deployed in large numbers to (a) facilitate easy creation of high-density sensor networks for monitoring large areas, (b) act as tagging devices for physical assets or commodities of interest, and (c) identify authorized personnel.
- Autonomous robotic vehicles can provide an automatic mechanism for tag locating and continuous tag mapping, as well as surveillance and specialized sensing capabilities during site patrolling or in case of an emergency. Besides offering specialized surveillance functionality, mobile robotic platforms can at any time poll information about intelligent sensors and update the database with sensor data history. They can also improve esti-

mates of sensor location, thus overcoming fundamental limitations in sensor-sensor communication and placement.
- The envisioned functionalities will be conveyed under a joint operational environment, offering complete, centralised situation awareness, supported by satellite communication technology. The iWATCH project addresses the deployment of existing hardware components with appropriate middleware, as well as the development of sensor-based security applications based on the resulting embedded network system and the capabilities offered therein. Finally, iWATCH will offer higher level integration at a centralized location, aiming at complete situation awareness among geographically distributed sites. The 'tactical surveillance and control picture' generated by each application site can be transmitted by reliable location-independent means such as space relays (satellites), and combined at a central station of choice (headquarters or control room) to form a unified operations environment. This environment, presented as an 'information wall', is the fundamental instrument for achieving situation awareness.

Finally, iWATCH brings together researchers from autonomous vehicles, intelligent sensors, networking, and signal processing and capitalizes on existing, cutting-edge expertise and technology in order to provide automated tracking, identification, mapping and inventory control for assets and persons as well as patrolling by autonomous mobile vehicles in large-scale premises.

Link:
http://www.ics.forth.gr/~trahania/

Please contact:
Panos Trahanias, ICS-FORTH, Greece
Tel: +30 2810 391 715
E-mail: trahania@ics.forth.gr

Securing JAFAR – An Architectural Framework for E-Commerce Applications

by Nicolas Guelfi, Cédric Pruski, Jan Reimen and Paul Sterges

Security is the key to the success of e-commerce applications. In the context of the 'FIDJI' project, the Software Engineering Competence Center (SE2C) at the University of Luxembourg has developed JAFAR, a J2EE (Java 2 Platform, Enterprise Edition) architectural framework for the development of secure e-commerce applications. Research is now focusing on the various security problems that are directly linked to the use and development of such applications.

Since the early nineties and the advent of distributed e-commerce, computer science has been confronted with problems related to security, in both the exchange and server-side handling of data. The key to success for e-commerce applications resides in gaining the trust of potential users. Applications must therefore be developed in order to offer adequate levels of security. The ADS (Architecture Engineering for Dependable Distributed Systems) team of the SE2C is interested in three major problems relating to security and trust management in the development and use of distributed e-commerce applications.

First, in the eyes of users, data confidentiality during an exchange between the client application and the server constitutes the most critical point of security. This phenomenon, which remains purely a network problem, increased in importance with the arrival of third-generation telecommunications networks and the appearance of the mobile commerce (m-commerce) paradigm. The use of an architectural framework such as JAFAR allows the developers of distributed applications to take into account and to more easily solve, at the software level, this problem of data exchange. This is managed in particular through the use of security solutions already implemented in the heart of the framework (eg communication interfaces supporting the TLS protocol, the use of certificates to encrypt the data etc). In addition, where the level of confidentiality demands it, the framework is enriched by the development of other components that integrate more powerful means of encryption. Security certificates also are important to authenticate the e-commerce site to the users. We are currently studying the various security solutions that can be used in the context of e-commerce and integrated in JAFAR, in particular solutions over UMTS networks.

Second, both the persistence and the availability of data are significant points in the success of e-commerce applications. Indeed, considering the company Amazon, an outage of its Web interface would have a very significant effect on its sales turnover. Thus, the applications must be designed in order to prevent remote attacks such as denial of service attacks, which target the availability of resources. During the development of JAFAR, we have studied the various possibilities for detecting this type of attack, as well as the measures to be taken when a major resource is no longer available. In particular, we are working on the decentralized aspect of the resources in order to reduce as far as possible the length of time for which they are unavailable. However, we also envisage integrating the adaptation and resilience aspects of the system in the near future, by means of other projects such as CORRECT, privileging the fault-tolerance as-

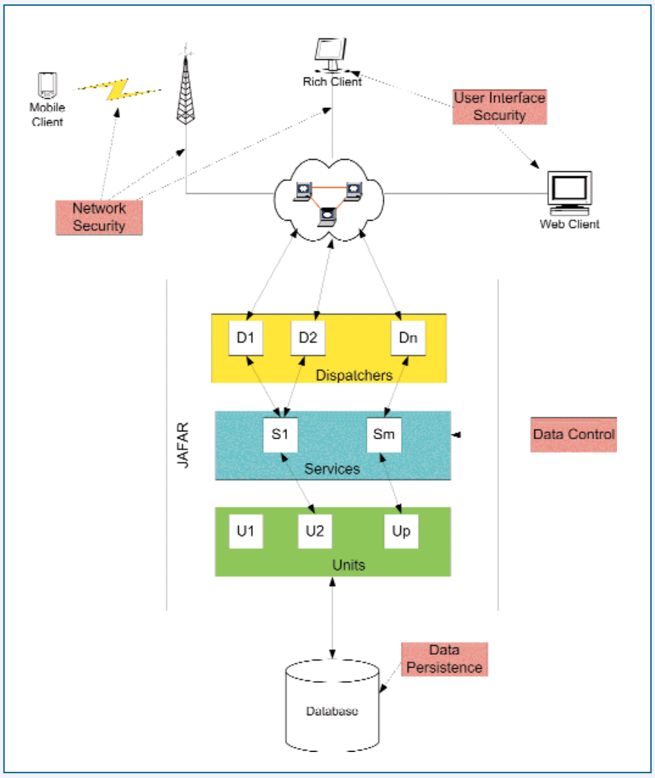

Architecture of JAFAR.

pect, or new projects awaiting validation from the European commission.

Lastly, since storage security is relative to the business application and its supporting storage infrastructure, a risk mitigation methodology is a sound way to strengthen storage availability, reliability and privacy. This is why storage of and access to stored data are the final two significant points that influence the design of e-commerce applications. The layer architecture of JAFAR is conceived in a rigorous way and determines the permissions granted to users after their identification. In addition, JAFAR was developed in order to offer a maximum of services to users while reducing the number of interfaces that make it possible to seize information or introduce erroneous data. The latter particularly includes code that is interpretable by the system, such as SQL or HTML, which can have disastrous effects on stored data. We are studying ways to define security rules which will be checked at run-time and which can be changed on a production system to react to observed patterns of abuse. Moreover, JAFAR integrates a module making it possible to rigorously control the data to be stored by detecting possible pieces of malicious code.

Links:
SE2C homepage: http://se2c.uni.lu

FIDJI Project: http://se2c.uni.lu/projects/fidji

CORRECT Project:
http://se2c.uni.lu/projects/correct

Please contact:
Nicolas Guelfi, University of Luxembourg
Tel: +352 420 101 251
E-mail: Nicolas.Guelfi@uni.lu

Distributed Reputation Systems for Internet-based Peer-to-Peer Systems and Mobile Ad-Hoc Networks

by Jochen Mundinger, Sonja Buchegger and Jean-Yves Le Boudec

Reputation systems are widely and successfully used in centralized scenarios. Will they work equally well, however, in decentralized scenarios such as Internet-based peer-to-peer systems and mobile ad hoc networks?

Reputation and recommender systems are widely used to provide a basis for the choice of transaction items and partners in online scenarios. They have proven useful in online auctioning systems such as eBay or online book stores such as Amazon, and can be viewed as a substitute for the word-of-mouth mechanisms observed in offline personal encounters.

As a result, it has recently been proposed that reputation systems be extended to decentralized and self-organized systems such as Internet-based peer-to-peer systems and mobile ad-hoc networks. The need here is predominantly to provide incentives for cooperation and to protect the network from node misbehaviour. However, reputation systems need to be distributed in decentralized scenarios and their potentially complex behaviour is not yet fully understood. This understanding is necessary to investigate whether distributed reputation systems might prove as useful as their counterparts for centralized systems. In a collaborative effort started in 2004 between EPFL, SIMS and the Statslab, we are therefore analysing the behavior of distributed reputation systems and evaluating whether and how they can best be used.

We have developed a reputation system to detect, discourage and stop node misbehaviour in Internet-based peer-to-peer sys-

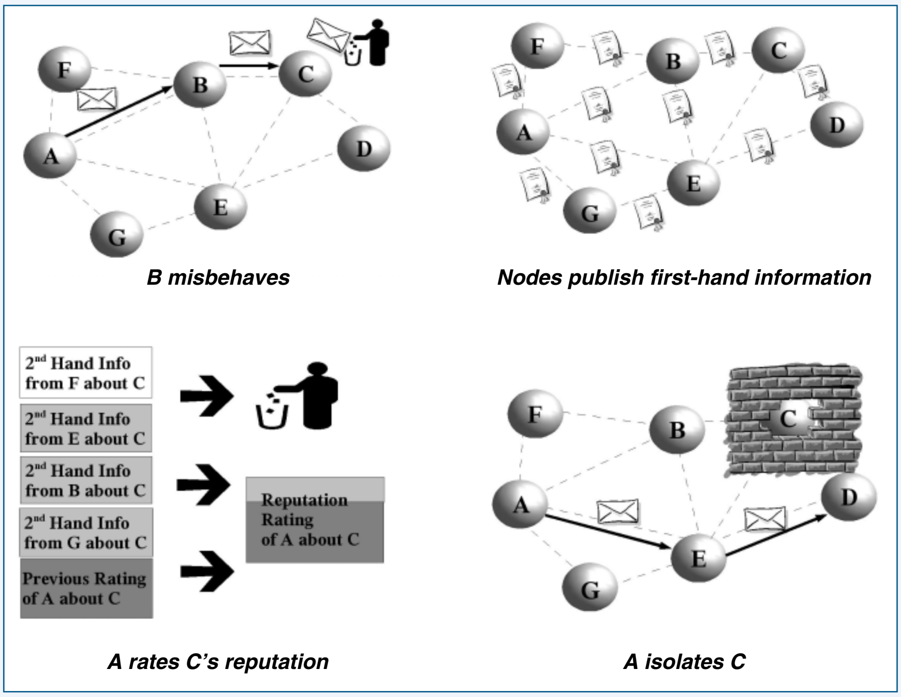
F
B
C
A
D
E
G
B misbehaves
Nodes publish first-hand information
2nd Hand Info from F about C
2nd Hand Info from E about C
2nd Hand Info from B about C
2nd Hand Info from G about C
Previous Rating of A about C
Reputation Rating of A about C
A rates C's reputation
A isolates C

Distributed Reputation System in a Mobile Ad-hoc Network.

tems and mobile ad hoc networks. We proposed a protocol called CONFIDANT (Cooperation Of Nodes — Fairness In Dynamic Ad hoc NeTworks) to cope with misbehaviour. Detection is based both on first-hand observation and on second-hand information provided by other nodes.

To be useful, reputation systems need to be reliable and robust against liars. They can, however, be tricked by the spread of false reputation ratings. On the other hand, simple solutions such as exclusively relying on one's own direct observations do not make use of all the information available.

Our fully distributed reputation system is based on a modified Bayesian estimation procedure and can cope with false information so as to effectively use second-hand information. Each node maintains a reputation rating and a trust rating for all other nodes it cares about. Reputation ratings capture the quality of the behaviour of a node as an actor in the network performing routing and forwarding tasks. From time to time reputation information is exchanged with others. However, second-hand information is only accepted if it is compatible with the current reputation rating or comes from a trusted node. To decide whether it is compatible, we introduce a simple mechanism called the deviation test. If the received estimation of misbehaviour deviates more than a threshold value from the current one, it is rejected, otherwise it is accepted. The trust rating is updated each time a deviation test has been performed, meaning compatible nodes are thus more trusted.

We use simulation to evaluate and demonstrate the performance. We found that CONFIDANT can keep the network performance high even when up to half of the network population misbehaves. We showed that our approach of using second-hand information significantly speeds up the detection of misbehaving nodes while keeping the number of false positives and negatives negligibly low.

Simulation results also suggest that the deviation test performs nearly as well on its own without the trust component. We have therefore analysed it in more detail and in a more general context. We introduced a mathematical model for a distributed reputation system based on the deviation test. We show that it exhibits a phase transition behaviour, ie critical parameter values exist, below which liars have no impact and above which liars do have a certain impact and corrupt the reputation system. The critical values are obtained via a mean-field approach and are confirmed by direct computation as well as simulation. We thus give guidelines for a good choice of parameters. We also provide insight into a fundamental design question of such systems, namely whether or not direct observations only should be passed on as second-hand information.

Our current focus is on combining our simulation-based approach with the analytical approach. In particular, we are applying the insights gained by the analysis to find and zoom in on relevant parts of the experiment space. We are interested in seeing whether and how the phase transition effects observed in the analytical work carry through to our simulations of a more complex system, such as a mobile ad hoc network implemented in GloMoSim. Based on our analytical results, we now know what sort of behaviour we need to look for in the simulations. This will then enable us to see how the guidelines for a good choice of parameters from the analytical approach translate into a good choice of parameters for the more realistic system. We thereby optimize the performance of the distributed reputation system.

In conclusion, we provide a methodology to identify potential sweet spots and irregular behaviour of complex distributed reputation systems in decentralized self-organized systems. We expect to complete the combination of the two approaches by the end of the year.

Links:
http://lcawww.epfl.ch/salem/IP4Web/ImmuneNetworking.htm
http://winet-coop.epfl.ch/
http://icapeople.epfl.ch/leboudec
http://www.sims.berkeley.edu/~sonja
http://www.statslab.cam.ac.uk/~jm288

Please contact:

Sonja Buchegger, SIMS, UC Berkeley, USA
Tel: +1 510 643 2251
E-mail: sonja@sims.berkeley.edu

Jochen Mundinger, Statistical Laboratory, University of Cambridge, UK
Tel: +44 1223 337952
E-mail: J.Mundinger@statslab.cam.ac.uk

Jean-Yves Le Boudec,
EPFL-IC-LCA, Switzerland
Tel: +41 21 69 36631
E-mail: jean-yves.leboudec@epfl.ch

Secure Wireless Ad-Hoc Networking

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

The lack of a fixed infrastructure in ad hoc networks forces ad hoc hosts to rely on each other in order to maintain network stability and functionality. It also introduces several problems relating to security. This project focuses on the security issues of ad hoc networks employed to meet specific emergency-preparedness requirements. Examples of such emergency networks include those deployed for disaster relief efforts following natural disasters or terrorist attacks, or for military operations.

Consider a disaster situation, a terrorist attack for example, in which a wireless network needs to be formed on an ad hoc basis, without the support of any fixed infrastructure, in order to interconnect all relevant computing and communication devices. The objective is to share information with the highest security possible, since no-one can guarantee that 'high-tech' terrorists/attackers won't try to disrupt or intercept the rescue efforts. However since neither a certification authority nor a secure communication channel exists, attackers have the ability to eavesdrop and tamper with messages transmitted over the air. Additionally, since no Identification Authority is present, group members may easily be impersonated.

In the above scenario, the main problem involves group members establishing a secure wireless network and at the same time eliminating outside threats. Moreover, in the case where a new node arrives and wishes to become a member of an existing group, how can this particular node participate in the group session without distracting the initial group? Having entered the group session, how can this node then gain the same privileges as other group members? And of course, in a more general case where a whole group is formed from scratch, how can the network be extended to incorporate it?

In order to answer the above questions, the Telecommunications and Networks Laboratory (TNL) of FORTH-ICS has introduced a new family of security protocols for wireless ad hoc networks. The security problems addressed by our system are:

- *Contributory key establishment.* An ad hoc network is set when a session key is agreed to by all network entities. This session key is generated through a process where all participating entities contribute equivalently.
- *Secure authentication.* Strong authentication must be derived from a small password.
- Resilience against attacks during the key formation:
  - *Forward authentication.* If a malicious party manages to compromise a

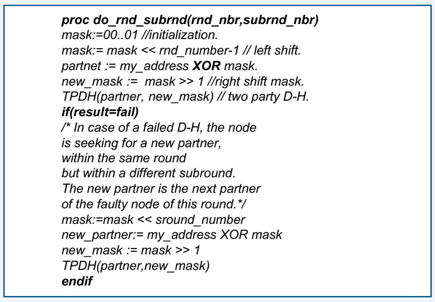

```
proc do_rnd_subrnd(rnd_nbr,subrnd_nbr)
mask:=00..01 //initialization.
mask:= mask << rnd_number-1 // left shift.
partnet := my_address XOR mask.
new_mask :=  mask >> 1 //right shift mask.
TPDH(partner, new_mask) // two party D-H.
if(result=fail)
/* In case of a failed D-H, the node
is seeking for a new partner,
within the same round
but within a different subround.
The new partner is the next partner
of the faulty node of this round.*/
mask:=mask << sround_number
new_partner:= my_address XOR mask
new_mask := mask >> 1
TPDH(partner,new_mask)
endif
```

Figure 1: The node algorithm for finding a partner and performing a DH two-party key exchange.

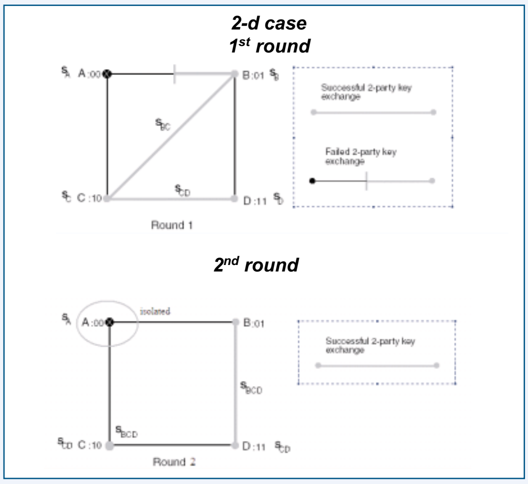

Figure 2: The 2-d case of the general d-cube protocol.

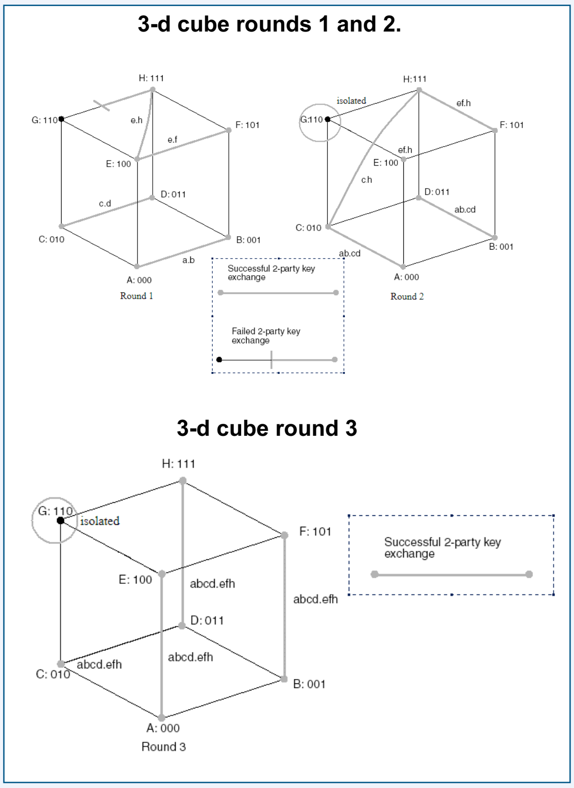

Figure 3: The 3-d case of the general d-cube protocol.

network node, he will be unable to participate in the network

- *Tolerance against passive attacks.* Even if an attacker inserts, deletes, or modifies the key-formation messages exchanged among legitimate entities, the security of the network will not be compromised.

Within this project, we have modified protocols for password authenticated multi-party Diffie-Helman (DH) key exchange, to make them more resistant against dictionary attacks. Figure 1 presents the node algorithm for finding a partner and performing a DH two-party key exchange, Figure 2 depicts the 2-d case, which will be described below, and Figure 3 depicts the 3-d case of the general d-cube protocol.

In the 2-d case it is considered that node A is a faulty partner. Therefore, in round 1, node B will fail to complete the Diffie-Hellman key exchange with A. However (see Figure 1), instead of remaining idle during this round, node B will perform a new DH key exchange with node C. This exchange will be successfully completed resulting in a new key and at the same time informing that is faulty. Meanwhile, node C has already performed a successful key exchange with node D creating key . Round 1 will be concluded with the creation of the keys and . On the second round, node B will perform a DH key exchange with node D and the key will be created. This key will be the common session key for the entities . Node C can calculate the common session key for itself, since it has all necessary information needed from the previous round. The main point in this scenario is that during the second round node won't have to perform any DH key exchange with the faulty node A. is being isolated because its neighbors in earlier rounds, than those foreseen by the d-cube algorithm, have performed extra DH key exchanges due to its behavior.

Additionally, we have also introduced new key establishment methods based on elliptic curve cryptography due to its lighter computational requirements, which is important for ad-hoc networks consisting mostly of thin clients. Finally we have proposed new faster protocols for the dynamic case, where the composition of the ad-hoc network changes in time with the arrivals and departures of nodes.

The current project was developed by TNL under the Information Security Program of FORTH-ICS.

Please contact:
Ioannis G. Askoxylakis, ICS-FORTH
Tel: +30 2810 391723
E-mail: asko@ics.forth.gr

Decentralized Identity for the Digital Business Ecosystem

by Jean-Marc Seigneur

Trustworthy decentralized identity mechanisms are promising to foster the Digital Business Ecosystem (DBE), an EU-funded FP6 IST Integrated Project. While progress has been made and driver SMEs are lobbying for more, such mechanisms still remain on the research agenda.

The Digital Business Ecosystem (DBE) funding consists of a €14 million three-year research project supported by the European Commission's 6th Framework Programme IST Thematic Priority. The project started in November 2003 and is planned to end in October 2006. Twenty partners from ten EU Member States are involved. Regional involvement ensures that the end objective is to benefit SMEs.

Thanks to the technical commons provided by the DBE, even micro-companies will soon have access to advanced information and communication technologies to grow their business. A minimal set-up example may be a simple broadband connection to the Internet combined with a running peer-to-peer version of the DBE SERVENT (SERVer/clieNT) to have access to the free common DBE services (such as business Web presence and networking). Thanks to the open-source DBE Studio, more customized services or revenue-generating owned services could just as easily be developed and served.

This example highlights a few visible aspects in favour of the sustainability of the DBE: the use of basic means to access the Internet; the provision of valuable open-source building blocks; and the unobtrusive sharing of unused computer resources of those who found value in running the basic building blocks. The research so far has focused on the scalability, availability and reusability of the technical commons. Advances have been made in peer-to-peer replication; firewall/NAT transparent transversal; service composition; business ontology modelling, learning and evolution; and self-organization of service proxy and super peers.

The interest from driver SMEs has been so significant that many of these driver companies have lobbied for a faster implementation of a more secure version of the technical commons. Indeed, without authentication of the interacting entities, the three main security properties – confidentiality, integrity and availability – can be trivially violated.

Unfortunately, the current state of the art in security for identity management is challenged by open large-scale decentralized environments, such as peer-to-peer on top of the Internet. The Identity Gang task force has recently been set up to discuss whether or not the new Microsoft Identity Meta-system proposal is a sound basis on which to build decentralized identity management. This proposal seems to scale down the scope of Microsoft's previous Passport effort with regard to identity management. It appears to indicate that federated identity management may never be globally adopted by SMEs due to their overheads.

In addition, in the spirit of a self-organized DBE, identity management in DBE should be an entity-centric identity management solution rather than a system-imposed identity management. Entity-centric identity management is inherently decentralized because any entity is free to choose how its identifiers and credentials are managed. The current proposal is to store IDBEs – the identifiers and their associated credentials – in CREdential Servers (CRES), similarly to the GRID MyProxy credential repository. In fact, the GRID initiative has similar security requirements to those of DBE, especially to coordinate resource sharing in dynamic, multi-institutional virtual organizations. The worldwide Grid community has put a lot of effort into security between decentralized virtual organizations. It therefore makes sense to reuse their work, which eventually consists of a comprehensive tool kit called the Grid Security Infrastructure (GSI). However, since securely using the GSI involves quite a steep learning curve, we envisage a simplified version of the GSI, with more convenient graphical user interfaces than the basic GSI command line tools. We are also working towards a more portable CRES than the MyProxy credential server.

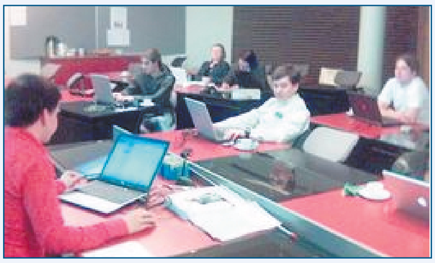
Photo: Henri Bergius / Nemein.

The first DBE driver companies workshop.

The CRES may be either local or remote and would be used to retrieve X509 Proxy Certificates credentials when needed. The advantage of a remote CRES is that users can use IDBEs on different computers without undertaking the risky task of moving the long-term credentials between these computers. If the CRES is managed by a professional CRES DBE service provider, the long-term credentials become better protected thanks to the knowledgeable security staff of the provider. The choice of the trusted Certificate Authorities (CAs) is left to the SERVENT owners. By default, the trusted CAs may be limited to CAs run by known DBE regional catalysts. However, our approach is open to external providers such as the current main Internet CAs or specific peers running a DBE CRES service. Thus, we allow the DBE peers to use a spectrum of technical trust in IDBEs: from self-signed, to certified by a web of computational trust or free CAs, such as CAcert, to current professional CAs including insurance and fraud protection services.

Links:

DBE website: http://www.digital-ecosystem.org

The DBE team at Trinity College Dublin (thanks to DSG DBE for the general discussions and David O'Callaghan for his comments with regard to Grid security): http://www.dsg.cs.tcd.ie/?category_id=-55

GRID Security Infrastructure: http://www.globus.org/security/

Computational trust: http://www.trustcomp.org/

The Identity Gang: http://www.identitygang.org/

Example DBE driver SME lobbying for more security: http://yukatan.fi/display/yukatan/2005/07/12/DBE+updates

Please contact:

Jean-Marc Seigneur,
Trinity College Dublin, Ireland
Tel: +353 1 608 1761
E-mail: Jean-Marc.Seigneur@trustcomp.org

Towards Privacy-Aware Identity Management

by Ernesto Damiani, Sabrina De Capitani di Vimercati
and Pierangela Samarati

The overall goal of the PRIME project (Privacy and Identity Management for Europe) is the development of a privacy-enhanced identity management system that allows users to control the release of their personal information. The PRIME architecture includes an Access Control component allowing the enforcement of protection requirements on personal identifiable information (PII).

Nowadays, a global information infrastructure connects remote parties worldwide through the use of large-scale networks, relying on application-level protocols and services such as the World Wide Web. Human activities are increasingly based on the use of remote resources and services, and on the interaction between remotely located parties that may (and sometimes should) know little about each other. Because of the vast amount of personal information thus available, concerns are growing regarding the privacy of users: effective information sharing and dissemination can take place only if users have some assurance that disclosure of sensitive information is not a risk. Digital identity management is therefore of paramount importance for supporting successful interaction. A comprehensive identity management solution should provide complete support for the definition and the life-cycle management of digital identities and profiles, as well as infrastructure for exchanging and validating this information.

Emerging scenarios of user-service interactions in the digital world are also pushing toward the development of powerful and flexible privacy-enhanced access control models and languages. The need for privacy means that access control policies and models must be rethought, and new forms of authorization specification and enforcement developed. In particular, two major issues exist: i) access control needs to operate even when interacting parties wish to remain anonymous or to disclose only specific attributes about themselves; ii) data collected during access control as well as data stored by the different parties may contain sensitive information on which privacy policies need to be applied.

In the context of the PRIME project, our main task is the development of the Access Control Decision Function (ACDF) module, together with the definition of a privacy-aware model and language for specifying and enforcing protection requirements on PII.

The access control component is based on a simple and expressive language whose main features are summarized as follows:

- Flexible and expressive access control rules. Access control rules make use of partial identities associated with users. It is also possible to specify access control rules relating to subjects accessing the information and to resources to be accessed, in terms of rich ontology-based metadata.
- Interactive enforcement. An access control component may not have all the information it needs to decide whether or not access should be granted. On the other side, requesters may not know in advance which information they need to present to get access. As a consequence, the access control process is a way of negotiating with the access requester the disclosure of additional personal information to achieve a final access decision.
- Client-side restrictions. In addition to traditional server-side access control rules, users should be able to place restrictions on the use of their personal information once released to a third party. For this purpose, we introduce the notion of release policies governing the release of properties, credentials and PII of the party.
- Anonymity and end-user control. The access control system enables full end-user control over the digital identity to be used. In other words, access control needs to operate even when interacting parties wish to remain anonymous or to disclose only specific attributes about themselves.
- Interchangeable policy format. Parties need to specify protection requirements on the data they make available using a format that is readable by both humans and machines, and is easy to inspect and interchange. The language therefore has a simple declarative form.

The ACDF module is under development and will be integrated with the PRIME architecture, which will be developed in collaboration by Compagnie IBM France, IBM Research GmbH (Switzerland), Unabhangiges Landeszentrum fuer Datenschutz (Germany), Technische University Dresden (Germany), Deutsche Lufthansa AG (Germany), Katholieke Universiteit Leuven (Belgium), T-Mobile Deutschland GmbH (Germany), Hewlett-Packard Ltd (UK), Karlstads Universitet (Sweden), Universita' degli Studi di Milano (Italy), Joint Research Centre (Italy), Centre National de la Recherche Scientifique (France), Johann Wolfgang Goethe Universitaet Frankfurt (Germany), Chaum LLC (USA), Rheinisch-Westfalische Technische Hochschule Aachen (Germany), Institut EURECOM (France), Erasmus Universiteit Rotterdam (The Netherlands), Stichting Katholieke Universiteit Brabant (The Netherlands), Fondazione Centro San Raffaele del Monte Tabor (Italy), and Swisscom AG (Switzerland).

Links:
PRIME project: http://www.prime-project.eu.org/

Security Group at the Dipartimento di Tecnologie dell'Informazione, Universita' degli Studi di Milano: http://seclab.dti.unimi.it

Please contact:
Pierangela Samarati
Universitá degli Studi di Milano, Italy
E-mail: samarati@dti.unimi.it

CareGrid: Autonomous Trust Domains for Healthcare Applications

by Naranker Dulay, Emil Lupu, Morris Sloman, Jean Bacon, David Ingram and Ken Moody

The overall goal of the CareGrid project will be to develop middleware for supporting decisions based on trust, privacy, security and context models. Health care will be used as the application domain, but the middleware developed will be applicable to other e-science applications.

Future large-scale health care will involve many different organizations cooperating in patient care, including hospitals, GPs, dentists, pharmacies, drug companies and insurance companies. With the advent of new wireless healthcare devices, it is becoming feasible to contemplate new applications that offer real-time health care to patients, and involve complex interactions between many services in many organizations.

Consider a simple scenario in which a patient with an acute heart condition subscribes to a monitoring service that provides wearable sensors and a small wireless controller. These devices send monitored information to the service centre and provide feedback, if necessary, to the patient from a medic. If an emergency is detected, the monitoring service calls an ambulance. The monitoring service needs access to patient cardiac history from the patient's GP and from the hospital where the patient had treatment, and so liaises with the emergency services and the hospital to which he/she will be taken for emergency treatment. Assume the monitoring service also provides anonymous monitoring records for medical research. Hospitals need to interact with the patient's GP and possibly social services in order to provide care for a patient following treatment. In a small hospital, there may not be sufficient local expertise to evaluate patient information such as X-rays and ECG readings, and so these need to be sent to a remote expert over the network. Perhaps the patient's usual consultant is unavailable and a new trusted one must be chosen: this is a form of trust-based choice of service. A consultant evaluating an X-ray or an ECG may wish to search for similar examples via a medical services grid but then the question of trust in the source of the examples arises. Issues of trust, privacy, security and context pervade this simple scenario.

CareGrid aims to provide middleware for organizing and coordinating trust, privacy and security decisions across collaborating entities using autonomous trust domains and context. Trust domains can be federated and/or grouped in hierarchical or P2P fashion. This requires protocols for group-membership and trust negotiation, as well as an overarching architecture that is self-managing. Examples include a body-area network monitoring the health of a patient, a team of care workers responsible for a patient, and a hospital or regional health authority. In addition to trust domains, specific components of the CareGrid architecture include the following:

- A trusted communication layer somewhat akin to SSL/TSL that supports trust negotiation, privacy control and evidence collection. Requests for trusted interactions are forwarded to the local trust domain, which is responsible for determining whether the interaction should succeed or fail. This can involve negotiation with the trust domain of the requester as well as other trust domains. The trust domain is also capable of providing a signed statement of the reasons for success or failure. If a request succeeds, the system will typically establish a new secure channel and trigger any necessary security adaptations, eg in the communications or access control systems.
- A language framework for specifying trust, privacy, security and management policies.
- A federated access control model suitable for expressing authorization policy for dynamic trust-domains and dynamically created security associations, in particular to dynamically change authorizations, or mandate changes to the security policy in response to changes in trust and context.
- An evidence service that collects, filters, synthesizes and anonymizes experience, risk, recommendation and reputation data that can be used as evidence for trust evaluation. Note that evidence may have to be archived for audit and statistical evaluations. The evidence service will include anonymization mechanisms to maintain the usefulness of evidence data to the greatest possible extent, while still honouring privacy requirements.
- A context management service that allows trust, privacy and security to be related to context, and for triggering trust-privacy-security adaptations when context changes. Examples of context include the location of a person or device, the time of day, environmental readings, physiological state (eg heart rate), patterns of past behaviour, user preferences and current roles. The context management service will support context schemas, context sensing and flexible context querying. Initial work will be performed at Imperial College on incorporating uncertainty into context values and defining functions over uncertain contexts.
- Mechanisms for protection against attacks on the trust-privacy-security-context infrastructure.

The CareGrid project is a collaborative project between groups at Imperial College London and the University of Cambridge. The two groups have common and complementary expertise in distributed and ubiquitous systems, including security and trust management. The project is funded by the UK's EPSRC and is due to start in October 2005.

Please contact:
Naranker Dulay, Imperial College London, UK
Tel: +44 20 7594 8288
E-mail: nd@doc.ic.ac.uk

Secure UPnP and Networked Health Care

by Kari Keinänen and Mika Pennanen

For today's rapidly growing mobile environments, VTT Information Technology's research is providing middleware for the development of networked health-care systems, and other applications for various end-user devices (PDAs, mobile phones etc) and wireless networks (WLAN, Bluetooth etc). The key point is the addition of network facilities to existing applications, rather than the development of new applications.

Resource discovery and communication are fundamental problems in the networking of devices and services. Universal Plug and Play (UPnP) is a widely accepted solution for discovering, controlling and monitoring networked appliances. Network installation becomes simple; furthermore, networks can be built in which one terminal controls all appliances and each appliance can be controlled by many different control points. However, UPnP does not specify sufficient security mechanisms. Secure UPnP was therefore developed to ensure that only authorised nodes can control and monitor devices.

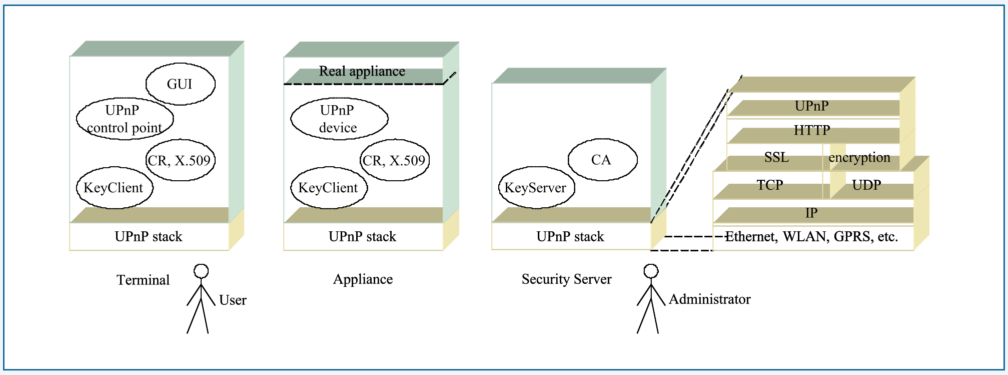

Figure 1: Secure UPnP architecture.

Security Solution

Our Secure UPnP provides authentication of hosts, data confidentiality and integrity, as well as key management. We employ well-known and proven security components, in particular Secure Sockets Layer (SSL) and X.509 certificates.

SSL is widely used, for example, to secure bank account access over the public Internet. We use SSL to secure all TCP traffic, which carries most of UPnP messages. To establish an SSL session, each node must have a X.509 certificate for authentication. Certificates are granted by a local Certificate Authority (CA) but only if the Administrator has accepted the new node. UPnP discovery phase uses UDP where it is not possible to use SSL, but we encrypt UDP data. The UDP encryption key is shared by the whole network and distributed using SSL.

Application Areas

Secure UPnP makes it possible to build secure networks that are easy to install and have multiple control terminals. A variety of physical networks can be used and shared with other applications. Application areas include health care in homes, hospitals, gyms and outdoor sports, home networks, building networks, industrial automation, sensor networks, and transport telematic networks.

The Networked Health Care system

VTT pioneered the concept of 'overall personal health-care information system'. Networked Health Care can be used with several health care instruments, eg scales, exercise cycles or fitness steppers to improve people's health. We work with several personal scenarios that aim to effectively control a person's weight and improve his/her overall fitness. We focus on the middleware layer and communication solutions. Users receive advice and information and the option of exercise training lessons with their own personal instructor, eg motivational virtual cycling environments with adjustable resistance. Providing detailed exercise feedback is a good way of encouraging people to manage their own personal health care.

The Networked Health Care system can be integrated with a number of appliances: for instance, an exercise cycle which lets you choose the speed of resistance, with the motivating virtual cycling environment working as a remote control. Personal health information such as heart rates, speed and calorie consumption/supply are collected directly from the Networked Health Care instruments (eg scales, exercise cycle) or from the user's profile (eg eating habits). Such information helps to monitor and control weight and fitness progression, to encourage physical exercise and to improve the user's health.

The system provides secure communication between instruments (see devices below) and only authenticated devices can join the Secure UPnP network. The network is a self-organizing system, which allows 'on-the-fly' associations between application entities and network resources without complex configuration.

Future work will concentrate on adding new appliances to the system. We also intend to provide remote controlling support between private networks.

Please contact:
Kari Keinänen, VTT, Finland
Tel: +358 9 456 5673
E-mail: kari.keinanen@vtt.fi

Mika Pennanen, VTT, Finland
Tel: +358 9 456 5623
E-mail: mika.pennanen@vtt.fi

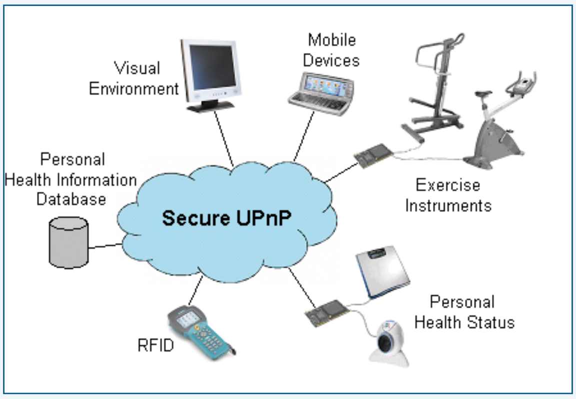

Figure 2: Networked Health Care system.

An Akogrimo Approach to Securing Virtual Organizations within Mobile GRID Computing Environments

by Thomas Kirkham, Giuseppe Cirillo, Julian Gallop, Damian Mac Randal, Brian Ritchie and Pierluigi Ritrovato

The Akogrimo Project (EU FP6-IST, 2004 to 2007) was introduced in a previous article, 'The Grid Goes Mobile' in ERCIM News (Issue 59). Akogrimo is aiming to radically advance the pervasiveness of Grid computing across Europe. To achieve this goal, and in addition to embracing layers and technologies which are intended to make up the so-called next-generation Grids (eg knowledge-related and semantics-driven Web services), Akogrimo aims to design and prototype a blueprint of a next-generation Grid that exploits and closely cooperates with evolving mobile Internet infrastructures based on IPv6. In this article, initial higher-level work on the security of virtual organizations and future plans for work in the security area are discussed.

Mobile Dynamic Virtual Organizations (MDVO)

Services shared by groups or individuals within distributed computing environments can be seen as virtual organisations (VOs). Within a traditional Grid these organizations can be seen as generally static and non-mobile. These traditional non-mobile VOs benefit from being able to register and link to services in permanent ways. Since service addresses and location details very rarely change, they can be made subject to common security measures associated with static networked topologies.

In Akogrimo, mobility is clearly a key concern. Since the VOs are both mobile and dynamic, several security issues are raised ranging from connection insecurity (wireless or otherwise) to Authentication, Authorization and Accounting (AAA) challenges. These security problems are present in particular during service discovery and re-discovery, since mobile services are prone to loss of connection, changes in bandwidth and so on.

Operative Virtual Organisations (OpVO)

Within Akogrimo this problem is approached by using a separate VO for the processing of services. This separate VO is referred to as an Operative Virtual Organization and is in existence for the lifetime of the particular workflow. In this model, user and service agents wrap a security layer around Grid users and services. The OpVO is linked to a Base VO via secure messaging and shared components, but it is essentially a temporary environment for the execution of services as opposed to the more permanent traditional view of the VO. A simple illustration of the basic security steps in the execution of a mobile Grid service can be seen in the figure.

The figure shows a service able to operate within VHE2, a Virtual Hosting Environment, wrapped by a service agent (SA), and able to be invoked by a workflow in the OpVO. Security is handled in the VHE where the service resides and also in the Base VO where the service is registered. The Base VO is central to the model and has the power to create and destroy an OpVO. This allows the existence of a hierarchy, the apex of which is a central point of security and policy enforcement for the workflow. Direct communication should not be possible between entities (eg users, services, resources etc) that belong to different administration domains, without going through the Base VO's security services. The execution of services within the OpVO and VHE reduces the workload on the Base VO, and it can be argued that this reduces the potential of a security breach occurring within the main Base VO, which could be running multiple OpVOs.

During workflow execution the Base VO in the model is largely used as a repository of VHE service details and security policy. The Workflow Manager links the VHEs to the Base VO for the purpose of discovery and authorization during the creation of an OpVO. This link calls services located in the VHEs by looking into the Base VO's service registry. As services are discovered and brought into the

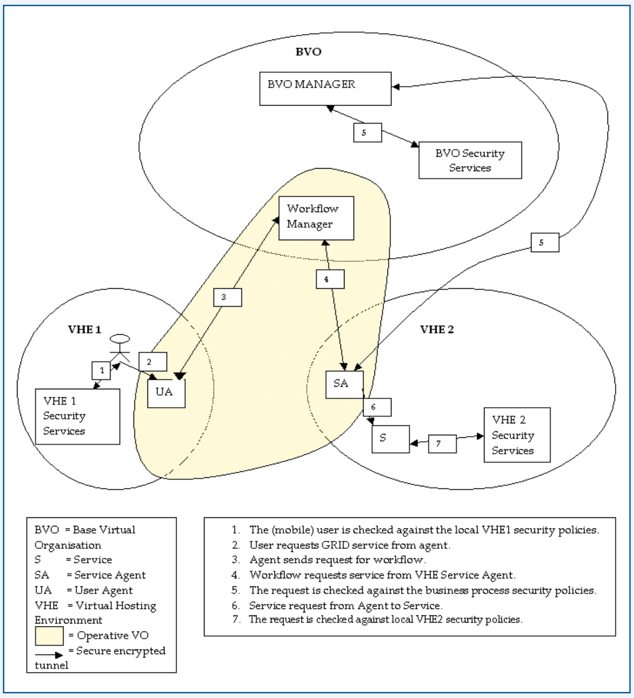

1. The (mobile) user is checked against the local VHE1 security policies.
2. User requests GRID service from agent.
3. Agent sends request for workflow.
4. Workflow requests service from VHE Service Agent.
5. The request is checked against the business process security policies.
6. Service request from Agent to Service.
7. The request is checked against local VHE2 security policies.

The Operative VO architecture and key security points.

OpVO they are authenticated and granted access via the secure exchange of tokens issued by the Base VO. If a service drops out or loses its connection after discovery, the workflow manager has the ability to send a request from the OpVO for a new service to be used; in this case the discovery and authentication process from the Base VO will be repeated. This re-authentication will help prevent security breaches such as ‘man in the middle’ attacks, which are popular on wireless connections. In addition, all traffic from the OpVO to the Base VO will be via an encrypted Web Service Secure Conversation link.

Conclusion and Future Work

In finding a solution to the problem of mobile Grid-based security, it can be argued that the use of Operative Virtual Organisations is the tip of the iceberg. The model is presented here in simplistic terms: plenty of scope exists for future investigation. The areas in which we are currently looking to improve security include greater monitoring, integration with security capabilities and constraints provided at the Network (mobile) layer (eg OASIS SAML standard for identity management and single sign-on), specifications for common security wrappers for integrating resources of administration domains having their own security policies and mechanisms, further details on the secure exchange of messaging using Web Service Secure Conversation and encryption, and the defining of policy in a more detailed fashion at all levels of the project.

Links:
CCLRC www.cclrc.ac.uk/
CRMPA: http://www.crmpa.it/
Akogrimo: http://www.mobileGrids.org/

Please contact:
Julian Gallop, CCLRC, UK
E-mail: j.r.gallop@rl.ac.uk

MOBIUS – Securing the Next Generation of Java-Based Global Computers

by Gilles Barthe

Global computers have the potential to realize the vision of ambient intelligence and to offer citizens a global and uniform access to services. Yet their success is conditioned by the development of appropriate security architectures that help establish trust and security in a global setting. Within the Global Computing 2 initiative, the FET Integrated Project ‘MOBIUS’ aims at providing the technology for securing next-generation global computers, building on the Proof Carrying Code paradigm.

Global computers aim at providing a global and uniform access to services through distributed computational infrastructures consisting of very large numbers of interacting devices. Prominent examples include the Internet, banking and telephone networks, digital video infrastructures, peer-to-peer and ad hoc networks.

While global computers may profoundly affect our quality of life, they will only become pervasive if novel security architectures are developed for bringing to users the level of reliability and security they expect for sensitive services. To realise this aim, the project will develop and combine type systems and program logics that can be used to ensure functionality and security policies for Java-enabled global computers.

Next Generation Global Computers

The next generation of global computers will emphasize the emergence of infrastructures with increasingly autonomous and heterogeneous devices:

- autonomy: devices will not be subjected to a global and uniform control, may belong to several global computers and may even move between different global computers
- heterogeneity: devices shall present significant differerences in their computational infrastructure (operating systems, communication protocols, libraries) and resources (memory, power autonomy, connectivity).

In order to accommodate these trends, the next generation of global computers will also require that devices are extensible with the computational infrastructure, platform or libraries needed to execute services as requested. This evolution will cause global computers to escape the scope of computational models which permeate mobile code, the Grid, or agents, and which impose a sharp separation between untrusted mobile applications, and the fixed and trusted computational infrastructure upon which they execute. At the same time, the evolution towards autonomy and extensibility will create new security threats that would not be found in current computational models, and thus any security architecture for global computing must comply with requirements that reach far beyond the limits of the current state –of –the art.

A Security Architecture based on Verifiable Evidence

The objective of the Mobius project is to develop a security architecture that meets the needs of global computers, by providing:

- innovative trust management, dispensing with centralized trust entities and allowing individual components to gain trust by providing verifiable evidence that they do not affect the security of the overall system
- static enforcement mechanisms, sufficiently flexible to cover the wide range of security concerns arising in global computing, and sufficiently resource-

aware and configurable to be applicable to the wide range of devices in global computers

- support for system component downloading, for compatibility with the view of a global computer as an evolving network of extensible devices.

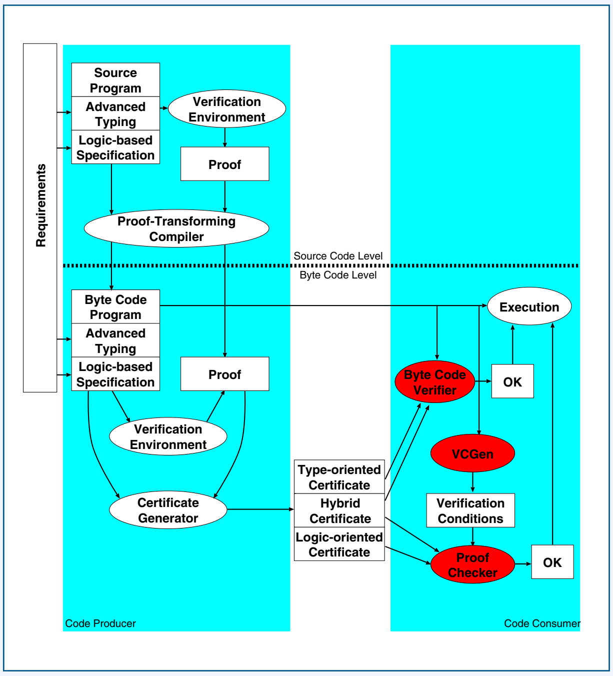

Modern verification environments based on program logics typically operate on source programs. The Mobius project proposes to combine these environments with type systems, which provide an automated means to enforce many basic policies, and use the resulting framework to cover a wide range of security policies for global computers. Evidence of programs adherence to their policy will be recorded by certificates: condensed, easily checkable formal proofs. In order to bring the benefits of source code verification to the code consumers, compilers will be enhanced to transform specifications and proofs for bytecode programs, yielding certificates that establish the correctness of bytecode programs and can be checked efficiently by code consumers.

The security architecture builds on ideas from Proof Carrying Code (PCC), and requires that mobile code is provided with a certificate, ie a condensed mathematical proof that the code is secure. In order to be applicable to global computers, the MOBIUS project will pioneer a PCC architecture that accommodates the distributed nature of global computing and allows enforcement of advanced policies, including both functional properties and advanced security properties such as non-interference or resource control. To ensure scalability, the MOBIUS project will also extend and combine two prime enabling technologies of PCC, ie type systems and program logics, and use the strength of these two techniques in hybrid certificates, to be verified through combined type checking and proof checking. Finally, the MOBIUS project will develop certificate translation as a means to bring to the code consumers the benefits of program verification, which is almost universally performed at source code level.

To maximize its impact, the MOBIUS project is focusing on Java-enabled global computers, and uses program logics that support the Java Modelling Language (JML). It will allow an implementation of the MOBIUS security architecture to be built on top of existing tools developed within the consortium. This will be evaluated on case studies from a range of application domains, covered by the consortium industrial partners, and by the End User Panel.

The project is coordinated by the French research institute INRIA, and is part of the pro-active initiative Global Computing 2, launched by the Future and Emerging Technologies unit of the IST programme.

Link:
Project web site: http://mobius.inria.fr

Please contact:
Gilles Barthe, Project coordinator:
Tel: +33 4 92 38 79 38
E-mail: Gilles.Barthe@sophia.inria.fr

Trust in Virtual Communities

by Marcin Czenko, Jeroen Doumen and Sandro Etalle

The objective of our research in the context of the BSIK Freeband project I-SHARE is to provide a sophisticated trust management framework for virtual communities.

Virtual Communities (VC) are a means of linking people with common interests, professional occupations, or resource-sharing habits. Several virtual communities already exist – Orkut, Kazaa, or Bittorrent to name the most prominent – and attract millions of Internet users. As a member of a virtual community one can access community resources that should be protected from outsiders. The challenges are to identify someone as a community member, since communities grow and shrink dynamically, and to establish trust between users on the basis of limited knowledge between parties. Existing virtual communities do little to protect their members from malicious internal and external users and as such are not acceptable for commercial use.

VC provide mechanisms that help to implement efficient, secure access to confidential data or protected resources. A member of a VC can access community resources that are not accessible to outsiders. For example, Apple could build a community of users who wish to download high-quality movies from iTunes' distributed secure servers. There should be no centralized access point, so that the system is more scalable and easier to maintain. Different community members may also have different privileges, including those of reselling music to other members of the iTunes community. The actual quality and the exact limitations on the amount of music that is accessible to community members depend on how much they wish to pay. In this scenario, the outsiders would be all the Internet users who do not pay for this access. As they are not members of the iTunes community, they cannot access any files there and cannot inject fake music into the network. Additionally, such controlled access to community resources helps not only to protect confidential information, but also to plan the use of bandwidth or load of service providers.

Trust management makes security transparent to the end users so that they no longer require substantial knowledge of the security mechanisms. Many existing approaches to trust management require a centralized architecture and thus do not fit well in the distributed nature of the Internet. Therefore, in our research, which started late 2004, we are combining distributed trust management and VCs into one solution that is suitable for both commercial and private use.

The first question we are trying to answer is how to formalize support for VCs in a trust management system. Secondly, because of the highly distributed nature of the Internet, we also need to investigate under which conditions the information necessary for secure user authentication is available. Finally, we aim to construct a suitable system design that provides security that is transparent to the end user.

Each community member must be able to prove membership of a community prior to being granted specific access permission. Such proof can be thought of as a distributed set of signed statements called credentials. Each credential consists of at least an issuer, a subject and a statement representing the actual meaning of the credential. The issuer and the subject of a credential both represent Internet users. A credential expresses the issuer's trust in the subject, giving the subject the permissions stated in the credential. The number of VCs one can join is unlimited; therefore one can be the subject of many credentials at the same time.

As a typical scenario, imagine that Alice would like to share the videos she took during her last holiday (see Figure 1). She does not want the whole world to see her footage, but rather wishes to restrict the potential audience to her friends. Alice can create a virtual community – 'friends' – which contains Alice's friends. To realise this, Alice issues credentials, one for each of her friends, saying that he or she is a member of the community 'Alice's friends'. In practice, Alice does not have to contact her friends personally, but the whole process can take place behind the scenes, using software agents acting on behalf of Alice and her friends. Alice's friends may not be even aware of the whole process and simply view Alice's film as if it was their own.

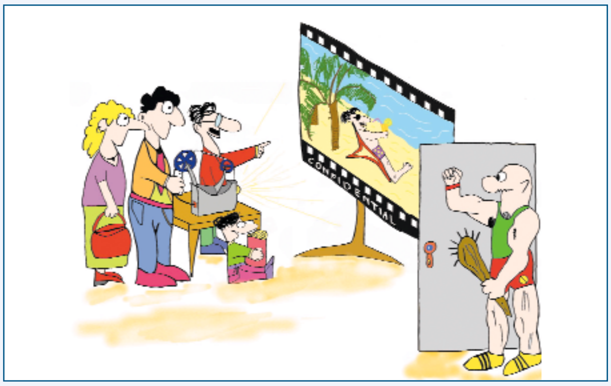

Virtual Communities allow for controlled access to shared resources.

When Bob wants to see Alice's films, his software agent contacts Alice's software agent to check whether the credentials provided by Bob's software agent are valid. Checking validity of credentials is not the only task a software agent must perform. If Alice's software agent detects that the credentials received from Bob's software agent are not sufficient to grant permission to see the pictures, Bob's software agent might be asked to provide additional credentials. As all credentials are stored in a distributed fashion, the job of Bob's agent is to track down the necessary credentials. The difficulty is in locating all the necessary credentials efficiently.

Imagine that some time later, Bob, encouraged by his good experience with Alice, also wants to share his videos. Bob then creates his own 'friends' community, and each member of the 'Bob's friends' community is allowed to see Bob's movies. Additionally, as Bob trusts Alice, he decides that all members of 'Alice's friends' community can also access his movies. In other words, Bob says that any friend of Alice is also a friend of his. We say that Bob delegates authority over his 'friends' community membership to Alice.

Even more sophisticated scenarios might arise. Therefore, our overall objective is to design and implement a sophisticated trust management framework that simplifies the protection of confidential resources in virtual communities. We use a so-called role-based trust management approach to model both trust and virtual communities. In our approach we use logic and logic programming to represent complex trust relationships that can appear in a real life experience. We also propose a new role-based trust management language with a formal declarative meaning based on the well-founded semantics for logic programs.

Our research is conducted as part of the I-SHARE project, which is part of the Freeband consortium funded by the Dutch national BSIK program. I-Share involves five project partners: Delft University of Technology (TU-Delft), Eindhoven University of Technology (TU/e), University of Twente (UT), Vrije Universiteit Amsterdam (VU) and Philips Research, Eindhoven.

Links:
Freeband Communication web site:
http://www.freeband.nl/index.cfm?language=en

I-SHARE Project web site:
http://www.freeband.nl/project.cfm?id=520

DIES group home page:
http://dies.cs.utwente.nl

Please contact:

Marcin Czenko,
University of Twente, The Netherlands
Tel: +31 53 489 3709
E-mail: Marcin.Czenko@utwente.nl

Jeroen Doumen,
University of Twente, The Netherlands
Tel: +31 53 489 2801
E-mail: Jeroen.Doumen@utwente.nl

Secure and Trusted Virtual Organization Management

by Adomas Svirskas, Alvaro Arenas, Michael Wilson
and Brian Matthews

TrustCoM is a European Integrated Project that aims to develop a framework for trust, security and contract management in dynamic Virtual Organizations.

The TrustCoM framework will enable the secure enactment of collaborative business processes in self-managed and dynamic value-chains of businesses and governments. One of the central aspects of TrustCoM research is VO (virtual organization) management, which aims to support dynamic virtual communities throughout their entire life cycle. Although proprietary implementations of VO management tools exist, secure tools based on interoperating open standards are not yet available. The open standards on which to build them are just being released as reliable implementations. The TrustCoM project provides an answer to these problems with its VO Management framework, based on the open Service-Oriented Architecture and open Web Services standards.

Current Grid-based VO management supports only the VO memberships function – listing VO members who are entitled to use VO resources. It does not support the management of the risks associated with VO membership through:

- the identification of potential VO partners through reputation management
- the roles defined in business process models that VO partners perform to limit resource access and reputation transfer
- the contractual or SLA obligations on the VO for security and privacy
- the enforcement of policies derived from contracts for quality and timeliness of business process enactment.

The VO management subsystem should provide the services necessary for maintaining the VO structures, monitoring members' performance, enforcing VO policies, assigning members to play certain roles and perform tasks and so forth. In essence, the VO management process involves ensuring that the members of a VO play by commonly agreed-upon rules and that members' behaviour is observable, thus allowing these rules to be enforced.

The TrustCoM VO management component will not only provide a membership function but will also provide life-cycle and context management functions. These will provide a generic VO management layer that records membership as well as addressing the management of the risks of VO membership.

TrustCoM is following the life-cycle model developed in the VO roadmap project (Camarinha-Matos and Afsarnabesh, 2003), including phases such as identification, formation, operation/evolution and dissolution. The identification phase deals with setting up the VO; this includes selection of potential business partners by using search engines or looking up registries. VO formation involves partnership formation, including the VO configuration by a VO manager (who distributes information such as policies, Service Level Agreements (SLAs) etc), and the binding of the selected candidate partners into the actual VO. After the formation phase, the VO can be considered to be ready to enter the operation phase where the identified and properly configured VO members perform according to their role. Membership and structure of VOs may evolve over time in response to changes of objectives or to adapt to new opportunities in the business environment. Finally, the dissolution phase is initiated when the objectives of the VO have been fulfilled.

Making the creation, operation and dissolution processes rapidly responsive requires both the appropriate legal mechanisms and dynamic management of the VO. There appears to be an obvious match between the business-driven desire to create and manage dynamic VOs, and the technological solution available in composable Web services. Although proprietary IT implementations of composable Web service tools exist, secure tools for VO management (VOM) based on interoperating open standards are not yet available, let alone those that also address legal issues. Consequently, IT based dynamic VOM is confined to closed communities that have adopted a single proprietary solution. Although proprietary implementations of VO management tools can operate either within single organizations or across cooperative organizations, they also pose substantial adoption costs, which in turn are only justified by long-term relationships within the closed community using the proprietary approach.

In addition, the TrustCoM VO management solution will use, where appropriate, declarative specifications of VO management processes. This approach will allow the publishing of and access to commonly understood and approved protocols of VO management and VO members' interactions. Having these interactions explicitly choreographed will reduce the complexity of end-point software components of VO members' software, thus increasing the robustness, sustainability and efficiency of VO management procedures. TrustCoM is using WS-CDL specification (a candidate W3C standard) to specify choreographies and emerging open-source tools, which allow modelling and validation of the protocols as well as end-point generation for WSDL-based and BPEL-enabled services.

Tools for managing the identified risks are being developed in the TrustCoM framework for the trust, business process management, contracts and policy components respectively. Putting them together to manage the risks of joining and operating within a VO is the TrustCoM's innovation for VO management. The TrustCoM VO management component represents a generic VO application layer that ties together these TrustCoM components, and it is upon these components that individual VOs can build and incorporate specific applications (eg aerospace modelling tools, remote learning tools). By defining the standard VO management protocols based on open standards, TrustCoM is contributing to the establishment of VO management patterns that may be useful beyound the scope of the project itself.

Link:
http://www.eu-trustcom.com/

Please contact:
Michael Wilson, CCLRC,
Tel: +44 1235 446619
E-mail: m.d.wilson@rl.ac.uk

Trusted Network on Wheels

by Matthias Gerlach

Vehicular ad hoc networks may become the largest ad hoc network ever deployed, and of fundamental importance for the safety and comfort of their users – the drivers and passengers. Clearly, deploying a revolutionary technology on such a large scale presents major challenges in the secure design of the system, its protocols and applications.

Within Germany's nationally funded Network on Wheels (NoW) project, which started in June 2004, there is a dedicated group of experts from various companies and research institutes looking at potential attacks on such networks and devising methods and mechanisms to protect them. This is the first occasion that security considerations have been applied from the start of network development. This article provides some insight into the current and future work on security in the NoW project, and demonstrates the seriousness with which security and privacy are regarded.

Vehicular Ad Hoc Networks

of vehicular ad hoc networks, devising both core technologies and possible applications for these networks. It is expected that in the near future, the current specification of WAVE (Wireless Access for Vehicular Environments), a 802.11 based technology, followed by specifications from the Car2Car Communication Consortium (C2CCC) will put into place basic technologies for applications such as advanced active safety, faster and better driver information systems and exciting applications for entertainment and information based on vehicle to vehicle (v2v) and vehicle to infrastructure (v2i) communications.

The challenges in this new type of network are the short contact times between different mobile nodes, a far from ideal radio environment that includes large vehicles which obstruct the radio path, and the sheer size of the network. In addition, management of the network can no longer rely on ubiquitous infrastructure access, as ad hoc networks allow for infrastructure-independent operation. These restrictions of the network make the design of security mechanisms very challenging, and the fact that the system is under constant development does not make the process any easier.

Security Requirements

Users, that is, the customers, will only buy and use a system they trust. From the users' point of view, the system can be trusted if it provides the orthogonal requirements of availability, privacy and correctness.

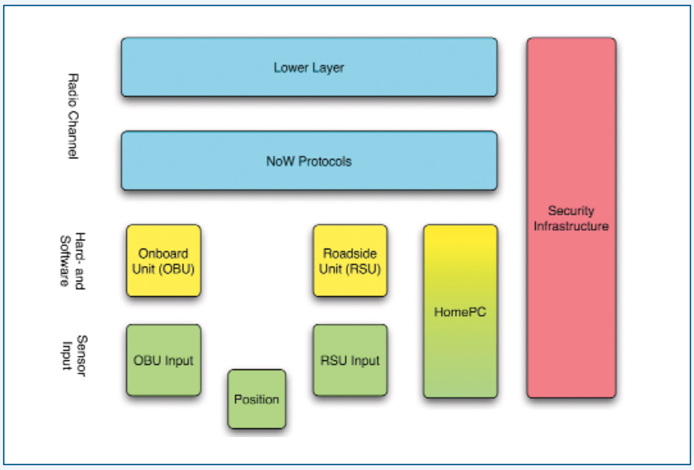

Figure 1: Generic system model in support for a heuristic security analysis.

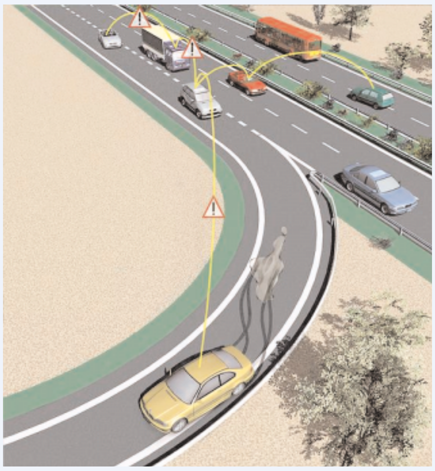

copyright BMW group

Figure 2: As soon as two or more vehicles are in radio communication range, they connect automatically and establish an ad hoc network.

Availability of the system implies that the system is robust even in the presence of malicious or faulty nodes, which due to the network size might be the general case. Note that this is not solely a security requirement but a common system requirement. Hence, security mechanisms can only provide a basis to enforce availability.

The privacy of users is an important asset in public networks. Basically, privacy requires untraceability of actions to a user and unlinkability of the actions of a node. These must be provided by the applications as far as possible, and be inherent to the internal functionalities of the communication system. The communication system should give away as little information as possible that could be used to violate the privacy of the users.

Finally, correctness in the security domain relates to secure communication. This boils down to the well-known security objectives of authenticity, freshness (which is, in fact authenticity in the time domain), integrity and non-repudiation. Another important aspect within such a system will be authorization in the different levels of the system, starting from the authorization to send messages on the radio interface to the use of certain application layer services provided by the network.

Approach in NoW

Specifying the security requirements is not enough, however. As the envisioned system is a complex one, we would like to identify potential attacks on the system in advance, in particular as we already know some of the base technologies. The generic system model depicted in the figure helps to specify the different subsystems' vulnerabilities and security requirements. Current work consists of detecting attacks on the different parts of the system and estimating both their impact and probability.

Starting from general attacks such as the insertion of false messages, system denial of service and privacy violations, attacks can be refined using attack trees. Attack trees represent a hierarchical organization of an attacker's goals in AND and OR conjunctions, which become more detailed the deeper down the tree you go. Attack trees can also be used to assess the impact of a system's vulnerabilities, so as to decide where the priority of the work must be. Within the Security Working Group of NoW, we are currently constructing attack trees for a variety of applications using different services of the communication subsystem. Both the attack trees and mechanisms to secure the network will be published and discussed openly as they are finalized. In the Security Working Group of NoW, we believe that security mechanisms should be publicly scrutinized before they become part of the actual system.

Links:
NoW project website:
http://www.network-on-wheels.de

C2CCC web site: http://www.car-to-car.org

Please contact:
Matthias Gerlach, Fraunhofer FOKUS
Tel: +49 30 3463 7132
E-mail: Matthias.gerlach@fokus.fraunhofer.de

A Graphical Delegation Solution for X.509 Attribute Certificates

by Isaac Agudo, Javier Lopez, Jose A. Montenegro

Delegation is becoming a major topic for distributed authorization. Several approaches have been proposed in order to provide delegation in distributed environments. However, none of these matches all the requirements of a flexible delegation method. Based on work from the project PRIVILEGE, a flexible solution based on the use of graph theory is proposed.

Delegation is a major goal when a real scalable distributed authorization system is needed. However, the uncontrolled use of delegation statements can become an important security threat; for instance, any user could improperly obtain over a resource the same privileges as the owner of that resource. Therefore, delegation solutions should include a mechanism to control the delegation of privileges as well as making use of suitable authorization statements.

One of the first tasks in the project PRIVILEGE, developed in the Computer Science Department at the University of Malaga, has been to study and put into perspective the delegation implications of standard schemes that have been proposed in the literature as solutions for distributed authorization problems. As such, we have realized that in PolicyMaker and Keynote schemes, the delegation statement does not exist; that is, any authorization statement can be delegated once and then again without any control. On the other hand, SDSI considers three different possibilities for controlling delegation, although SPKI reduced it to a Boolean condition. Such a Boolean parameter is only a modest mechanism to control the depth of delegation.

There are more theoretical solutions that use logic programming for the representation of authorization and delegation statements. In logic programming, those statements are represented as predicates, and decisions are based on formulae verification. Although logic programming offers a powerful mechanism to represent authorization and access control decisions, it has important drawbacks: it is difficult to understand and has obscure transcription. Moreover, no logic solu-

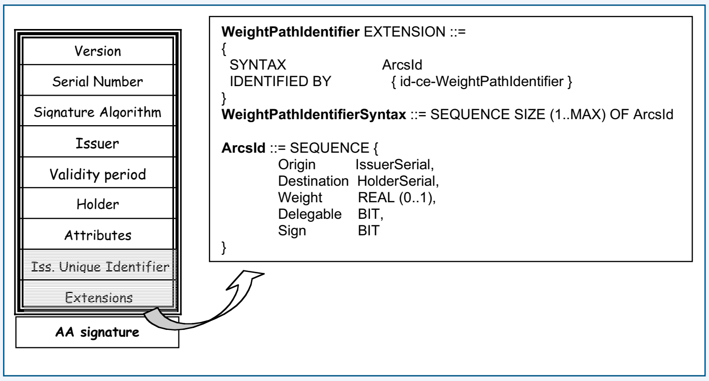

```
WeightPathIdentifier EXTENSION ::=
{
 SYNTAX                    ArcsId
 IDENTIFIED BY             { id-ce-WeightPathIdentifier }
}
WeightPathIdentifierSyntax ::= SEQUENCE SIZE (1..MAX) OF ArcsId

ArcsId ::= SEQUENCE {
        Origin        IssuerSerial,
        Destination   HolderSerial,
        Weight        REAL (0..1),
        Delegable     BIT,
        Sign          BIT
}
```

Figure 1: Attribute certificate and weight path identifier extension.

tion has been integrated in any standard authorization framework.

Additionally, there are graphical solutions that are thought to be less powerful but more expressive and easily understood. A graphical solution may be based on the use of directed graphs to model authorization and delegation processes. Basically, this maps each credential in the system to a directed edge in a graph. Edges go from the issuer of the authorization or delegation statement to the subject who is authorized or granted privileges. Usually, the root of the tree is the owner of the resource under consideration. It is therefore possible to study the relations between entities in the system in a graphical way.

The project PRIVILEGE focuses on the use of X.509 Attribute Certificates. Therefore, it includes a practical implementation of a Privilege Management Infrastructure (PMI). As part of our work, we have developed a mechanism to perform a controlled delegation that uses the extension fields of the attribute certificates. Our proposal is based on graphical solutions, attaching extra information to every edge in the graph. In particular, we include an index, a real number in the interval [0,1], that measures the level of confidence of the issuer on the issued certificate. Moreover, we distinguish between positive and negative statements. Positive ones grant the right encoded in the certificate and negative ones deny it. In order to encode it, we use the variable sign (see Figure 1). We also add another Boolean variable, delegation, to define whether the certificate can be chained, ie delegated.

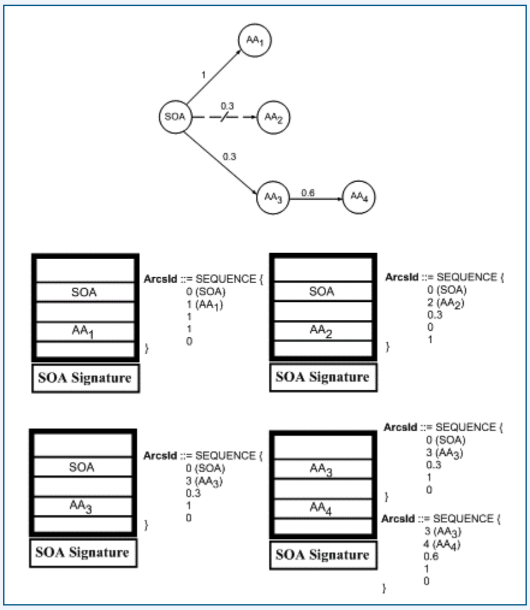

```
ArcsId ::= SEQUENCE {
        0 (SOA)
        1 (AA1)
        1
        1
        0
}
```

```
ArcsId ::= SEQUENCE {
        0 (SOA)
        2 (AA2)
        0.3
        0
        1
}
```

```
ArcsId ::= SEQUENCE {
        0 (SOA)
        3 (AA3)
        0.3
        1
        0
}
```

```
ArcsId ::= SEQUENCE {
        0 (SOA)
        3 (AA3)
        0.3
        1
        0
}
ArcsId ::= SEQUENCE {
        3 (AA3)
        4 (AA4)
        0.6
        1
        0
}
```

Figure 2: Design of statements and corresponding certificate chains.

We add this information directly in the certificate, by using the extensions field. This field allows us to include additional information into the attribute certificate. Although the X.509 standard provides several predefined extension categories, we focus on the delegation extension category, which defines different extension fields. Among them, the ITU-T recommendation includes:

Authority attribute identifier: In privilege delegation, an Attribute Authority (AA) that delegates privileges shall itself have at least the same privilege and the authority to delegate that privilege. An AA that is delegating privilege to another AA or to an end-entity may place this extension in the AA or end-entity certificate that it issues. The extension is a back pointer to the certificate in which the issuer of the certificate containing the extension was assigned its corresponding privilege. The extension can be used by a privilege verifier to ensure that the issuing AA had sufficient privilege to be able to delegate to the holder of the certificate containing this extension.

That extension is suitable for our purposes. However, it does not define the weight associated to the edge between the issuer and the holder of the certificate. Therefore, we define our own extension, in ASN.1 (see Figure 1), based on the authority attribute identifier.

This new extension determines a sequence between the source of authority (SOA) and the holder. Each sequence includes another sequence, ArcsId, in which is included the information of the edges in the graph, weight of the edge, origin node, and Boolean information about statements, delegation and sign. The destination node must match the serial number of the attribute certificate.

As the reader can infer, the design of authorization and delegation statements in a graphical mode can be converted automatically into X.509 attribute certificate chains. The example included in Figure 2 shows the graphical design of delegation statements (normal line) and authorization statements (dotted line) and its equivalent representation using attribute certificates. Every attribute certificate stores, in the extensions field, the graphical information.

PRIVILEGE will finish at the end of 2006, and is scheduled to have a complete practical implementation of the delegation solution, further elaborating on the complexity of management of delegation chains in fully distributed authorization systems.

The Research Project PRIVILEGE (TIC2003-8184-C02) has been funded by the Spanish Ministry of Science and Technology.

Please contact:
Javier Lopez, University of Malaga, Spain
Tel: +34-952-131327
E-mail:jlm@lcc.uma.es

Security and Trust Management Extensions to the PERMIS X.509 Privilege Management Infrastructure

by David Chadwick

Authorization in vitual organizations (VOs) and multi-organization federations is difficult to set up and manage. Having a pan-VO role- or attribute-based access control infrastructure can ease the burden, providing that trust relationships between the various entities can be safely managed. This is the problem domain that the PERMIS authorization system (www.openpermis.org) has been addressing for several years. Currently, three research projects are adding significant new capabilities to it, in the form of dynamic delegation of authority, separation of duties, and reputation management of the participants.

The PERMIS Trust Model

The resource (target) owner is the authorization root of trust for all resources under his/her control. This is termed the Source of Authority (SOA) in an X.509 Privilege Management Infrastructure (PMI). The SOA creates his/her policy and stores it in a digitally signed policy attribute certificate (AC). When the PERMIS authorization engine is initialized, it receives the distinguished name of the SOA and the location of the LDAP directory where it will find the SOA's policy. PERMIS reads the policy from the SOA's LDAP entry and checks its signature. PERMIS can now be assured that it has the correct policy to trust, and that it has not been tampered with.

The SOA's policy, which follows the classical Role Based Access Control (RBAC) model, comprises two parts:

- a Role Assignment Policy (RAP) which specifies who is trusted to assign which roles (in the form of X.509 role ACs) to which groups of user
- a Target Access Policy (TAP) specifying which roles are needed to access which target resources under which conditions.

The RAP enables static delegation of authority because the SOA names one or more (possibly remotely located) managers who are trusted to assign roles. Since the names of these remote managers are included inside the digitally signed policy AC, they cannot be unknowingly tampered with; therefore PERMIS is able to trust these remote managers to assign X.509 role ACs to groups of users. Any role ACs that PERMIS is passed or retrieves in getcreds that do not conform to the RAP are simply discarded. In this way the SOA can be assured that his/her delegation policy is being rigorously enforced by PERMIS.

Current Research Projects

DyVOSE, run jointly with the e-Science centre at the University of Glasgow, is adding dynamic delegation of authority to PERMIS. This feature will allow the SOA to indicate whether he/she trusts remote managers to further dynamically delegate their roles to other users in the same domain as themselves. Dynamic delegation of authority is supported in the X.509 PMI model through an appropriate certificate extension, and the SOA can set an integer in his/her PERMIS policy to indicate the length of the delegation chain that can be trusted. Once this is fully implemented, the SOA will not need to update his/her RAP policy with the names of additional remote managers who can be trusted, as is currently the case. Instead, as long as these additional managers have a valid delegation path to a remote manager in the PERMIS policy, then any ACs issued by them will be trusted.

DyCom is combing PERMIS with GRASP to create a fine-grained access control infrastructure for Grids, and is also adding separation of duties to the PERMIS trust model. Separation of duties will ensure that a user with mutually exclusive roles is not allowed to perform conflicting tasks. This requires PERMIS to keep a record of past and present authorized actions, so that future conflicting ones can be denied. In an offshoot of this project, we have developed a secure

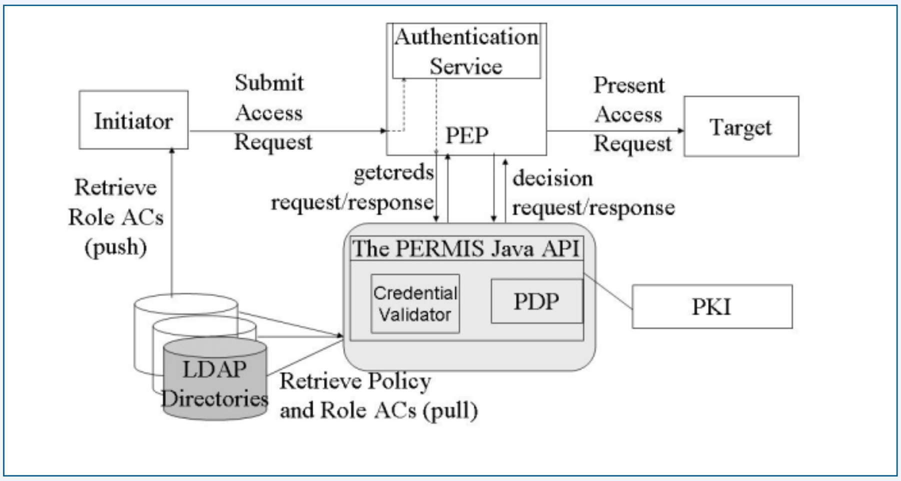

PERMIS authorization system.

audit Web service (SAWS) as a general purpose audit tool.

As part of the EC TrustCoM integrated project, we have built a reputation management system capable of recording the reputations of users (for example, as performed by eBay). The next step is to link this to the PERMIS decision engine so that access control decisions can be based on the current reputation of a user (which is related to their trustworthiness). Currently users are either trusted or not to access a target resource, based on their X.509 ACs. Once reputations are included in the decision-making however, users' permissions may be removed if their reputation drops below a certain value. In addition, the TrustCoM project is defining standard protocols for credential validation (ie calls to getcreds) and the making of policy decisions (ie calls to decision). It is likely that WS-TRUST and XACML respectively will be used for these. PERMIS will be enhanced to support these protocols once they have been finalized by the consortium.

Link:
http://www.eu-trustcom.com

Please contact:
David W. Chadwick, University of Kent, UK
Tel: +44 1227 82 3221
E-mail: D.W.Chadwick@kent.ac.uk

Richard Sinnott, (for DyVOSE)
University of Kent, UK
E-mail: ros@dcs.gla.ac.uk,

Damian Mac Randal (for DyCom)
CCLRC, UK
E-mail: D.F.Mac.Randal@rl.ac.uk

Theo Dimitrakos (for TrustCoM), BT, UK
E-mail: theo.dimitrakos@bt.com

Interactive Access Control and Trust Negotiation for Autonomic Communication

by Hristo Koshutanski and Fabio Massacci

Recent advances in Internet technology and globalization of peer-to-peer communications offer organizations and individuals an open environment for rapid and dynamic resource integration. In such an environment, federations of heterogeneous systems are formed with no central authority and no unified security infrastructure. Considering this level of openness each server is responsible for the management and enforcement of its own security policies with a high degree of autonomy.

Dynamic coalitions and autonomic communication add new challenges: a truly autonomic network is born when nodes are no longer within the boundary of a single enterprise, which could deploy its policies on each and every node and guarantee interoperability. An autonomic network is characterized by the self-management and self-configuration of its constituent nodes. In an autonomic network, nodes are partners that offer services and lightly integrate their efforts into one (hopefully coherent) network.

Policy-based network access and management already requires a paradigm shift in the access control mechanism: from identity-based access control to trust management and negotiation, but even this is not enough for cross-organizational autonomic communication.

In an autonomic communication scenario, a client may have all the necessary credentials to access a service but may be unaware of this. Equally, it is unrealistic to assume that servers will publish their security policies on the Web so that clients can perform policy combinations and evaluations themselves. Rather, it should be possible for a server to ask a client, on the fly, for additional credentials: the client may then choose whether or not to disclose them. The server then re-evaluates the client's request taking the newly submitted credentials into consideration, and iterates the process until a final decision (of 'grant' or 'deny') is reached. We call this modality interactive access control.

While some of these challenges can be solved by using policy-based self-management of networks, this is not universally the case. Indeed, if we abstract away the details of the policy implementation, we can observe that the only reasoning service that is actually used by policy-based self-management approaches is deduction: given a policy and a set of additional facts, find out all consequences (actions or obligations) of the policy and the facts. We simply look at whether granting the request can be deduced from the policy and the current facts.

Autonomic communication needs another reasoning service: abduction. Loosely speaking, we could say that abduction is deduction in reverse: given a policy and a request to access a service, we want to know what credentials/events are missing that would grant access.

The contribution of the framework is in the way we bootstrap from the basic reasoning services of deduction, abduction and consistency checking, a comprehensive interactive access-control algorithm that computes on the fly the missing credentials needed for a client to get access. We extended the algorithm to cope with arbitrary access policies so that in cases of inconsistency it performs a recovery step and finds a set of excessing credentials banning the client to get a solution for the desired resource. Following this, a strengthened version of the algorithm is devised that is resistant to DoS attacks. We have modelled a fully fledged ac-

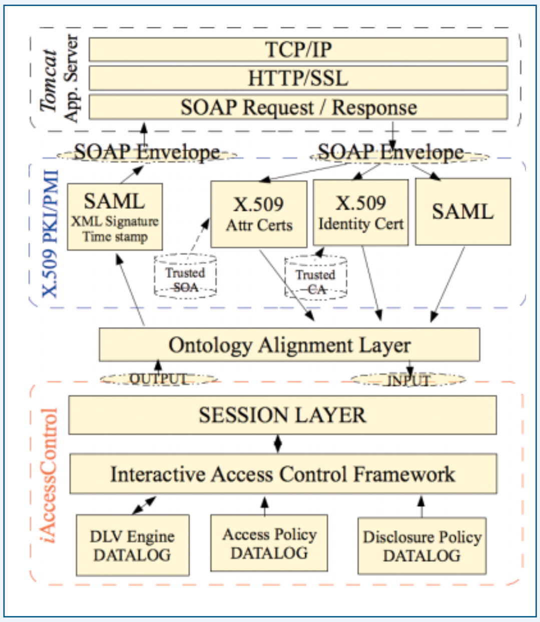

Figure 1: Interactive Access-Control Engine Architecture.

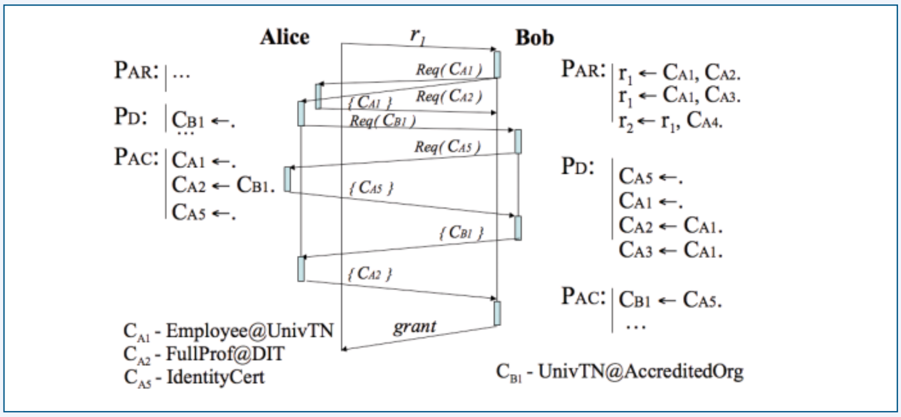

Figure 2: Example of Interoperability of the Negotiation Protocol.

P_{AR} denotes the policy for access to resources, P_{AC} denotes the policy for access to credentials and PD the policy for disclosure of (foreign) credentials. Credentials used in the policies are in the following notations: Alice's local credentials are marked with subscript 'A' and Bob's with 'B', respectively. Bob's access policy P_{AR} says that access to resource r_1 is granted if $\{C_{A1},C_{A2}\}$ or, alternatively, $\{C_{A1},C_{A3}\}$ are presented by Alice. Access to r_2 is granted if Alice satisfies the requirements for access to r_1 and presents C_{A4}.

Analogously, we read Bob's disclosure policy P_D as meaning that to disclose the need for credential C_{A2} there should exist already-disclosed credential C_{A1}, which by default is always disclosable. In contrast, the need for credential C_{A4} is never disclosed but is expected by Bob's access policy PAR when r_2 is requested.

The real interactions start when Alice requests r_1 from Bob. Then, suppose that the set $\{C_{A1},C_{A2}\}$ is minimal with respect to the other alternative $\{C_{A1},C_{A3}\}$, and say that C_{A2} contains a role lower in the hierarchy than the role in C_{A3}. Then, Bob replies with two counter requests to Alice. Alice, in her turn, runs the two requests in new threads and replies to Bob, according to her policy for access to sensitive credentials P_{AC}, with a counter-request for C_{B1} and the disclosure of C_{A2}.

The negotiation process continues, as shown in the figure, until Bob discloses all credentials requested by Alice and Alice, in her turn, discloses all credentials requested by Bob so that at the end the desired resource is granted.

cess-control framework and shown its correctness and completeness.

Based on the interactive access-control algorithm, we introduce a trust negotiation protocol that runs on both client- and server-sides. It automatically interoperates and negotiates missing credentials until either a final decision of 'grant' is taken and the negotiation is successfully completed, or one of the parties fails to negotiate the requirements and the service request is denied. Figure 2 shows a message-flow example of the negotiation protocol.

We have implemented the interactive access-control modality as a Web Service. For this purpose we used X.509 PKI/PMI and OASIS SAML standards as a unified way of conveying credentials and defining authorization statements respectively. We integrated the two standards with the W3C SOAP protocol so that our access control engine can be used and invoked in a platform-independent manner. We have done a Java wrapper for the DLV system that implements the interactive access-control algorithm. The DLV system is used as a core engine of the basic functionalities of deduction, abduction and consistency checking. The architecture of the access control engine is shown in Figure 1.

Future work will look at characterizing the complexity of the framework and extending it to cope with stateful systems and especially with the open issues of revocation of credentials.

Links:

X.509 PKI/PMI:
http://www.ietf.org/html.charters/pkix-charter.html

OASIS SAML:
http://www.oasis-open.org/committees/security

W3C SOAP: http://www.w3.org/TR/soap

DLV: http://www.dlvsystem.com

Please contact:
Hristo Koshutanski and Fabio Massacci,
University of Trento, Italy
E-mail: hristo@dit.unitn.it, massacci@dit.unitn.it

Access-Control Policy Administration in XACML

by Erik Rissanen and Babak Sadighi Firozabadi

In recent years, researchers at SICS have been looking at managing large numbers of access permissions in a dynamic and decentralized network. The main results of our work are a framework and a calculus, called privilege calculus, for access permissions and their administration.

The eXtensible Access Control Markup Language, XACML, is a very effective and now widely adopted standard language for expressing access control policies. The specification of XACML includes the language, its semantics and a framework for making access control decisions based on XACML policies. However, XACML is currently lacking an access control model for the policy itself.

The current XACML model of policy administration puts the access control of policy administration outside the policy model. To control who may edit the policy, mechanisms such as access control at the operating system level must be used. In large distributed systems, such mechanisms may prove difficult to manage. There may be a need to manage the policies in parts of the system not under the control and within the trust of a specific Policy Decision Point, for instance from a mobile device. The rights to change the policy may themselves be highly dynamic. Consequently, there is a need for the policy itself to have an access-control policy model. Our research has been focused on these issues.

The Policy-Based Reasoning group at SICS has for several years been performing research on how best to manage large numbers of access permissions in a dynamic and decentralized network. The main results of our research are a framework and a calculus, called privilege calculus. In the framework, we distinguish between access permissions and administrative permissions, both referred to as privileges. Privilege calculus allows us to reason about privileges and their administration. The core mechanism of privilege calculus is constrained delegation, which allows constraints to be put on the creation of privileges, access permissions or administrative permissions.

Recently, a number of XACML Technical Committee (TC) members have discussed the need for adding administrative support to XACML. The discussed ideas are very similar to the delegation mechanism of privilege calculus. We are now looking into the possibility of extending the current XACML specification and implementing our delegation model in SUN's open-source XACML implementation. Our work will be part of two projects – the TrustCom EU FP6 project and Decentralized Authorization Management in Network-Based Defence – in which we investigate the use of XACML as a policy language for distributed services in highly dynamic and decentralized networks.

Adding delegation to XACML involves defining new forms of policy that can express administrative rights, and a new processing model that can verify that delegations have been performed in an authorized manner. The new features of XACML help users to implement flexible decentralized access control management, for instance in the setting up of a large organization or joint business venture. This will reduce the administration costs of the organizations and make them more flexible. Having these features available in a standard access control language will make their use simpler and more widely adopted.

Please contact:
Babak Sadighi Firozabadi or Erik Rissanen,
SICS, Sweden
Tel: +46 8 633 1500
E-mail: babak@sics.se, mirty@sics.se

Contributions of Team Automata in Security

by Maurice ter Beek, Gabriele Lenzini, and Marinella Petrocchi

Researchers from two CNR Institutes in Pisa are studying ways in which a formal model of team automata can be exploited to specify and analyse security-related issues.

The Information Security group at the Institute for Informatics and Telematics (IIT-CNR) has both practical and theoretical experience with many aspects of security. The Formal Methods and Tools (FM&&T) group at the Institute of Information Science and Technologies (ISTI-CNR) has experience in research on formal methods for the specification, design and verification of computer systems. Recently, researchers from these two groups teamed up to investigate how a formal model of team automata can contribute to the specification and analysis of security issues. This cooperation will be continued in the context of an EU-funded project on Software Engineering for Service-Oriented Overlay Computers (SENSORIA).

Team automata form a mathematical framework introduced in 1997 by C.A. Ellis to model components of groupware systems and their interconnections. Their usefulness however extends to modelling collaborations between system components in general (for an overview, see [TA]).

A team automaton is composed of component automata that distinguish input, output and internal actions. Input actions are not under the automaton's control, but are triggered by the environment, which can include other automata. Output and internal actions are under its control, but only the output actions are observable by other automata. Input and output actions together constitute the external actions and form the interaction interface between the automaton and its environment; internal actions do not participate in any interactions.

In composing a team automaton, the crux is to define the way in which those originally independent components interact. Their interactions are formulated in terms of synchronizations of shared actions, a method for modelling collaboration among system components that is well known in the literature. A component automaton does not necessarily participate in every synchronization of an action it shares. Hence there is no such thing as the unique team automaton over a set of component automata. Rather, a whole range of team automata, distinguishable only by their transition relation, can be constructed from a given set of components. It is this freedom to choose a transition relation that sets the team automata framework apart from most other automata-based models, most of which use a single and very strict method for choosing the transition relation of an automaton composed over a set of automata - in effect resulting in composite automata that are uniquely defined by their constituents.

In a series of papers (see [TA]) we have shown how team automata can adequately be used to model (and sometimes verify) various access control policies, multicast/broadcast communication protocols and general (cryptographic) communication protocols.

To begin with, we have demonstrated the model usage and utility for capturing information security and protection structures, as well as critical coordinations between these structures. On the basis of a spatial access metaphor, various known access-control strategies have been given a rigorous formal description in terms of synchronizations in team automata. Moreover, we have initiated to validate some of the resulting specifications with the model checker Spin.

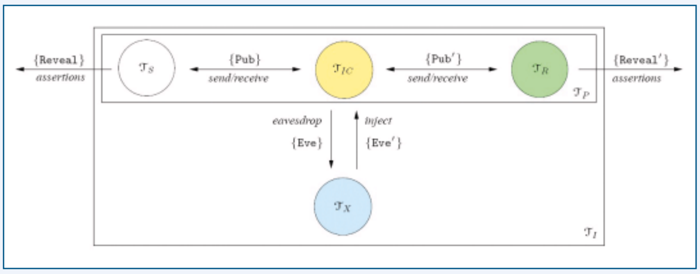

An insecure communication scenario for team automata.

Later we have initiated the use of team automata for the security analysis of multicast and broadcast communication. For this purpose, we have performed a case study in which team automata were used to model an instance of a particular stream signature protocol. The one-to-many and one-to-all communications, which are so typical of multicast and broadcast communications, were captured by team automata in a native way as synchronizations between the set of component automata constituting a team automaton. We have also developed a framework for security analysis with team automata, which has required three basic formal steps.

First, we defined an insecure communication scenario based on the addition of a so-called 'most general intruder' to a team automaton model of a secure communication protocol. The intruder was modelled as an active agent able to influence communication among honest agents. This insecure scenario can be used to analyse some security properties of cryptographic communication protocols involving two roles – an initiator and a responder. Rather than occurring directly, all communication is assumed to flow through an insecure channel. This insecure channel may release some messages to an intruder, which in its turn can either listen to or modify (fake) the messages passing through this channel. When verifying security properties for cryptographic communication protocols, it is indeed quite common to include an additional Dolev-Yao-style intruder that is supposed to be malicious and whose aim is to subvert the protocol's correct behaviour. A protocol specification is consequently considered secure with reference to a security property if it satisfies this property despite the presence of the intruder. Abstracting from the cryptographic details concerning the operations according to which messages can be encrypted, decrypted, etc, the insecure scenario is informally described by the team automata interactions sketched in the figure.

Second, a well-established theory for defining and verifying a variety of security properties was reformulated in terms of team automata and subsequently, a compositional analysis strategy was described for it. Under appropriate assumptions, this can be used to verify some security properties in the communication protocol modelled by the scenario.

Third, this framework was applied to show that integrity is guaranteed for the particular setting of the case study. This shows the effectiveness of our approach for a realistic stream signature protocol, thus facilitating an easy comparison for those familiar with other approaches. In fact, an approach that uses an automata-based formalism for the specification and verification of properties in the field of security is not unique, but has become very popular in recent years.

Finally, very recently, team automata have been used to model and verify a protocol aiming at privacy in communication among mobile agents. This was the first attempt to use team automata for the analysis of privacy properties.

Links:
[TA]: http://fmt.isti.cnr.it/~mtbeek/TA.html
FM&&T: http://fmt.isti.cnr.it/
IIT: http://www.iit.cnr.it/

Please contact:
Maurice H. ter Beek, ISTI-CNR, Italy
Tel: +39 050 315 3471
E-mail: maurice.terbeek@isti.cnr.it

Using Probabilistic I/O Automata to Improve the Analysis of Cryptographic Protocols

by Ran Canetti, Ling Cheung, Dilsun Kaynar, Moses Liskov, Nancy Lynch, Olivier Pereira and Roberto Segala

Modelling cryptographic protocols and analysing their security is a tricky business. On the one hand, valid modelling and analysis must address the concurrency aspects of asynchronous distributed systems, with potentially adversarial scheduling of events. On the other hand, realistic analysis must accommodate the fact that, in most interesting cases, it is impossible to completely prevent successful attacks against the protocol. Instead, we can only bound the success probability of attacks that use a bounded amount of computational resources, based on underlying computational hardness assumptions.

Cryptographic modelling and analysis is typically complex, involving many subtleties and details, even when the analysed protocols are simple. Furthermore, analysis is handwritten and often tedious to verify. These factors make the security analysis of cryptographic protocols susceptible to errors and omissions.

This project demonstrates how to cast cryptographic security analysis of distributed protocols within the Probabilistic I/O Automata (PIOA) framework of Lynch, Segala and Vaandrager. This framework provides standard tools for arguing rigorously about the concurrency and scheduling aspects of protocols. It supports reasoning with multiple levels of abstraction and has a well-defined notion of composition. Consequently, using the PIOA framework can help in making cryptographic analysis more precise and less susceptible to errors.

In the context of this project, we aim to develop general techniques applicable to the analysis of a wide range of security protocols that exhibit various levels of complexity in terms of adversarial behavior or the use of cryptographic primitives. As a first step, we are currently analysing a relatively simple protocol, the two-party Oblivious Transfer (OT) protocol, in the presence of a semi-honest adversary (essentially, an eavesdropper). The particular OT protocol we are studying is the classic protocol by Even, Goldreich and Lempel, which uses trapdoor permutations (and hard-core predicates for them) as the underlying cryptographic primitive. For the underlying cryptographic notion of security, we start from Canetti's Universally Composable Security.

In spite of the relative simplicity of the investigated case, the exercise is non-trivial and requires addressing a number of fundamental issues. These include modelling resource-bounded computations, resource-bounded adversarial behaviour and scheduling, combining non-deterministic and probabilistic choices, and modelling computational hardness assumptions. Some of these are beyond the reach of the existing semantic theory of probabilistic I/O automata. This project therefore involves using not only probabilistic I/O automata in security protocol analysis but also extending the basic theory to address the requirements of such analysis.

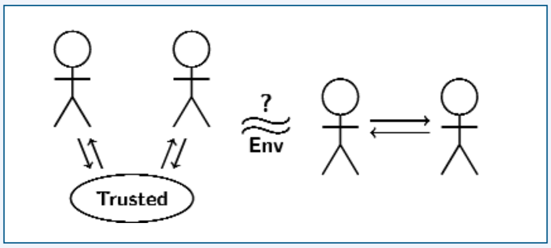

The environment tries to distinguish the ideal-world system, representing the specification of the protocol, from the real-world system running the actual protocol.

Approach

One common approach to simplifying cryptographic protocol analysis and improving its correctness is to model cryptographic primitives as 'symbolic operations', or 'ideal boxes', which represent the security properties of the primitives in an idealized way that involves no error probabilities or computational issues. This approach is quite promising; however, it does not completely remove the need for cryptographic analysis of protocols. Rather, it only proves the security of the overall protocol by assuming the security of the cryptographic primitives in use. One still has to prove the security of these primitives in a fully fledged cryptographic model with all its subtleties. Our project proposes an alternative (in fact, complementary) approach to making cryptographic protocol analysis more systematic and rigorous, and thus less susceptible to errors.

We benefit from the powerful proof techniques that have traditionally been used within I/O automaton frameworks. In particular, we express the system at multiple levels of abstraction, where the highest level in the abstraction hierarchy represents the specification of the protocol and the lowest represents the real-world system running the protocol. We demonstrate simulation relations between these levels that allow us to conclude that the real-world system implements the abstract specification. The composition operation for PIOAs makes it possible to separate the specification of correctness and security requirements of the protocol, and to express the real-world system naturally, as a composition of logically separate units interacting with another.

Future Work

We would like to assess the techniques we have developed so far in the analysis of more complex protocols. This complexity arise through more powerful adversaries, more complex interaction patterns between the components of the protocol, or more subtle uses of cryptographic primitives.

One limitation of our current approach is that the existing scheduling mechanism is oblivious to execution histories and hence significantly limits the capabilities of the adversarial components. We will investigate how our assumptions about task-PIOAs and scheduling mechanisms can be relaxed to enable the analysis of a larger class of protocols.

Once our modelling and proof techniques become more general and established, we can focus on mechanizing our proofs with interactive theorem-provers, or even automate some or all of the proof steps.

Link:
Technical report:
http://www.lcs.mit.edu/specpub.php?id=1792

Please contact:
Olivier Pereira, Université catholique de Louvain (FNRS), Belgium
Tel: +32 10 47 91 63
E-mail: pereira@dice.ucl.ac.be

Cryptographic Security by Swamping Adversaries with Quantum Information

by Ronald Cramer and Serge Fehr

In the ongoing quest for cyber security, quantum mechanics plays an increasingly important role. The Cryptology and Information Security Group at CWI in Amsterdam recently designed a new practical scheme for so-called Oblivious Transfer. This is a cryptographic tool known to be sufficient to secure any cooperation among mutually distrusted parties. The new scheme is based on quantum mechanics and the technological difficulty of storing photons without disturbing their quantum state. The basic idea behind the scheme is to swamp an adversary with more quantum information than they can possibly store. According to quantum information theory, the uncertainty a potential adversary has about the total amount of information can be exploited by the honest parties in order to secure private information.

The security of most of the cryptographic schemes in use today is based on the assumption that some computational problem is difficult to solve. For example, the widely used RSA encryption scheme is based on the difficulty of factoring large integers. Part of the research done in the Cryptology and Information Security Group at CWI is devoted to the study of models that allow for cryptographic schemes based on assumptions of a different nature.

One approach is to use quantum-mechanical effects, and to try to base the security solely on the laws of quantum physics, without restricting the adversary's power. This has proven to be fruitful in the context of Secure Communication, where two mutually trusting communicating parties require security against an outside adversary. However, it falls short of providing security for mutually distrusting communicating parties against each other. This problem is known as Secure Cooperation. A simple but already non-trivial Secure Cooperation is the generation of a random bit, solely by interactive (asynchronous) communication, such that each party is convinced that the bit is not biased against his preference. Another example is the so-called Millionaires' Problem, where two parties would like to find out who is richer, but in such a way that the wealth of each is not revealed to the other. In general, using Secure Cooperation techniques allows mutually distrusting parties to evaluate any given function on their private inputs – revealing the function value but protecting the individual private inputs.

A fundamental form of Secure Cooperation is Oblivious Transfer. A sender transmits two bits to a receiver, who may choose which one he would like to receive. This is executed in such a way that the sender does not learn the receiver's choice and the receiver only learns the chosen bit and not the other one. Oblivious Transfer is complete for Secure Cooperation in that, at least in principle, it can be used as a building block to achieve any Secure Cooperation.

Another approach is to design schemes in such a way that a possible adversary is literally swamped with data. He will have to store them all in order to break the scheme, and to obtain classified information and/or provoke a malfunction of the scheme. However honest users of the scheme who are not trying to break it, need only store a small part of the information. Such a scheme is then secure under the assumption that a potential adversary does not have overwhelming memory, even though he might have infinite computing power.

A particularly interesting feature of this so-called Bounded-Storage Model is 'everlasting security': if the adversary fails to store the mountain of data during the execution of the scheme, then his chance of breaking it is lost forever together with the data he failed to store. That would never change, even if they should obtain larger storage capacity sometime in the future. This is in sharp contrast to conventional computational cryptography, in which schemes may be broken 'in retrospect', once an adversary has gained sufficient computing power. On the other hand, while basing cryptography on computational assumptions allows for very efficient schemes, a major disadvantage of the Bounded-Storage Model is the immense amount of data that must be communicated (even for the honest users) in order to be able to swamp a potential adversary's memory.

Recently, in collaboration with researchers from Aarhus University (Denmark), CWI's Cryptology and Information Security Group has put forward and studied a variant of the Bounded-Storage Model which is based on swamping the adversary's quantum-memory: the Bounded-Quantum-Storage Model. Quantum mechanics provide a good platform for such an approach: on the one hand, according to Heisenberg's Uncertainty Principle, converting quantum information to classical information by measuring irreversibly destroys some of that information, and the challenge is to arrange it so that the adversary cannot afford this loss while honest users can. On the other hand, in contrast to classical information, storing quantum information is very difficult (at this point essentially impossible), and thus it requires very little to swamp the adversary's quantum memory.

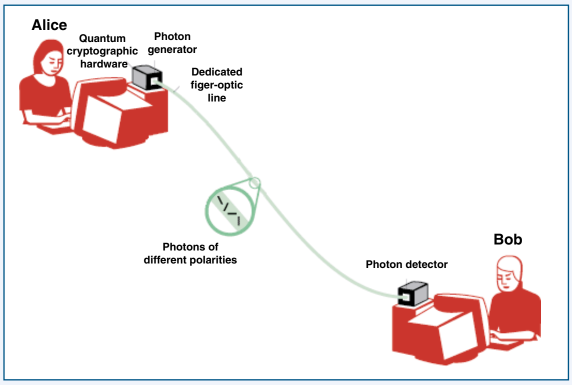
Alice
Quantum cryptographic hardware
Photon generator
Dedicated figer-optic line
Photons of different polarities
Bob
Photon detector

Using quantum communication to achieve cryptographic security.

In this model the CWI researchers have constructed a scheme for Oblivious Transfer, and thus indirectly for any Secure Cooperation. The scheme involves the transmission of a stream of single photons (or other atomic particles), and it is proven secure under the sole assumptions that Heisenberg's Uncertainty Principle holds and that a potential adversary cannot store more than a large fraction of the transmitted photons, even though he may have infinite computing power and classical memory.

Apart from quantum cryptography, CWI's Cryptology and Information Security Research Group focuses on all mathematical aspects of cryptology, such as algebraic secret sharing and secure multi-party computation, computational number theory (eg the Number Field Sieve Project for factoring RSA challenge numbers), information-theory-based cryptography, public-key cryptography in general (in particular, chosen cipher-text security for encryption schemes), cryptographic protocols, and formal security analysis.

Link:
http://www.cwi.nl/crypto

Please contact:

Ronald Cramer, CWI, The Netherlands
Tel: +31 20 592 4166
E-mail: Ronald.Cramer@cwi.nl

Serge Fehr, CWI, The Netherlands
Tel: +31 20 592 4257
E-mail: Serge.Fehr@cwi.nl

Building a Stochastic Model for Security and Trust Assessment Evaluation

by Karin Sallhammar, Svein Johan Knapskog and Bjarne Emil Helvik

The ICT systems of today are complex inventions and we rely on their existence in almost all aspects of our everyday life. It is therefore crucial that they can provide the services we need whenever we require them. Due to the interconnection of networked systems, attacks are becoming increasingly sophisticated and can be performed remotely. To what degree can we trust that a system will perform its intended task in a secure and reliable manner?

The new paradigms of ubiquitous computing and high capacity data transfer have opened up the Internet as the main area for information interchange and electronic commerce. Attacks against the computer networks used by modern society and economics for communication and finance can therefore threaten the economical and physical well-being of people and organizations. To allow continuous risk estimation of today´s ICT systems, there is an urgent need for models providing probabilistic measures of operational security.

Stochastic Modelling

During the last decade, significant research has been performed on applying traditional dependability techniques to quantify the security attributes of ICT systems. In particular, stochastic modelling techniques such as Markov chains

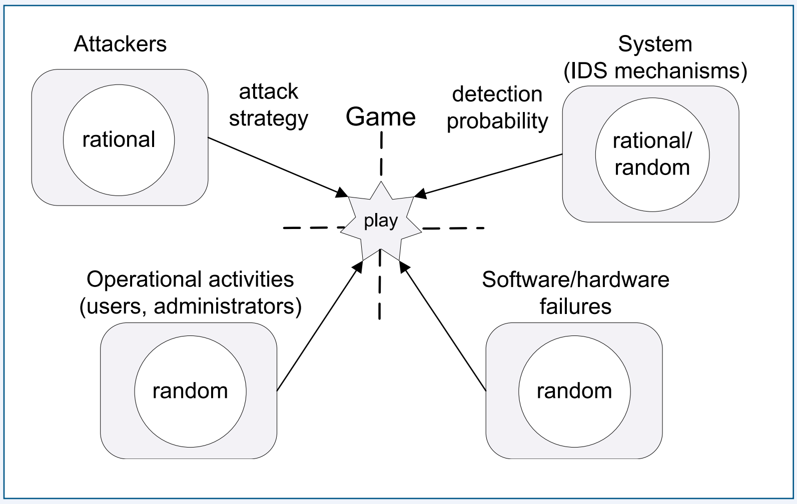

The interactions between the attacker and the system modelled as a stochastic game.

or stochastic Petri nets have been identified as promising approaches. In a dependability context, a system will continuously be vulnerable to failures of software and hardware, which may transfer the system from a good state into a corresponding failed state. Usually, these methods do not consider failures due to malicious acts. However, by using an analogy between a system failure and a security breach, it is possible to model an intrusion attempt as one or more state changes that transfer the system into a security breach state, ie a state which deviates from the specified security policy. The use of a stochastic model, which combines security-related attacks with traditional dependability fault sources has a wide range of application:

- to quantify security: by using the steady-state probabilities of the stochastic model, one can calculate operational measures such as the 'mean time to security compromise' for the system
- for trade-off analysis: for example, one may evaluate the possible effect of security countermeasures before implementing them
- as a method to help administrators find optimal defence strategies and to calculate the expected loss associated with different strategies.

However, attacks may not always be well characterized by models of random nature. Most attackers will act with intent and will consider the possible consequences (satisfaction, profit and status versus the effort and risk of their actions) before they act. One of the remaining challenges is therefore how to incorporate intelligent attacker behaviour into the stochastic models.

The Game Model

At the Q2S centre at NTNU, Norway, we are developing a stochastic model that can be used for assessing the security and trustworthiness of ICT systems. Our model considers all aspects that may affect the security or dependability attributes of the system, including:

- normal user behaviour
- administrative activities
- random software and hardware failures
- intentional attacks.

To incorporate intentional attacks in the model, the attacker behaviour must be predicted. By using a stochastic game model, we can compute the expected attacker behaviour for a number of different attacker profiles.

The game model in the figure is based on a reward/cost concept. This assumes that attackers will consider the reward of successful actions versus the possible cost of detection before they act, and that they will always try to maximize the expected outcome of the attack. The dynamics of the states of the stochastic games form a Markov chain, under the assumption that attackers, users and administrators do not change their behaviour over time. Having solved the stochastic game, the expected attacker behaviour can then be reflected in the transitions between states in the system model, by weighting the transition rates according to a probability distribution. In the final step, the corresponding stochastic process is used to calculate security measures of the system, in a similar manner to the common availability and reliability analysis of ICT systems.

Previous research has shown that stochastic models can be used to model and analyse the trustworthiness of ICT systems in terms of both security and dependability attributes. Our current research indicates that game theory is a suitable tool for incorporating the expected attacker behaviour in such models. However, verifying the method's ability to predict real-life attacks will require further research, including validation of the model against empirical data.

Link:
http://www.q2s.ntnu.no

Please contact:
Karin Sallhammar
Centre for Quantifiable Quality of Service (Q2S), NTNU, Norway.
E-mail: karin.sallhammar@q2s.ntnu.no

Flexible, TCP/IP-Based and Platform-Independent Management of Biometric Data for Access Control Systems

by Burkhard Stiller

The use of biometric data for access control, eg fingerprint scans, offers a simplification of room access control or a key-management function. A detailed cost analysis of existing systems shows that in addition to specific security requirements, the loss of physical keys or magnetic access cards can increase the overall costs dramatically. A large employee turnover, a large number of access points, such as doors and gates, and various user- as well as key-groups compounds the situation. A biometric data management approach would solve these problems and additionally, would ensure that keys could not be passed on to different personnel.

The collaborative project 'Biometric Access Control', run with the ITIS e.V. (located at the University of Federal Armed Forces, Munich, Germany) and Biometronix GmbH (Munich, Germany), has developed a new, platform-independent, flexible and TCP/IP-based Biometric Local Area Network Control Center (BioLANCC) in support of both access-control systems and single sign-on solutions.

The following example illustrates the motivation behind the current work on networked biometric data control centres. Consider a holding company with its headquarters and human resources department in Munich, and two external research laboratories in Budapest and Zürich. Since the main administration server in Munich hosts all the access rights of employees, visitors, temporary employees and cleaning staff, Munich has an overview of the entire company. Access rights may be changed (employee relocation), withdrawn (employee retirement), or added (new employment). Of course, the two research labs in Budapest and Zürich also need to be able to change access rights according to their local requirements. In addition, many different types of access-control hardware, such as fingerprint scanners, need to be grouped according to labs, research groups, or buildings. Such an organization demonstrates the administrative and organizational gains to be achieved with BioLANCC.

Thus, the development of BioLANCC was driven by a detailed organizational requirements analysis, key security requirements, and the need for an operation on multiple hardware devices as well as on different operating systems. The logically centralized management approach includes distributed access and maintenance of single users' access. In particular, the fast and efficient support of access rights for temporary employees, or students in an academic environment, is a beneficial feature not found in any other existing system. Remote and centralized control have been integrated based on a networked TCP/IP solution.

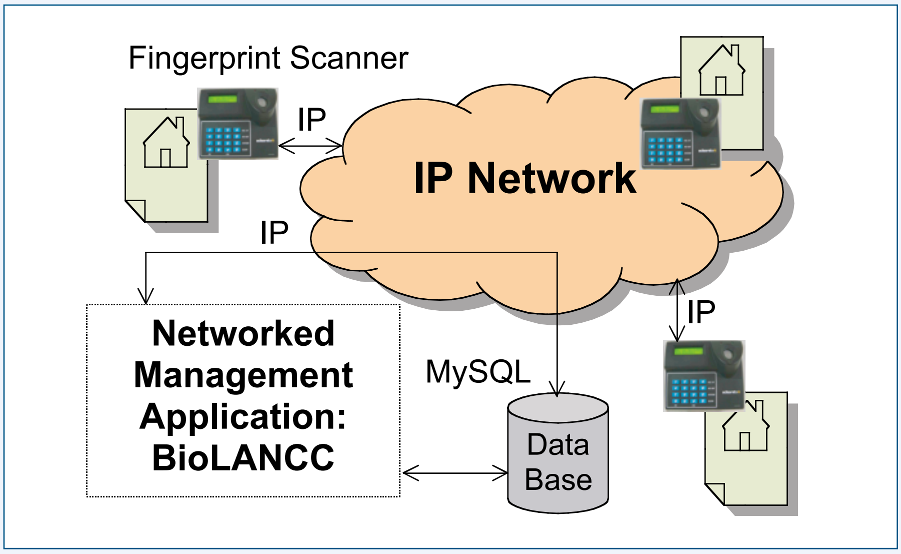

The interactions between the attacker and the system modelled as a stochastic game.

Figure 1 depicts the system's architecture. The implementation basis is determined by a two-tier, Java-programmed application, which accesses an SQL database. The initial use of MySQL is in the process of being replaced by a Postgress database. In addition, a three-tier version is under development, in which multiple thin management clients will be supported. BioLANCC's specific features include the standardized use of an IP network for communications between all biometric access devices and the control centre, and full platform independency; currently supported are Windows, Linux, Solaris, and Mac OS X. Other features are the integration of multiple hardware devices (based on an open device interface and the BioAPI), a flexible user and device-group management, a time-zone manager for users and devices, and a user-selectable, graphical interface for all management operations (drag and drop). Finally, a reporting module customizable by the user allows for the supervision and backtracking of correct access as well as the investigation of, for example, error reports and failed authentications.

BioLANCC has been successfully implemented at the Department of Computer Science of the University of Federal Armed Forces in Munich. The system includes eighteen fingerprint scanners and caters for several hundred permanent and temporary employees, including professors, research assistants and students. At this stage, fingerprint scanners (V20 devices) from Identix are in use; this installation is currently the largest biometric access control system of its type in Europe. During the course of the project, additional requirements were identified and are being integrated. This covers the BioAPI mentioned above, as well as different types of biometric hardware devices (eg iris scanners). A device wizard offers the possibility of integrating a new access point easily into the system. Multilingual support is also at hand. Finally, its continued effective operation in a harsh academic environment demonstrates BioLANCC's robustness.

Links:
http://de.biometronix.com/index.php?biolancc
http://us.biometronix.com/index.php?biolancc

Please contact:
Burkhard Stiller,
University of Zurich, Switzerland
and ITIS e.V. Munich, Germany
E-mail: stiller@ifi.unizh.ch

BT's Security Research

by Theo Dimitrakos

BT's Security Research Centre was formed in April 2004 under the BT Group Chief Technology Office as part of the ongoing expansion of security research within BT. The goal of the centre is to deliver world-class security research to support BT's growing business in mobility, broadband and ICT (Information and Communications Technology). The new research centre, led by Bryan Littlefair, plans to grow to 25-30 staff, many PhD-qualified. This team is in addition to the many engineers working on development projects as part of the BT security programme.

The Research Challenge

Many emerging security challenges are the result of disappearing boundaries:

- innovations in mobile technology and the convergence of personal communications and computing devices mean that the ability to access corporate data is no longer limited to making a physical connection to the corporate network
- flexible working practices and networked personal electronics are blurring the boundary between work and leisure time, making work an activity rather than a place
- converged multi-service IP networks are removing the distinction between voice/video and data, and computation and communication
- service-oriented technology such as Web services, Grid computing and utility computing are being used to interconnect computing resources and databases, and to connect business processes across enterprise boundaries, heralding the era of the virtual organization.

Such changes and advances are not insecure in themselves, but security has traditionally been about defining and defending boundaries. In a world without clear boundaries, new approaches will be needed.

The research programme at BT is addressing these challenges by conducting research in three broad areas:

- securing the converged network
- securing the virtual organization
- security management frameworks.

The results of this research will be critical to BT's ability to secure its own business against attack, and will help the company develop its position in the market for security and compliance services.

Securing the Converged Network

To improve the security of BT's networks, the company is researching ways of identifying the true source of any packet of data sent across the Internet. Attackers often 'spoof' IP addresses to obscure their identity and location, so a reliable means of working out exactly where messages came from, regardless of the claimed IP address, would be a strong deterrent.

BT is also researching ways of structuring networks that could enhance security and performance, and improved ways of monitoring the data generated by Internet intrusion detection systems. Simulations are used to help researchers improve their understanding of the security issues and protection needs of IP and MPLS networks.

Research is also being conducted into improved ways of validating the identities of employees and customers. Among other options, BT is exploring the use of speaker verification technologies to 'recognize' people based on the characteristic acoustic features of their voice – that is, recognizing not what is said but who is saying it.

In the context of the EU-funded project Ambient Networks (http://www.ambient-networks.org), the team is working towards creating next generation of network solutions for mobile and wireless systems beyond 3G. It will enable scalable and affordable wireless networking while providing rich and easy-to-use communication services for all. It is geared towards increasing competition and cooperation in an environment populated by a multitude of user devices, wireless technologies, network operators and business players.

A collaboration between BT and Imperial College London is exploring the issue of providing overall dynamic protection against cyber-based attacks.

Securing the Virtual Organization

Trust is essential in any form of communication and especially in securing relationships between enterprises. BT's researchers are exploring the development of systems that can enable the appropriate sharing of information assets while protecting them from external software attacks and physical theft. It is also exploring how new technology might make intranets more flexible – something closer to a secure online shared workspace than a physically separate network.

The team is also working as a partner in the EU-funded project TrustCoM (http://www.eu-trustcom.com), which is working to establish a way of representing trust and contractual relationships that enables the creation of collaborative business processes between organizations. The ideas are being tested against the needs of collaborative engineering projects, and those of small and medium-sized enterprises that could work together to deliver services to clients as a 'virtual supplier'.

The team is also conducting research in the context of the EU-funded project GUIDE (http://www.guide-project.org/), which aims at creating an architecture for secure and interoperable e-government electronic identity services and transactions for Europe. The project's approach is multi-disciplinary and includes technology, and procedural and policy development across Europe.

In addition, the twin topics of Web services and Grid security are being researched to develop secure ways of integrating services and resources across enterprise and organizational boundaries.

Security Management Framework

BT's researchers are developing a long-term roadmap that outlines how governance and engineering issues will need to be addressed in the future.

Partners

BT's research programme includes both projects undertaken entirely by the company's own researchers and those undertaken together with research partners from leading universities and research institutes, corporate research centres, and end-user organizations. Some projects are sponsored by strategic EU and UK government programmes, while BT's corporate research partners include Microsoft, IBM, SAP and BAE Systems. The company also has a growing number of links with HP Laboratories, Bristol, complementing the commercial alliance between BT and HP in the ICT marketplace. Finally, BT's Security Research Centre has appointed Dr Theo Dimitrakos as a representative on the industrial advisory board of the newly established ERCIM technical working group on Trust & Security.

Acknowledgement

The following BT colleagues have contributed to this security research overview: Theo Dimitrakos, Mark Drew, Andy Jones, Paul Kearney, Bryan Littlefair, Mark Pawlewski.

Please contact:
Bryan Littlefair, Head of Security Research
Security Research Centre, BT Group CTO, UK
Tel: +44 (0) 1473 649 427
E-mail: bryan.littlefair@bt.com

Theo Dimitrakos, ERCIM Trust & Security
Working Group representative,
Security Research Centre, BT Group CTO, UK
Tel: +44 (0) 1473 646706
E-mail: theo.dimitrakos@bt.com

Bringing the Common Critera to Business Enterprise

by Christophe Ponsard, Philippe Massonet and Jean-François Molderez

Security is a primary concern for companies, many of which are quickly becoming aware of it but are still ignorant of how to handle it correctly. The Common Criteria (CC) provides a systematic approach with a methodology and a set of reusable protection profiles. In practice however, they have proven difficult to adopt and costly to use. In order to overcome these obstacles, CETIC (Centre of Excellence in Information Technology and Communication) is promoting an adapted usage of CC, and has developed supporting tools based on an existing requirements-engineering toolkit.

Information security has become a critical aspect in the operations of many companies, from multinationals to SMEs. Failing to address it adequately can have dramatic consequences: loss of access to data can freeze an organization, sensitive information might be sold to competitors, and so on. In an ever more connected world, ever more exposed to malicious users, many companies are exposed to such a painful experience before becoming aware of the importance of designing for security. Such design starts by correctly specifying what is to be protected, against what threat, and in which environment.

Over the years, a number of standards that were developed separately finally converged to the ‘Common Criteria’ (CC), an internationally agreed method for gathering, organizing and assessing the security requirements for an IT product. CC provides a scale of eight evaluation assurance levels (EALs) in line with the effort and techniques required to develop a secure IT product. The resulting process is similar to a typical requirements engineering (RE) process (see figure). It includes an analysis of the environment, the definition of high-level objectives to address potential threats and their refinement into concrete requirements. Certification requires a strong rationale to be provided for the artefacts under evaluation.

Besides the fact that CC is an international standard, it is also based on sound RE practice and includes flexible adaptation mechanisms. Designers are free to reuse and/or adapt predefined requirements, or to add specific and customized requirements. They can also build requirements models at various abstraction levels, thus leaving details for later implementation stages.

Connecting with Business Enterprise

Despite the inherent qualities of CC that makes it a good vehicle to address security concerns, its use remains focused on specific projects, generally addressed by large companies and mainly for the purpose of product certification.

Based on discussions with Belgian companies of various sizes, CC was perceived as difficult to understand, learn and apply in the context of a small structure, and definitely not worth using for systems in which security is less critical. In order to overcome those difficulties, a number of improvements have been identified. For example, companies often ignore the fact that several protection profiles are already freely available; these can provide an exhaustive checklist of security requirements for various systems such as cash machine, firewalls, e-purse systems and so on. The conformance between the system and the profile can also be managed at the right assurance level, achieving a good balance between system criticality and available resources. It is also possible to apply only part of the process. This is the case in the Electronic Money System Security Objectives: based on a comprehensive risk analysis for e-money systems, this document develops a list of security objectives that should be fulfilled in order to cover these risks/threats in a given environment.

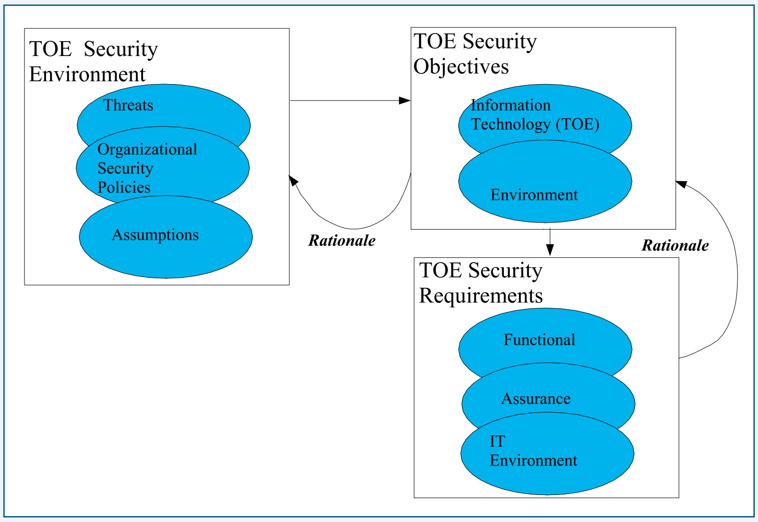

An overview of the CC process.

An important part of the problem is also that security is a ‘vertical concern’, implying all system levels from hardware to middleware to software applications. A successful security design should therefore rely on an interdisciplinary approach, involving people from fields such as software quality, distributed systems and electronic systems, the three main fields in which CETIC is active.

Finally, tool support is also crucial. So far there exist few tools for the Common Criteria (and more generally in the RE field) that have a real methodological support. The most widely used tool is a word processor, possibly complemented with a traceability tool. An ideal tool should provide a rich meta-model capturing relevant security concepts such as the environment, threats, agents (user or malicious), security objectives and requirements. It should support the CC library, query it and provide operations (refinement, iteration, selection etc) on it. Most importantly, it should capture the rationale between all the required CC artefacts in order to automate document generation as far as possible.

Further work

Our current goal is to adapt an existing goal-oriented requirements engineering framework called Objectiver to the security domain. This tool provides services like diagram edition, query support and semi-automated document generation. Being based on meta-modelling, it can easily be tuned to support the Common Criteria. We also plan additional support to navigate within the CC, manage rationales and generate protection profiles and security target documents.

CETIC is currently working on this area in collaboration with Belgian companies and universities, in the context of various European projects. Based on our expertise in requirements engineering for critical systems and in tool development, CETIC is working on the specification of domain-specific protection profiles and the development of an adapted tool support to ease the use of the Common Criteria. Security in domains such as the GRID and ambient intelligence are being investigated, as well as the application of CC at a high assurance level with the use of formal methods.

Link:
Objectiver: http://www.objectiver.com

Please contact:
Philippe Massonet,
CETIC Research Center, Belgium
Tel: +32 71 91 98 27
E-mail: phm@cetic.be

Articles in this Section

Born of INRIA — Twenty Years Creating Spin-Off Companies

by Laurent Kott

The foundation of a company is an opportunity to make the most of the enthusiasm, experience and dynamism of a research team. Naturally, the creation of a company with solid growth potential requires that a complete team be assembled, highly competent in many other areas than science and technology.

INRIA thus had to set up a whole process. This process begins with the conviction that certain research results are very likely to spawn innovations on the market. The right person or persons to head up the project must then be found, a motivated and enthusiastic team must be built and the initial setting up and transition problems must be resolved. In addition, a business model capable of convincing and bringing in investors must be developed, along with a good understanding of the market.

The main reason behind the birth of INRIA-Transfert was a realization of the urgent need to help finance the creation of innovative companies. We decided to start by setting up a subsidiary with the aim of establishing a clear boundary between the public subsidies allocated to INRIA and the funds dedicated to creating start-ups. Hence, a concern for clarity and transparency was paramount in the creation of INRIA-Transfert. Once it had been set up, we realized it could also play another role: providing guidance to project leaders to help them build their companies better and faster.

The only dedicated structure in the French ICST sector, INRIA-Transfert benefits from the experience accumulated by INRIA. The institute had already contributed to creation of start-ups over the years, as and when opportunities arose. The name INRIA-Transfert was a way to capitalize on this past experience and the institute's renown. Indeed, thanks to the positive image associated with INRIA's reputation as a centre of scientific and technological competencies and the fact that it is well known to financiers and investors, its subsidiary is able to offer this function of guidance not only to its own research scientists but also to people associated with other research institutes, both public and private. This is a deliberate and conscious choice: from our point of view, encouraging company creation is a concrete way of ensuring that research carried out in public (and private) research centres is brought to a successful conclusion and transformed into products and services through the creation of start-ups.

We believe that having a network of innovative companies in this very important sector is in the interests of both France and INRIA. It is a way of transferring our approach and being attentive to what is happening in the real world. My enthusiasm and conviction does not stop me from being realistic: I'd like to stress the fact that turning an idea into a successful new company depends 80% on prevailing economic conditions. However, since the creation of INRIA-Transfert, this opportunity is clearly one of the career paths offered to INRIA researchers. We have proved our capacity to foster the creation of start-ups, directly or indirectly, through the funds managed by I-Source Gestion. INRIA-Transfert is a challenge that we are in the process of winning.

INRIA-Transfert Assessment: A Sign of Recognition

After a year or more of guidance and evaluation carried out in collaboration with the entrepreneurs involved, the best company creation or development projects are awarded one of two special certifications. They are a concrete representation of our evaluation as of the date on which they are awarded, depending on the company's state of development. They reflect a decision and a commitment, associated with a certain acceptance of risk. It's also a way of giving these start-ups greater credibility.

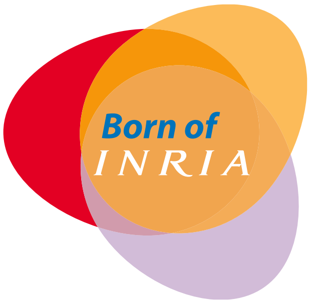

Pertinence IT® certification confirms the usefulness of the technology and the economic viability of the start-up based on a number of evaluation criteria. It is awarded by INRIA-Transfert in consultation with its team of experts and the company's pilot customers, and illustrates their opinion of the company's economic value and its ability to survive in the marketplace.

Croissance IT® (Growth in English) certification is awarded to companies that have already obtained Pertinence IT® certification. It indicates that the start-up's products, team, technology and sales strategy put it in a favourable position in a market with high growth potential. This certification is awarded by INRIA-Transfert, its experts, and institutional investors that already hold part of the company's capital or are likely to invest in it.

Start-Ups Bearing the INRIA Trademark

As of 19 April 2005, companies stemming from INRIA have the right to use the slogan 'Born of INRIA'. Brought together under the same slogan, modelled on the well-known 'guaranteed vintage' label, 'Born of INRIA' start-ups will be able to capitalize on the institute's name and renown. It is also a way of highlighting the strong relationship between each of these companies and the institute, either through the people responsible for their creation or through the technology they apply which was originally developed by INRIA.

More than ever, INRIA remains convinced that company creation in the field of ICST is one of the most efficient and long-lasting means of technology transfer. INRIA wishes to express its admiration and gratitude to the research scientists, the engineers and their collaborators who embark on such ventures. Success or failure partly depends on the project proponents, and on general conditions. Let us hope that confidence returns to economic and financial circles, because it has never deserted scientists and technology experts.

The technology companies stemming from INRIA

AM^2 Systems
Semantic-web technology supplier
www.am2systems.com

Athys Technologies
Behaviour detection and measurement
www.blueeyevideo.com

CAPS entreprise
Code optimisation for embedded sytems
www.caps-entreprise.com

Cogenit
Tests and engineering in telecommunications
www.cose.fr

Cose
Data acquisition for scientific computing
www.cose.fr

Diatelic
Telemedicine system for patients undergoing CAPD
www.diatelic.com

Distene
Solution for pre- and post-processing to optimize the simulation-based design chain
www.distene.com

Dolphin Integration
Products and services in microelectronics
www.dolphin.fr

eliKya
Knowledge acquisition aid
www.elikya.com

Enginest Software
Planning and scheduling tools
www.enginest.com

Ergomatic Consultants
Computer system ergonomy

Esterel technologies
Reliable system design software
www.esterel-technologies.com

Finoptech
Computation intensive solutions for financial engineering
www.finoptech.com/

GeometryFactory
Geometric Software Components
www.geometryfactory.com

Gene-IT
Bio-informatics and genome
www.gene-it.com

Icatis
Software for clustering and Grid computing
www.icatis.com

ILOG
C++ and Java software components
www.ilog.fr

Istar
Air and space image processing
www.istar.fr

Jalios
Content Management solution
www.jalios.com

Kelkoo
Internet buying guide
www.kelkoo.com

KEENEO
Software editor for behaviour recognition and intelligent video surveillance
www.keeneo.com

LTU technologies
Software solutions that analyze and describe the content of images
www.LTUtech.com

Lorin
Real time calculators
www.lorin.fr

Medience
A software vendor for data integration
www.medience.fr

N2NSoft
Simulation and optimization of large IP networks
www.n2nsoft.com

Noesis
Computer vision and image processing
www.noesisvision.com

PolySpace Technologies
Testing and validation tools
www.polyspace.com

Probayes
Mastering uncertainty
www.probayes.com

QuantifiCare
Services to provide accurate and objective measurements of disease evolution
http://www.quantificare.com

RaisePartner
Quantitative analysis solutions for finance
www.raisepartner.com

Realviz
Digital special effects and animated computer images
www.realviz.com

Robosoft
Mobile robots and associated peripherals
www.robosoft.fr

RT Consultant
Ergonomics and applied human factors in aeronautics and information technology
www.rt-consultant.com

ScalAgent
Mediation Software Systems
www.scalagent.com

TAK Imaging
Development in microelectronics and multimedia
www.takimaging.com

Tak'Asic
Development in microelectronics and multimedia
www.takasic.fr

Time-AT
Conception, integration and forework of vision, Industrial computing and telecommunications products
www.timeat.fr

Trusted Logic
Open and secure architectures for embedded systems
www.trusted-logic.fr

UDcast
One-way broadcast on the Internet
www.udcast.com

VSP-Technology
Walking around a model in real time
www.vsp-tech.com

W2G Technologies
A new breed of computing architecture
www.web2grid.com

Xyleme
Enabling intelligent access to XML content
www.xyleme.com

Zeliade Systems
XML technology to the processing of financial derivatives products
www.zeliade.com

Links:
http://www.inria.fr/valorisation/20ansdestartup/index.en.html

http://www.inria-transfert.fr/english/index.html

Please contact:
Laurent Kott, INRIA-Transfert, France
E-mail: Laurent.Kott@inria.fr

The GridML Project: Next-Generation Data Mining on High-Performance, Parallel Distributed Systems

by Csaba Szepesvári

The aim of the GridML project is to develop a prototype of an advanced, next-generation data-mining system, capable of more efficiently exploiting the vast computational resources of distributed, parallel computational environments. Researchers will look not only into the problem of predicting the performance of the algorithms separately, but also at the whole process of data mining, where the cost of making such predictions is also taken into account.

GridML (Grid-based Machine Learning) is a research project supported by the Economic Competitiveness Operative Programme (GVOP) of the Hungarian Government. The project began in January 2005 and will run for two years. The GridML Consortium consists of four partners: SZTAKI (coordinator) with the participation of the Machine Learning Group and the Laboratory of Parallel and Distributed Systems, the Research Group on Artificial Intelligence from the University of Szeged, AAM Consulting Inc, and Magyar Telekom Ltd (T-Systems).

There exist numerous applications areas for data mining. Examples include marketing, chemistry, CRM, environmental control and security, just to mention a few.

One characteristic feature of data mining is that it works with extremely flexible models, often involving dozens, sometimes hundreds of tunable parameters. The promise is that such flexible models are capable of capturing the complex relationships hidden in data. The price one must pay is a significant increase in the computational cost of fitting such models to the data. However, experience has shown that the quality of the resulting solutions makes this approach worthwhile.

Often however, the process is expensive not only in terms of computational resources but also in a monetary sense. This is primarily due to the high human resource costs, which follow from the current practice of assigning several data-mining experts for the full duration of the project. The tasks of these experts range from defining the problem, through obtaining and cleaning the data, to selecting the right model and tuning the models' and learning algorithms' parameters. Out of these, the latter two contribute most significantly to the total time spent on the project.

Researchers plan to apply bandit theory to optimizing the execution of data-mining algorithms on Grids.

Due to the high cost of experienced experts, many data-mining projects are abandoned or never see the light of day. Another problem with the current approach is that the quality of the solutions obtained depends largely on the preferences or the experience of the experts involved. This significantly decreases the predictability of the results of the mining process, making companies uncertain of whether it is worth starting the project. It is not only industrial data mining that suffers from the issue of ad hoc parameter and model selection: researchers also face this problem due to the lack of sufficient computational power.

One possibility to overcome the computational bottleneck is to employ a large number of computers in parallel, connected into a cluster or Grid. Projects devoted to migrating data-mining algorithms to such distributed environments have existed for some time. The two most widespread approaches are either to parallelize the algorithms themselves or to let many instances of the algorithms (same or different) run in parallel.

It would be naive to think that distributed data mining is a universal cure. With the increase of computational resources, researchers and data-mining experts see new opportunities that would not previously have been feasible, and the available resources an soon become insufficient. Hence, the central issue investigated in our project is how to better utilize the available resources, with the ultimate goal of the project being to develop a prototype data-mining system that organizes computations in such a way that given a fixed computational budget it produces the best possible models in a reliable manner.

As such, the main research theme of the project belongs to the domain of meta-learning, a relatively new field concerned with the relationship of tasks or domains and learning algorithms. However, unlike most approaches to meta-learning, in our project the emphasis is on the online performance of the algorithms. Most previous methods produce an initial ranking of the various algorithms without giving further advice on how best to utilize the resources. In contrast, our approach not only provides an initial ranking, but also continuously monitors the performance of the various methods throughout the mining process, and then feeds the result back to the resource allocation strategy. The resource allocation problem is treated as an online learning problem, the best-known example of which is the so-called multi-armed bandit problem where a gambler faces the problem of deciding which arm of a K-slot machine to pull to maximize his total reward in a series of trials. The payoffs of the machines are random and initially unknown. After trying each machine at least once, the gambler has to decide if a machine with a high estimated payoff should be chosen or one that has a lower estimated payoff but which has been tried less often.

In addition to researching the performance of various advanced online resource allocation schemes for optimizing distributed data mining, another goal of the project is to create a prototype system capable of meeting several technical requirements such as accountability, security and the automated documentation of the mining process.

Link:
http://zaphod.aml.sztaki.hu/wg/index.pl/gridml

Please contact:
Csaba Szepesvári, SZTAKI, Hungary
Tel: +36 1 279 6262
E-mail: szcsaba@sztaki.hu

Crossing the Rubicon: from Haskell to .NET through COM

by Beatriz Alarcón and Salvador Lucas

.NET is an emerging Microsoft project which promotes a new framework for Software Development that emphasizes the use of Internet resources and the interaction between components written in different programming languages. Whereas functional programming languages such as Haskell are well suited for developing tools to analyse, verify and transform programs, typical Haskell compilers do not provide sophisticated capabilities such as support for XML-Web services, assisted GUI development and so on, which occur frequently in most .NET development frameworks. This article reports on our experience in the integration of Haskell applications into the .NET framework.

The CLR (Common Language Runtime) is the heart of the .NET framework; together with the FCL (Framework Class Library), a consistent, object-oriented set of classes, they both provide a platform supporting multiple programming languages. The idea is that the developers should use the language they are most familiar with or the one best suited to the task at hand. The (.NET) language compilers should then convert source code into Intermediate Language (IL) code, which is used by CLR to convert language-independent IL code into the machine code of the device on which it is intended to run. In this way, different languages can be combined to write a single application.

While this description of the basics of the CLR may suggest that making a programming language into .NET is a simple task, this is, of course, not the case. There exist many different kinds of programming languages inspired by quite different programming paradigms, which exploit and exhibit quite different features. Functional programming languages, for instance, deal with higher-order functions, fully parametric polymorphism in data structures and function definitions and lazy evaluation of expressions. These features are typically missing in imperative (object-oriented) languages like C, C++, or C#, which are the ‘natural’ .NET evolution of previous forms. Fortunately, however, a number of different attempts have successfully fitted functional languages into the .NET platform, namely, SML.NET and Mondrian.

Haskell is a general purpose, purely functional programming language with an impressive number of libraries, extensions, compilers and projects for current and future developments (see the Haskell Communities URL below). During recent years, we have used Haskell to develop several tools that implement the compilers, termination analyses, and declarative debugging techniques available for programming and specification languages like Maude, CafeOBJ, OBJ and XHTML. All these tools target different but complementary purposes, and programming languages would greatly benefit from an integrated framework in which full interaction and combination is possible. As remarked above, .NET has excellent capabilities in this regard. Although the Haskell community is also interested in this, the functional language Haskell has yet not been ported into .NET.

Interoperability (ie allowing a program on one system to access programs and data on another system) is a general problem in Software Engineering and a number of solutions (namely middleware systems) have been devised. One of the main goals of the new research project SELF (Software Engineering and Lightweight Formalisms) is the exploration of middleware translators and interfaces, from the existing programming languages and development frameworks (like C# and .NET) to the formalisms or languages which underlie the program analysis tools (eg Haskell). Fortunately, well-known middleware techniques like Microsoft's COM (Component Object Model) have received some support from both Haskell and .NET communities. The figure summarizes how it works: our starting point is HaskellDirect (HDirect) which provides a Foreign Function Interface (FFI) for Haskell based in the standard IDL (Interface Definition Language). This permits the specification of a programming interface which does not depend on any concrete programming language. With HDirect it is possible to build COM components in Haskell that are then imported and used by a .NET application by means of a COM DLL (Dynamically Linked Library). In the figure, an input IDL file (Example.idl) containing the appropriate interface information given by the programmer is processed by HDirect. A file (ExampleProxy.hs) is automatically generated. Together with Example.hs (which contains the Haskell code of the COM component), the HDirect library Com.lhs and a Run Time Support (RTS) C module, the COM module is obtained after processing everything with the Glasgow Haskell Compiler (GHC). The output is a COM DLL that requires the use of a Runtime Callable Wrapper (RCW) to become part of a .NET application. The RCW is part of the .NET SDK. We are now able to see the application written in native (unmanaged, in the .NET terminology) Haskell code as managed code, which can be executed by the CLR together with the other components of the .NET Framework.

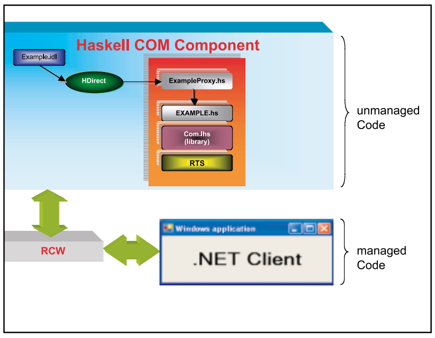

We are using these ideas to make our software tools available within the .NET framework. Moreover, the COM DLLs that we are producing are a good basis for easily extending the original capabilities of our Haskell applications.

Links:
Haskell: http://www.haskell.org

Haskell Communities:
http://www.haskell.org/communities

SELF: http://self.lcc.uma.es/

Please contact:
Salvador Lucas, Universidad Politécnica de Valencia / SpaRCIM
E-mail: slucas@dsic.upv.es
http://www.dsic.upv.es/~slucas

A Tookit for Analysis, Testing and Restructuring of Web Applications

by Filippo Ricca and Paolo Tonella

We have defined and implemented a toolkit to support the quality of Web applications. The validity of this toolkit has been assessed by extensive empirical work.

A number of studies show that the quality of Web applications is often poor or unsatisfactory; end users often make the same claim. Quality demands on these software systems are thus increasing.

Since 1999, in the SSI division at ITC-irst (a public research centre of the Autonomous Province of Trento, Italy), we have been investigating, defining and applying a range of techniques for the analysis, testing and restructuring of Web applications. Our purpose is to assess the quality of Web applications during their development and evolution. The results of the analyses can be used to check whether a Web application satisfies suitable constraints, and to detect possible anomalies. The goal of testing is to improve the quality of the Web application and expose possible failures, that is, deviations of the application from the intended behaviour. The restructuring activity aims to improve certain quality factors of the Web application without changing its external behavior.

The approach we have adopted is based on reverse engineering. Unlike the more traditional forward engineering approach, we move from the assumption that a Web application already exists. Our starting point is therefore the actual implementation of the Web application (Web pages, server programs, forms, frames etc), and our techniques work on the abstract models that we derive from the implementation. This approach is also based on the observation that several well-established methods for the analysis, testing and restructuring of traditional software systems already exist, and these can be adapted to Web applications (see Figure 1).

Our work is divided into three parts:

- the development of models and views of Web applications that can accommodate static as well as dynamically generated entities
- the application of a variety of analyses, restructuring and testing techniques to these models
- empirical studies on real-world applications.

As a necessary component of our experimental work, two research prototype tools, ReWeb and TestWeb (see Figure 2), consisting of a number of software modules, have been built to implement a range of analyses/restructuring techniques and to support application developers in testing activities.

The main contributions of our work extend in several directions:

- modelling and model extraction of dynamic Web applications

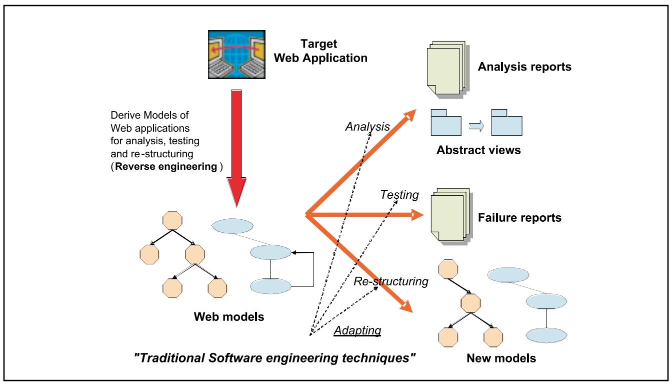

Figure 1: Reverse engineering through the application of software engineering techniques.

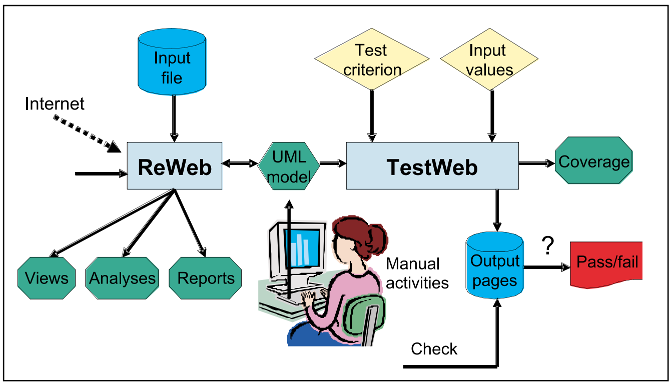

Figure 2: ReWeb and TestWeb.

- structural and evolution analysis
- Web application slicing
- structural and statistical testing
- multilingual Web site restructuring
- migration of static Web sites to dynamic Web applications
- Web application understanding: producing abstract views.

Many interesting developments and extensions of the proposed techniques and tools are possible. We are currently working on Web application testing; in particular we are collecting data in order to define a fault model for Web applications. Our goal is to verify the fault detection abilities of the various testing techniques.

At present, we are completely rewriting ReWeb and TestWeb using HttpUnit (a Java framework which provides a library for implementing a Spider and automated test scripts for Web applications). This restructuring was necessary for many reasons. The original Spider did not use a fully fledged parser, and support for cookies and Javascript was limited with the consequence that the models extracted were in some cases partial. Moreover, the testing process was overly complex and the test cases were not expressed in the Java language.

The new versions of ReWeb and TestWeb will be available for research purposes soon. Further information can be found on our Web site.

Links:
http://star.itc.it/

Please contact:

Filippo Ricca, ITC-irst, Trento, Italy
Tel: +39 0461 314 522
E-mail: ricca@itc.it

Paolo Tonella, ITC-irst, Trento, Italy
Tel: +39 0461 314 524
E-mail: tonella@itc.it

Telephony over IP: Experience and Challenges

by Laurent Burgy, Charles Consel, Fabien Latry, Nicolas Palix and Laurent Réveillère

Telephony over IP (ToIP) makes it possible to program telephony platforms that have long been kept closed and proprietary. As new services are developed that enrich telephony with complex features such as database access, Web service invocation, and agenda look-up, the telephony platform is increasingly exposed to software bugs. Because telephony is heavily relied on, the success of its evolution to IP critically depends on the reliability of services.

SIP (Session Initiation Protocol) — a protocol for Telephony over IP — is at the forefront of the rapid emergence of ToIP. It was standardized by the IETF and adopted by the ITU; it is used as the underlying signalling protocol for 3G/UMTS mobile systems and various communication tools running under Unix and Windows. Besides telephony, SIP supports other forms of communication: Internet conferencing, presence, event notification and instant messaging. SIP implementations are available for a variety of devices including desktop computers, wired and wireless phones and PDAs.

A SIP platform is based on a client-server model and consists of a signalling server performing telephony-related operations. In addition, service logic is executed by an application server and can access resources from the computer network such as Web resources, databases and agendas. This evolution brings a host of new functionalities to the domain of telephony, which in turn enables telephony to be customized with respect to the preferences, trends and expectations of ever more demanding users.

Challenges for Service Development

The open nature of an SIP platform makes it possible for a wide range of developers to write telephony services. Yet, telephony is as heavily relied on as water and electricity. As telephony platforms become programmable by any developer, they become less reliable. A lack of reliability may disrupt or even crash a service or the entire platform. Furthermore, to create a service, a developer should have extensive expertise in the telephony domain, SIP and related protocols, distributed programming, networking, and SIP APIs.

The TelIP Project

The Phoenix research group at INRIA Futurs/LaBRI/ENSEIRB in Bordeaux is developing a software engineering approach to producing programs whose reliability is guaranteed with respect to critical properties of a given domain. This approach relies on the development of Domain-Specific Languages (DSL),which ease program development without sacrificing safety. Phoenix has been studying the DSL approach in the context of communication services.

We have deployed an SIP platform named TelIP at ENSEIRB, a graduate engineering school of information and communication technologies to which the group is affiliated. TelIP interfaces with the PABX (Private Automatic Branch eXchange) of ENSEIRB, allowing IP terminals to call anyone either inside or outside the school. Because ENSEIRB provides Wi-Fi, terminals can be wired or wireless, enabling longer-range mobility than conventional wireless telephony systems. TelIP already has two dozen users at ENSEIRB and the number is growing fast.

The group has developed an application server allowing telephony services to be deployed on TelIP. The server functionalities include the binding of a user to one or more services. This binding allows an incoming or outgoing call to be handled by the service logic of the corresponding user, enabling a wide variety of customizations.

To ease the development of telephony services, we have designed and implemented a DSL known as SPL (Session Processing Language). This language offers domain-specific constructs and extensions that abstract over the intricacies of the underlying technologies. By design, SPL guarantees critical properties, far beyond the reach of general-purpose languages such as Java or C++. Examples of errors detected in services include call loss, incorrect state transitions, and unbounded resource usage.

To enable non-programmers to define services, a graphical, simplified version of SPL has been developed. This version of SPL offers intuitive visual constructs and menus that permit users to quickly develop a wide variety of services ranging from simple redirection to agenda-dependent call handling.

To validate SPL, rich telephony services are being developed. In particular, one service is dedicated to handling calls to the telecommunications department of ENSEIRB. If the department secretary is already on the phone, this service queues incoming calls. An estimated waiting time is periodically computed and played to the waiting callers. Furthermore, the availability of the secretary is automatically updated, thanks to the user status managed by SIP. If the secretary on duty is unavailable, calls are redirected to an alternative secretary. In the context of a conventional telephony platform, such a service would typically need to be written by a third-party application developer who has been certified by the PABX manufacturer. As a result of this development model, telephony service creation has been anecdotal and the cost has been prohibitive.

In the future, the goal is to significantly increase the base of TelIP users. ENSEIRB consists of more than 700 students and 140 faculty and staff members. Enabling users with different skills and interests to succeed in developing the services they envision is a real challenge. Doing so without compromising the robustness of the platform is a key step toward validating the platform and the language.

Links:
http://phoenix.labri.fr
http://phoenix.labri.fr/documentation/sip

Please contact:
Charles Consel
INRIA Futurs/LaBRI/ENSEIRB, France
Tel: +33 5 40 0035 19
E-mail: Charles.Consel@inria.fr

Adaptive Scene and Traffic Control in Networked Multimedia Systems

by Leif Arne Rønningen

The Distributed Multimedia Plays (DMP) Systems Architecture provides three-dimensional multiview video and sound collaboration between performers over packet networks. To guarantee an end-to-end time delay under twenty milliseconds, and to obtain high network resource utilization, the perceived quality of audio-visual content is allowed to vary with the traffic in the network. Typical applications are in the next generation of televisions (Multimedia Home Platform (MHP) extended with DMP), games, education, and virtual meetings that have a realistic feel. The architecture is also suitable for use in creating virtual collaborations for jazz sessions, music lessons and distributed opera.

Research at the Norwegian University of Science and Technology is building on the concept of DMP, which was first proposed in 1996 as an extension to MHP. A collaborative effort is underway between researchers in the departments of Psychology, Telematics and Art and Media Studies, with research looking at the design and testing of visual sequences that support a comparison between realistic meetings and virtual collaborative meetings using three-dimensional, high-

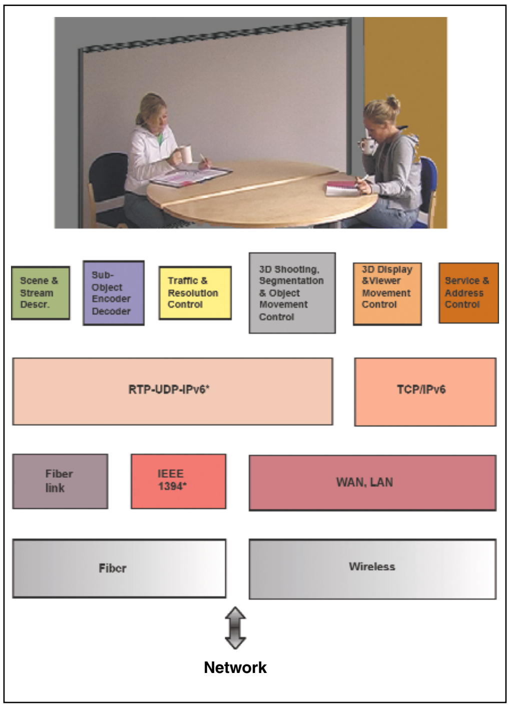
Scene & Stream Descr.
Sub-Object Encoder Decoder
Traffic & Resolution Control
3D Shooting, Segmentation & Object Movement Control
3D Display &Viewer Movement Control
Service & Address Control
RTP-UDP-IPv6*
TCP/IPv6
Fiber link
IEEE 1394*
WAN, LAN
Fiber
Wireless
Network

Figure 1: The DMP Terminal System.

resolution display.. The temporal and spatial resolution will be varied for individual objects in scenes. Eye movements and object focus are tracked when subjects watch selected video clips. Practical applications for this developing technology could be in use before 2015.

A DMP terminal system is shown in Figure 1. The two upper layers must always be present, while the two lower layers represent wireless and fibre alternatives. Network nodes include resolution and traffic control in addition to normal routing functions.

Scene Characteristics

To give the virtual collaboration a natural feel, the size, form and position of objects should be near natural size, the displayed picture shall be flicker-free, and individual pixels should not be visible at a viewing distance of fifty centimetres. The normal comfortable distance between people generally varies between fifty centimeters and several metres, but may also be several hundred metres. The natural viewing area of human eyes is used as the 'frame' of the scene.

Sub-objects and 2D Interlace

In the RL algorithm - an integrated resolution and traffic control algorithm - each object is divided into two, four or more sub-objects, which are sent in separate streams. This is called the 2D interlace.

Video Segmentation and Multiview Shooting

Object recognition and tracking is carried out by analysing video sequences, shot by multiview cameras. Some of the most important factors for DMP are face and eye recognition and tracking. Note that the motion estimation of object-oriented scene objects is moved away from the coder (compressor) to the object- and eye-tracking systems.

Compression and Coding

Compression algorithms can make use of the Discrete Cosine Transform together with Huffman coding, or wavelet transforms that represent data in the time-frequency space. The Wavelet transform makes level scalability easy.

The RL Resolution and Traffic Control Algorithm

A detailed description of the RL-C algorithm can be found in the reference. Simulations show that the end-to-end delay can be guaranteed. However, formal tests must still be conducted to show when and to what extent the time-varying resolution of audio-visual content is acceptable.

Service and Scene Setup, and Address Allocation

SMIL (Synchronized Multimedia Integration Language) is used for scene composition. The SIP/TCP/IPv6 protocols are used for set-up, and the SIP URL identifies services, scenes and subscenes. Parameters in the combined RTP/UDP/IPv6 protocol header identify the same for content transfer. IPv6 addresses will be allocated in a private way, and the SIP and RTP protocols are adapted for this application.

Variable-Resolution Display

A viewer focuses on one scene object at a time, and most viewers focus on the same object at any given time. An eye-tracking system identifies these objects, and orders the encoder on the shooting side to shoot and represent these objects with high resolution.

A Virtual Dinner Scenario

Researcher A in Trondheim enters his dining room, sits on his sofa and requests a Virtual Dinner with researcher B in Padova, also sitting on his sofa. This interaction generates different levels of traffic from A sent to B. The system 'finds' two faces and a plate with food for researcher B. A and B talk for about 30 seconds (7 Gbps), then researcher A arises (0.5-1 sec), goes out for a plate of food (5-10 sec), is absent (1-2 min), returns (5-10 sec) and sits down (0.5-2 sec). This increases traffic to nearly 60 Gbps, which then drops to the background of 127 Mbps. The system tracks the plate and the food. They start eating and talking, during which their faces, arms, hands, plates and food dominate the data rate, about 8 Gbps.

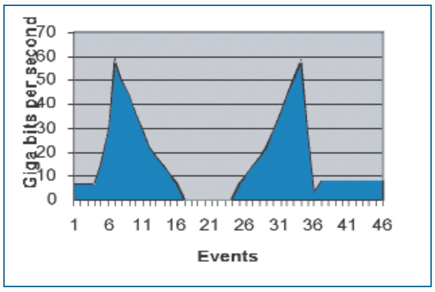

Figure 2: Virtual Dinner, traffic generated by researcher A (time-scale not correct)

Researcher B stands up and walks across the room. After a few minutes they need to talk to researcher C in Poznan, and establish a three party DMP. After eating, B leaves the room. A asks the system to disconnect B. Meanwhile, researcher C has to leave home and go to his office, but he wants to continue the session with A while travelling to the office.

PhD funding has been provided from the Research Council of Norway (NFR).

Link:
http://www.item.ntnu.no/~leifarne/

Please contact:
Leif Arne Rønningen
Norwegian University of Science and Technology (NTNU), Norway
Tel: +47 9003 0473
E-mail: leifarne@item.ntnu.no

INEX: Evaluating XML Retrieval Effectiveness

by Mounia Lalmas

The Initiative for the Evaluation of XML Retrieval (INEX) is an international campaign involving more than fifty organizations worldwide. It provides a means of evaluating retrieval systems that provide access to XML content.

Documents today contain a mixture of textual, multimedia, and metadata information. One way to format this mixed content is according to the eXtensible Mark-up Language (XML). In contrast to HTML, which is mainly layout-oriented, XML follows the fundamental concept of separating the logical structure of a document from its layout. XML thus offers the opportunity to exploit the logical structure of documents to allow more precise searching. Providing effective access to XML-based content has become a key research issue.

Effective access to XML repositories is the core of XML retrieval research. XML retrieval systems aim to exploit the logical structure of documents to retrieve, in response to a user's query, document components (ie XML elements) rather than whole documents. Implementing this more focused form of retrieval means that an XML retrieval system must not only find relevant information, but also determine the appropriate level of element granularity to return to the user. Evaluating the effectiveness of these systems therefore requires test-beds utilizing criteria that take into account the imposed structural aspects.

In 2002, INEX started to address these issues. An infrastructure was established, and a large XML test collection and appropriate scoring methods were provided for the evaluation of content-oriented XML retrieval systems.

Evaluation is carried out using test collections assembled specifically for evaluating particular retrieval tasks. A test collection consists of a document collection, a set of user requests (ie topics) and relevance assessments. The characteristics of traditional test collections have been adjusted to appropriately evaluate XML retrieval effectiveness: the document collection comprises documents marked up in XML, the topics specify requests relating to both content and structure, and the relevance assessments are made at element level. In addition, relevance is measured such that it appropriately quantifies the system's ability to return the correct granularity of XML elements.

Ad Hoc Retrieval

The main retrieval task to be performed in INEX is ad hoc retrieval. This can be described as a simulation of how a library might be used, and involves the searching of a static set of documents using a set of topics. In INEX, the library consists of XML documents, the queries may contain both content and structural conditions, and in response to a query, arbitrary XML elements may be retrieved.

In 2005, we identified two ad hoc retrieval sub-tasks that depend on how structural constraints are expressed. In the content-only sub-task, it is left to the retrieval system to identify the most appropriate XML elements to return to the user. Three different strategies have been defined, depending on the preferred output format of an XML retrieval system. In a focused strategy, we assume that a user would prefer the single element most relevant to the topic; in a thorough strategy, we assume that a user would prefer all highly relevant elements; and in a fetch and browse strategy, we assume that a user is interested in highly relevant elements contained within highly relevant documents. An extension of this sub-task is when a user adds structural constraints to the topic to narrow down the number of returned elements.

In the content-and-structure sub-task, the structural constraints are explicitly stated in the topics and can refer to where to look for the relevant elements (ie support elements) and what types of elements to retrieve (ie target elements). Structural constraints can be interpreted as either strict or vague, and these interpretations can be applied to both support and target elements, giving a total of four strategies for this sub-task.

Tracks

INEX has separate tracks looking at specific issues in XML retrieval. The interactive track aims to investigate the behaviour of users when interacting with XML documents, and to develop effective user-based approaches to XML retrieval. The heterogeneous track is concerned with cases where an XML collection comprises documents from different sources and based on multiple document structures. The multimedia track aims to provide an evaluation platform for XML retrieval systems that include text, images, speech, and video. The document-mining track aims at defining techniques for effectively mining information from XML documents, with a focus on classification and clustering.

INEX is led by Norbert Fuhr from the University of Duisburg-Essen, Germany, and Mounia Lalmas from Queen Mary University of London, United Kingdom, and is partly funded by the DELOS Network of Excellence on Digital Libraries administrated by ERCIM.

Links:
INEX: http://inex.is.informatik.uni-duisburg.de/
DELOS: http://www.delos.info/

Please contact:
Mounia Lalmas
Queen Mary University of London, UK
E-mail: mounia@dcs.qmul.ac.uk

Providing Access to Biodiversity Data in the National Biodiversity Network

by Charles Hussey and Steve Wilkinson

The National Biodiversity Network (NBN) is a partnership working to build the UK's first network for sharing biodiversity information. A number of services and tools are being developed to mobilise the large quantity of data present in the UK.

A report by the Co-ordinating Commission for Biological Recording, in 1995, recognised that there were at least 60,000 active biological recorders in the UK, 2,000 organisations were involved and 60 million species records had been accumulated. This report was instrumental in the formation of the NBN in 2000. The NBN includes the UK Countryside agencies, the Wildlife Trusts, Environment Agency, Natural History Museum (NHM), Joint Nature Conservation Committee (JNCC), Biological Records Centre (BRC) and other partners. The NBN currently provides free access online to over 18 million species records and these are also contributed to the Global Biodiversity Information Facility (GBIF).

An important part of the work of the NBN Trust is ensuring that access to data is free at point of use, whilst letting data providers exercise control over the level of detail. The NBN Trust has built up considerable experience in the field of Intellectual Property and database rights. They have drafted template agreements for data providers, data collators and end users, as well as putting in place the technical web components to manage the process. A data exchange standard, incorporating an XML Data Type Definition, has been established. The NBN is also responsible for the GBIF UK National Node.

Delivering Data

The NBN's main portal to data is through the NBN Gateway, which was launched as an operational system in June 2004. The construction and management of the Gateway is a joint venture by JNCC and BRC. Special mapping displays have been developed for this project. Polygon information allows generation of species lists within defined boundaries, such as protected sites. Extensive use is made of map layers, including Ordnance Survey maps and land cover. 136 different datasets containing over 18.6 million records have been integrated to allow searching by species, designated site or by 10 km square.

Gathering Data

A software application for use in biological recording is managed by JNCC. The current version is known as Recorder 2002. This is the latest in a pedigree that stretches back to the days of DOS, when the software was written in Advanced Revelation. The software has been developed with public funding and is available at very low cost through a network of resellers. The software is structured around surveys, sites and site visits and supports 'virtual recording cards' which mimic manual systems familiar to the recording community in the UK. Features include mapping using polygons and Ordnance Survey map tiles, bibliographic references and image handling. The application has been designed to accommodate additional modules. Recently the Musée national d'histoire naturelle, Luxembourg has extended Recorder with a comprehensive Collections Management module and a multi-lingual thesaurus.

Taxonomy is Complex

To properly account for what is found in a country requires a maintained taxonomy that tracks changes in nomenclature and classification. Ongoing taxonomic research may lead to splitting or lumping of species. Species may be reassigned to different genera and revisions made to higher classifications. In particular, names written into legislation may quickly become out of date and need be mapped to the current name. Within the NBN, this work is undertaken by the Species Dictionary project, which provides the taxonomy underpinning both the Gateway and Recorder. The Dictionary is managed by the NHM and has the ultimate aim of being the master list of what is found in the UK. It currently provides complete coverage for 70% of the 109 taxonomic groups recorded from the UK. It contains over 410,000 records from over 250 different checklists representing 190,000 separate variants of names. A Name Server is being developed that provides links between names and this will, in due course, be accessible through a web service to external applications.

Many of the current online biodiversity resources presume some knowledge of scientific names. The NHM received a grant from the New Opportunities Fund to create a website for use by the general public that would allow them to search using familiar common names - in particular, informal names for higher groups (such as "ants" or "conifers"). The site provides a finding aid for resources and images relating to UK wildlife. Common names are mapped to equivalent scientific names and presented in a display that encourages users to explore relationships between organisms. The default display is based on a classification showing broader and narrower terms but users can switch to a dynamic 'spider-diagram' display developed using software from http://www.touchgraph.com.

Links:
NBN: http://www.nbn.org.uk
NBN Gateway: http://www.searchnbn.net
NBN Species Dictionary:
http://www.nhm.ac.uk/nbn/
Nature Navigator:
http://www.nhm.ac.uk/naturenavigator

Please contact:
Charles Hussey, Species Dictionary Project Manager, Natural History Museum, London, UK
Tel: +44 (0) 207 942 5213
E-mail: c.hussey@nhm.ac.uk

Steve Wilkinson, Manager for NBN Gateway and Recorder projects,JNCC, Peterborough, UK
Tel: +44 (0) 1733 866 865

sponsored by ERCIM

ECDL 2005 –European Conference on Research and Advanced Technology for Digital Libraries

by Andreas Rauber

Some 400 researchers and practitioners met in Vienna, Austria for the 9th European Conference on Research and Advanced Technology for Digital Libraries (ECDL 2005) 18-23 September 2005.

The conference offered a fascinating program, starting with a set of six tutorials on Sunday, covering general concepts and formal frameworks for digital libraries, context-enhanced digital library services, thesauri and ontologies in digital libraries, building digital library collections with Greenstone and DL interoperability standards. This was followed by a dense 2 1/2 day program of presentations.

The first keynote speaker was Neil Beagrie from the British Library and the Joint Information System Council, UK, who addressed the potential of personal digital libraries, considering specifically the massive amounts of data that are and will be collected by the population at large as part of their day-to-day life, and the resulting potential and consequences for personal information systems.

The second keynote by Erich Neuhold, former head of Fraunhofer-IPSI, now Vienna University of Technology, picked up this topic and pointed out future directions in DL development, focusing particularly on the promises and challenges of integrating and combining DL, peer-to-peer and GRID technology.

Costantino Thanos, DELOS Scientific Coordinator (left) is handing over the DELOS Research Exchange Award for the best paper presented by a young author, to Christos Tryfonopoulos for his paper “LibraRing: An Architecture for Distributed Digital Libraries Based on DHTs”.

The quality of the presentations, held in two parallel sessions, was excellent, with only 32% of all papers submitted to the conference having been accepted for presentation. In addition to the presentations ECDL featured two panels. Donatella Castelli discussed with a prominent group of panelists one of the key messages addressed in the second keynote, namely whether digital libraries on the GRID will turn out to be hell or heaven, whereas Tamara Sumner tried to answer the question whether e-science needs digital libraries. The whole range of current digital library research as well as project activities were presented in a poster session, featuring both research posters, application demos, as well as current project activities of the DELOS Network of Excellence on Digital Libraries. Furthermore, a large group of PhD students had the chance to present and discuss their work not only within the ECDL Doctoral Consortium, but also with the whole ECDL community as part of the poster session.

Following the main conference program, ECDL was rounded off with a set of well-attended workshops. CLEF, the cross-language evaluation forum, drew more than 100 participants during its two and a half days of presentations and evaluation sessions. IWAW, the ECDL Workshop on Web Archiving and Digital Preservation was extended to fill an intensive two-day program this year, with about 60 participants.

Both the NKOS Workshop on Mapping Knowledge organization Systems, as well as HDL, Healthcare in Digital Libraries, ran in their successful 1-day form, soliciting intensive discussion.

ECDL 2005 prooved to be a splendid opportunity for people to meet and discuss current core aspects of digital library research. We are looking forward to continue these discussions at the 10th aniversary of ECDL, which will be held September 17-24 2006 in Alicante, Spain.

ECDL participants enjoy the ECDL gala dinner.

Link:
http://www.ecdl2005.org
http://www.ecdl2006.org

Please contact:
Andreas Rauber, Vienna University of Technology, Austria
Tel: +43 1 58801 18826
E-mail: rauber@ifs.tuwien.ac.at

Cross-Language Evaluation Forum — CLEF 2005

by Carol Peters

The results of the sixth campaign of the Cross-Language Evaluation Forum were presented at a two-and-a-half day workshop held in Vienna, Austria, 21-23 September, immediately following the ninth European Conference on Digital Libraries. The workshop was attended by well over 100 researchers and system developers from academia and industry.

The main objectives of the Cross-Language Evaluation Forum (CLEF) are to stimulate the development of mono- and multilingual information retrieval systems for European languages and to contribute to the building of a research community in the multidisciplinary area of multilingual information access. These objectives are realised through the organisation of annual evaluation campaigns and workshops. Each campaign offers a series of evaluation tracks designed to test different aspects of mono- and cross-language system performance. The focus is diversified to include different kinds of text retrieval across languages and on different kinds of media (ie not just plain text but collections containing images and speech as well). In addition, attention is given to issues that regard system usability and user satisfaction with tasks to measure the effectiveness of interactive systems.

The scope of CLEF has gradually expanded over the years. In CLEF 2005 eight tracks were offered to evaluate the performance of systems for:

- mono-, bi- and multilingual document retrieval on news collections (Ad-hoc)
- mono- and cross-language structured scientific data (domain-specific)
- interactive cross-language retrieval (iCLEF)
- multiple language question answering (QA@CLEF)
- cross-language retrieval on image collections (ImageCLEF)
- cross-language speech retrieval (CL-SR)
- multilingual web retrieval (WebCLEF)
- cross-language geographic retrieval (GeoCLEF).

In order to cover all these activities, the CLEF test collection has been expanded considerably: the main multilingual comparable corpus now contains over 2 million news documents in twelve European languages- new entries this year were Hungarian and Bulgarian. Sub-collections from this corpus were used by the Ad-Hoc, QA, iCLEF and GeoCLEF tracks. The collection used to test domain-specific system performance consists of the GIRT-4 database of English and German social science documents and the Russian Social Science Corpus. ImageCLEF used a number of different collections: an archive of historic photographs provided by St Andrews University, Scotland, and several sets of medical images with French, English and German case notes and annotations made available by University Hospitals, Geneva, and by Aachen University of Technology. The cross-language speech retrieval track (CL-SR) used speech transcriptions from the Malach collection of spontaneous conversational speech derived from the Shoah archives. Finally, WebCLEF used the EuroGOV corpus, a multilingual collection of about two million webpages crawled from European governmental sites.

Participation in this year's campaign was considerably up with respect to the previous year with 74 groups submitting results for one or more of the different tracks as opposed to 54 groups in CLEF 2005: 43 from Europe, 19 from North America, ten from Asia and one each from South America and Australia. The introduction of the Speech, Image and Question Answering tracks in previous years and of the GeoCLEF track this year means that the growing CLEF community is increasingly multidisciplinary, with expertise in diverse areas such as natural language processing, speech recognition and analysis, geographic information systems, image processing, medical informatics, etc..

The campaign culminated in the workshop held in Vienna, Austria, 21-23 September. In addition to presentations by participants in the campaign, Noriko Kando from the National Institute of Informatics, Tokyo, gave an invited talk on the activities of the NTCIR evaluation initiative for Asian languages. Breakout sessions gave participants a chance to discuss ideas and results in detail. The final session was dedicated to proposals for activities for CLEF 2006.

The presentations given at the CLEF Workshops and detailed reports on the experiments of CLEF 2005 and previous years can be found on the CLEF website at http://www.clef-campaign.org/ The preliminary agenda for CLEF 2006 will be available from mid-November.

CLEF is an activity of the DELOS Network of Excellence for Digital Libraries.

Link:
http://www.clef-campaign.org

Please contact:
Carol Peters, ISTI-CNR, Italy
CLEF Coordinator
Tel: +39 050 3152897
E-mail: carol.peters isti.cnr.it

CLEF steering committee.

HCI International 2005

by Constantine Stephanidis

The 11th International Conference on Human-Computer Interaction (HCI International 2005) was held in Las Vegas, Nevada, USA, 22-27 July 2005, jointly with the Symposium on Human Interface (Japan) 2005, the 6th International Conference on Engineering Psychology and Cognitive Ergonomics, the 3rd International Conference on Universal Access in Human-Computer Interaction, the 1st International Conference on Virtual Reality, the 1st International Conference on Usability and Internationalization, the 1st International Conference on Online Communities and Social Computing, and the 1st International Conference on Augmented Cognition, under the auspices of nine distinguished international boards of more than 200 members from 30 countries.

The ERCIM Working Group (WG) 'User Interfaces for All', through its dedicated and prolific work, actively supports since 2001 the Universal Access in Human-Computer Interaction Conference, held this year for the third time. The UAHCI Conference alternates with the ERCIM biannual workshop, and many members of the WG have participated in the Conference in Las Vegas.

HCII 2005 was one of the biggest ever organised in the fields related to Human Computer Interaction and Information Society Technologies, and attracted more than 2100 participants from 60 countries, representing the research and academic communities, as well as the industry.

The Conference Programme was organised into nine thematic areas, namely Ergonomics and Health Aspects of Work with Computers, Human Interface and the Management of Information, Human-Computer Interaction, Engineering Psychology and Cognitive Ergonomics, Universal Access in Human-Computer Interaction, Virtual Reality, Usability and Internationalization, Online Communities and Social Computing, and Augmented Cognition. The programme featured an opening session, 251 parallel paper sessions, 22 tutorials, 271 poster presentations, eleven demonstrations and five Special Interest Groups. Dr. Gerald M. Edelman, Nobel laureate (The Neurosciences Institute, San Diego, California, USA) was the keynote speaker. His keynote address was entitled 'From Brain Dynamics to Consciousness: A Prelude to the Future of Brain-Based Devices'. During the opening session, a movie premier was also shown, entitled 'The Future of Augmented Cognition - An Alexander Singer Film', introduced by Dr. Dylan Schmorrow, DARPA, USA.

The Proceedings have been published on CD-ROM by Mira Digital Publishing (ISBN 0-8058-5807-5), and are structured in eleven volumes, entitled: 'Engineering Psychology, Health and Computer System Design', 'The management of Information: E-Business, the Web, and Mobile Computing', 'Human-Computer Interfaces: Concepts, New Ideas, Better Usability, and Applications', 'Theories Models and Processes in HCI', 'Emergent Application Domains in HCI', 'Human Factors Issues in Human-Computer Interaction', 'Universal Access in HCI: Exploring New Interaction Environments', 'Universal Access in HCI: Exploring New Dimensions of Diversity', 'Advances in Virtual Environments Technology: Musings on Design, Evaluation, & Applications', 'Internationalization, Online Communities and Social Computing: Design and Evaluation', and 'Foundations of Augmented Cognition', and a 'Posters' section.

The exhibition of the HCII 2005 Conference.

The 12th International Conference on Human-Computer Interaction (HCI International 2007) will be held in Beijing, PR China, 22-27 July 2007, jointly with the Symposium on Human Interface (Japan) 2007, the 7th International Conference on Engineering Psychology and Cognitive Ergonomics, the 4th International Conference on Universal Access in Human-Computer Interaction, the 2nd International Conference on Virtual Reality, the 2nd International Conference on Usability and Internationalization, the 2nd International Conference on Online Communities and Social Computing, the 2nd International Conference on Augmented Cognition, and the 1st International Conference on Human Digital Modelling. The call for papers is available at http://www.hcii2007.org.

Link:
http://www.hci-international.org/

Please contact:
Constantine Stephanidis, ICS-FORTH, Greece
Tel: +30 2810 391741
E-mail: cs@ics.forth.gr

sponsored by ERCIM

CASSIS: Construction and Analysis of Safe, Secure and Interoperable Smart Devices

by Gilles Barthe, Benjamin Grégoire, Marieke Huismanand Jean-Louis Lanet

Following last year edition in Marseilles, the second edition of the International Workshop on the Construction and Analysis of Safe, Secure and Interoperable Smart Devices was held in Nice, 8-11 March 2005.

The aim of the CASSIS workshop is to bring together experts from the smart devices industry and academic researchers, with a view to stimulate research on formal methods and security, and to encourage the smart device industry to adopt innovative solutions drawn from academic research. In order to address the different issues raised by the evolution of smart devices, the workshop consisted of seven thematic sessions:

- *Session 1: Research trends in smart devices.* The session was organized by Jean-Jacques Vandewalle, from Gemplus. The session was dedicated to providing perspectives on possible evolutions of smart devices. The keynote speaker was Gilles Privat, from France Telecom R&D.
- *Session 2: Web services.* The session was organized by Cédric Fournet and Andy Gordon, from Microsoft Research Cambridge. This session focused on security issues for web services, including trust and identity management, and formal and automatic verification of web services deployments. The session was followed by a panel discussion on web services security, chaired by Andy Gordon. The keynote speaker was Cédric Fournet.
- *Session 3: Virtual machine technology.* This session covered new developments in Java technologies and in for developing generic, adaptable and maintainable platforms for smart devices. The keynote speaker was Sophia Drossopoulou, from Imperial College London.
- *Session 4: Security.* This session was devoted to security issues from a wider perspective, and addressed issues such as elecronic voting, Internet threat analysis, privacy and language-based security. The keynote speaker was Dan Wallach, from Rice University, Texas.
- *Session 5: Validation and formal methods.* This session, organized by Thomas Jensen, from IRISA Rennes, focused on verification techniques for Java-like applications, including run-time verification, program analyses, and interactive verification. The keynote speaker was Klaus Havelund, from Kestrel Technology at NASA Ames Research Center.
- *Session 6: Proof-Carrying Code.* The session was organized by Adriana Compagnoni. It was devoted to the presentation of Proof-Carrying Code architectures, and of their applications to advanced security policies about resouce control and information flow. The keynote speaker was George Necula, from the University of California at Berkeley.
- *Session 7: Embedded devices.* The final session was organized by Traian Muntean, from Marseilles University, and Jean-Louis Lanet, now at Gemplus. The session focused on technology issues that arise from the evolution of embedded devices into networked mobile devices. The keynote speaker was Rajesh Gupta, from the University of California at Irvine.

The workshop was attended by over 70 participants. The organizers would like to thank the session organizers, speakers and participants for contributing to make CASSIS 2005 a stimulating and enjoyable event. The organizers would also like to acknowledge financial support from ERCIM, Gemplus International S.A, and Oberthur Card Systems. A special thanks goes to the support teams at INRIA Sophia-Antipolis, and in particular to Nathalie Bellesso and Monique Simonetti for their help in organizational matters.

Link:
CASSIS'05: http://www-sop.inria.fr/everest/events/cassis05/index.html

Please contact:
Gilles Barthe, INRIA
Tel: +33 4 92 38 79 38
E-mail: Gilles.Barthe@sophia.inria.fr

PHISE'05 - 1st Workshop on Philosophical Foundations of Information Systems Engineering

by Cesar J. Acuña and Belen Vela

Esperanza Marcos from Rey Juan Carlos University, Spain, and Roel Wieringa from the University of Twente, The Netherlands, have successfully organized the 1st International Workshop of Philosophical Foundations on Information Systems Engineering (PHISE'05), held in conjunction with the 17th CAISE in Porto, Portugal.

Empirical and formal sciences are older than engineering sciences and, for this reason, they have their own branch of philosophy, the philosophy of science. Engineering sciences have not fully developed a philosophical branch working towards conceptual clarity. The PHISE workshop tries to cover different aspects related to philosophical foundations of the research in information system engineering and aims to promote an interdisciplinary forum as a result of the confluence of three disciplines: philosophy of science and technology, software engineering and information systems. From the 21 submissions, twelve papers were accepted as regular papers and six as short papers. The selected papers were structured in four sessions:

- foundations of information systems engineering,
- research methods in information systems engineering,
- models and ontologies in information systems engineering
- miscellany.

In addition Dr. Heinz K. Klein from the Department of Management Information Systems (MIS) of the Temple University of Philadelphia, USA, present the invited lecture: 'Methodological Differences between the Cultural and Natural Sciences'.

Link:
http://kybele.escet.urjc.es/PHISE05/

Please contact:
Esperanza Marcos, SpaRCIM, Spain
E-mail: esperanza.marcos@urjc.es

Proposed Retention of Electronic Communications Data in Europe

Current proposals for a legal requirement for the retention of electronic communications data, which would include details of mobile phone calls and internet traffic, will affect the privacy of every EU citizen. A recent draft framework proposing the retention of electronic data by communications service providers for up to three years was rejected by the European Parliament. A separate proposal by the European Commission for a new Directive on the retention of electronic communications data has now been considered by the European Data Protection Supervisor ('the EDPS'), Peter Hustinx, who delivered his opinion on the proposal on 26 September 2005.

The provisions of the proposed Directive were considered by the EDPS with regard to data protection legislation and the protection of privacy given by Article 8 of the European Convention for the Protection of Human Rights and Fundamental Freedoms (the ECHR). The EDPS has suggested amendments to the proposal that would provide a legal basis for the retention of data by service providers and access to the data by law enforcement agencies of the Member States.

The EDPS recognises that access to certain traffic and location data by law enforcement agencies can assist in the combat of terrorism and other serious crime. However, the necessity of the obligation to retain data must be demonstrable and proportionate. The EDPS is not convinced of the absolute necessity of the retention of traffic and location data for law enforcement purposes, or that it is in itself an effective response, unless additional safeguards apply including that:

- the data should be transmitted from the service providers to the authorities, rather than allowing them direct access
- access should be prohibited for data-mining activities or for routine 'fishing operations'
- access by the authorities should only be for the investigation of certain serious criminal offences, and should be subject to judicial control
- data protection principles must be expressly provided for, such as the exercise of rights by the data subject, the need for data quality and security, and limitations as to lawful purposes of retention of data.

The new Directive provides that the obligation to retain data should apply for a maximum period of 6 months in the case of Internet communications data, and for up to one year for the retention of other communications traffic data. The retention of all data for such periods will mean the creation of huge databases, and require technological measures to ensure that data is erased after the relevant retention period, the provision of a verifiable audit trail of actions, and an adequate and effective search engine to allow for targeted searching for specific data.

The costs of data retention and related technological measures such as effective search engines will be considerable. Providers must be offered compensation as an incentive to install the technological and procedural security measures to retain the data and control access. Member states will bear the cost of compensating the service providers. There are interesting cross-border jurisdictional issues, for example in cross-border phone calls, crossing a border during a phone call, and the use of a provider in another country than that where the individual resides.

The Council of Europe is scheduled to meet on 12th October to discuss the issue of data retention further. The full text of the Opinion of the European Data Protection Supervisor on the Commission's Proposal can be found at http://www.edps.eu.int/12_en_opinions.htm.

By Judy Beck, CCLRC, UK
Editor: Heather Weaver, CCLRC, UK

Heather Weaver, would like to make it clear that the article on Digital Rights Management (DRM), which appeared in the Euro Legal column of edition 62 of ERCIM News, was an extract from a paper written by Renato Iannella and the URL included in the article (http://www.dlib.org/dlib/june01/iannella/06ianella.html) will take you to the full paper.

As the debate on open access continues to gain pace, Renato has kindly agreed to contribute an article on his latest research into DRM for the January edition of ERCIM News which will look into Open Access issues in the European Scene.We look forward to that.

ERCIM NEWS

ERCIM News is the magazine of ERCIM. Published quarterly, the newsletter reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of 10,500 copies. The printed version of ERCIM News has a production cost of 8 Euro per copy. It is available free of charge for certain groups.

Advertising
For current advertising rates and conditions, see http://www.ercim.org/publication/ERCIM_News/ or contact office@ercim.org

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

ERCIM News online edition
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: Jérôme Chailloux, ERCIM Manager

Central Editor:
Peter Kunz, ERCIM office
peter.kunz@ercim.org

Local Editors:

AARIT:	n.a.
CCLRC:	Martin Prime M.J.Prime@rl.ac.uk
CRCIM:	Michal Haindl haindl@utia.cas.cz
CWI:	Annette Kik Annette.Kik@cwi.nl
CNR:	Carol Peters carol.peters@isti.cnr.it
FORTH:	Margherita Antona antona@ics.forth.gr
Fraunhofer ICT Group:	Michael Krapp michael.krapp@scai.fraunhofer.de
FNR:	Patrik Hitzelberger hitzelbe@crpgl.lu
FWO/FNRS:	Benoît Michel michel@tele.ucl.ac.be
INRIA:	Bernard Hidoine bernard.hidoine@inria.fr
Irish Universities Consortium:	Ray Walshe ray@computing.dcu.ie
NTNU:	Truls Gjestland truls.gjestland@ime.ntnu.no
SARIT:	Harry Rudin hrudin@smile.ch
SICS:	Kersti Hedman kersti@sics.se
SpaRCIM:	Salvador Lucas slucas@dsic.upv.es
SRCIM:	Gabriela Andrejkova andrejk@kosice.upjs.sk
SZTAKI:	Erzsébet Csuhaj-Varjú csuhaj@sztaki.hu
VTT:	Pia-Maria Linden-Linna pia-maria.linden-linna@vtt.fi
W3C:	Marie-Claire Forgue mcf@w3.org

Subscription
Subscribe to ERCIM News 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/

SICS — **Staffan Truvé** is the new Managing Director of SICS and new member of ERCIM's Board of Directors. Staffan Truvé has been appointed Managing Director of SICS as of 1st July 2005. The position is his third as Managing Director within the same group: Staffan Truvé already leads the Interactive Institute as well as the umbrella organization SITI, where SICS and Interactive Institute are included together with Viktoria and S:ta Anna Institutes. Consequently he will also join ERCIM's Board of Directors.

Staffan has 20 years of experience in research and high-tech startups. His research background is in compiler technology, computer architecture and computer vision. In 1994 he co-founded CR&T (Carlstedt Research and Technology), a research and development organization in the areas of computer science and computer engineering. He is also the co-founder of Spotfire, AppGate, Pilotfish, Gatespace Telematics and CADSIM. Staffan is the chairman of Gatespace Telematics and acts as a technical/business advisor to startups, venture capital funds, universities and government agencies. Staffan is 42 years old, and holds a PhD in Computer Science from Chalmers University and a BA in Business Administration from Göteborg University. He has been a visiting Fulbright scholar at MIT and maintains a strong academic network in Sweden and in the US.

The appointment of Staffan Truvé as Managing Director takes SICS one step closer to leading the co-ordination of the four IT research institutes in Sweden.

FhG — **The Drug Discovery application**, where scientists carry out in silico docking, was initiated and implemented by Fraunhofer Institute for Algorithms and Scientific Computing (SCAI) in Germany and the Corpuscular Physics Laboratory (IN2P3) of Clermont-Ferrand in France. It has been running on the EGEE (Enabling Grids for E-Science) production service since December 2004. In silico docking enables researchers to compute the probability that potential drugs will dock with a target protein and is one of the most promising approaches to speed-up and reduce the cost to develop new drugs to treat diseases such as malaria.

The application is part of the WISDOM data challenge which demonstrated how grid computing can help drug discovery research by speeding up the whole process. The sheer amount of data generated indicates the potential benefits of grid computing for drug discovery. The next step is to analyse the huge quantity of results provided by this experiment which will require significant data mining effort.

http://public.eu-egee.org/news/fullstory.php?news_id=53

SpaRCIM — **Spanish Awards in Informatics** — during the First Spanish Conference on Informatics (CEDI 2005) in Granada, 13-16 September 2005, two new national awards have been granted to prominent Spanish researchers. The 'José García Santesmases Award' for the most outstanding professional career was presented to Prof. Antonio Vaquero from the Complutense University of Madrid (UCM). Antonio Vaquero is one of the founders of Informatics in Spain, both as an academic discipline and as a very active research field. The 'Aritmel Award' for the researcher developing the most significant scientific contributions to informatics engineering was given exaequo to Manuel Hermenegildo, from the Technical University of Madrid (UPM) and Mateo Valero, from the Technical University of Catalonia (UPC). Mateo Valero has strongly contributed to the development of high performance computing. Manuel Hermenegildo is well-known for his many scientific contributions in the areas of software engineering and programming languages. Finally, the 'Mare Nostrum Award' was given to the companies Telvent and iSOCO, and the Ramón Llull Award gone to the Junta de Extremadura.

http://cedi2005.ugr.es/premios.shtml

INRIA — **Call for proposals for 2006-2007 Cooperative Research Initiatives**. INRIA launched a new call for proposals targeted at the scientific community in France and in Europe. This program aims to foster the rapid emergence of certain topics, encourage initiatives by young research scientists and support synergies between teams having different and complementary skills. These initiatives also are a means for INRIA to open itself up to new research partnerships. External collaborations are largely open in to INRIA's priviledged European partners such as the members of ERCIM and transfrontier projects are encouraged.

The allocated resources support research activities for two years. The present call for proposals concerns a sum that is not entirely defined yet but that should be close to 500 KEuros per year over two years. Around ten proposals could be selected after this call.

Cooperative research initiatives are intended to support already begun scientific activities that are not connected to a short term transfer project. It is not necessary to set up industrial partnerships or development initiatives in this context. This however could happen at a later stage. The potential applications remains an important element in the assessment of proposals.

Concerning the definition of themes for the initiatives, emphasis is placed on the priority challenges set by the INRIA's scientific policy. All proposals will however be examined according to their merits. Cooperative research initiatives must make it possible to launch research initiatives that could not be undertaken otherwise, or with difficulty. Proposals that include more than one geographical site have priority once they are scientifically justified.

The application deadline is 10 November 2005.

http://www.inria.fr/recherche/arc/arc2006.en.html

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

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

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

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

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

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

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

Fonds National de la Recherche
6, rue Antoine de Saint-Exupéry, B.P. 1777
L-1017 Luxembourg-Kirchberg
Tel: +352 26 19 25-1, Fax +352 26 1925 35
http:// www.fnr.lu

FWO
Egmontstraat 5
B-1000 Brussels, Belgium
Tel: +32 2 512.9110
http://www.fwo.be/

FNRS
rue d'Egmont 5
B-1000 Brussels, Belgium
Tel: +32 2 504 92 11
http://www.fnrs.be/

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

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

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

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

Spanish Research Consortium for Informatics
and Mathematics c/o Esperanza Marcos, Rey Juan Carlos
University, C/ Tulipan s/n, 28933-Móstoles, Madrid, Spain,
Tel: +34 91 664 74 91, Fax: 34 91 664 74 90
http://www.sparcim.org

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

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

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

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

Irish Universities Consortium
c/o School of Computing, Dublin City University
Glasnevin, Dublin 9, Ireland
Tel: +3531 7005636, Fax: +3531 7005442
http://ercim.computing.dcu.ie/

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

✂- -

Order Form

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

Send, fax or email this form to:
ERCIM NEWS
2004 route des Lucioles
BP 93
F-06902 Sophia Antipolis Cedex
Fax: +33 4 9238 5011
E-mail: office@ercim.org

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

I wish to subscribe to the

❑ ***printed edtion*** ❑ ***online edition (email required)***

Name: ....

Organisation/Company: ....

Address: ....

....

Post Code: ....

City: ....

Country ....

E-mail: ....

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

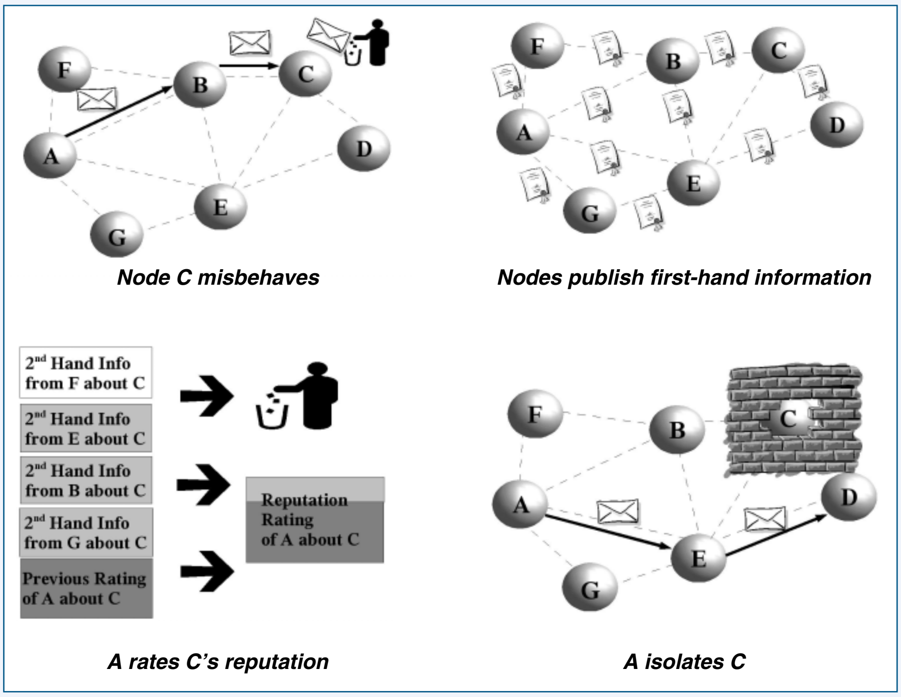
F
B
C
A
D
E
G
2nd Hand Info from F about C
2nd Hand Info from E about C
2nd Hand Info from B about C
2nd Hand Info from G about C
Previous Rating of A about C
Reputation Rating of A about C

Node C misbehaves

Nodes publish first-hand information

A rates C's reputation

A isolates C

Distributed Reputation System in a Mobile Ad-hoc Network.

tems and mobile ad hoc networks. We proposed a protocol called CONFIDANT (Cooperation Of Nodes — Fairness In Dynamic Ad hoc NeTworks) to cope with misbehaviour. Detection is based both on first-hand observation and on second-hand information provided by other nodes.

To be useful, reputation systems need to be reliable and robust against liars. They can, however, be tricked by the spread of false reputation ratings. On the other hand, simple solutions such as exclusively relying on one's own direct observations do not make use of all the information available.

Our fully distributed reputation system is based on a modified Bayesian estimation procedure and can cope with false information so as to effectively use second-hand information. Each node maintains a reputation rating and a trust rating for all other nodes it cares about. Reputation ratings capture the quality of the behaviour of a node as an actor in the network performing routing and forwarding tasks. From time to time reputation information is exchanged with others. However, second-hand information is only accepted if it is compatible with the current reputation rating or comes from a trusted node. To decide whether it is compatible, we introduce a simple mechanism called the deviation test. If the received estimation of misbehaviour deviates more than a threshold value from the current one, it is rejected, otherwise it is accepted. The trust rating is updated each time a deviation test has been performed, meaning compatible nodes are thus more trusted.

We use simulation to evaluate and demonstrate the performance. We found that CONFIDANT can keep the network performance high even when up to half of the network population misbehaves. We showed that our approach of using second-hand information significantly speeds up the detection of misbehaving nodes while keeping the number of false positives and negatives negligibly low.

Simulation results also suggest that the deviation test performs nearly as well on its own without the trust component. We have therefore analysed it in more detail and in a more general context. We introduced a mathematical model for a distributed reputation system based on the deviation test. We showed that it exhibits a phase transition behaviour, ie critical parameter values exist, below which liars have no impact and above which liars do have a certain impact and corrupt the reputation system. The critical values are obtained via a mean-field approach and are confirmed by direct computation as well as simulation. We thus give guidelines for a good choice of parameters. We also provide insight into a fundamental design question of such systems, namely whether or not direct observations only should be passed on as second-hand information.

Our current focus is on combining our simulation-based approach with the analytical approach. In particular, we are applying the insights gained by the analysis to find and zoom in on relevant parts of the experiment space. We are interested in seeing whether and how the phase transition effects observed in the analytical work carry through to our simulations of a more complex system, such as a mobile ad hoc network implemented in GloMoSim. Based on our analytical results, we now know what sort of behaviour we need to look for in the simulations. This will then enable us to see how the guidelines for a good choice of parameters from the analytical approach translate into a good choice of parameters for the more realistic system. We thereby optimize the performance of the distributed reputation system.

In conclusion, we provide a methodology to identify potential sweet spots and irregular behaviour of complex distributed reputation systems in decentralized self-organized systems. We expect to complete the combination of the two approaches by the end of the year.

Links:
http://lcawww.epfl.ch/salem/IP4Web/ImmuneNetworking.htm

http://www.sims.berkeley.edu/~sonja/CONFIDANT.html

http://winet-coop.epfl.ch/

http://icapeople.epfl.ch/leboudec

http://www.statslab.cam.ac.uk/~jm288

Please contact:
Sonja Buchegger, SIMS, UC Berkeley, USA
Tel: +1 510 643 2251
E-mail: sonja@sims.berkeley.edu

Jochen Mundinger, Statistical Laboratory, University of Cambridge, UK
Tel: +44 1223 337952
E-mail: J.Mundinger@statslab.cam.ac.uk

Jean-Yves Le Boudec,
EPFL-IC-LCA, Switzerland
Tel: +41 21 69 36631
E-mail: jean-yves.leboudec@epfl.ch