

ERCIM NEWS

www.ercim.eu

Special theme:

Mobile Computing

Also in this issue:

Keynote

*Strategies of the European ICST
Public Research Organisations
towards Horizon2020
by Domenico Laforenza*

Joint ERCIM Actions

*University of Southampton joins
ERCIM*

Research and Innovation

*VMC: A Tool for the Analysis of
Variability in Software Product Lines*

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. 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 about 8,500 copies..

*ERCIM News is published by ERCIM EEIG
BP 93, F-06902 Sophia Antipolis Cedex, France
Tel: +33 4 9238 5010, E-mail: contact@ercim.eu
Director: Jérôme Chailloux
ISSN 0926-4981*

Editorial Board:

Central editor:

Peter Kunz, ERCIM office (peter.kunz@ercim.eu)

Local Editors:

Austria: Erwin Schoitsch, (erwin.schoitsch@ait.ac.at)

Belgium: Benoît Michel (benoit.michel@uclouvain.be)

Cyprus: George Papadopoulos (george@cs.ucy.ac.cy)

Czech Republic: Michal Haindl (haindl@utia.cas.cz)

France: Thierry Priol (thierry.priol@inria.fr)

Germany: Michael Krapp (michael.krapp@scai.fraunhofer.de)

Greece: Eleni Orphanoudakis (eleni@ics.forth.gr),

Artemios Voyiatzis (bogart@isi.gr)

Hungary: Erzsébet Csuhaj-Varjú (csuhaj@inf.elte.hu)

Italy: Carol Peters (carol.peters@isti.cnr.it)

Luxembourg: Thomas Tamisier (tamisier@lippmann.lu)

Norway: Truls Gjestland (truls.gjestland@ime.ntnu.no)

Poland: Hung Son Nguyen (son@mimuw.edu.pl)

Portugal: Joaquim Jorge (jorgej@ist.utl.pt)

Spain: Silvia Abrahão (sabrahao@dsic.upv.es)

Sweden: Kersti Hedman (kersti@sics.se)

Switzerland: Harry Rudin (hrudin@smile.ch)

The Netherlands: Annette Kik (Annette.Kik@cwi.nl)

United Kingdom: Martin Prime (Martin.Prime@stfc.ac.uk)

W3C: Marie-Claire Fogue (mcjf@w3.org)

Contributions

Contributions should be submitted to the local editor of your country

Copyright Notice

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

Advertising

For current advertising rates and conditions, see <http://ercim-news.ercim.eu/> or contact peter.kunz@ercim.eu

ERCIM News online edition

The online edition is published at <http://ercim-news.ercim.eu/>

Subscription

Subscribe to ERCIM News by sending email to en-subscriptions@ercim.eu or by filling out the form at the ERCIM News website: <http://ercim-news.ercim.eu/>

Next issue

July 2013, Special theme: Intelligent Cars

Strategies of the European ICST Public Research Organisations towards Horizon 2020



*Domenico Laforenza
Director of the Institute for
Informatics and Telematics (IIT),
Italian National Research
Council (CNR), and ERCIM
Vice-President*

The scientific and professional societies in ICST in Europe have a key responsibility in helping to understand and to shape the digital future. They should collectively identify the vision, needs and priorities, and offer their expertise to the entire discipline and to society as a whole.

A round table on the “Role and Strategies of the European ICST Public Research Organisations towards Horizon 2020”, was held in October 2012 in the framework of the ERCIM Fall meetings (<http://www.ercim.eu/fall-meetings-2012>) at INRIA Sophia Antipolis. The panel addressed how European public research organisations working in Information and Communication Sciences and Technologies (ICST) can contribute to the success of Horizon 2020 (H2020), and in particular to the strategic direction of the “Excellence in Science” agenda. Members of the panel were: Domenico Laforenza (CNR, ERCIM Vicepresident – panel moderator), Keith Jeffery (STFC-RAL, ERCIM President), Michel Cosnard (CEO and Chairman, INRIA), Jan van Leeuwen (Utrecht University, Chair of the European Forum for ICST), Jos Baeten (CEO, CWI), Matthias Jarke (Chairman, ICT Group, Fraunhofer Gesellschaft).

The initial presentation addressed the new challenges introduced by H2020. After an overview of the objectives and structure of H2020, focusing on the three pillars of the H2020 strategy (Excellent Science, Industrial Leadership and Societal Challenges), DG “Connect” (the new Directorate General for Communications Networks, Content and Technology created by the EC in July 2012 in order to manage the Digital Agenda for Europe) was presented.

The main question raised in the panel was: How can European public ICST research organisations contribute to shaping the H2020 work programmes? In the past, each organisation interacted independently with the EC; however, the lack of a single voice has seriously impacted on their ability to be heard by EC decision makers. The development of common viewpoints and strategies for ICST in Europe and, whenever appropriate or needed, a common representation of these viewpoints and strategies at the international level are the foundational principles of the European Forum for Information and Communication Sciences and Technologies. EFICST (<http://www.eficst.eu/>) was established in November 2011 by the joint action of seven leading organisations and societies in ICST in Europe: ACM Europe (ACM Europe Council), European Association for Programming Languages and Systems (EAPLS), European Association of Software Science and Technology (EASST), European Association for Theoretical Computer Science (EATCS), European Coordinating Committee for Artificial Intelligence (ECCAI), European Research Consortium for Informatics and Mathematics (ERCIM), and INFORIE (Informatics Europe).

The Forum is intended to be an open platform for cooperation among the scientific ICT societies in Europe.

As reported in an EFICST white paper (Shaping the Digital Future of Europe) under preparation, the advances in information and communication science and technology continue to dramatically impact on our economies and our society. Computational thinking and virtualization are revolutionizing both science and technology. Driven by the need for automation and the exciting opportunities that are emerging, the development of ICST must be accompanied by a sound and far-reaching vision of the way we will live, work, and do business in the years to come. This requires a permanent dialogue between ICST professionals and the many other stakeholders involved, aimed at understanding where future innovations will take us. In its strategy for 2013-2015, the European Forum for ICST is defining a set of actions intended to meet these challenges in view of the crucial role of ICST as a key driver of innovation and change. The strategy focuses on three areas: Software as the key enabling technology of Europe’s digital future; ICT as the catalyst of industrial and societal innovation, and Informatics as a scientific discipline in vocational and professional education.

It was observed in the panel that although software is worldwide considered a strategic key enabling technology, unfortunately “software” is not included in the list of the Key Enabling Technologies (micro- and nano-electronics; photonics; nanotechnologies; advanced materials; biotechnology; advanced manufacturing and processing) reported in official EC H2020 documents. In order to raise awareness that the lack of recognition of the strategic importance of software technology will lead to a significant reduction in global European competitiveness, an interesting ISTAG report, titled “The Missing KET: Toward a Strategic Agenda for Software Technologies in Europe”, has been released. The objective of this report is to create a sense of urgency in the European software industry and awareness of software as the prime industrial differentiator and basis for innovation.

An additional topic discussed during the panel was how to reinforce relationships with the Knowledge and Innovation Communities (KICs) of the European Institute for Innovation and Technology (EIT), and in particular with the EIT KIC ICT Labs, in order to stimulate innovation through a more rigorous and dynamic link with higher education, research and business.

The final message of the panel was that European digital future is in the hands of many stakeholders, driven by different goals and needs. It is important to achieve coordinated thinking and actions so that society can benefit in the best possible way and not run the risk of missed opportunities or unwanted effects.

Domenico Laforenza

KEYNOTE

- 3 Strategies of the European ICST Public Research Organisations towards Horizon 2020**
by Domenico Laforenza

JOINT ERCIM ACTIONS

- 6 Cor Baayen Award 2013 Call for Nominations**
- 6 ERCIM Postdoc Fellowship Programme: Last Round of the Co-funded “ABCDE” Project**
- 7 University of Southampton joins ERCIM**
by Inés Teresa-Palacio

SPECIAL THEME

The special theme section “Mobile Computing” has been coordinated by Edgar Weippl, SBA Research, AARIT, Austria and Pietro Manzoni, Universitat Politècnica de València, SpARCIM

- 8 Introduction to the Special Theme**
by Edgar Weippl and Pietro Manzoni

Invited articles

- 10 Reality Mining at the Convergence of Cloud Computing and Mobile Computing**
by Matthias Steinbauer, Ismail Khalil and Gabriele Kotsis
- 11 Positioning Terminals in Mobile Computing Networks**
by Francisco Barcelo-Arroyo, Israel Martin-Escalona and Marc Ciurana-Adell
- 12 DarkDroid -- Exposing the Dark Side of Malicious Mobile Applications**
by Engin Kirda

Computing services

- 14 Automatic Offloading of Mobile Applications Using Evolutionary Algorithms**
by Gianluigi Folino and Francesco Sergio Pisani
- 15 Scheduling Data Mining Applications in Mobile Computing Environments**
by Carmela Comito, Deborah Falcone, Domenico Talia and Paolo Trunfio
- 17 On the Benefits of a Poly-Cultural Sensor Setup: Controlling Embedded Sensors with a Smart Phone**
by Rolf Adelsberger and Gerhard Tröster
- 18 SmartLab: Empowering Mobile Computing Research through an Open Smartphone Cloud**
by Georgios Larkou, Panayiotis Andreou, Andreas Konstantinides and Demetrios Zeinalipour-Yazti
- 19 Phone Accessories as an Interface**
by Mattias Jacobsson, Stina Nylander and Ylva Fernaeus
- 21 Boosting Performance of Wireless Networks with Concurrent Access by Smart Traffic Splitting**
by Gerard Hoekstra and Rob van der Mei

Platforms

- 22 CoMobility: A Mobile Platform for Transport Sharing**
by Carlos E. Cuesta, Paloma Cáceres, Belén Vela and José María Cavero
- 23 Quality of Service Information System: Get to Know the Performance of Your Mobile Network Operator Anywhere-Anytime**
by Katarzyna Wac
- 25 Mobile Service Platforms Based on Opportunistic Computing: The SCAMPI Project**
by Marco Conti, Franca Delmastro and Andrea Passarella

27 TravelDashboard - a Framework for the Delivery of Personalized Mobility Services to Urban Travellers
by Licia Capra, Pierre Chatel, Animesh Pathak and Roberto Speicys Cardoso

28 APISENSE: Crowd-Sensing Made Easy
by Nicolas Haderer, Romain Rouvoy, Christophe Ribeiro and Lionel Seinturier

Security and privacy

30 Mobilities: Analyzing Privacy Leaks in Smartphones
Jagdish Prasad Acharya, Franck Baudot, Claude Castelluccia, Geoffrey Delcroix and Vincent Roca

31 Privacy-Preserving Interest-Cast for Android Smartphones
by Gianpiero Costantino, Fabio Martinelli and Paolo Santi

32 Market-Based Security for Mobile Devices
by Gabriele Costa, Alessio Merlo and Luca Verderame

34 Revealing Social Links Between Owners of Wi-Fi Enabled Smartphones
by Mathieu Cunche, Mohamed Ali Kaafar and Roksana Boreli

35 Secure Collaboration for Smartphones
by Abdessamad Imine and Michaël Rusinowitch

Localization

36 Indoor User Localization Using Mobile Devices
by Jonáš Ševčík

37 Airplace: Indoor Geolocation on Smartphones Through WiFi Fingerprinting
by Christos Laoudias, Georgios Larkou, Demetrios Zeinalipour-Yazti and Christos G. Panayiotou

39 Knowledge Representation and Management in Indoor Mobile Environments
by Imad Afyouni, Cyril Ray and Christophe Claramunt

Applications

41 Mobile Real Time Applications for Enhancing Public Transport User Experience - The MOVE-ME Project
by João Falcão e Cunha, Teresa Galvão and Jeremy Pitt

42 U-AirPoll: Mobile Distributed and Collaborative Air Pollution Measurement
by Marino Linaje and Luis Miguel Dominguez-Peinado

43 ARGO Sentinel: The Mobile App for Reporting Oil Spillages at Sea
by Massimo Martinelli, Davide Moroni and Ovidio Salvetti

45 Mobile Devices to Improve Breast Cancer Information Management
by Damià Segrelles, Maite Giménez and Ignacio Blanquer

RESEARCH AND INNOVATION

This section features news about research activities and innovative developments from European research institutes

46 Secure and Privacy-Aware Mobile Identity Management
by Fabio Martinelli

47 Social Electricity: When Awareness About Electricity Becomes Social
by Andreas Kamlaris, George Taliadoros and Andreas Pitsillides

49 New Interaction Paradigms in Energy Management
by Paulo Carreira and Alfredo Ferreira

50 VMC: A Tool for the Analysis of Variability in Software Product Lines
by Maurice ter Beek, Stefania Gnesi and Franco Mazzanti

52 Wood Variety Recognition on Mobile Devices
by Pavel Vácha and Michal Haindl

53 The New SHIELD Architectural Framework
by Mariana Esposito, Andrea Fiaschetti, Francesco Flammini

54 A Biofeedback System for Self Empowerment and Improved Quality of Life
by Johanna Mercurio

55 Software and Hardware-Intensive Activities for Supporting Creative Learning
by Michail Giannakos and Letizia Jaccheri

EVENTS

56 W3C at the Mobile World Congress 2013
by Marie-Claire Forgue

56 PROMISE Winter School 2013 on Bridging between Information Retrieval and Databases
by Nicola Ferro

57 Announcements

IN BRIEF

59 CWI researcher Floor Sietsma youngest PhD in the Netherlands

59 Warsaw Center of Mathematics and Computer Science Established

Cor Baayen Award 2013 Call for Nominations

The Cor Baayen Award is awarded each year to a promising young researcher in computer science and applied mathematics. The award was created in 1995 to honour the first ERCIM President. The award consists of a cheque for 5000 Euro together with an award certificate. The selected fellow will be invited to the ERCIM meetings in autumn. A short article about the winner, together with the list of all candidates nominated, will be published in ERCIM News.

Conditions

Nominees must have carried out their work in one of the "ERCIM countries": Austria, Belgium, Cyprus, Czech Republic, Finland, France, Germany, Greece, Hungary, Italy, Luxembourg, Norway, Poland, Portugal, Spain, Sweden, Switzerland, The Netherlands and the United Kingdom. Nominees must have been awarded their PhD (or equivalent) after 30 April 2010. A person can only be nominated once for the Cor Baayen Award.

Submitting a nomination

Nominations should be made by a staff member of an ERCIM member institute. Self nominations are not accepted. They should concern nominees having performed their research in any research institution from the country of the nominating institution. Nominations must be submitted online

Selection

The selection of the Cor Baayen award winner is the responsibility of the ERCIM Human Capital Task Group, who will consult expert opinion when reaching their decision.

Deadline

Deadline for the 2013 award nominations is 30 April 2013

Further Information:

<http://www.ercim.eu/activity/cor-baayen-award>.

About Cor Baayen

The Cor Baayen Award is named after the first president of ERCIM and the ERCIM 'president d'honneur'. Cor Baayen played an important role in ERCIM's foundation.



Cor Baayen was scientific director of the Centrum Wiskunde & Informatica (CWI) in the Netherlands, from 1980 to 1994. He joined the institute in 1959 as a researcher of the department of pure mathematics, where he became group leader in 1965. As scientific director, Cor Baayen convinced the government to include CWI in a program that stimulated Dutch computer science research. He initiated or stimulated several new research areas including cryptography, computer algebra and performance analysis.

55 Post-Doctoral Positions at Inria in 2013

In 2013, Inria is offering 55 government-funded post-doctoral positions among others, each lasting about 12 or 24 months, for holders of a PhD or other doctorate. Post-doctoral positions offer the opportunity to pursue rewarding research work in the field of computer science and mathematics in an internationally recognized working environment with links to industry. Applicants must have defended their thesis in 2012 or 2013. The offered positions focus on the institute's priority research areas. However, high-quality applications to work on non-priority subjects will also be considered. Positions are open to candidates from all over the world. This campaign begun in January and ends on 14 June.

For subjects and available positions, see:

<http://www.inria.fr/en/institute/recruitment/offers/post-doctoral-research-fellowships/campaign-2013>

ERCIM Postdoc Fellowship Programme: Final Round of the Co-funded "ABCDE" Project

The final round of the successful "ABCDE" programme co-funded by the EC Marie Curie Actions is open. For this round, the fellowships are of 12-month duration to be spent in one ERCIM member institute. Topics cover most disciplines in Computer Science, Information Technology, and Applied Mathematics and applicants can also propose their own research agenda. The deadline for applications is 30 April 2013.

Conditions

Applicants must:

- have obtained a PhD degree during the last 8 years (prior to the application deadline) or be in the last year of thesis work with an outstanding academic record.
- be fluent in English.
- be discharged or get deferment from military service.
- have completed the PhD before starting the grant (proof will be requested).
- start the fellowship before May 2014.

An opportunity for ERCIM members

Since the start of the programme in 2010, some 130 fellowships have been co-funded by the European Commission's Marie-Curie Actions. Participating institutes appreciate the high quality of the candidates and that part of the costs of hosting a fellow is refunded. For this last round, co-funding is still available for about 50 fellowships and all ERCIM member institutes are entitled to benefit from this programme.

More information and application form:

<http://fellowship.ercim.eu/>

University of Southampton joins ERCIM

by Inés Teresa-Palacio

The University of Southampton is one of the top 15 research universities in the UK and is ranked among the top 1% of universities in the world (2012 QS World University Rankings). It is a founder member of the prestigious Russell Group of leading research universities in the UK and is part of the Worldwide Universities Network, engaging in collaboration and exchange agreements with some of the world's leading universities. The University of Southampton has around 22,000 students of which over 5,000 are international from more than 130 different countries. It has eight faculties covering a wide range of subject areas: from humanities, art and social sciences, to biology, chemistry, electronics and computer science, engineering, health sciences, medicine, mathematics, oceanography, environmental sciences, optoelectronics, physics and astronomy.

Electronics and Computer Science (ECS), which is part of the Faculty of Physical Sciences and Engineering, is one of the world's largest and most successful electronics and computer science departments, with more than 60 years of development in leading edge technology. It has a global reputation for computer science, agent technologies, web science, biometrics, open access, digital libraries, nanofabrication, wireless communications, energy harvesting and pervasive computing. ECS contains six research groups, an IT Innovation Centre, 100 academic staff, 120 research staff and 270 PhD students. ECS also plays a major role in two of the University's Doctoral Training Centres (DTCs) – the Institute for Complex Systems Simulation and Web Science. Both DTCs offer four-year postgraduate programmes in these specialist areas enabling students to work effectively across disciplines and helping them to develop the skills they need to address some of society's biggest problems.

Southampton is ranked first in the UK for Electronics and Electrical Engineering, and we are consistently placed in the top 10 places in the UK for Computer Science and IT. ECS researchers achieved exceptional success in the 2008 Research Assessment Exercise (RAE). Computer Science was ranked joint second in the UK for the quality of its research, with 85 per cent of its research work receiving either the top 4* rating (defined as 'world leading') or the 3* rating ('internationally excellent'). In Electronics and Electrical Engineering (in which ECS was assessed jointly with the University's Optoelectronics Research Centre), ECS (and the ORC) came second in the 'medals' tables, with 42 researchers rated as achieving research of either world-leading or internationally excellent quality. We are also one of only a handful of universities to receive the prestigious title Regius Professor, an honour bestowed by The Queen. The rare professorship marks our excellence in the field of Computer Science and reflects our exceptionally high quality of teaching and research.

ECS benefits from state-of-the-art facilities, including the Mountbatten Complex (one of the world's leading clean-

UNIVERSITY OF
Southampton



Mountbatten Complex

room laboratory complexes for materials and device research in diverse fields ranging from electronics to photonics and bio-nanotechnology) and the High Voltage Laboratory (one of only a handful of similar facilities in Europe that provide the high-level testing and research facilities required by electricity supply companies).

"I am very excited that the University of Southampton is joining the ERCIM network," says Professor Dame Wendy Hall who is representing the university in the ERCIM General Assembly. "Over the years we have hugely benefited from the international networks that we have been part of, including the World Wide University Network and the Web Science Trust international network of research laboratories. More recently we have formed a strategic partnership with the Fraunhofer-Gesellschaft through the SoFWIREd (Southampton Fraunhofer Web Science, Internet Research & Development) project. It is a natural step for us to join the ERCIM network, which represents some of the best computing research laboratories in Europe and enables the development of networks of excellence to tackle the grand challenges in our subject today, which require big teams working collaboratively across disciplinary and geographical boundaries".

Links:

<http://www.southampton.ac.uk/>

<http://www.ecs.soton.ac.uk/>

Please contact:

Professor Dame Wendy Hall, ERCIM General Assembly representative

Leslie Carr, ERCIM General Assembly substitute
University of Southampton, UK

E-mail: wh@ecs.soton.ac.uk, lac@ecs.soton.ac.uk

Introduction to the Special Theme

by Edgar Weippl and Pietro Manzoni

A few years ago, the IT industry forecast that consumer demand for data would surpass that of the voice market. At the time, few could conceptualize such an outlook. These days the data market is at least 70% larger than the mobile market, and still growing. Mobile technology has advanced in leaps and bounds over recent years, such that experts predict that, within the next few years, mobile computing will be strictly bound to cloud computing. Mobile cloud computing is set to impact and transform the mobile communication landscape and the whole computing infrastructure. Computing offloading, for example, is one of the main features of mobile computing to improve the battery life of mobile devices and to improve the performance of applications. However, there are many associated issues to solve, including efficient and dynamic offloading in a variable environment.

Mobile users may face certain problems, such as congestion owing to wireless bandwidth limitations, network disconnection, and signal attenuation caused by mobile users' mobility. Although many researchers have proposed solutions to optimally and efficiently allocate bandwidth, bandwidth limitation is still a major concern because the number of mobile and cloud users is increasing dramatically. 4G network and Femtocell are emerging as promising technologies that are revolutionizing bandwidth optimization, helping to overcome the traditional limitations. Efficient network access management not only improves link performance for mobile users but also optimizes bandwidth usage. Cognitive radio is likely to offer a solution to achieve wireless access.

Today's mobile phones harness the power of a computer, thus making sensitive information more widely available (eg information about user behaviour, GPS position, and personal data). Mobile phones are often fully integrated into social applications such as social networks, e-mail clients, messengers or data-harvesting analytics scripts on websites, enabling not only a holistic analysis of an individual, but also paving the way for new attack vectors operating against individuals and the companies they work for. Malware targeting mobile clients, especially those on the more open Android-system, have become very common.

Articles in the special theme

The articles presented in this issue reflect the many faces of our special theme, "Mobile computing", and address topics including: Mobile service platforms and new ways of establishing networks, privacy concerns and security challenges, location services and collaborative data capture.

The invited article by **Engin Kirda** summarizes a research project that focuses on Android malware detection. **Matthias Steinbauer et al** discuss effects of the convergence of cloud and mobile computing. **Francisco Barcelo-Arroyo et al** present results related to the important topic of providing indoor localization through the combined use of communication networks.

Different kinds of computing services form the focus of the first six articles:

Folino and Pisani describe a framework for generating decision tree-based models that use evolutionary algorithms to take automatic decisions regarding the offloading of mobile applications into a cloud environment. **Larkou et al** present SmartLab, an open IaaS cloud of smartphones that improves the efficiency of systems-oriented mobile computing research.

Mobile device networks also offer enormous potential for data mining. **Comito et al** have defined a distributed architecture with an energy-aware scheduling strategy that assigns data mining tasks in a peer-to-peer network of mobile devices.

Adelsberger and Tröster present a means of using smartphones to synchronize and control sensors wirelessly, more effectively, and in a more energy-efficient way than is usually the case for wireless synchronizing of multiple data streams.

With mobile Internet on the rise, the demand for wireless networks is growing at an unprecedented pace. **Hoekstra and van der Mei** describe new methods, developed in the Netherlands, that use smart algorithms to split traffic over the numerous overlapping networks in the country, thereby increasing wireless speed.

Moving towards the viewpoint of the user, **Jacobsson et al** describe a design concept for changing setups and user interface styles of smartphones by physically attaching phone shells or accessories such as jewellery or headphones to the device, enabling quick shifts between, for example, business and leisure modes.

The user perspective is also present in different mobile service platforms. The near ubiquity of smartphones makes them ideally suited for transport planning services. **Capra et al** introduce the TravelDashboard project, which will allow customizable trip planning according to personal preferences, limited not only to time tables and real-time delay information but also incorporating user-generated content with information on issues such as crowdedness of buses. In the same vein, **Cuesta et al's** CoMobility platform also allows travel planning but integrates carpooling with public transportation, allowing users to plan their trips so as to optimize cost, energy consumption, and CO₂ emissions.

Conti et al discuss opportunistic computing, a self-organizing dynamic networking paradigm that combines pervasive environmental network devices with mobile devices to allow communication and services. They present the CAMEO middleware platform, which focuses on the management and elaboration of context information in such opportunistic networks.

The pervasiveness of mobile devices has raised interest in collecting their sensor data via "crowd sensing". **Haderer et al** address the needs of various research communities with their APISENSE platform, which gives researchers an online environment in which they can set up an experiment without in-depth technical knowledge while ensuring the privacy and security of the participants collecting the data.

Privacy and security are, indeed, important issues as consumers and companies enter the world of mobile computing with unprecedented enthusiasm, and a number of contributions in this issue are dedicated to them.

One reason for the widespread use of mobile devices is, of course, the wide availability of Wi-Fi networks. Transmitting data over unsecured networks has well-known risks, but **Cunche et al.** go further and show how the automatic search for networks enabled on most smartphones can be used to not only fingerprint individual devices, but also to identify social links between their owners.

Networked calendars and other collaborative applications are very popular but pose considerable security challenges. **Imine and Rusinowitch** have developed a decentralized and secure shared calendar that is independent of third-party servers, instead allowing users to share their calendar events in dynamic groups in ad-hoc networks.

Constantino et al present an implementation of the cryptographic FairPlay framework for Android smartphones. It protects users' privacy in opportunistic networks by ensuring that information is exchanged with other users' devices via Bluetooth only if they have matching interest profiles, which can be determined without sending sensitive information in plaintext.

Achara et al have examined the information that can be gained from smartphones to educate users about risks. The Mobilities project investigates both the Android and iOS operating systems and apps for these platforms for potential privacy leaks and has found that many apps access information not necessary for their operation.

One way of making apps more secure can be found in **Costa et al's** proposal for a security-enabled app marketplace, where applications are analyzed to ensure that they comply with the security policy and can be installed without affecting the device's security configuration.

While GPS-based localization offers a multitude of location-based services (eg, navigation) in outdoor environments, indoor mobile environments require different approaches, some of which are presented in the next three contributions.

Laoudias et al's Airplace indoor positioning system uses a radiomap built from received signal strength (RSS) fingerprints. This radiomap is transmitted to the user's Android smartphone upon entering the building and allows the device to determine its position based on the signal strength it receives from surrounding wireless access points without revealing its personal state.

RSS fingerprinting is also one of the three techniques that are combined in the user localization method presented by **Ševčík**, the others being dead reckoning and sequential Monte Carlo filtering. The prototype displays the position of the user on a floor plan. Further developments are planned to allow navigation to a selected point on the map and augmented reality.

Afyouni et al focus on the representation and management of spatial data and location-dependent query processing required for the development of efficient and flexible context-aware indoor navigation systems. They developed a hierarchical data model of an indoor environment and algorithms that utilize it to process location-dependent queries continuously and effectively.

Finally, we have four contributions dedicated to applications. From bus travel to oil spills, they present four innovative uses of mobile computing combined with crowdsourcing. While **Falcao e Cunha et al** present their MOVE-ME multimodal travel planning project with user contribution, **Segrelles et al's** TRENCADIS allows the secure sharing, organization and searching of medical images on mobile devices and has been prototyped for breast cancer diagnosis and treatment.

The papers by **Trigueros et al. and Martinelli et al** both present applications for environmental monitoring with a strong focus on crowdsourcing: Trigueros and Peinado's U-AirPoll is a novel approach to distributed and collaborative air quality measurement, while Martinelli et al.'s ARGO Sentinel allows volunteers at sea to immediately report sighted oil spillages.

From localization and transportation services to sensor control and environmental monitoring, we hope you enjoy the contributions in this issue, which show the wide applicability of mobile computing and give us a taste of things to come.

Regarding the future, we can see that the area of mobile computing, within the wider area of the Information and Communications Technologies (ICT) is a very promising and strategic discipline. The European Commission, through its new program "Horizon 2020", has presented an €80 billion package for research and innovation funding, as part of the process to create sustainable growth and new jobs in Europe.

In particular, and strongly related to mobile computing, ICT in Horizon 2020 will be a crucial actor and will support the development of solutions for industrial leadership, by supporting the development of the next generation of computing thanks to advanced computing systems and technologies supported by enhanced network infrastructures, technologies and services for the future Internet, including content technologies and information management.

The final objective is to provide answers to societal challenges such as health, demographic change and wellbeing (eg, e-health, self management of health, improved diagnostics, improved surveillance, health data collection, active ageing, assisted living); secure, clean and efficient energy (eg, Smart cities; Energy efficient buildings; smart electricity grids; smart metering); smart, green and integrated transport (eg, smart transport equipment, infrastructures and services; innovative transport management systems; safety aspects); climate action, resource efficiency and raw materials (eg, ICT for increased resource efficiency; earth observation and monitoring); and finally inclusive, innovative and secure societies (eg, digital inclusion; social innovation platforms; e-government services; e-skills and e-learning; e-culture; cyber security; ensuring privacy and protection of human rights on-line).

Please contact:

Edgar Weippl
SBA Research/AARIT, Austria
E-mail: EWeippl@sba-research.org

Pietro Manzoni
Universitat Politècnica de Valencia, Spain
E-mail: pmanzoni@disca.upv.es

Reality Mining at the Convergence of Cloud Computing and Mobile Computing

by Matthias Steinbauer, Ismail Khalil and Gabriele Kotsis

During any social interaction, nonverbal social signals convey just as much information as the conversation itself. While transmitting and analysing conversation is quite a common task for machines, the transmission and analysis of social signals is not. The convergence of cloud computing and mobile computing leads to a situation where insight into social systems is possible, thus paving the way for exciting new applications.

Over the last decade, we have witnessed an increasing usage of mobile devices for capturing, analysing, and predicting human behaviour in everyday activities.

Most modern mobile devices are equipped with a plethora of sensors that capture every aspect of the user's physical context represented by attributes such as time, location, light, sound, weather, temperature, or even physiological state. Combined with social computing applications such as blogs, email, instant messaging, social networking (Facebook, Twitter, Linked,

formalized procedures such as workflows, recurring sequences of actions (such as routine tasks), types of motion (such as walking, running and standing), tasks (such as having lunch, washing dishes, and driving) and goals (for example, socializing, hiring, selling, keeping fit or simply having fun) [1].

Reality mining is the collection and analysis of machine-sensed environmental data pertaining to human social behaviour, with the goal of identifying predictable patterns of behaviour,

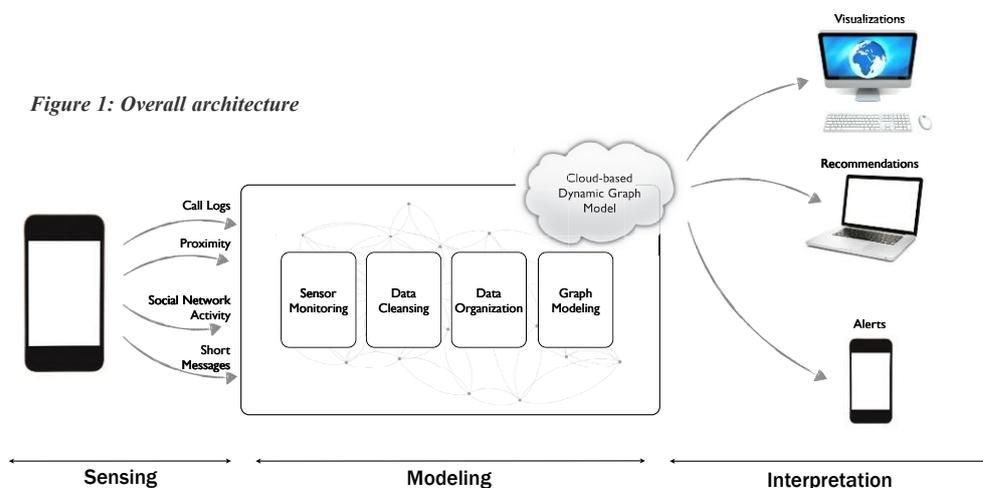
the social network from which they originate.

Such clues can be used to detect communication bottlenecks in organizations (such as when a single individual connects departments) and behavioural stereotypes applying to each member of a work group in order to determine the roles of individuals within groups (who is a leader in which group, who is an expert in which group).

In our research we focus on cloud-based reality mining. We are aiming to provide tools and techniques

that allow work groups to analyse their social network in near real time. Our objective is to improve group performance by providing a holistic overview of the group, its activities, situations and goals, in order to improve the group's overall performance.

The overall architecture of our approach is displayed in Figure 1. Reality mining comprises three phases: Sensing, modelling and interpretation.



Google+), Wikis, and social bookmarking, computers are able to capture the social context of the users in terms of interpersonal relationships and roles.

These physical and social contexts describe the existence of a relationship between two entities. Although the structure and nature of such relationships can be interpreted as a semantic network that can be used as the basis for understanding the meaning of an interaction, it fails to reflect the dynamics of a relationship over time. The dynamic patterns of interaction are essential in

including how computers can learn to extract social clues from social systems [2].

In reality mining, techniques inherited from data mining and data analysis are applied to data generated by human interactions, ie phone call logs, e-mail messages, Bluetooth proximity logs and cell tower logs, etc.

Reality mining analyses traces left by mobile devices, social networks and communication systems in the environment to extract social clues about

The sensing part is concerned with the recording, storing and transmission of sensor data that are generated by mobile phone sensors. These data are sent to a cloud based dynamic graph model for further processing.

The modelling phase, which shows the data processing pipeline in the cloud, consists of a sensor monitoring component in which raw sensor data are retrieved from roaming mobile devices. A REST service allows mobile clients to post their data to the service. In the data

cleansing component, the raw data are cleaned of duplicates and unnecessary content. This step also allows us to clean data from blocked phone numbers and e-mail addresses.

The data organization component performs a first alignment of the raw data to the dynamic graph model. Here we look up user profiles, and conversation topics or create them on demand if they did not exist.

Finally, in the graph modelling component, the data are inserted into the graph model as vertices and edges in a single transaction.

The final phase is the data interpretation in which the graph model is used to generate visualizations, recommendations and alerts.

Current research on reality mining tends to focus on one of two aspects: (1) implementing mechanisms and tools

that run on mobile devices or (2) implementing methods to analyze data sets in the cloud with low latency.

Our research focuses on dynamic graph models. These models keep their history as the graph changes. In this way we are able to analyse various phenomena, for instance changes in the structure of social networks. On top of these dynamic models we use queries that adapt their results whenever the graph changes in order to reflect the new model.

We are working on integrating methods learned from this research into the design of collaborative applications and in applications of computer supported cooperative work.

Mobile devices have become an important platform for understanding social dynamics and influence, because of their pervasiveness, sensing capabilities and computational powers. The conver-

gence of mobile and cloud computing is forming the breeding ground for real world applications in a research field known as reality mining. This is undoubtedly an interesting field, which is gaining new momentum with the convergence of mobile and cloud computing.

References:

[1] B. Begole: "Ubiquitous computing for business", Upper Saddle River, NJ: Financial Times Press; 2011

[2] Nathan Eagle and Alex Pentland: "Reality mining: sensing complex social systems", Personal and Ubiquitous Computing, 2006.

Please contact:

Matthias Steinbauer, Ismail Khalil, Gabriele Kotsis
Institute of Telecooperation, Johannes Kepler University Linz, Austria
E-mail: {matthias.steinbauer, ismail.khalil, gabriele.kotsis}@jku.at

Positioning Terminals in Mobile Computing Networks

by Francisco Barcelo-Arroyo, Israel Martin-Escalona and Marc Ciurana-Adell

The ability to pinpoint a terminal's position is useful for many applications of mobile computing and for network optimization (eg handovers, tariffs, resource management). A range of techniques are available to obtain a terminal's position [1]. GPS, for example, is used externally to the network and achieves good accuracy outdoors, with the trade off of increased energy consumption. Communication devices, however, are frequently used indoors, connecting to private networks, such as WLAN. Since GPS is inaccurate indoors owing to signal blockage and multipath errors, further research on indoors localization through communication networks is required. Mobile computing is linked to indoors positioning in applications such as: aged care, remote health control and security of buildings such as hospitals.

The fingerprinting technique is used extensively for WLAN positioning. The terminal collects the received signal strength from several access points and, during a precalibration phase, compares the achieved vector to the vectors previously recorded along with their positions. This technique does not involve modifications to the hardware. Other techniques use the time of flight (ie the time needed by a signal to travel between two nodes) to estimate the distances to several access points at known positions and then apply a trilateration process. The time of flight is more consistent than the signal strength. But in order to avoid modifications to the terminal's hardware, the time of flight must

be obtained from communication messages by using only software.

Recent research at the Technical University of Catalonia has led to a procedure to measure distances between terminals [2]. This procedure is aimed at obtaining the time of flight after adding timestamps to messages sent and to the corresponding acknowledgements received. The round trip time (RTT) is computed as the difference between both timestamps, and the distance between the nodes is inferred by considering that the trip occurs at the speed of the radio signal. The software must interact with the link layer of the protocol stack of the device. A simple

approach is to use a network interface capable of providing time measurements made by the hardware, but the timestamps performed do not have an acceptable accuracy for many location applications (eg the characteristics of IEEE 802.11 lead to a resolution of one microsecond corresponding to 300 metres in distance.)

A more sophisticated approach is presented in Figure 1. The protocol stack of the terminal is updated by introducing two software layers that are registered in each terminal. The registration process is run once and replaces the network interrupt handler (responsible for handling the events related with the net-

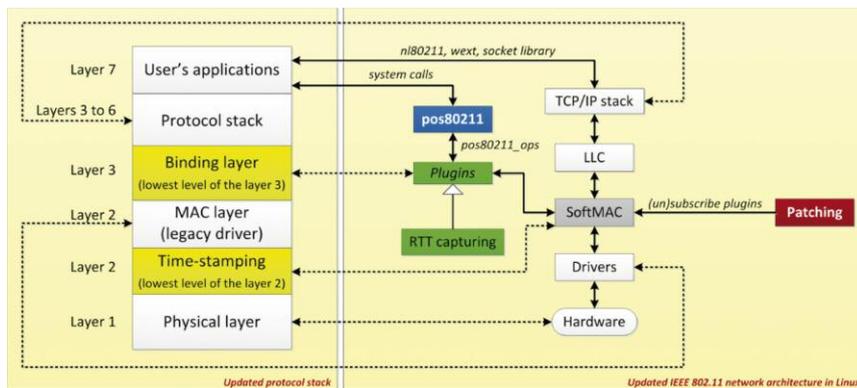


Figure 1: Proposed Linux implementation of the measurement system.

work interface) with a new one. This new handler analyses the traffic transmitted and received and adds timestamps to certain messages so that the RTT can be finally computed. Figure 1 also shows the network architecture including the interrupt handler in a Linux-based device. The sources of the Linux kernel have been patched to allow location metrics (ie RTT) to be observed. These changes alter the mac80211 subsystem, which implements most of the common MAC features in Linux. The goals of these changes are 1) to allow the location-related capabilities to be registered and released and 2) to add timestamps to the messages exchanged between the terminal and the access point. The capabilities are implemented as plugins, so that

each works as standalone. This allows isolation of the bugs and extension of the capabilities without impacting those that are already working. An RTT plugin has been developed in order to calculate the RTT between a node and an IEEE 802.11 access point. This plugin is responsible for most of the tasks developed by the interrupt handler. Specifically, it filters the traffic not suitable for location purposes and matches the transmission and reception messages involved in an RTT, so that the RTT can finally be computed. The interaction between the user's applications and the RTT plugin is done by means of system calls to a new module named pos80211 [3]. This module provides the computed and buffered RTTs to the user's applications.

Currently, the proposed implementation has been prototyped and tested at distances shorter than 30 metres with good results. Future work includes more testing (eg different scenarios, longer distances) and developing new plugins for hyperbolic, instead of circular, trilateration. The current plugin provides the RTT which is useful for circular trilateration, while time differences are appropriate for hyperbolic. The use of this proposal in ad hoc networks is also investigated.

References:

- [1] Y. Gu, A. Lo, I. Niemegeers: "A survey of indoor positioning systems for wireless personal networks", IEEE Commun. Surveys & Tutorials, vol. 11, no. 1, pp. 13-32, 1Q 2009
- [2] F. Barcelo-Arroyo, M. Ciurana, I. Martin-Escalona: "Process and system for calculating distances between wireless nodes", U.S. Patent 8 289 963, October 26, 2012
- [3] J. Corbet, A. Rubini, G. Kroah-Hartman: "Linux Device Drivers", O'Reilly Media, Third edition, 2005

Link: <http://grxca.upc.edu>

Please contact:

Francisco Barcelo-Arroyo, Universitat Politecnica de Catalunya, Spain
Tel: +34 934016010
E-mail: barcelo@entel.upc.edu

DarkDroid – Exposing the Dark Side of Malicious Mobile Applications

by Engin Kirda

The Android operating system is a burgeoning platform for deploying mobile applications to users, with more than 550,000 activations per day and an approximate 75% share of the global smartphone market that eclipses the once-dominant Apple iOS [1,3]. This trend is expected to continue, considering that Android's liberal licensing structure, open development environment, wide adoption across multiple hardware manufacturers and carriers, and modern end-user experience make it an attractive platform for both civilian and military use.

The Android security model is a particularly compelling reason for security-conscious organizations to adopt Android as a platform for developing and distributing mobile applications. In this security model, applications must declare a set of permissions that describes the set of privileged actions that might be performed during execution. Examples of such actions include accessing camera data, determining the location of a

device, or placing phone calls. Prior to installation, the user must explicitly approve the set of permissions requested by the application. During execution, the Android OS is responsible for ensuring that the application only performs those privileged actions that have been approved by the user.

This containment-based approach to preventing malicious behavior has the

benefit of potentially high scalability, since Android applications are not necessarily subject to a manual review process (as is the case for iOS applications distributed in Apple's App Store). Unfortunately, the discovery of Android-based malware in the wild has prompted concerns that the Android security model is inadequate, especially if Android devices are to be adopted in high security environments (eg, the US

Department of Defense, the European Commission, etc.) [2] A major drawback of the platform's approach to security is its reliance on user security awareness to detect potentially malicious applications. Specifically, most users are not possessed of the requisite motivation or security knowledge to determine whether an application legitimately requires a declared set of permissions, or to determine whether an attack has occurred at runtime. Therefore, there is a clear and pressing need to supplement the Android security model with the means to detect malicious code and remove the applications that contain such code from Android marketplaces.

Given their centralized nature, Android application marketplaces are well-suited as a venue for deploying static analyses to detect potentially malicious behavior in Android applications. Static analysis has the significant advantage that the entire program code can be analyzed, and thus, it is possible to guarantee the absence of certain, malicious behaviors. This is crucial for governmental stakeholders who might place a significant amount of trust into their applications. Furthermore, marketplaces also present an opportunity to deploy application instrumentation frameworks to remove malicious behavior or harden vulnerable applications against known classes of attacks prior to application installation.

DarkDroid is a joint project between Northeastern University in Boston and University of California, Santa Barbara, that is being sponsored by DARPA since 2012. The goal of DarkDroid is to research static mechanisms for proving the non-existence of malicious code in Android applications. We use scalable and precise static analyses on Dalvik bytecode to discover attacks ranging from privacy violations to unauthorized modifications to critical sensor data. Using the results of our analyses, our tools can then excise the malicious code or harden other apps against attack.

Static code analysis has come a long way; it is nowadays an accepted technique that is routinely used for examining large source code repositories for the presence of bugs and security vulnerabilities. In most scenarios, the goal of a static analyzer is to find as many program flaws as possible rather than to

guarantee the complete absence of bugs. While this unsound approach makes the analysis scalable and results in few false positives, it fails in the presence of malicious code. The reason is that malware authors write their programs to specifically exploit the difficult corner cases that an unsound analysis misses. Hence, we require a precise static analysis that accurately handles such difficult cases.

In the case of Dalvik (Android) applications, we identify two main challenges for precise analysis: First, programs make use of complex data structures (such as hash tables) and polymorphic classes with virtual methods. Second, Dalvik applications interact with the user in non-trivial ways, and multiple applications can collude (exchange messages) to achieve a single, malicious goal. We propose novel static analysis techniques that improve the precision of the analysis (and data-flow tracking) when facing complex data structures, user interaction (through intents), and virtual method calls. This guarantees that we will identify and capture malicious code with a low number of false positives, even when malware authors attempt to obfuscate their actions.

Another significant problem is that the absence of a particular class of malicious code does not imply that the application is not malware. In particular, it is crucial that the analysis guarantees the absence of a broad and diverse range of malicious behaviors for the analysis to be useful. In DarkDroid, we are working on novel ways to characterize

malicious activity. First, we leverage data-flow analysis to capture malicious code that aims to breach confidentiality. Moreover, we extend data-flow tracking with a fine-grain value set analysis to distinguish cases where certain data elements are legitimately transmitted from cases where unintended information flows occur. Finally, we introduce code patterns that represent attacks against integrity and availability of data, such as patterns that indicate that the program tries to draw an unnecessary amount of energy from the battery.

If successful, the impact of this project will be significant. In particular, DarkDroid will guarantee the absence of broad classes of malicious code from Android applications. Given the anticipated low false positive rates and the detailed output to render human analysis easy, we believe that it is possible to detect malicious applications before they are deployed on an Android marketplace.

References:

- [1] GOOGLE, INC. 2011 Q2 Earnings Conference Call Slides, July 2011 <http://kwz.me/01>
- [2] MCAFEE, INC. McAfee Q2 2011 Threats Report Shows Significant Growth for Malware on Mobile Platforms, August 2011, <http://kwz.me/02>
- [3] GOOGLE, INC. Advanced sign-in security for your Google account, February 2011, <http://kwz.me/05>

Please contact:

Engin Kirda
 Northeastern University, USA
 E-mail: ek@ccs.neu.edu

DarkDroid: Exposing the Dark Side of Android Marketplaces

University of California, Santa Barbara: Christopher Kruegel, Giovanni Vigna
 Northeastern University: Engin Kirda, William Robertson

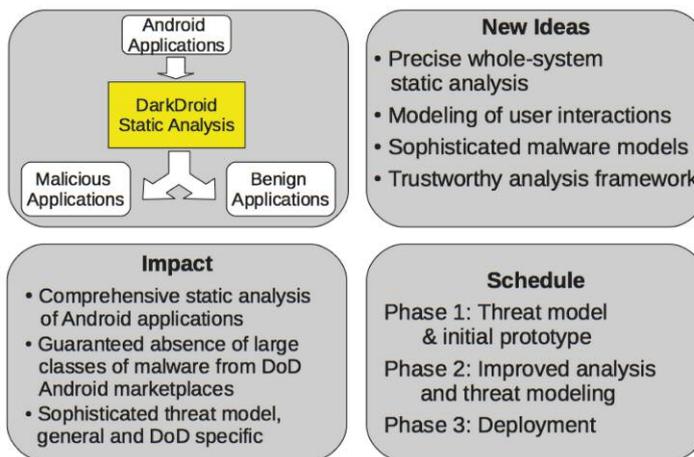


Figure 1: Summary of the DarkDroid Project

Automatic Offloading of Mobile Applications Using Evolutionary Algorithms

by Gianluigi Folino and Francesco Sergio Pisani

We present a framework for generating decision-tree based models that take automatic decisions on the offloading of mobile applications onto a cloud computing platform using an algorithm based on the genetic programming (GP) approach.

A key problem of modern smartphones is the limited life of their batteries. The introduction of larger screens and the large usage and availability of cpu-consuming and network-based mobile applications has aggravated this problem. The offloading of computation on cloud computing platforms can considerably extend battery duration. However, it is important to be able to verify not only whether offloading guarantees real advantages with respect to the computing power needed for data transfer but also if user requirements are satisfied, with respect to the quality of service and the costs of using the clouds. All the issues involved in the offloading decision, such as network disconnections and variability, data privacy and security, variations in load of the server, etc. need to be evaluated carefully.

At ICAR-CNR, we have designed a framework for the automatic offloading of mobile application using a genetic programming approach, which attempts to address the issues listed above. The framework comprises two parts: a module that simulates the entire offloading process, and an inference engine that builds an automatic decision model to handle the offloading process.

The simulator and the inference engine both apply a taxonomy that defines four main categories concerning the offloading process: user, network, data and application. The simulator evaluates the performance of the offloading process of mobile applications on the basis of user requirements, of the conditions of the network, of the hardware/software features of the mobile device and of the characteristics of the application. The inference engine is used to generate decision tree based models that take decisions concerning the offloading process on the basis of the parameters contained in the categories defined by the taxonomy. This is based on a genetic programming tool that generates the models using the parameters

defined by the taxonomy and driven by a function of fitness, giving different weights to the costs, time, energy and quality of service depending on the priorities assigned.

Taxonomy

A taxonomy of parameters and properties has been defined and is used to take decisions in order to build the model that decides the offloading strategy. The taxonomy only considers aspects that influence the offloading process and is based on four different categories: Application (parameters associated with the application itself), User (parameters assigned according to the user needs), Network (parameters con-

cerning the type and the state of the network), and Device (parameters reflecting the hardware/software features of the devices). The decision model built by the GP tool will take the decision whether to offload or not on the basis of the parameters associated with these categories.

cerning the type and the state of the network), and Device (parameters reflecting the hardware/software features of the devices). The decision model built by the GP tool will take the decision whether to offload or not on the basis of the parameters associated with these categories. The two main modules depict the inference engine, consisting of a Genetic Programming module which develops a population of models capable of deciding the possible offloading of a mobile application, and the simulation module comprising the well-known GreenCloud simulator (simulating the cloud part of the offloading process)

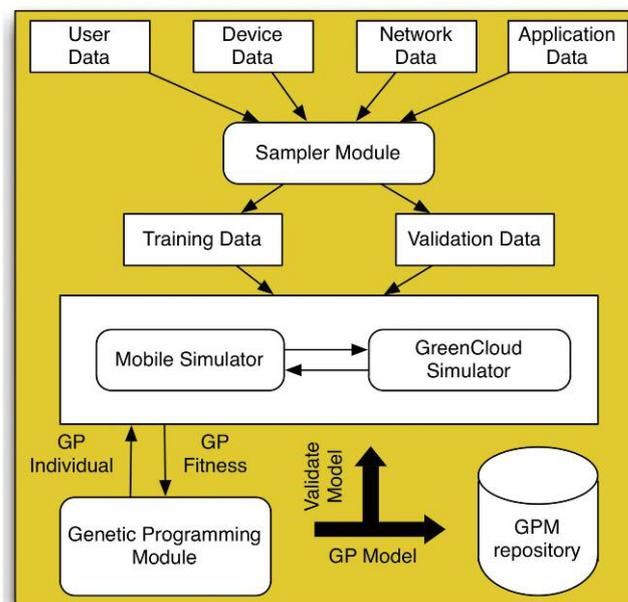


Figure 1: The software architecture of the system

cerning the type and the state of the network), and Device (parameters reflecting the hardware/software features of the devices). The decision model built by the GP tool will take the decision whether to offload or not on the basis of the parameters associated with these categories.

Architecture

Figure 1 shows the software architecture of the framework. The first four modules contain sets of data for each of

and an ad hoc mobile simulator modeling the mobile device's behaviour. In practice, each model generated by the GP module, will be input to the simulator module and a weighted fitness function will be computed that evaluates the performance of the model on the basis of the energy wasted, time consumed, costs and QoS.

At the end of the process, the best model (or the best models) will constitute the rules adopted by the offloading engine,

which will decide whether an application must be offloaded, according to the conditions assigned (user requirements, bandwidth, characteristics of the mobile device and so on). All these models must be validated using the simulation engine with the validation dataset. If the result of this evaluation is above a pre-defined threshold, the model will be added to a model repository for future use.

Conclusions and Perspective

This work presents an automatic approach to generate decision-taking

models for the offloading of mobile applications on the basis of user requirements, conditions of the network, the hardware/software features of the mobile device and the characteristics of the application. The system constitutes a general framework for testing offload algorithms and includes a mobile simulator, which computes the energy wasted in the process of offloading. Ongoing and future activities involve testing the framework with real datasets and verifying whether the obtained models improve battery performance in real environments.

Reference:

K. Kumar and Y.-H. Lu: "Cloud computing for mobile users: Can offloading computation save energy?", *IEEE Computer*, 43(4), 2010.

Links:

<http://www.genetic-programming.org/>
<http://www.icar.cnr.it/>

Please contact:

Gianluigi Folino
ICAR-CNR, Italy
E-mail: folino@icar.cnr.it

Scheduling Data Mining Applications in Mobile Computing Environments

by Carmela Comito, Deborah Falcone, Domenico Talia and Paolo Trunfio

Data mining is emerging as a promising topic in mobile computing environments. We have defined a distributed architecture in which mobile devices cooperate in a peer-to-peer style to perform a data mining process, tackling the problem of energy capacity shortage by distributing the energy consumption among the available devices. An energy-aware (EA) scheduling strategy assigns data mining tasks over a network of mobile devices optimizing energy usage. The main design principle is to find a task allocation that prolongs the network residual life by balancing the energy load among the devices.

The wide availability and growing computing power of mobile devices has opened the way for data analysis and mining in mobile scenarios [1]. Mobile applications exploiting data mining techniques have appeared on the market in recent years. Examples include smartphone-based systems for body-health monitoring, vehicle control, and wireless security systems. An important aspect that must be addressed is that of ensuring energy efficiency, as most mobile devices have battery power which would last only a few hours. Data mining tasks in mobile environments should be allocated and scheduled so as to minimize the energy consumption of low-capacity mobile devices.

We have worked in this direction by defining a distributed architecture in which mobile devices cooperate in a peer-to-peer style to perform data mining tasks, tackling the problem of energy capacity shortage by distributing the energy consumption among the available devices. Efficient resource allocation and energy management is achieved through clustering of mobile devices, as shown in Figure 1. With this approach, mobile nodes can be assigned

different roles, such as cluster-head or cluster member. A cluster-head serves as the local coordinator for its cluster, performing intra-cluster transmission arrangement and data forwarding, while a cluster member is a non-cluster-head node without any inter-cluster links.

To evenly exploit all available resources, a proper distribution of data mining tasks among clusters and individual devices is crucial. To this end, we defined an energy-aware (EA) distributed task scheduling strategy whose goal is to find a task allocation that prolongs the network residual life. The EA scheduler implements a two-phase heuristic-based algorithm. When an assignment decision has to be made for a task, the first phase, denoted local assignment phase, is responsible for local task arbitration: it considers the energy consumption of task execution on the different devices within the local cluster. The algorithm tries to minimize the total energy consumed in the cluster by assigning the task to that device that permits the cluster residual life to be extended. If the first phase is not feasible, the second phase, denoted global assignment phase, is responsible for

task arbitration among clusters: the task will be assigned to the most suitable device, in the network of clusters, that maximizes the network residual life.

We performed an evaluation of the EA allocation strategy using a network simulator, which allowed us to assess its effectiveness on a set of data mining tasks [2]. The simulation mainly aims to study the behaviour of the scheduler with respect to network residual life, number of alive devices, and number of completed tasks at the end of the simulation. To assess the effectiveness of the EA strategy, we compared its performance with that achieved by the well-known round-robin (RR) scheduling algorithm. As a first step, the simulator builds a network composed of 100 mobile devices, and lets them group into clusters. Then, an initial energy capacity ranging from 3,000 to 11,000 Joules is assigned to each device, following a normal distribution. After the initial setup, mobile devices start generating a set of data mining tasks to be executed; these are allocated to the available nodes according to the EA strategy. For the purpose of our simulation, we considered three reference data

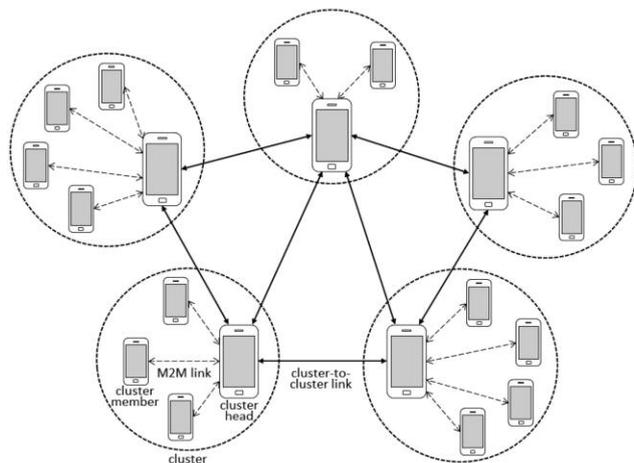


Figure 1: Reference architecture

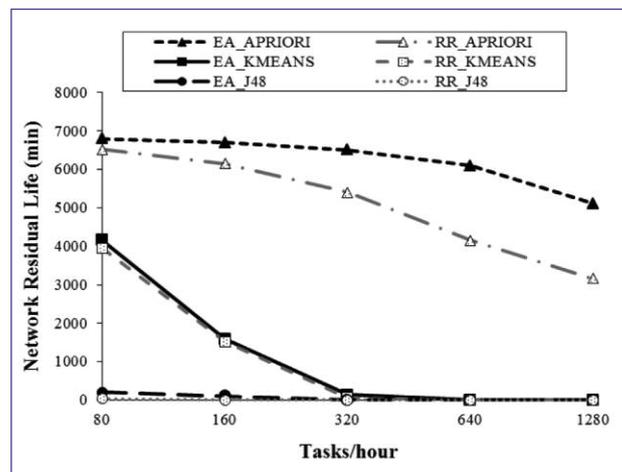


Figure 2: Network residual life

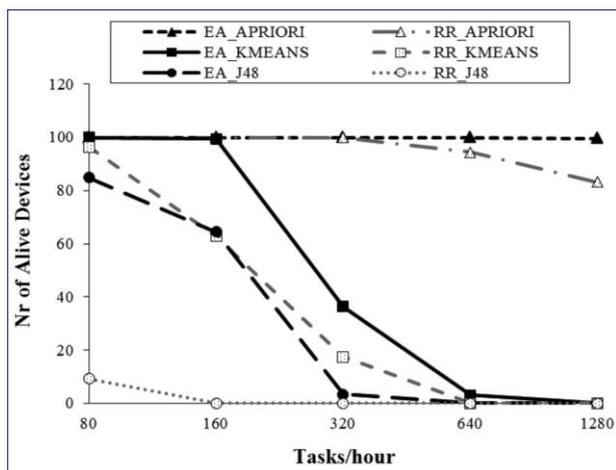


Figure 3: Number of alive devices

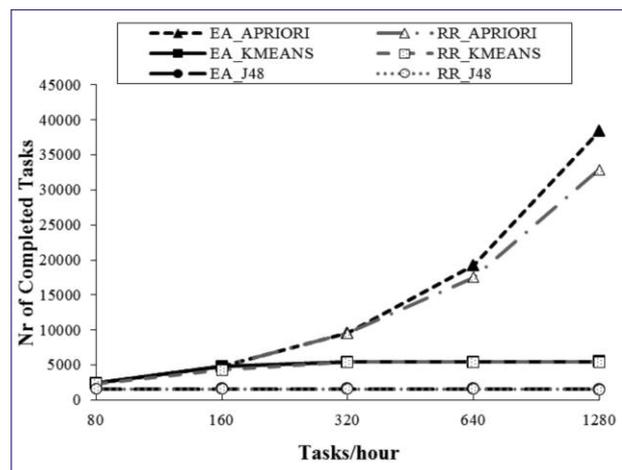


Figure 4: Number of completed tasks

mining algorithms from the Weka project [3], namely J48, K-means and Apriori.

For each reference algorithm, we ran a set of tasks with a fixed dataset size (200 kB) and task arrival rate λ varying from 80 to 1280 tasks per hour. Figure 2 shows the network residual life measured at the end of the experiments for the three algorithms, using EA and RR. As expected, increasing the task arrival rate, the network residual life tends to zero both for EA and RR. However, for the lightest of the three data mining algorithms (Apriori), the residual life does not reach zero and the difference between EA and RR increases with λ in favor of EA. Figure 3 shows the number of alive devices for the three algorithms, using EA and RR. Also in this case, the results demonstrate that the number of alive devices with EA is greater than that achieved by RR. Finally, Figure 4 compares the performance of EA and RR in terms of completed tasks for the three algorithms. With Apriori, both EA

and RR are able to complete more tasks as λ increases, but EA ensures better performance. With J48 and K-means, with a given task arrival rate, the number of completed tasks cannot increase because the network residual life is zero, as shown in Figure 2. However, even in these cases, there is an advantage for EA compared to RR with λ lower than 320 tasks/hour.

In summary, the experimental results showed that a significant improvement can be achieved using our EA scheduler compared to the time-based RR scheduler. Our algorithm: i) is effective in extending the network residual life by reducing the energy consumption, without limiting the number of completed tasks; ii) keeps alive most of the mobile devices, in all the experiments performed, thanks to its energy load balancing strategy.

Link:
<http://grid.dimes.unical.it>

References:

- [1] D. Talia, P. Trunfio: "Service-oriented distributed knowledge discovery", Chapman and Hall/CRC, 2012
- [2] C. Comito et al: "A distributed allocation strategy for data mining tasks in mobile environments", in proc. of the 6th International Symposium on Intelligent Distributed Computing, Studies in Computational Intelligence 446: 231-240, 2012
- [3] M. Hall et al: "The WEKA data mining software: an update". SIGKDD Explorations 11(1): 10-18, 2009.

Please contact:

Carmela Comito and Domenico Talia
 ICAR-CNR and DIMES, University of Calabria, Italy
 E-mail:
 {ccomito,talia}@dimes.unical.it

Deborah Falcone and Paolo Trunfio
 DIMES, University of Calabria, Italy
 E-mail:
 {dfalcone|trunfio}@dimes.unical.it

On the Benefits of a Poly-Cultural Sensor Setup: Controlling Embedded Sensors with a Smart Phone

by Rolf Adelsberger and Gerhard Tröster

We present our wireless solution to device synchronization, control and real-time feedback for small and light-weight Inertial Measurement Units, (IMUs). The controller is implemented on a smartphone. We base a custom application layer on the protocol layer implemented on a low-power radio chip by ANT+TM. Our system is capable of synchronized (< 15 μsec) control of an arbitrary number of sensors (ETHOS), streaming data for real-time visualization and it reduces power consumption compared to other approaches.

Multi-device sensor systems often introduce the problem of synchronizing multiple data streams, as sub-millisecond accuracy is often needed. Synchronizing multiple sensors at run-time is a complex task. After data acquisition, offline synchronization by correlation analysis between two data streams solves the problem. Two IMUs, for instance, can be synchronized if there are shared features in some of the data signals. However, for statistically independent data sources this approach is not feasible. In this situation, a physical connection can help by providing a synchronization signal. Synchronization between two identical devices via wireless communication consumes a lot of power and requires an advanced network protocol. Our solution uses a single smartphone to control an arbitrary number of IMUs (ETHOS Sensors [1]).

Controlling Sensors via ANT+TM

We use small, light-weight IMUs that feature on-device storage [1]. An ANT+TM-capable network chip takes care of the wireless communication. ANT+TM defines a data source, ie a sensor, as a master device. An aggregator, such as a smartphone, collects data and is dubbed slave in this setting. The network chip uses a TDMA-like adaptive isochronous network technology for communication.

An embedded device's energy scarcity is exacerbated by the requirements of wireless communication. Our sensor nodes draw current from 300 mAh Li-Ion batteries; hence, power awareness was crucial during the development. Deployment of the devices in master mode is suboptimal for two main reasons: Constantly sending data, even if not used by an aggregator, wastes energy: not only does the cost of wireless communication add to the power budget, but also the increased load on the operating system uses more power

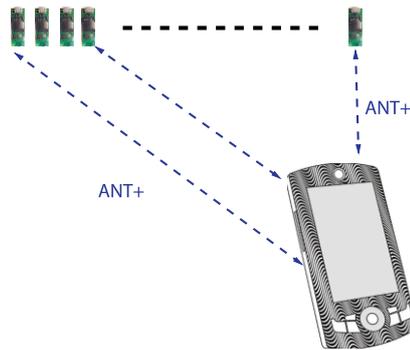


Figure 1: System setup: A smartphone communicating to an arbitrary number of sensors; (size ratio of sensors to smart phone is accurate).

since a node cannot go to sleep. Further, frequent collisions on the physical layer at high transmission rates with multiple master devices would not only reduce the throughput of the network: Retransmissions also impact the real-time data processing of an aggregator and thus aggravate the problem of

power scarcity. These considerations motivated our solution. We configured the sensors to slaves: they do not emit data packages on their own, but listen to new packages on the channel. On the other hand, our aggregator device - an X-Peria S by Sony - assumes the role of a master: in idle mode it emits heartbeat packages that are received by any sensor device tuned to a common radio frequency.

In addition to this network-layer connection, we also implemented a simple application layer: A user of the system can control all or specific nodes with a set of commands. Each node is configured to a node ID - a simplistic equivalent of an IP-address. The master emits either broadcast or addressed packages. A sensor node only reacts to a network package if that package either is a broadcast package or if the destination ID of the package matches the node's ID. There is one command to start sam-

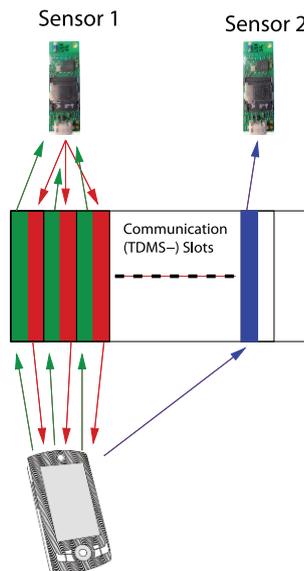


Figure 2a: Communication with ANT+TM-. A smartphone requests data from sensor 1 (green slots) which answers in the same slot (red). Later, the smartphone sends a command to sensor 2 (blue slot).

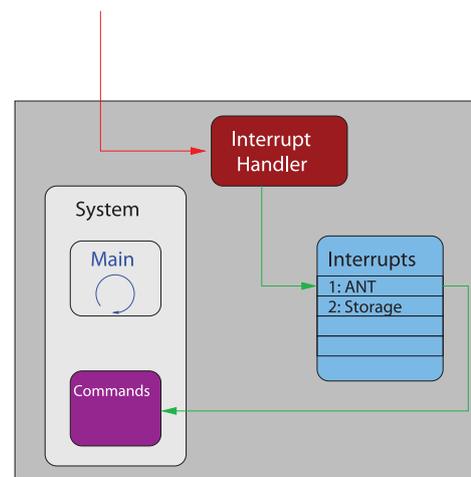


Figure 2b: Network interrupt priority. A smartphone requests data from sensor 1 (green slots) which answers in the same slot (red). Later, the smartphone sends a command to sensor 2 (blue slot).

pling. The same operating system (firmware) runs on each sensor node. Network communication is handled by an interrupt service routine with highest priority.

In fact, every other software interrupt is overwritten by the network-stack interrupts. This way, we can ensure that each node receives the commands at virtually the same time: the micro-controller of our sensor boards features fixed-overhead interrupt handlers.

For 100 iterations we started the sampling on 10 rigidly connected sensor nodes, exposed all of them to a common peak in acceleration and calculated the maximal inter-node offset between any

two sensor nodes. The mean offset was 12 μ sec. A user can select a sensor to stream its data to the aggregator: the GET DATA command requests data packages from an addressed sensor node. The ANT+TM-protocol allows slave devices to reply to a master device in the same time slot. We exploit this feature to implement real-time data streaming: in streaming mode the master sends addressed packages to one device which replies with packages containing the requested sensor data.

Conclusions

We have shown that multiple types of wearable devices can be combined into a sensor system to create something more powerful than a mono-cultural system.

By carefully programming the software modules wireless synchronization can be as accurate as 12 μ sec.

Link:

<http://www.ife.ee.ethz.ch/>

Reference:

[1] H. Harms et al: "ETHOS: Miniature orientation sensor for wearable human motion analysis" in proc. of IEEE Sensors conference, IEEE, 2010.

Please contact:

Rolf Adelsberger
Federal Institute of Technology Zurich, ETHZ, Switzerland
E-mail: rolf.adelsberger@ife.ee.ethz.ch

SmartLab: Empowering Mobile Computing Research through an Open Smartphone Cloud

by Georgios Larkou, Panayiotis Andreou, Andreas Konstantinides and Demetrios Zeinalipour-Yazti

SmartLab is a first-of-a-kind open cloud of smartphones that enables a new line of systems-oriented mobile computing research.

The year 2011 marked the beginning of the post personal computer (PC) era, as the number of smartphone shipments exceeded, for the first time in history, the sales of all types of PCs combined (ie, notebooks, tablets, netbooks and desktops). According to IDC's "Worldwide Mobile Phone 2012-2016 Forecast and Analysis", Android is projected to dominate the future of the smartphone industry with a share exceeding 53% of all devices shipped in 2016. Currently, an Android smartphone provides access to more than 650,000 applications and those apps introduce unprecedented possibilities, knowledge and power to humankind.

Re-programming smartphones and instrumenting them for application testing and data gathering is currently a tedious, time-consuming process that poses significant logistical challenges. To this end, we have implemented and demonstrated SmartLab [1], a first-of-a-kind open Infrastructure-as-a-Service (IaaS) cloud that enables fine-grained control over both real and virtual smartphones via an intuitive web-based interface. Our current infrastructure is ideal for scenarios that require fine-grained and low-level control over real smart-

phones, eg, OS, Networking, DB & storage, security, peer-to-peer protocols, but also for scenarios that require the engagement of physical sensors and geo-location scenarios. Our preliminary edition has been utilized extensively in-house for our research and teaching activities [2,3] but it has also been open to selected research groups around the globe. SmartLab provides a diverse, high-availability platform that can be utilized by the mobile computing research community to engage more effectively in systems-oriented research on smartphones.

SmartLab's hardware comprises 40+ Android smartphones and our data centre that encompasses over 16TB of



Figure 1: Part of the Smartlab smartphone fleet connected locally to our datacentre

RAID-5 / SSD storage on an IBM X3550, as well as 220GB of main memory on 5 IBM / HP multiprocessor rackables. SmartLab supports a variety of connection modalities. In particular, most of our devices are connected to the server with USB hubs. Additionally, a number of smartphones are connected from within our DMSL research lab (in the same building) through the network on virtual servers that reside on our data centre. The given mode is particularly promising for scenarios we want to apply to scale our testbed outside the department. Finally, a few devices within the department are also wirelessly connected using WiFi.

SmartLab supports four modes of user interaction with the smartphones: i) Remote Control Terminals (RCT), a web-based remote screen terminal for Android, developed in-house using Ajax that mimics touch screen clicks and gestures among other functionalities; ii) Remote Shells (RS), a web-based shell developed in-house with Ajax that enables a wide variety of UNIX commands to be issued to the Android Linux kernels; iii) Remote Scripting Environment (RSE), which

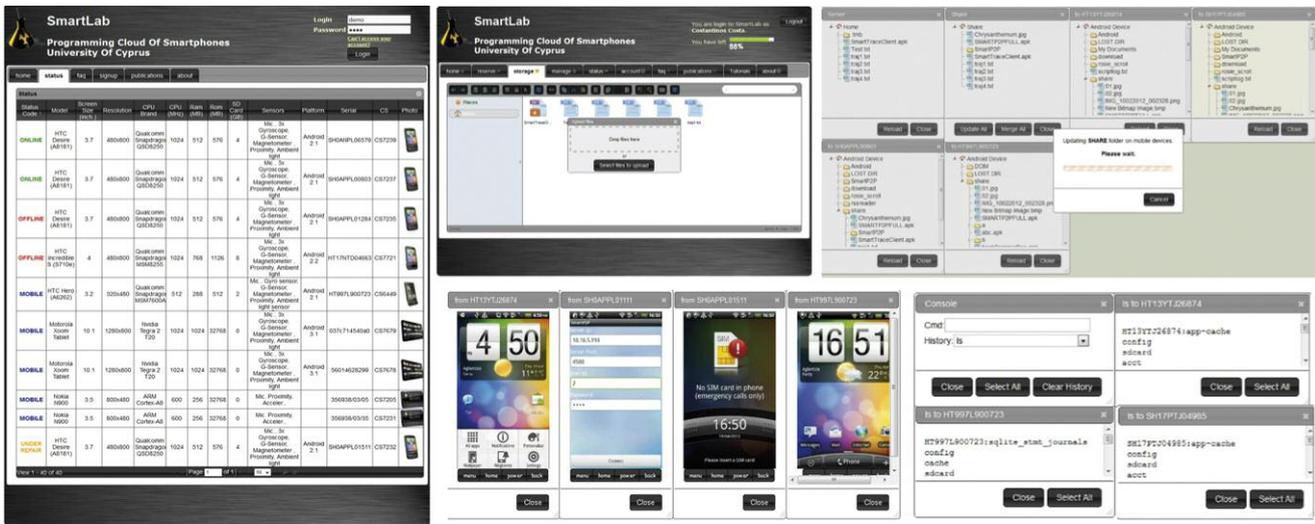


Figure 2: Screenshot of the SmartLab web-based user interface

allows users to author Android MonkeyRunner automation scripts (written in python) and upload them to the devices to perform automated tasks; and iv) Remote Debug Tools (RDT), which provide web-based debugging extensions to the Android Debug Bridge (ADB). Through these tools, SmartLab facilitates research in smartphone network programming environments, communication protocols, system design and applications.

In the near future, we are planning a number of exciting extensions for SmartLab. Firstly, we are integrating a prototype Apache HBase installation within our data centre, to store sensor readings in a tabular and scalable (ie, column-oriented) format. The given store can be utilized to store billions of

sensor readings that can be fed to our GPS/Sensor Mockup subsystem. Additionally, we are currently working with local telecommunication companies in order to obtain 3G data time for our mobile fleet, and local transportation companies, which will be moving our devices around within a city. This will allow testing of algorithms, protocols and applications within a real mobile urban environment, thus providing a first-of-kind open mobile programming cloud. Furthermore, we are developing a Web 2.0 JSON-based API of our testbed using the YII framework. In particular, this effort will allow users to access the subsystems of our testbed in a programmable manner (ie, Web 2.0 JSON interactions). Finally, we are also looking into mechanisms for supporting security and privacy experiments.

Links:

<http://smartlab.cs.ucy.ac.cy/>
<http://dmsl.cs.ucy.ac.cy/>

References:

- [1] A. Konstantinidis et al: "Demo: a programming cloud of smartphones", ACM Mobisys '12
- [2] D. Zeinalipour-Yazti et al: "Crowdsourced Trace Similarity with Smartphones", IEEE TKDE, 2012
- [3] G. Chatzimiloudis, et al: "Crowdsourcing with Smartphones", IEEE Internet Computing, Volume 16, 2012.

Please contact:

Demetrios Zeinalipour-Yazti
 Department of Computer Science
 University of Cyprus
 Tel: +357 22 892755
 E-mail: dzeina@cs.ucy.ac.cy

Phone Accessories as an Interface

by Mattias Jacobsson, Stina Nylander and Ylva Fernaeus

Mobile ActDresses is a design concept that utilizes existing practices of accessorizing, customization and manipulation of a physical mobile device to predict and control the behaviour of its software. Existing smartphone technology can be augmented by taking inspiration from existing practices in relation to costume, jewellery, personalization and fashion.

There are various physical means by which people personalize their digital devices: by placing stickers on laptops; by using customized cases; and by attaching mascots and charms to smartphones, for instance [1]. Another area of customization is the use of personalized digital themes and the growing amateur practice of making small and personal

mobile applications. The mobile phone as such is thus not merely a technological device but also an object for personal expression. Inspired by the trend of personalizing devices, we have been investigating the potential of using physical clothing, accessories and labelling as an alternative means of controlling mobile interactive systems by

developing a prototype system: Mobile ActDresses.

The design space of Mobile ActDresses ranges conceptually, from single on/off mode switchers to more complex configurations with combinations of active labels and accessories. Our prototype enables people to attach physical acces-

sories to their smartphone and, at the same time, change its digital functionality or appearance. Below we outline examples of different strategies for implementing and deploying this design concept based on existing standards on mobile phones currently on the market.

Various implementation experiments

There are many forms of technology available for mobile phone handsets that could potentially be used to implement the concept of ActDresses [2]. However, wireless protocols such as Bluetooth and WiFi require active transmitters, and also have a relatively long range of communication, features that are incompatible with having signs in the immediate physical context of the device that they control. Using the phone's camera to read barcodes on an object, for instance, requires explicit reading and is not ideal given that it conflicts with the immediate physical context requirement.

Another emerging wireless method is Near Field Communication (NFC), which enables the exchange of data between devices at a distance of up to ten centimetres from one another. Drawbacks in this case are that the reading must be constantly active on the device and that most devices currently do not support simultaneous reading of multiple tags. As with the wireless protocols, there is a range of possible wired or direct contact solutions for realizing the concept. Examples include experimental solutions such as Pin & play, iButtons, conductive stickers, resistors, USB, or even the built in memory cards.

There are thus a number of partially complete technical options currently available. Our conclusion is that it will be crucial to have a dedicated channel, whether it is wireless or wired. The key point is that physical 'sockets' restrict-positioning of tags whereas a wireless solution can be designed to be both free and 'socketed'.

The design prototype

In our first prototype (Figure 1), we equipped physical shells with small but strong Neodym (NdFeB) magnets positioned at different locations. The absolute distance between the magnet and the magnetometer in the phone is sensed, and can be used to trigger events in software. In this case a simple service application was developed that



Figure 1: Mobile phone prototype using magnets on shells to control themes and other content.

changes the theme on a mobile Android phone according to the style of the shell. Furthermore, we used the headphones/handsfree jack as an example of a wired ActDress together with jewellery (Figure 2).

A general design concept for expression

To conclude, our goal with this prototype was to explore how the concept of ActDresses could be extended to digital artifacts in general [3]. Our inspiration comes from how clothes are worn by people to serve a range of communicative functions, indicating appropriate behaviours, group belongings, and expected interactions, for example. Similarly, physical accessories attached to a device could be used as a resource to indicate what mode the device is currently in, and what behaviours and interactions could be expected.

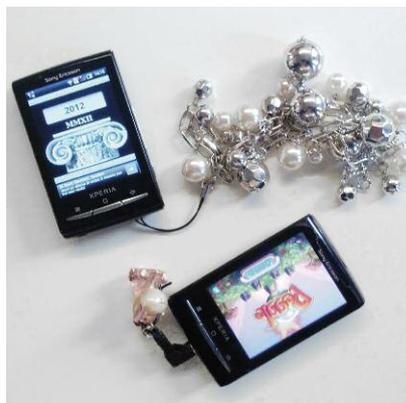


Figure 2: Headphone and jewellery hack to control applications and other content.

A day in Jill's Mobile Life

Jill rarely leaves the house without her mobile phone, which she customizes physically as well as digitally to match her own clothes of the day. Thereby the phone itself works as a fashion item that she seeks to match with her outfit. Upon entering her office, Jill attaches the company shell to her mobile phone handset, which enables the phone to let her into the building, as well as acting as a company identity marker and label. Plus, it goes well with her work outfit. The phone is now set into a mode that automatically loads her work contacts as her primary address book in the phone. While the shell is on, all charges on the phone get placed on the company, rather than on her personal, phone bill. When leaving the office she removes the shell on her phone, which then replaces her office applications with her favorite spare time applications on the front screen. Later in the evening Jill goes out for a drink and attaches her 'after work'-charm that shows her VIP status at one of the local clubs. Not only does it enable rebates at the club when paying with the phones digital wallet, it often becomes a conversation piece, and therefore a marketing tool for that specific club.

Link: www.sics.se/projects/actdresses

References:

- [1] J. E. Katz, S. Sugiyama: "Mobile phones as fashion statements: evidence from student surveys in the US and Japan", *New Media Society*, 8(2), 321–337, 2006 .
doi:10.1177/1461444806061950
- [2] M. Jacobsson, Y. Farnaes, R. Tieben, R.: "The look, the feel and the action", in *proc. of the 8th ACM Conference on Designing Interactive Systems - DIS'10*, p. 132, New York, USA: ACM Press, 2010
doi:10.1145/1858171.1858196
- [3] Y. Farnaes, Y., M. Jacobsson: "Comics, robots, fashion and programming: outlining the concept of actDresses" (pp. 3–8). Cambridge, UK: ACM, 2009
doi:10.1145/1517664.1517669

Please contact:

Mattias Jacobsson, SICS Swedish ICT
Tel: +46 703 128557
E-mail: majac@sics.se

Boosting Performance of Wireless Networks with Concurrent Access by Smart Traffic Splitting

by Gerard Hoekstra and Rob van der Mei

The spectacular growth of mobile internet on smartphones and tablet computers has boosted the demand for fast wireless networks. By 2015, mobile data exchange is expected to be 26 times larger than it was in 2010 [1]. In the first six months of 2012, the mobile internet traffic volume in the Netherlands was 21% higher than that recorded in the last six months of 2011 [2]. A highly promising means to meet the increasing demand is to take advantage of the fact that many geographical areas are covered by a multitude of overlapping networks. This phenomenon is referred to as concurrent access (CA). CWI and technology company Thales Nederland B.V. have developed new methods to make efficient use of the possibilities of CA by developing and implementing smart algorithms to split traffic over the multitude of wireless access networks.

Most locations in developed regions are covered by a multitude of overlapping wireless networks from several network providers that have 3G/4G networks deployed, as well as hotspots and local Wi-Fi-networks. Usually, users can only access one of these networks at a time. But what if several of the available networks could be used simultaneously? Would splitting data over two wireless networks also make mobile communication twice as fast? And would it increase the overall performance of the communication? The results of research by CWI and Thales exceeded expectations: simultaneous use of two networks did not just double the speed of wireless communication, but could increase it up to a factor of ten.

High speed gain

The key factor in this speed gain is the strong fluctuation in available network capacity. In this context, one may use temporal under-utilization in one network to compensate for congestion in the other network. Even if both networks are fully utilized, there are still very high fluctuations in supply and

demand of network capacity on a micro scale. The main idea is to split traffic streams over the different access networks in such a way that the user-perceived performance is optimal, for example by assigning packets to the network with the lowest utilization. This smart assignment of packets to parallel wireless access networks not only leads to an increase in the network throughput, but also leads to great enhancements of response-time performance, availability and robustness of the applications.

Removing peaks in capacity demand is very effective in reducing congestion. As in several logistic problems, speed is not linearly, but rather exponentially, related to available capacity. If the number of cars on a road could be decreased by 10% during rush hour, traffic jams would not decrease in length by 10%, but would be halved or even eliminated. The same happens in wireless networks: the download times of individual users can be up to ten times shorter when capacity peaks are flattened by using two networks at the same time.

Sharing networks

The developed techniques can be realized in the short term, simply by devising a smartphone that can combine any mobile networks, hotspots or Wi-Fi-networks, and a piece of software based on our models. But for a more effective implementation, network providers need to adapt their software and hardware to this new mode of operation and more importantly, reach an agreement with competing providers to share networks.

Application areas

Many application areas, such as emergency and security services and the military, could greatly profit from fast and reliable networking. Such organizations could use their own networks to profit from these new techniques.

Developing the model

The mathematical challenge of modeling the use of multiple networks at the same time lies in devising accurate models that determine how the file should be split. The complication lies in the fact that the networks involved are coupled: the download starts at the

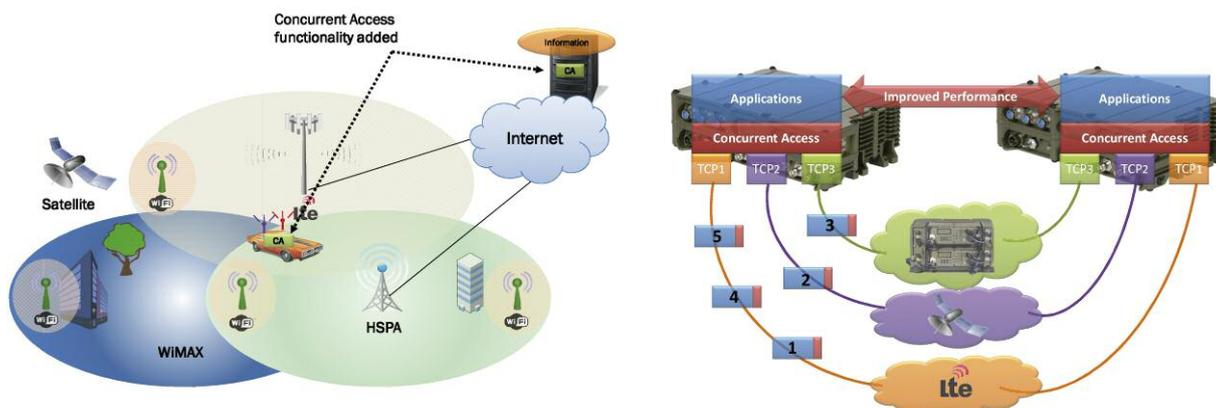


Figure 1: Users covered by multiple, overlapping wireless networks from different technologies and/or providers experience improved performance by using concurrent access.

same time in all networks, with transported data through each network depending on the original file size and the statistical (stochastic) variation in available capacity in each network.

We first developed an accurate, experimentally verified model of file downloads in a single network. This model is based on queuing; abstractions used in computer science to model waiting lines. The behaviour of multiple networks was modelled by a network consisting of several of these queuing abstractions. The performance of these

networks and the efficiency of the splitting algorithms was analysed and optimized using Processor Sharing models and optimization techniques like Markov Decision Processes and Bayesian analysis. By assuming perfect splitting of files, the researchers found a theoretical solution for this model. Experimental results subsequently showed that this theoretical solution indeed provides nearly optimal performance in practice, with a difference between modelled and practical outcomes of no more than a few percent. More details can be found in [3].

References:

- [1] TNO Monitor draadloze technologieën 2011, TNO, August 2011
- [2] OPTA, Markt cijfers tweede kwartaal 2012, 06-11-2012
- [3] G.J. Hoekstra: "Concurrent Access and Traffic Control Methods in Wireless Communication Networks", Ph.D. thesis, VU University Amsterdam, 2012 <http://kwz.me/OS>

Please contact:

Rob van der Mei, CWI, The Netherlands
Tel: +31 20 592 4129
E-mail: mei@cwi.nl

CoMobility: A Mobile Platform for Transport Sharing

by Carlos E. Cuesta, Paloma Cáceres, Belén Vela and José María Cavero

Mobile systems are becoming ubiquitous, which, in combination with wide-range service-oriented architectures, offers enormous potential for a range of uses. In this context, we have developed the CoMobility platform, a system designed to integrate carpooling and the use of public transport networks, with the goal of reducing energy consumption and CO2 emissions. CoMobility defines a service-oriented platform to help mobile users plan their use of transport, including sharing, with the purpose of saving both energy and money.

Mobile systems are not only useful in their ability to provide pervasive access to computing systems (ie enabling computer access anywhere), but they also provide the inputs of a mobile and dynamic environment into a computational system; it is now possible to perform computations that, until recently, were simply impossible. A classic example is geolocation; it is now easy to provide the physical location of a user, and to use this data for a variety of purposes.

In particular, such data can be used in the context of transportation. Transport is an important issue in modern societies, especially in big cities, and represents a significant proportion of global energy consumption. In many cases, energy consumption can be significantly reduced by an efficient use of transport media. Not only is choosing the right path within the transport network important, but also being able to share a vehicle (including private vehicles) with others. There are several initiatives assisting people in sharing transport, the most popular being carpooling: more than one person sharing a car.

However, no existing solution combines private transport sharing with the use of

public transport, which would make it more flexible. We have designed an IT platform, called CoMobility, to assist in intermodal transport sharing, integrating the use of carpooling with public transport, as well as other private transport media.

To take advantage of this approach, individuals must be convinced of the benefits of reducing the number of private cars, and of the need of a new model of transport. For this purpose,

our CoMobility platform has "customized" analytics on savings and energy consumption, to make individuals aware of the benefits of this new way of travelling. These data are obtained by combining private data from carpooling, and open data from public transport networks, and from energy-aware institutions.

Our system requires the integration of two types of data set: private data (provided by private providers and con-

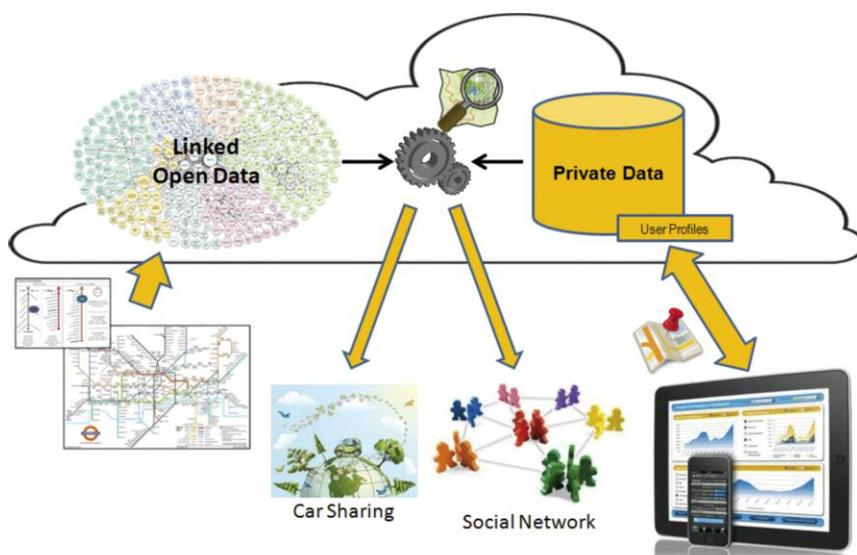


Figure 1: Conceptual architecture of the CoMobility platform

sumers of the transport system) and open data (ie data from public transport networks and public institutions). The format of public data within the open data initiatives prevents non-experts from using them directly, and thus it requires additional semantics, as provided by “Linked Open Data” [1][2] .

In summary, CoMobility provides a systematic approach to (i) accessing open, integrated and semantically annotated transportation data and street maps, (ii) combining them with private data, and (iii) supplying mechanisms to allow the actors to share and search these data. Therefore, the CoMobility conceptual architecture provides the means to perform the following tasks (also depicted in Figure 1).

First, the platform can identify, select, extract and integrate data from different and heterogeneous sources, stemming from the transportation, geographical and energy domains. Second, data from public institutions is obtained automatically in the form of open data. Third, these data are annotated as linked data, and a set of heuristics generate links between data items from different

sources without human intervention. Fourth, these data are integrated with private data provided by users themselves. And finally, CoMobility provides intuitive and customized data analytics and visualization, allowing individuals to become aware of the environmental impact of their transport choices.

The CoMobility platform is provided on the Internet “as a service”, where both public transport information and data provided by users themselves are stored and accessed “in the cloud”. The cloud approach is necessary as scalability is one of the most important requirements of this kind of wide-range service architecture. The platform needs to access a great amount of data, which is also stored in the cloud – both the private data of carpoolers, and the public data accessed in a linked open data approach. Users are able to access their information in several formats, particularly in mobile devices (currently, Android devices) and web applications. Through these devices, they are able to plan their paths in the transport network, moving from a shared car to the underground, and from there to a bus

line; and at the same time receiving an estimation of the saving of both money and energy.

With this system, we hope to encourage social and individual change towards a new - more efficient and environmentally friendly - model of transport.

Link:

VorTIC3 Research Group (Rey Juan Carlos University).
<http://www.vortic3.com>

Reference:

- [1] Linked Open Data.
<http://linkeddata.org/>
- [2]T. Heath, C. Bizer: “Linked Data: Evolving the Web into a Global Data Space”, Morgan & Claypool, 2011.

Please contact:

Carlos E. Cuesta, Paloma Cáceres, Belén Vela and José María Cavero
VorTIC3 Research Group, Rey Juan Carlos University, Madrid, Spain
E-mail: carlos.cuesta@urjc.es,
paloma.caceres@urjc.es,
belen.vela@urjc.es,
josemaria.cavero@urjc.es

Quality of Service Information System: Get to Know the Performance of Your Mobile Network Operator Anywhere-Anytime

by Katarzyna Wac

Quality of Service Information System (QoSIS) focuses on measurement-based performance evaluation of wireless access networks provided by diverse mobile network operators in diverse locations and times. We have developed an Android OS mobile application that uses measurement data provided by real mobile users living in the Geneva area to predict the networks' expected performance. Measurement data, and therefore predictions, are available for Swiss operators: Swisscom, Sunrise, Orange CH, as well as French operators: SFR, Bouygtel and Virgin.

The effectiveness of any mobile service depends on the quality of service (QoS) provided by the wireless access network it uses. However, the QoS is often unknown, as public and private wireless access network providers, mobile network operators (MNOs) for instance, tend to not disclose detailed, real-world QoS-information. For marketing purposes, these providers usually advertise only the best data rate values for their networks.

According to the 4G vision, in the near future wireless access networks of various providers - employing different wireless access technologies - will be ubiquitously available for mobile service users. Also, a seamless handover between these networks will support users' mobility. Ideally, users should have a priori knowledge about the QoS provided by different networks. Based on that knowledge, a mobile device, on behalf of its user, could make an informed choice about which wire-

less access network provider and technology to use for the preferred mobile services of the user.

To date, there are no unbiased, external providers of such information in the mobile business landscape. Users, who gain personal experience of various mobile services, may eventually use their acquired knowledge to manually reconfigure their devices and share their knowledge with family and friends. However, there is no service platform

enabling users to collaboratively share their collected knowledge. We are developing a Quality of Service Information System (QoSIS) to fill this gap. QoSIS distributes predictions to mobile devices about the QoS provided by the different wireless access networks available at a given geographical location and time. These predictions are then used by a mobile device, on behalf of the user, to choose a wireless access network provider and technology to be used. The knowledge furnished by QoSIS also allows mobile service providers to adapt their service delivery to the predicted QoS, thus increasing the service quality and improving user experience.

Mobile devices, on behalf of their users, can contribute to the QoSIS database in a collaborative information-sharing manner by submitting collected data about the QoS provided during their mobile service use, given the selected wireless access network provider and technology. The principal dimensions of the QoSIS database are: geographical location, time, wireless access network provider and wireless access technology. QoSIS is based on a QoS prediction engine. Based on machine learning algorithms, the engine builds a heuristic from which to derive predictions.

We have assessed the feasibility of deriving predictions in a case study based on delay measurement data collected from the mobile device of a health telemonitoring service user. We have shown that it is feasible to predict the value of delay for a device, based on its own measurement data or on data collected by another device being used at the same location and time and using the same wireless access network provider and technology. It is also feasible to predict the value of delay for a device, based on delay measurement data collected by both devices [1].

In contrast to existing approaches, the proposed QoSIS is entirely user-driven: mobile service users collaboratively create the QoS-information that would provide the basis for QoS prediction for other users. QoSIS therefore implements the Mobile Web 2.0 paradigm. In line with this approach, we have also assessed the business scenarios for an enterprise based on QoSIS services - provisionally named "QoSIS.net". This

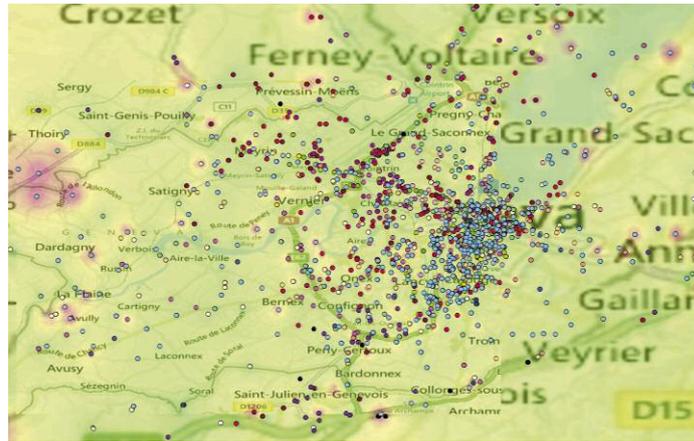


Figure 1: QoSIS service coverage in Geneva area (background colors indicate performance level, dots indicate separate network cells) (status: end of 2012)



Figure 2: Most frequent users' mobility paths for which QoSIS service is available (dots indicate separate network cells) (status: end of 2012)

enterprise could provide a QoS prediction service to its customers: mobile service providers and MNOs, as well as mobile service users. As part of the case studies, we have investigated business-to-business (B2B) and business-to-consumer (B2C) scenarios for QoSIS.net. We have shown that these scenarios are beneficial for all parties in terms of increased revenue, increases in mobile service quality and improvement of mobile user's experience [2].

The Quality of Life technologies group at University of Geneva, Institute of Services Science, is developing the QoSIS.net project. Future activities include deployment of QoSIS.net in the mHealth area – benefiting mobile patients and their caregivers by providing QoS-assurance for health telemonitoring and treatment services [3].

Links:

<http://www.qosis.com/>
<http://www.qol.unige.ch/>

References:

- [1] K. Wac, M. Hilario, B.J. van Beijnum, et al: "QoS-Predictions Service: QoS Support for Proactive Mobile Services", Wireless Network Traffic and Quality of Service Support: Trends and Standards, IGI Global Publisher, 2010
- [2] K. Wac: "Collaborative Sharing of Quality of Service-Information for Mobile Service Users", PhD thesis, Information Systems, June 2009, University of Geneva, Switzerland.
- [3] K. Wac, M. Bargh, B.J. van Beijnum, et al: "Power- and Delay-Awareness of Health Telemonitoring Services: the MobiHealth System Case Study", IEEE JSAC, Special Issue on Wireless and Pervasive Communications in Healthcare, 27(4): 525-536, IEEE Press, May 2009.

Please contact:

Katarzyna Wac
 University of Geneva, Switzerland
 Tel: +41 22 379 0242
 E-mail: Katarzyna.Wac@unige.ch

Mobile Service Platforms Based on Opportunistic Computing: The SCAMPI Project

by Marco Conti, Franca Delmastro and Andrea Passarella

Pervasive networking devices, including mobile devices, generate an environment saturated by heterogeneous hardware and software resources. Novel mobile computing paradigms allow this environment to be organized and orchestrated. "Opportunistic computing" is a new approach that allows applications to take advantage of self-organizing services built, in a dynamic way, out of the mobile resources that are available in pervasive environments.

Modern mobile devices (smartphones, tablets, etc.) bundle a number of sensing, computing and networking resources, such as cameras, microphones, wireless interfaces and memory. The environment is becoming increasingly saturated with pervasive devices (fixed cameras and sensor networks, for instance) that also feature computation, sensing and networking capabilities. As a consequence, the environment features a multitude of heterogeneous resources with dynamic availability due to factors such as the users' mobility. If orchestrated and managed through novel, appropriate computing paradigms, the availability of such resources has the potential to support innovative applications. The new mobile computing paradigm, "opportunistic computing" [1] aims to address this vision.

Opportunistic computing assumes the existence of a heterogeneous set of hardware and software resources contributed by users' devices and by the devices available in the environment. It composes and makes them available to the users' applications in a dynamic way, based on the current, situated needs of the applications and the mobility patterns of the users. Opportunistic computing is an evolution of opportunistic networking, a self-organizing networking paradigm that enables communication in dynamic pervasive networks. Each contact between mobile nodes is seen as an opportunity to forward content towards final destinations, such that end-to-end communication is supported even when simultaneous multi-hop paths between sources and destinations are not available. Opportunistic computing generalizes this approach, and views contacts between nodes as opportunities to exploit each other's resources, represented as service components. Ultimately, thanks to the dynamic composition of these components, applications can enjoy far richer

functionalities with respect to what is available in each individual node.

Developing the opportunistic computing concept and designing service platforms for future self-organizing pervasive networks is the goal of the EU SCAMPI project (Service Platform for Social Aware Mobile and Pervasive Computing [2]), which started in 2010 under the FP7 FIRE (Future Internet Research and Experimentation) initiative. The conceptual view on opportunistic computing of SCAMPI is described by the logical architecture of Figure 1.

The social layer at the bottom captures properties of users' movements and their social networks. The structure of social relationships between users can be used to derive very good predictors of mobility patterns and of communication opportunities [1]. Contact opportunities between users naturally translate into opportunities for utilization of their resources. In the intermediate layer a resource can be a CPU, a shared portion of memory, a sensor, a network connec-

tion, a piece of content, a function implemented as a piece of code, etc. A link between two resources in this layer represents the fact that these resources have a probability of "encountering" each other through opportunistic contacts, typically as a side effect of users' mobility. The opportunistic service layer (OSL) is responsible for abstracting resources into service components. It provides a functional description of resources, their limits and capabilities, and is responsible for orchestrating the composition of multiple components into a single service, as required by the applications. Note that this may require passing input/output parameters between different nodes providing different components, which is achieved through opportunistic networking techniques.

One issue that we are investigating in SCAMPI is how information collected dynamically on users' devices can be used to describe their context, and ultimately to optimize service provisioning in opportunistic networks. To this end, we designed and implemented

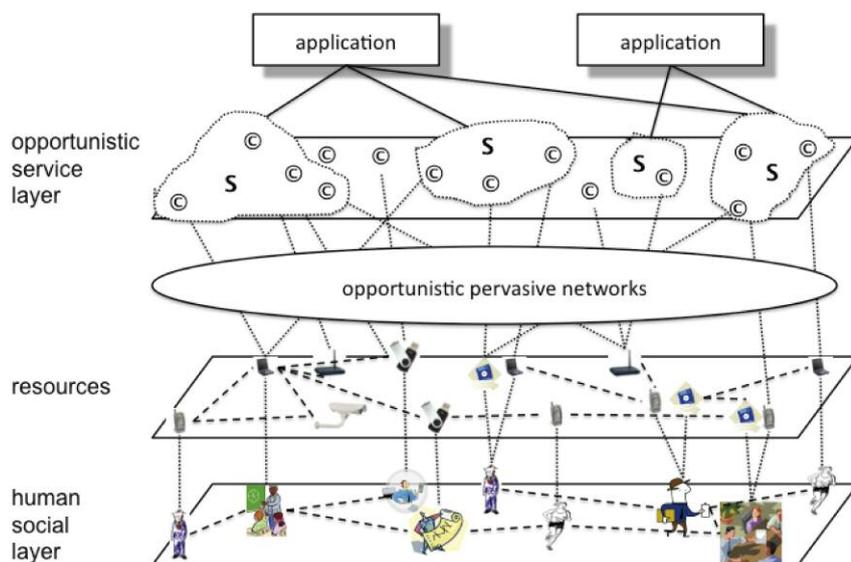


Figure 1: Logical architecture of opportunistic computing

CAMEO, a middleware platform focused on the management and elaboration of context information for opportunistic computing environments [3]. The architecture of CAMEO is described in Figure 2, and is composed of two main building blocks. The Local Resource Management Framework (LRM-Fw), aims at implementing features related strictly to the interaction with the local resources of the device, both hardware (eg, embedded sensors, capacity, battery, wireless interfaces) and software (eg, communication primitives and programming libraries). The Context-Aware Framework (CA-Fw), aims at storing, elaborating and disseminating all the context information, and gathering a view on the resources available on other devices, which can then be composed according to the applications' needs.

In order to test the functionality of CAMEO with real applications, we have developed a Tourist-MSN (Mobile Social Networking) application [3], designed to improve the tourist experience, for example, during a visit to a city, by collecting and sharing useful information and content possibly dynamically elaborated and enriched by other users themselves. Figure 3 provides some snapshots of the application interface. Let us consider, for instance, a couple visiting Rome; before leaving, they plan to visit several attractions in few days trying to optimize their time. While they are moving around the city they encounter other people that have just visited some of those attractions they are interested in. Such users may provide useful information and/or multimedia content related to the attractions they are visiting. Tourist-MSN can exploit the opportunistic computing functionalities to dynamically reschedule the tourists' visit based on the current conditions (eg, expected waiting times at queues). Moreover, it can allow its users to share each other's resources in order to cooperatively generate multimedia content starting from elements (pictures, clips, video, etc.) they generate during their visit. More generally, through CAMEO, applications like Tourist-MSN can provide the users with functionalities such as (i) identifying users in the social context interested in a specific content, post or discussion; (ii) disseminating selected contents to interested users; (iii) generating ratings of

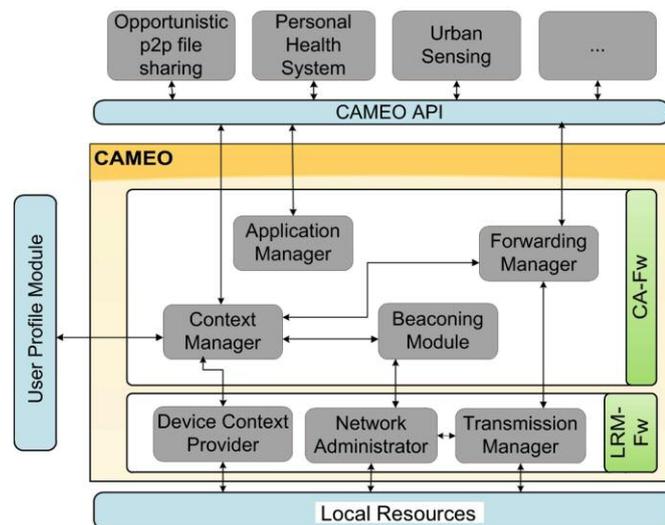


Figure 2: The CAMEO software architecture.



Figure 3: Screenshot of Tourist-MSN application exploiting the CAMEO middleware.

available contents depending on the local user's interests; (iv) establishing discussion forums with other users; (v) cooperatively annotating content and enriching it thanks to multimedia editing functionalities contributed by devices available in the environment.

Link: <http://www.ict-scampi.eu>

References:

- [1] M. Conti et al: "From Opportunistic Networks to Opportunistic Computing", IEEE Communications Magazine, Vol. 48, Issue 9, 2010
- [2] M. Pitkänen et al: "SCAMPI: service platform for social aware mobile and pervasive computing", Computer Communication Review 42(4): 503-508 (2012)
- [3] V. Arnaboldi, M. Conti, F. Delmastro: "Implementation of CAMEO: a Context-Aware Middleware for Opportunistic Mobile Social Networks", IEEE WoWMoM 2011, 20-23 June 2011.

Please contact:

Marco Conti
IIT-CNR, Italy
Tel: +39 050 315 3062
E-mail: marco.conti@iit.cnr.it

Franca Delmastro
IIT-CNR, Italy
Tel: +39 050 315 2405
E-mail: franca.delmastro@iit.cnr.it

Andrea Passarella
IIT-CNR, Italy
Tel: +39 050 315 3269
E-mail: andrea.passarella@iit.cnr.it

TravelDashboard - a Framework for the Delivery of Personalized Mobility Services to Urban Travellers

by Licia Capra, Pierre Chatel, Animesh Pathak and Roberto Speicys Cardoso

European researchers and companies join hands to address the woes of the urban traveller.

With over 70% of the world's population expected to be living in cities by 2050, the support of citizens' mobility within the urban environment is a priority for municipalities worldwide [1]. Although public multi-modal transit systems are necessary to better manage mobility, they alone are not sufficient. Citizens must be offered personalized travel information to make their journeys more efficient and enjoyable. Notably, such information should not only be objective (eg, bus timetable, live bus tracking), but crucially personalized – since every passenger's preferences and interests differ [2] (eg, crowdedness of trains, heat of tube platforms, sociability of the coaches).

Goal and objectives

Owing to the recent proliferation of smartphones, the relatively new field of mobile participatory sensing [3] could be leveraged towards providing a more fine-grained and up-to-date view of a city's transportation system. A pan-European team including partners from Alcatel/Lucent, Ambientic, Inria, Systematic, Thales, Transport for London (TfL) and University College London (UCL), have pooled their

resources for TravelDashboard; a project to produce an open source middleware platform, enriched with personalized mobility services for urban travellers, evaluated via real-life demonstrators' assessments, and accompanied by novel business models. The partners will tackle:

1. Development of a self-adaptive data collection middleware, that gathers both passive sensory information (eg, bus location from GPS) and active user-generated content (eg, road hazard reports, crowdedness of journey, etc.), and disseminates transport network status updates to travellers.
2. Development of personalized mobility services, including: user preferences about various aspects of mobility (eg, punctuality, crowdedness, and sociability of transport systems); personalized information aggregators, to combine the data streams that are of interest to a user in their current context.
3. Development of a demonstrator for real-life assessment, and deployment with actual end-users.
4. Elicitation of viable Business Models, based on the services whose rapid

development and deployment are supported by the middleware.

The middleware core is being worked on by Inria, Ambientic, and Thales; the planned added-value services are being developed by UCL and Alcatel; the demonstrator will be built and deployed under the experience of Ambientic, thanks also to key input from TfL; finally, Systematic brings the know-how to elicit business models for innovative technological services, for Ambientic to deliver new urban traveller services.

From research to innovation

This 12 month project (January to December, 2013) is funded by the European Institute of Innovation and Technology (EIT)'s "ICT Labs" knowledge and innovation community (KIC) under the "Intelligent Mobility and Transportation Systems" action line. The core goal of EIT ICT labs is to spur innovation in European ICT by leveraging results from ongoing research projects – as presented below for TravelDashboard. This activity expands on them by conducting domain-focused research, concretized through demonstrators, and by developing business models specifically tailored to the urban mobile participatory sensing domain. More precisely:

- Inria leverages its work on large scale mobile participatory sensing as part of the EC FP7 CHOReOS project on choreographies for the future internet, applying it to the urban transport domain, while Thales expands on the transportation-related use cases that it is developing as part of CHOReOS as well as the French ANR SocEDA project on large distributed social Event Driven Architecture (EDA) platform, focusing on user mobility and mobile device embedding aspects.
- UCL builds upon the recommender system it has been developing as part of the FP7 i-Tour project on intelligent urban transport systems, adding qualitative metrics such as crowded-

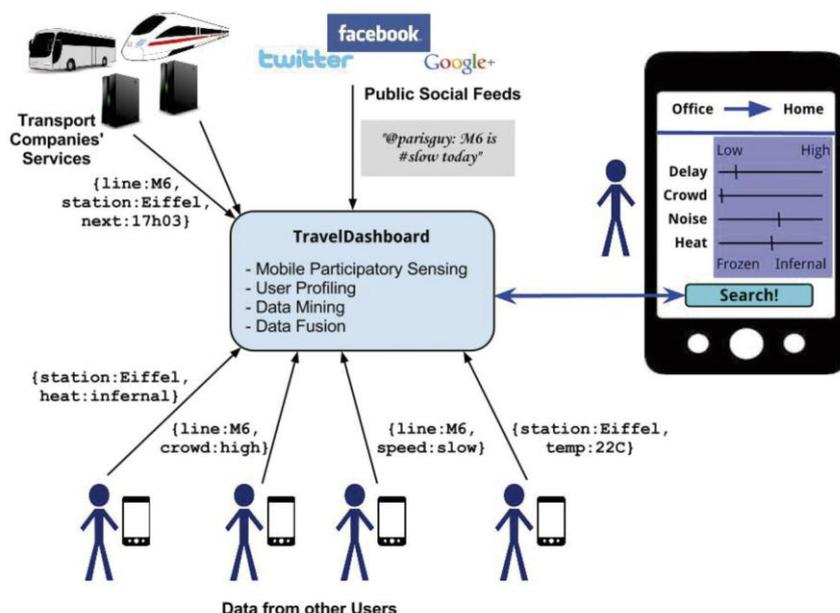


Figure 1: The TravelDashboard vision of urban transport

- ness to its information filtering and data mining services.
- Alcatel-Lucent is involved in an ongoing project on Human Mobility Pattern Mining, which combines different sources of user mobility traces to understand how communities are formed. In TravelDashboard, their researchers perform an empirical study to understand mobility patterns and characteristics (time, mode, route and price) of different user groups, and identify contextual factors that dictate user mobility in order to improve their user profiling system.
 - Founded on its work on multi-platform multimedia streaming service in the FP7 CONNECT project, the SME Ambientic leads the work on enabling access to data streams generated by mobile urban users, as well as the development of a mobile application for urban transport.

To enable the transfer from research to innovation, Systematic analyzes the business value of the TravelDashboard solution and identifies relevant business

models and strategies for the development and provisioning of personalized mobility services in the urban crowd-sourcing context.

Impact: A Smart City for Smart Travellers

The TravelDashboard framework will be realized by adhering to the principles of Service-Oriented Architectures. At its core is the CHOReOS service-oriented middleware, atop which each partner is then responsible for the development of its own services. This core will be released as open-source software, thus promoting the rapid and easy development of applications in support of highly personalized urban travel experiences. Both municipalities and transport operators can use this platform to build smart applications to help citizens navigate their cities. One demonstrator application will be built and deployed during this project, catering to the London public transport infrastructure, contributing to the development of business models and strategies for the provisioning of personalized urban mobility services.

Link:

TravelDashboard project: <http://www-roc.inria.fr/arles/traveldashboard>

References:

- [1] P. Bosnell: "Human Behaviours to Moving People More Intelligently", 2006
- [2] Rail Safety and Standards Board Ltd: "Topic Note on Travel Behaviour and Behavioural Change", 2010
- [3] N. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. Campbell: "A survey of mobile phone sensing", *Communications Magazine*, IEEE, 2010.

Please contact:

Pierre Chatel
Thales Group, France
Tel: +33 169415565
E-mail: pierre.chatel@thalesgroup.com

APISENSE: Crowd-Sensing Made Easy

by Nicolas Haderer, Romain Rouvoy, Christophe Ribeiro and Lionel Seinturier

The rapid emergence of mobile devices, such as TabletPC and smartphones, equipped with a rich array of sensors, enables a new means of acquiring sensor data, known as crowd-sensing. Crowd-sensing is currently receiving a lot of attention, not only from industry but also from various research communities interested in collecting a new class of data over a much larger population than was previously possible.

"Crowd-sensing" refers to the involvement of a large, diffuse group of participants in the task of retrieving reliable data from the field. This approach has been used in multiple research studies, including traffic and road monitoring, social networking and environmental monitoring.

However, developing a dedicated sensing application to collect a specific dataset over a target population presents a real challenge [1]. Indeed, when building an efficient application involving a large community of participants, a number of vital issues must be taken into account, including incentive mechanisms, recruitment models, energy limitations and privacy concerns. A thorough expertise in the area of mobile device technology is essential to address these issues. Consequently, it

can be difficult for scientists, inexperienced in this field, to collect realistic datasets for their studies. But more importantly, the ad hoc applications that are developed may neglect privacy and security concerns, resulting in the disclosure of sensitive user information, or even the adoption of unethical or illegal approaches.

The state-of-the-art therefore lacks reusable solutions for safely collecting and exploiting crowd activity traces, which are usually difficult to set up and are tied to specific data representations and device configurations.

In this context, the Inria ADAM project team has developed APISENSE [2], a platform targeting multiple research communities, and providing a lightweight solution to build and deploy

crowd-sensing applications for collecting experimental datasets.

The APISENSE Platform

The APISENSE platform distinguishes between two roles. The first, called scientist, is typically a researcher who wants to define and deploy an experiment over a large population of mobile users. The platform therefore provides an online environment (as a software as a service) allowing the researcher to describe experimental requirements via a scripting language, to deploy the experiment scripts over a subset of participants and to connect other services to the platform in order to extract and process the collected datasets (eg, visualization, analysis, replay). Technically, the server-side infrastructure of APISENSE is built on the principles of Cloud computing in order to offer a

```

location.onLocationChanged
distance : "10 m"
time : "5 min"
(event) -> {
  longitude : event.longitude
  latitude : event.latitude
  signal_strength : gsm.dbm()
}

```

Figure 1: Description of a sensing task (in CoffeeScript)

modular service-oriented architecture, which can be customized to suit the scientist's requirements. The second role is the mobile phone user, identified as a participant. The APISENSE platform provides the participant with a mobile application allowing experiments to be downloaded and executed in a dedicated sandbox, and the collected datasets to be automatically uploaded on the APISENSE server.

The APISENSE Language

To facilitate its adoption by scientists, APISENSE adopts standard scripting languages and provides an extension of the JavaScript, CoffeeScript and Python languages, as an efficient means of describing an experiment without any specific knowledge of mobile device programming models (eg, Android SDK).

Figure 1 shows how simple it is to define a sensing task to collect a specific dataset in the wild and automatically build an open data map from collected data. The sensing task is defined using CoffeeScript language, which is triggered whenever the location of a participant changes by a distance of 10 m in a period of five minutes. When these conditions are met, the script automatically builds a trace containing the location of the participant and attaches GSM networks characteristics. The sensing task can be defined by the scientist using a web interface and a single click makes it available to participants. Figure 2 depicts the resulting open map

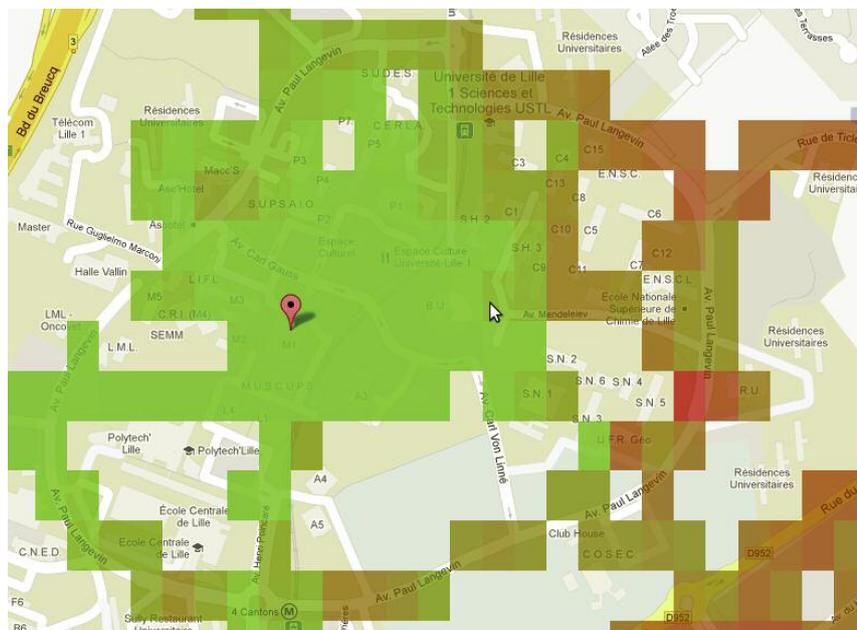


Figure 2: GSM open map generated from collected datasets

generated from datasets collected by only three participants and displaying the GSM signal strength of a given network operator in the area of the University Lille 1.

The scripting library supports a wide range of features to define data collected during a sensing experiment, including traditional sensors proposed by smartphone technology, such as GPS, Bluetooth, accelerometer, compass, phone call, SMS, application status (installed, running), for opportunistic sensing activities and also a graphical user interaction framework to describe user surveys in the case of participatory sensing activities.

Ongoing Work

Our current work focuses on the deployment of various large scale sensing experiments for building a collaborative method i) to identify Android malware in the wild and ii) to deploy micro-seismic mobile stations for earthquake monitoring.

Additionally, APISENSE collaborates with the MetroScope initiative, supported by Inria, to provide a new generation of instruments for observing Internet usage. By providing tools, datasets, and analyses that make sense for different disciplines, APISENSE not only contributes to the theoretical knowledge of the Internet, but also introduces a greater level of transparency by involving citizens and giving them direct feedback.

Links

APISENSE website:
<http://apisense.fr>
 ADAM project-team:
<http://adam.lille.inria.fr>
 MetroScope project:
<http://metroscope.eu>

References:

- [1] N.D. Lane et al: "A survey of mobile phone sensing", Communications Magazine 48(9), IEEE. September 2010
- [2] N. Haderer et al: "A preliminary investigation of user incentives to leverage crowd-sensing activities", in 2nd International IEEE Workshop on Hot Topics in Pervasive Computing (PerHot), March 2013.

Please contact:

Romain Rouvoy
 ADAM project-team,
 University Lille 1 - LIFL,
 Inria Lille - Nord Europe, France
 Tel : +33 3 59 35 87 77
 E-mail: Romain.Rouvoy@inria.fr

Mobilitics: Analyzing Privacy Leaks in Smartphones

by Jagdish Prasad Achara, Franck Baudot, Claude Castelluccia, Geoffrey Delcroix and Vincent Roca

Who, do you think, is aware of almost everything you do? Well, it's probably right there in your pocket, if you own a smartphone and carry it with you. In order to evaluate the actual privacy risks of smartphones and to raise public awareness of these risks, the CNIL (French data protection authority) and the Inria (French public science and technology institution dedicated to computational sciences) Privatics team started working together in 2012 as part of the Mobilitics project.

It is no surprise, given smartphones' convenience and utility, to see their wide adoption worldwide. Today, there are 1.08 billion smartphone users out of a total of five billion mobile phone users worldwide, and the ratio is constantly increasing. Smartphones are used not only to communicate but also to browse the web and run various internet-enabled Apps. As a result, they contain a lot of information about the cyber activities of their owners, and therefore users' interests and behaviours. Furthermore, smartphones are also equipped with GPS, NFC and Bluetooth units, along with a digital camera and are almost always connected to the Internet; thereby revealing a lot of information about the physical activities of their owners. On top of this, smartphones are very personal to the user and are barely turned off.

For the aforementioned reasons, combined with the fact that users tend to carry smartphones wherever they go, they are an ideal target for marketers who want to profile users to profit from their personal data. Some studies even suggest that the main business model for some developers (in the case of free Apps, for example) is based on the collection of personal data. As a result, many Apps might be leaking personal information to third parties, such as Analytics and Advertising (A&A) companies.

A few insights

The goals of the Mobilitics project are to investigate smartphone Operating Systems (OSs) and Apps for potential privacy leaks and to inform their users about the privacy risks. The project currently targets two OSs, namely Android and iOS, because they cover almost 75% of the whole smartphone OS market share.

As part of this project, we have developed a software solution (an Android version with similar functionalities is under development) for iOS to capture access to private information by various Apps. When an App makes a call to the iOS API to access a broad list of a user's personal data, eg Contacts, Location, Device Name, UDID, Calendar, Reminders, Photos, Notes and Accounts, our software logs this event for later analysis. Note that some Apps do actually need to access personal data to provide the desired service. These applications do not breach user privacy if they only process

and use the personal data to provide the desired service and don't transmit the data to remote third parties. In order to detect personal information leakage, we therefore also monitor whether the accessed personal data is sent to a third-party, perhaps without the user permission. Additionally, we are also developing a visualization tool to help people understand the privacy implications by aggregating, interpreting and displaying all private data stored and/or sent by various Apps.

For instance, our iOS tool reveals that many Apps are accessing the Unique Device ID multiple times (in the order of hundreds), which implies that it is probably being used for online tracking of the user. Some Apps are also, surprisingly, accessing the user's device name although there does not seem to be any obvious reason to do so. The name of the device is set during the initial device setup and often contains the real name of the user. Moreover, even if the user does not set it to his or her real name, it might easily be used for tracking purposes since the device owner does not generally modify it after the initial setup.

Our software solution was developed in 2012 for iOS 5.x before Apple launched iOS 6 in September 2012. iOS 5.x didn't seek the user's permission for private data access except for location information. In iOS 6,

Apple decided to change its strategy and introduced a new privacy-specific setting giving the user control over whether an App can access private data: the user is prompted the first time an App tries to access Contacts, Location, Reminders, Photos, Calendar and Social Networking accounts and later,

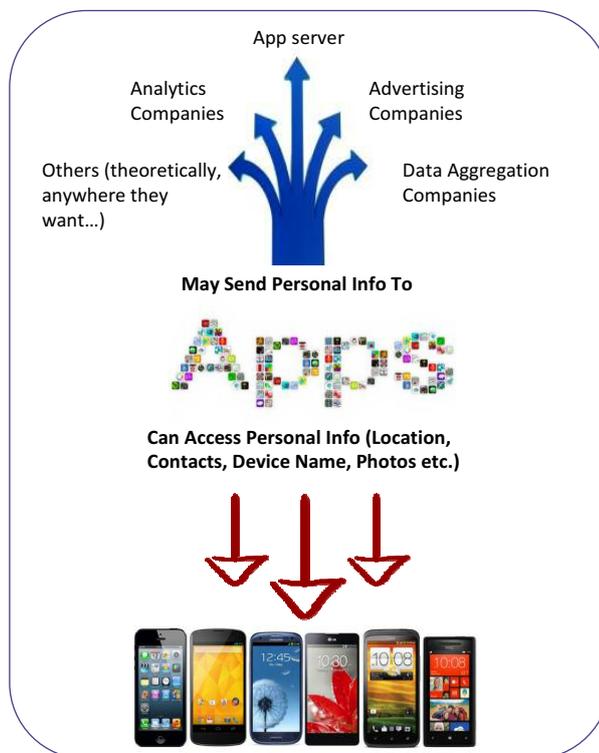


Figure 1: Android and iOS currently don't provide any mechanism to let users know how their personal information is being used by various Apps. Will it be used locally on the device or sent to remote servers? Being aware of it, users can probably make better decision whether to allow/deny access to their personal information for a particular App.

iOS remembers and follows the user preferences.

In our opinion, this is a decent step by Apple towards making iOS privacy-friendly. However, several questions still remain open: is the list of private data included in their privacy-settings sufficient? Is an authorization that does not consider any behavioural analysis sufficient? For instance, accessing the device location upon App installation, to enable a per-country personalization, is not comparable to accessing the location every five minutes. Also, does the App keep the personal information locally for internal purposes, or is it communicated to external servers? If the latter, where exactly are these servers? Moreover, A&A libraries included by the App developer also have access to the same set of user's private data as the App itself. However, a user giving access to his or her Contacts doesn't necessarily indicate

consent for these data to be shared with A&A companies. Might this pave the path for privacy invasion? These are the questions that Mobilities will attempt to answer.

Conclusion

Our preliminary results and the various scandals that occurred in 2012 show that privacy considerations are of utmost importance if we want to continue using these devices with serenity. We believe that smartphones can't be, in the long run, black boxes to their owners because nobody wants these great devices in our pockets to be the ultimate spy.

Mobilities is a CNIL-Inria project that involved the following participants: Jagdish Prasad Achara, Franck Baudot, Claude Castelluccia, Geoffrey Delcroix, James Douglas Lefruit, Gwendal Le Grand, Stéphane Petitcolas, and Vincent Roca

Links:

<https://team.inria.fr/privatics/mobilities/>
<http://www.go-gulf.com/blog/smartphone/>
<http://kwz.me/04>
<http://blogs.wsj.com/wtk-mobile/>
<http://kwz.me/00>

References:

- [1] W. Enck et al: "TaintDroid: An Information-Flow Tracking System for Realtime Privacy Monitoring on Smartphones", in proc. of OSDI 2010
- [2] M. Egele et al: "PiOS: Detecting Privacy Leaks in iOS Applications", in proc. of NDSS 2011.

Please contact:

Jagdish Prasad Achara, Claude Castelluccia, Vincent Roca
Inria, France
Tel : +33476615215
E-mail:
{Jagdish.Achara,Claude.Castelluccia, Vincent.Roca}@inria.fr

Privacy-Preserving Interest-Cast for Android Smartphones

by Gianpiero Costantino, Fabio Martinelli and Paolo Santi

We present an implementation of the FairPlay framework for secure two-party function computation on Android smartphones, which we call MobileFairPlay. Our application was developed to preserve the users' privacy within opportunistic networks considering the interest-casting model. Our tests show that the running times of the protocol on several Android phones, are very reasonable (up to five seconds in the worst case).

Crowded places present an opportunity for people to share personal information. In addition to sharing information through traditional, web-based platforms and applications such as Facebook and Twitter, the availability of short range radio interfaces in smartphones, tablet PCs, etc. allows individuals to share information with one another through direct, opportunistic communication (typically using the Bluetooth or WiFi interface).

This model of store-carry and forwarding data to others is known as opportunistic networking (OppNets). A common feature of these approaches is that, before making a decision about whether to share information with an individual, users have to exchange some sensitive information, such as history of past encounters [1], interest profiles, etc.

Given that the person encountered is generally a stranger, this exchange of sensitive information (which occurs in plain text in the approaches mentioned) is likely to be deemed unacceptable by the user in real-world scenarios, owing to privacy concerns.

To address this issue, we present a feasible implementation of a cryptographic framework for secure multi-party computation (the FairPlay framework proposed in [2]) targeted to the interest-cast model and running on the Android mobile platform. Our application, "Mobile-FairPlay" [3], has been developed with the aims of: 1) finding people in the user's (Alice) neighbourhood through a Bluetooth scan operation, 2) connecting to another user (Bob) and determining whether Bob and Alice have similar interest profiles without

disclosing sensitive information, and 3) sharing messages between Alice's and Bob's devices in the event that their profiles are similar.

When Alice and Bob have established a new connection, Bob, who received the connection, randomly selects different topics to verify their similarity with respect to these interests. Then, they start matching interests using the secure framework implemented in the App. During this execution, both Bob and Alice use their own value for the selected topic, extracted from the interest profile. However, these values are not sent to the other participants in plain, but are encoded in the garbled Boolean circuits exchanged through MobileFairPlay. At the end of the handshaking phase, Alice and Bob only know the result of the jointly computed

matching, without knowing the specific interest values of the other party. In the case of a positive comparison, the user who received the incoming connection can start sending their own information using real files.

The execution time of the interest matching performed with our application ranges between 2 to 2.5 seconds for a single topic comparison and between 3.5 to 4.5 seconds for four topic comparisons. The running time is dominated by the cryptographic part run during the matching, and depends mainly on the smartphone's hardware. Times reported here are obtained using Samsung Galaxy S2 and Samsung Galaxy S and other slower devices. However, by running more recent smartphones the computation time can be even lower.

We are currently developing a different version of this application that uses past users' locations to determine whether they have spent time in common places.

This matching is also performed whilst preserving the users' privacy. Common past locations are used to understand similarity in users' behaviour, and taking advantage of this can help users to increase their social contacts.

The application was released late spring 2012, and is available at the link below.

Acknowledgment

It took over six months to develop this application, and our thanks go to Dario Amoroso, who prepared his master degree thesis on this topic.

Link:

<http://www.iit.cnr.it/staff/gianpiero.costantino/>

References:

- [1] E. Daly and M. Haahr: "Social network analysis for routing in disconnected delay-tolerant manets", in ACM Mo- biHoc
- [2] D. Malkhi et al: "Fairplay: a secure two-party computation system", in proc. of the 13th conf. on USENIX Sec. Symp.
- [3] G. Costantino et al: "An implementation of secure two-party computation for smartphones with application to privacy-preserving interest-cast", PST 2012.

Please contact:

Gianpiero Costantino,
Fabio Martinelli and Paolo Santi
IIT-CNR, Pisa, Italy
E-mail: gianpiero.costantino@iit.cnr.it,
fabio.martinelli@iit.cnr.it,
paolo.santi@iit.cnr.it

Market-Based Security for Mobile Devices

by Gabriele Costa, Alessio Merlo and Luca Verderame

We present a security-enabled application marketplace that provides formal security guarantees to the existing mobile software distribution paradigm. Our proposal allows users and organizations to apply fine-grained security policies on top of the existing market-based software deployment with no need for invasive customization of devices and without compromising the system scalability and usability.

Most modern mobile Operating Systems (OSs) rely on application marketplaces, also called app stores, to allow users to download and install software packages that extend their devices with new functionalities. The success of this software distribution paradigm is confirmed by the gargantuan number of mobile applications released and installed every day.

Despite the enormous advantages in terms of customizability, access to many third party software packages raises several security issues. One major issue is how to provide fine grained, strong security guarantees without compromising usability and scalability. Most mobile OSs don't apply security-relevant settings automatically, critical choices about security aspects have to be made by the user. For instance, Android requires that users accept the application contract, or "manifest", at

installation time. Also, the marketplace collects positive and negative feedback that serves as a reputation system that customers can consider before installing new software.

Nevertheless, these mechanisms do not offer proper protection to users. Indeed, manifest descriptions are extremely coarse grained, eg, "this app will use the filesystem", and even expert users struggle to find a relationship between a manifest and their security needs. Recent studies have demonstrated that users generally disregard security-relevant warnings upon installation of applications, thus ignoring potential security risks.

Furthermore, once a user has installed an application, it can run freely and independently of the context, ie the device configuration, under which it is used. For instance, current mobile OSs

have no mechanism to prevent a user from playing games or executing applications that may spread sensitive information during business hours. Consequently, organizations that could benefit from their employees having a smartphone, eg, because the company could avoid purchasing multiple specific devices, may be reluctant to use them.

To cope with these issues, we recently proposed (see [1,2]) a novel approach that provides formal security guarantees to the users of mobile devices. Furthermore, the system is built on top of the existing mobile applications distribution paradigm, which makes the whole system scalable and reusable.

In practice, a mobile application marketplace works as follows: Owners of mobile devices register to a marketplace by means of access credentials.

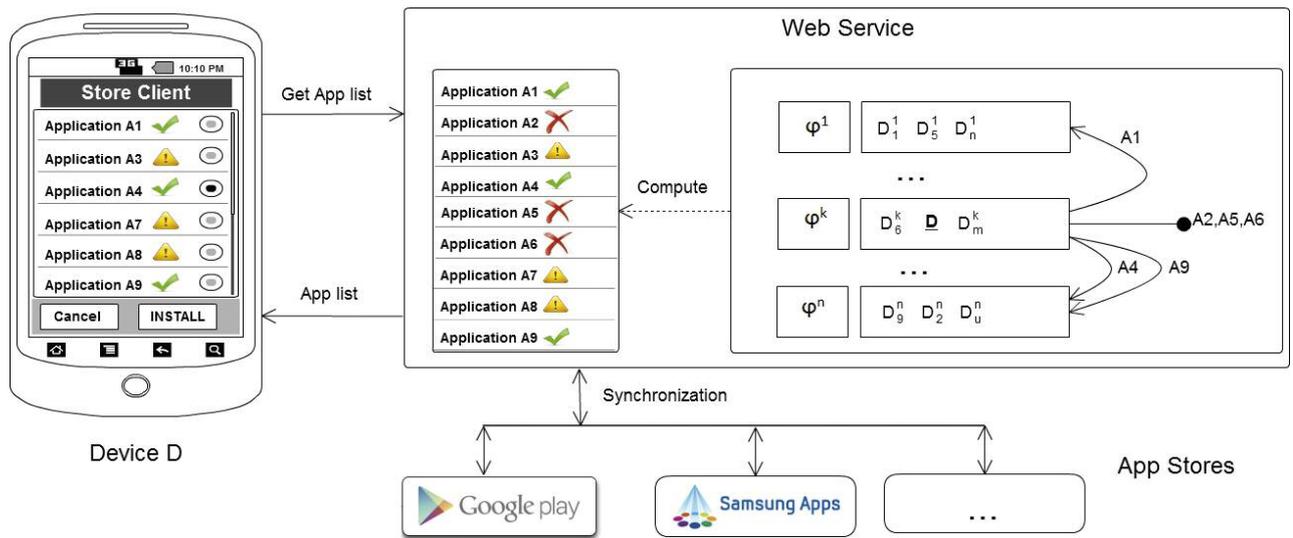


Figure 1: Architecture of the secure marketplace.

The marketplace stores device information and configuration, eg, hardware profile, OS version and installed applications. When the user decides to install a new application, the marketplace checks the compatibility with the device configuration and requests the user to provide the application with the requested access rights. If the user disagrees, the installation is cancelled, otherwise the marketplace sends the software package and the device installs it.

We propose the use of a security enabled marketplace that acts as a security proxy, providing a formal security guarantee. A secure marketplace exploits a security policies management system and verifies whether a certain application can be installed on a device without affecting the security configuration. Application packages are analyzed to extract security contracts that, in turn, extend and enrich the manifest. Contracts are safe representations of the application’s behaviour, and denote all the possible sequences of system access operations that the software can perform (including possible synchronizations with other installed applications).

The secure marketplace then applies a model checking procedure to verify whether the device configuration composed with the application contract is still policy-compliant. If it is, the application is labeled as safe. Otherwise, the

application can be modified by instrumenting its code with security checks, guaranteeing that it respects the policy. In both cases, the secure marketplace generates a safety proof that the user can verify, through automatic proof checking, before installation.

After installation confirmation, the secure marketplace updates the policy management system with the new security state of the device.

Our proposed system could work with most of the common usage contexts for mobile devices. Indeed, it applies equally to both private customers, ie, users wanting to apply security controls to their own devices, and companies, ie, organizations in which employees and affiliated persons must respect precise security restrictions.

We are currently developing a real implementation of our secure marketplace that customers can use to mediate the access to standard marketplaces, eg, Google play.

We plan to release a secure marketplace implementation for Android OS in mid-2013. Moreover, we are currently investigating potential extensions of our proposal; including the adoption of security monitoring to allow the user to install applications that statically violate the policy. In this way, application executions are guaranteed to respect the security policy.

References:

- [1] A. Armando, G. Costa, A. Merlo: “Formal Modeling and Reasoning about the Android Security Framework”, in proc. of TGC 2012, Newcastle upon Tyne, UK, 2012 <http://www.ai-lab.it/merlo/publications/AndroidModel.pdf>
- [2] A. Armando, G. Costa, A. Merlo, L. Verderame: “Bring Your Own Device, Securely”, in proc. of SAC 2013, Coimbra, Portugal, 2013 <http://www.ai-lab.it/merlo/publications/BYODroid-extended.pdf>

Please contact:

Gabriele Costa
 University of Genova, Italy
 Tel: +39 0103536545
 E-mail: gabriele.costa@unige.it

Alessio Merlo
 E-Campus University, Italy
 Tel: +39 0103532344
 E-mail: alessio.merlo@unicampus.it

Revealing Social Links Between Owners of Wi-Fi Enabled Smartphones

by Mathieu Cunche, Mohamed Ali Kaafar and Roksana Boreli

Wi-Fi technology, available in the vast majority of mobile phones, tablets, laptops and other computing devices that we use in our daily lives, has enabled widespread use of new applications and services. This technology, however, has a number of issues related to privacy loss, exacerbated by its ubiquity. Our research shows how the information freely transmitted by the Wi-Fi protocol can be used to identify links between people, ie whether they are family, friends, colleagues etc.

Wi-Fi protocol includes potential sources of personal information leaks. Wi-Fi enabled devices commonly use active discovery mode to find the available Wi-Fi access points (APs). This mechanism includes broadcast of the Wi-Fi network's names to which the mobile device has previously been connected, in plain text, which may be easily observed and captured by any Wi-Fi device monitoring the control traffic. The combination of the network names associated with any single mobile device can be considered as a Wi-Fi fingerprint which can be used to identify the user to whom the mobile device belongs. Our research [1] investigates how it is possible to exploit these Wi-Fi fingerprints to identify links between users, ie owners of the mobile devices broadcasting such network names.

From June to October 2012, we collected data using a laptop running Wi-Fi monitoring tools in the city of Sydney. Overall, Wi-Fi fingerprints of more than 8,000 devices were collected and we found that some devices were revealing their associations to more than 80 Wi-Fi networks. In order to test our hypothesis (social links can be inferred based on the Wi-Fi fingerprint) we collected the Wi-Fi fingerprint, as well as the existing social links, from a group of volunteers.

Identifying linked individuals

Our approach is based on the similarity between Wi-Fi fingerprints, which is equated to the likelihood of the corresponding users being linked. When computing the similarity between two Wi-Fi fingerprints, two dimensions need to be considered:

- The number of network names in common. Indeed, sharing a network is an indication of the existence of a link, eg friends and family that share multiple Wi-Fi networks.



Figure 1: A smartphone broadcasting in plaintext its Wi-Fi fingerprint, ie names of networks to which it has previously been connected.

- The rarity of the network names in common. Some network names are very common and sharing them does not imply a link between the users. This is the case for public network names, for instance McDonalds Free Wi-Fi, or default network names such as NETGEAR and linksys. On the other hand, uncommon network names such as Griffin Family Network or Orange-3EF50 are likely to be associated with a strong link between the users of these networks.

Utilizing a carefully designed similarity metric, we have been able to infer the existence of social links with a high confidence: 80% of the links were detected with an error rate of 7%.

Who should worry about it?

Owners of smartphones are particularly exposed to this threat, as these devices are carried on persons throughout the day, connecting to multiple Wi-Fi networks and also broadcasting their connection history.

What can be done to prevent the linking?

There are a number of industry and research initiatives aiming to address the Wi-Fi related privacy issues. The

deployment of new technology, ie privacy preserving discovery services, would necessitate software modifications in currently deployed APs and devices. The obvious solution - to disable active discovery mode - comes at the expense of performance and usability, ie it would take longer for the Wi-Fi capable device to find and connect to an available AP. As a possible first step, users should be encouraged to remove obsolete connection history entries, which may lower the similarity metric and thus reduce the ease of linkage.

Links:

<http://www.inria.fr/en/teams/privatics>
http://www.nicta.com.au/research/projects/trusted_networking

Reference:

[1] M. Cunche, M.A. Kaafar, and R. Boreli: "I know who you will meet this evening! linking wireless devices using wi-fi probe requests", WoWMoM 2012 IEEE International Symposium, p. 1-9, June 2012.

Please contact:

Mathieu Cunche
 INSA-Lyon and Inria, France

Secure Collaboration for Smartphones

by Abdessamad Imine and Michaël Rusinowitch

Designing secure collaborative applications has become a hot topic in the area of mobile devices, posing challenging problems such as data management and security.

Smartphones (or mobile computing devices) are proliferating globally. According to a study by comScore [1], by December 2012, more than 125 million people in the USA alone owned smartphones. These mobile devices provide a variety of networking options such as GSM, GPRS, Bluetooth and Wi-Fi. Therefore, smartphones are becoming the device of choice for people to collaborate with family members, friends and business colleagues/customers. They enable us to communicate not only by telephone, by email and short messages, to play games, to share information, organize videoconferences, and coordinate business events.

The adoption of collaboration via smartphones relies on critical factors such as network coverage and security concerns. Collaboration may involve mobile users located in the same limited geographical area, yet connected via wireless channels, which can be less reliable and offer lower data rates than traditional wired ones. Moreover, the amount of personal/corporate data stored in mobile devices has increased. Sharing such data requires preventive security and privacy mechanisms to protect and regulate access to this shared data by other users.

The CASSIS project team at Inria Nancy Grand-Est has designed and developed a decentralized and secure shared calendar, which can be easily deployed on ad-hoc networks of smartphones [2]. Unlike shared calendars such as Google Calendar, our collaborative application ensures privacy and confidentiality of calendar events since it is independent of any third-party or big company servers.

The calendar enables users to share their calendar events with selected users in a dynamic group (users can join or leave the group at any time), and is as responsive as a personal calendar. This high responsiveness is achieved by replicating the shared calendar at each participating user. The consistency of these replicas is maintained in a decentralized way (without a central coordination point) using Operational Transformation (OT) approach.

To prevent unauthorized access by illegal users, our shared calendar is endowed with a flexible access control mechanism where each user can define an authorization policy for controlling access to its proper calendar events. Furthermore, all authorization policies are replicated at the local memory of

each user in order to ensure low latency for access rights checking. Thus, a user will own two copies: the shared calendar and the access data-structure. It is clear that this replication allows for high data availability since when users want to read or update the shared calendar, this manipulation will be granted or denied by controlling only the local copy of the access data-structure. Due to the out-of-order execution of the shared calendar's updates and the authorization policy's updates, we have used an optimistic approach that tolerates momentary violation of access rights but then ensures the copies to be restored in valid states with respect to the stabilized access control policy.

Figure 1 shows the components of our decentralized and secure shared calendar: a user locally manipulates his or her copies (the shared calendar and its authorization policy), and the remote updates from other users (arriving by means of ad-hoc networks) are processed. When a user intends to update his or her local copy of the shared calendar, this update will be granted or denied by checking the local copy of the policy. Once granted and executed, the local update is processed by coordination and access control

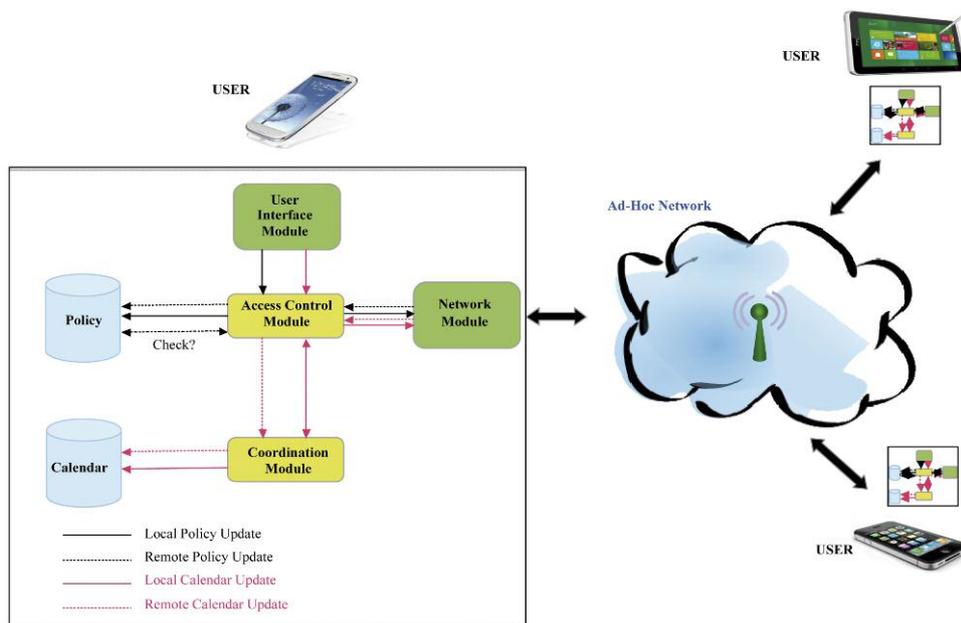


Figure 1: Architecture of our shared and secure calendar.

modules in order to be wrapped by some meta-information, which are necessary to enforce the calendar consistency and its authorization policy, and next it is broadcast to other users. As for the remote calendar update, the access control module checks whether or not this update is authorized and, if granted, sends it to the coordination module for further processing to maintain the calendar consistency before its immediate execution on the local copy of the shared calendar. Note that individual users administer their own proper calendar events. Thus, a user can define an authorization policy on which access rights for other users can be specified.

When an administrator modifies a local policy by adding or removing authorizations, this modification is sent to other users in order to update their local policy copies.

Our research work is funded by the European project FP7 NESSoS.

Links:

<http://www.loria.fr/~imine/tools/home.htm>
NESSoS: <http://www.nessos-project.eu>

References:

- [1] comScore, comScore reports December 2012, U. S. smartphone subscriber market share, press release, February 6, 2013
[2] J. Achara, A. Imine and M. Rusinowitch: "DeSCal - Decentralized Shared Calendar for P2P and Ad-Hoc Networks", in IEEE ISPDC 2011, Cluj-Napoca, Romania, 2011, pages 223-231.

Please contact:

Abdessamad Imine
University of Lorraine and Inria
Nancy-Grand Est
E-mail: Abdessamad.Imine@inria.fr

Indoor User Localization Using Mobile Devices

by Jonáš Ševčík

The ability to easily access your location along with additional relevant information has emerged over the last six years with the rapid development of the smartphone industry. The trend of using mobile phones to get a sense of location is now one of the main features of these devices. Most outdoor tracking is done by the Global Positioning System (GPS), but what are one's options if one is indoors, where you cannot use the GPS signal?

Navigation throughout vast complex buildings, eg, hospitals or school buildings, can be difficult even if the building is densely equipped with navigation signs. For acquiring location coordinates, the GPS can be used. This technique is commonly used in transportation and these days also in mobile phones to calculate coordinates on a virtual map. Although GPS tracking can be useful outdoors, it cannot be used indoors, because there is no direct visibility to GPS satellites.

Fortunately, there are other means for current smartphones to perform local-

ization. They are equipped with various sensors such as magnetometer, linear accelerometer, gyroscope, and wireless network adapters. Even a low-end smartphone contains at least a magnetometer and an accelerometer. Therefore, there is now a unique opportunity to realize a solution for indoor localization.

At Masaryk University in Brno we recognized that opportunity, and have been developing an indoor localization for mobile devices. For our project we have selected the Google Android platform, which offers well documented platform

tools enabling easy access to sensor data. In addition, developments for Android are mainly based on the Java programming language. Therefore, various already programmed libraries can be used.

Our localization system is based on algorithms published in [1, 2, 3]. The developed solution consists of three localization techniques merged together to provide more accurate results. The first technique is indoor tracking based on wireless networks. As presented in [1], we use the Received Signal Strength (RSS) fingerprinting to create

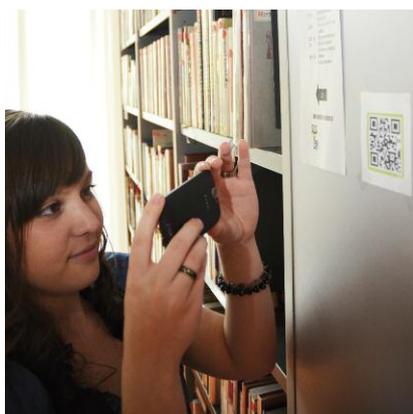


Figure 1: Recalibrating location in a library.

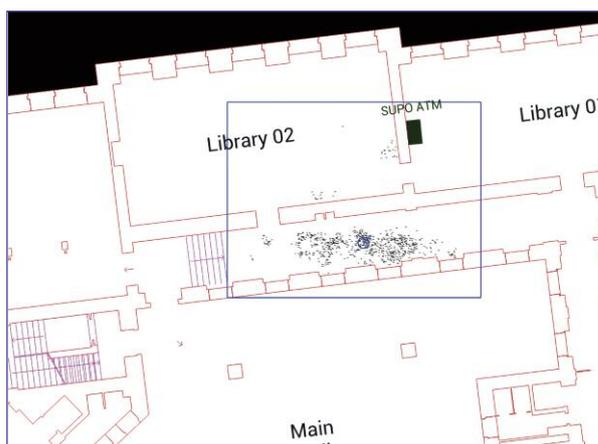


Figure 2: Map view with filter data.

a database of unified access point identifiers and their received signal strength which is mapped to location coordinates. RSS measurements at an unknown location are then compared to the signal strength maps to estimate the receiver's location.

The second technique is dead reckoning – a process of calculating the current position by using previously determined coordinates, which are advanced by known speed and course. This technique is well known from air navigation. We use dead reckoning in an inertial navigation system, where the initial location is determined via the former technique. Alternatively, coordinates can be encoded into a barcode, which when placed into the position of those coordinates, can be scanned by a camera and then decoded. Position tracking is done by calculating a number of steps. Step length is estimated using a neural network as presented in [2]. The course of steps is determined by a gyrocompass. Since dead reckoning is subject to cumulative errors, we use barcodes for recalibrating a position (Figure 1).

The third technique is Sequential Monte Carlo filtering [3], which uses particles evenly spread in the probable location

(as shown in Figure 2 by the blue rectangle) determined by a RSS fingerprint database. These particles are set in motion with events generated by step detection. Then these particles whose hypothetical motion leads through impassable obstacles, eg walls, are eliminated. This improves location estimation.

The indoor localization prototype is capable of displaying the position of a user on a floor map. The basic view offers display of walls, doors, staircases, and elevators. Additional information can be displayed using a set of layers, which contains points of interest such as room numbers, positions of vending machines, ATMs etc. Layers can be applied according to the needs of the user. This system is useful in unknown environments, providing users with crucial information about their location and helping them track their movement.

In a further development, we are planning to implement navigation to a selected point on the map. We will also extend view modes by introducing an augmented reality, which will allow users to perceive camera images enhanced by information about their surroundings.

The indoor localization prototype demonstrates that current smartphones are capable of difficult real-time calculations and that it is possible to implement a sophisticated localization system without using any additional hardware. This system provides guidance crucial for orientation in unknown buildings, where standard maps cannot be used.

References:

- [1] A. Chen et al: "An algorithm for fast, model-free tracking indoors", *ACM SIGMOBILE Mobile Computing and Communications Review*, ISSN 1559-1662, vol. 11 no. 3, 2006
- [2] S. Cho: "MEMS Based Pedestrian Navigation System", *Journal of Navigation*, ISSN 0373-4633, vol. 59, pp. 135–153, Royal Institute of Navigation, 2006
- [3] M. Arulampalam et al: "A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking", *IEEE Transactions on Signal Processing*, ISSN 1053-587X, vol. 50 no. 2, 2002.

Please contact:

Jonáš Ševčík
Masaryk University - CRCIM, Czech Republic
E-mail: jonas.sevcik@mail.muni.cz

Airplace: Indoor Geolocation on Smartphones Through WiFi Fingerprinting

by Christos Laoudias, Georgios Larkou, Demetrios Zeinalipour-Yazti and Christos G. Panayiotou

The wide availability of location data is undoubtedly reshaping the entire smartphone ecosystem with the advent of innovative location-aware services and applications. To accelerate their adoption by the public, the Airplace positioning system is a location-enabling platform that aims to deliver highly reliable and accurate location information right where the highest demand is anticipated in the near future: inside buildings.

Location information has become an integral part of the mobile user's daily life and Google is already reporting that a third of searches on portable electronic devices refer to spatially relevant content, while 94% of smartphone users have searched for local information. This is mainly due to the proliferation of powerful smartphones featuring on-board GPS sensors. According to Tomi Ahonen Consulting, there are already 1.2 billion smartphone devices in use worldwide, and 770 million of these are

equipped with GPS. As one might expect, the most successful Location-Based Service (LBS) is vehicle navigation and pedestrian guidance to points of interest in outdoor areas.

There has, however, been a recent increase in interest in location-aware content specific to large indoor environments, such as shopping malls, museums, exhibition centres, conference venues and airports. Statistics by Strategy Analytics indicate that people

spend 80-90% of their time inside buildings, while 70% of cellular calls and 80% of data connections originate from indoors. This has triggered an increasing interest in indoor applications, such as in-building navigation, asset tracking, elderly support for Ambient and Assisted Living (AAL) and others. Since GPS has low availability indoors due to the blockage or attenuation of the satellite signals, alternative geolocation solutions are required.

At the University of Cyprus, we appreciate the benefit of indoor LBS and our goal is to facilitate their wide acceptance. Thus, we focus on the provision of GPS-free location data inside buildings by leveraging location-oriented information collected with the smartphone's built-in sensors, such as WiFi, accelerometer, gyroscope and digital compass. The Airplace indoor positioning system [1] is a collaborative research effort between the KIOS Research Center and the Department of Computer Science at the University of Cyprus that was initiated a few years ago. Airplace is developed on popular Android smartphones and the unknown user location is determined with the use of Received Signal Strength (RSS) data from the surrounding WiFi infrastructure, ie, wireless Access Points (AP).

This approach is very convenient and extremely suitable for indoor applications running on the users' smartphones. Firstly, the existing WiFi APs, already deployed in indoor environments for wireless connectivity, are exploited thus avoiding the installation of costly dedicated positioning equipment, such as custom transmitters and antennas. Secondly, we rely on RSS data extracted through passive scanning of the WiFi beacon packets transmitted by neighbouring APs as part of the standard network functionality. Thus, commercial mobile devices, which are usually equipped with WiFi adapters, can be used without any hardware modifications and only a software agent is needed for monitoring the RSS values.

Airplace is based on the RSS fingerprinting approach to address the challenging indoor signal propagation conditions. A number of RSS fingerprints (ie, vectors of RSS measurements recorded from APs in the vicinity of the user) are collected a priori to build the "radiomap". Location is then estimated by finding the best match between the observed fingerprint and fingerprints in the radiomap using state-of-the-art algorithms developed in-house [2]. The prototype system follows a mobile-based network-assisted architecture in order to reduce the communication overhead, while at the same time addresses security concerns and respects user privacy [3]. In a typical positioning scenario, when a user enters an indoor environment, covered



Figure 1: Screenshot of the Airplace indoor positioning interface for Android



Figure 2: Screenshot of the Airplace indoor mapping interface for Android

by several WLAN APs, the user's smartphone obtains the RSS radiomap from the local distribution server in a single communication round and is thereafter able to self-locate independently using only local knowledge and, more importantly, without revealing its personal state.

The Airplace system won the Best Demo award at the 13th IEEE International Conference on Mobile Data Management (MDM'12) and the results of our research are also being exploited for the development of an enhanced prototype system via a technology transfer agreement with a large hardware and software company in Taiwan. In particular, we recently built and demonstrated a hybrid positioning solution that fuses the WiFi location with multi-sensory inertial data that are widely available on Android smartphones to deliver fine-grain location information. Releasing a "big-data" Web2.0-oriented geolocation service that will be scalable to large indoor environments is another direction of this research. In this context, some of our future activities include the investigation of algorithms for reliable floor determination in multi-storey buildings, supporting diverse mobile devices without compromising the positioning accuracy, developing an efficient map-matching component and building a user-friendly indoor navigation application using Google Maps.

References:

- [1] C. Laoudias et al: "The Airplace Indoor Positioning Platform for Android Smartphones", in proc. IEEE MDM'12
- [2] C. Laoudias et al: "Localization using radial basis function networks and signal strength fingerprints in WLAN," in proc. IEEE GLOBECOM, 2009
- [3] G. Chatzimiloudis et al: "Crowdsourcing with Smartphones", IEEE Internet Computing, Volume 16, 2012.

Links:

<http://www2.ucy.ac.cy/~laoudias/pages/platform.html>
<http://www.kios.ucy.ac.cy/>
<http://dmsl.cs.ucy.ac.cy/>

Please contact:

Christos G. Panayiotou
 KIOS Research Center for Intelligent Systems and Networks
 University of Cyprus
 Tel: +357 22 892298
 E-mail: christosp@ucy.ac.cy

Demetrios Zeinalipour-Yazti
 Data Management Systems Laboratory (DMSL), Department of Computer Science, University of Cyprus
 Tel: +357 22 892755
 E-mail: dzeina@cs.ucy.ac.cy

Knowledge Representation and Management in Indoor Mobile Environments

by Imad Afyouni, Cyril Ray and Christophe Claramunt

Despite the continuous development and improvements made in mobile computing and the variety of technologies that can be used to enable ambient indoor environments, there are still many research challenges in this area that need to be addressed. Spatial data representation and management as well as location-dependent query processing are among current issues to be coped with so that an efficient and sufficiently flexible indoor context-aware navigation system can be designed.

A key enabler of mobile computing, apart from the networking infrastructure, is the successful integration of semantics and advanced reasoning techniques into mobile information systems. These features have been recently characterized under the Ambient Intelligence (AmI) technologies. The range of applications in ambient information systems is progressively evolving from large to small scale environments. In particular, there is a growing need for applications that assist humans in their navigation-related activities in indoor spaces (eg, airports, museums, office buildings). A successful integration of indoor knowledge representation and ambient systems still requires the development of appropriate spatial data structures and data management facilities. We believe that this is mandatory for the delivery of intelligent and context-aware systems applied to indoor spaces.

Location-dependent queries and more generally context-dependent queries appear to have considerable impact for the development of different categories

of such location-based and context-aware services. The context-dependent character of these queries means that any change in the context (eg, changes in the locations and profiles of the objects that are involved in the query) may significantly affect the answer. As these queries are time-sensitive and location-dependent, they may be valid only for a given period of time and within a given area. Therefore, these queries are expected to be processed as continuous queries, meaning that the system should continually keep the answers up-to-date, until the query is explicitly cancelled by the user. A unique combination of challenges arises with location- and context-aware services and queries in indoor environments, as researchers must be able to represent different kinds of location-dependent and user-centred queries (typically path search and range queries) in a flexible manner, and to take into account additional context information, time-dependency, and the hierarchical layout of the indoor environment.

Indoor space modelling is the core element to be investigated for advanced services [1]. We postulate that modelling and designing a hierarchically organized context-dependent indoor data model is the best way to satisfy a wide range of location-based services. Such a data model can be viewed as a tree structure in which location information is represented at different levels of abstraction: (1) a fine-grained graph embedded within an occupancy grid (Figure 1); (2) an exit hierarchy (Figure 2); and (3) a location hierarchy (Figure 3). Moreover, such a model is able to represent: (i) static/moving features of interest, (ii) their spatial properties, and (iii) the behaviours or actions that emerge from them. The hierarchical design alleviates performance and scalability issues in location dependent query processing. In addition, time-dependent functions that compute the network distance and the travel time are introduced. Furthermore, a classification of user profiles is possible in order to perform an offline filtering of the multilevel data model, thus reducing the

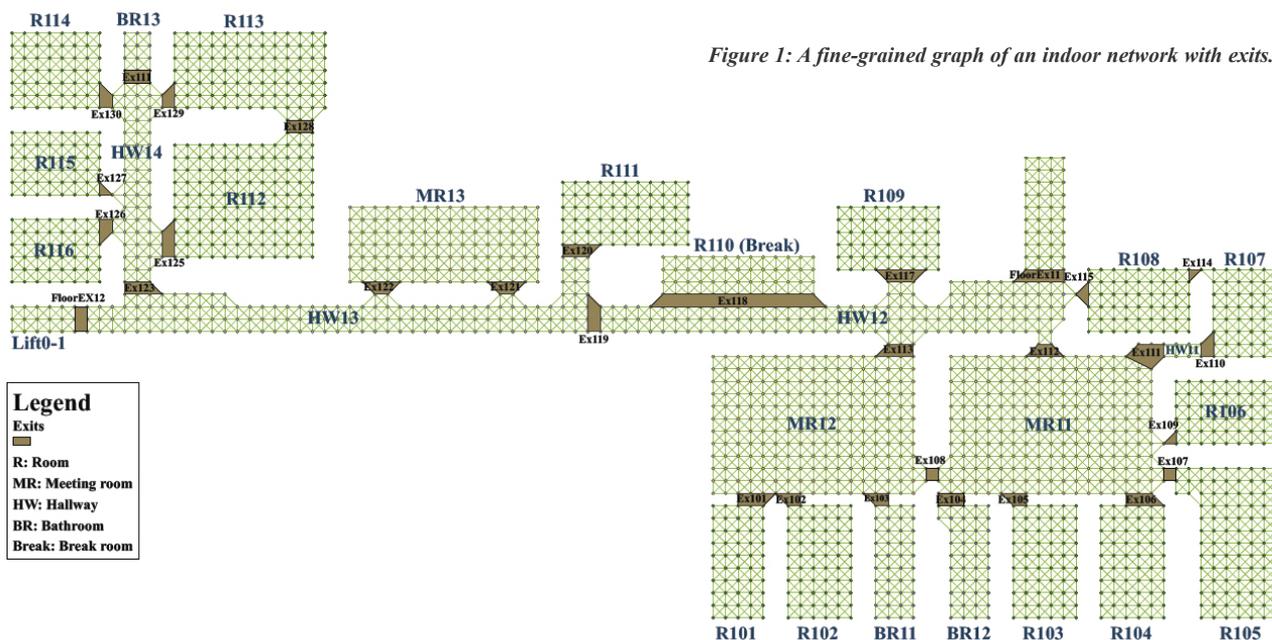


Figure 1: A fine-grained graph of an indoor network with exits.

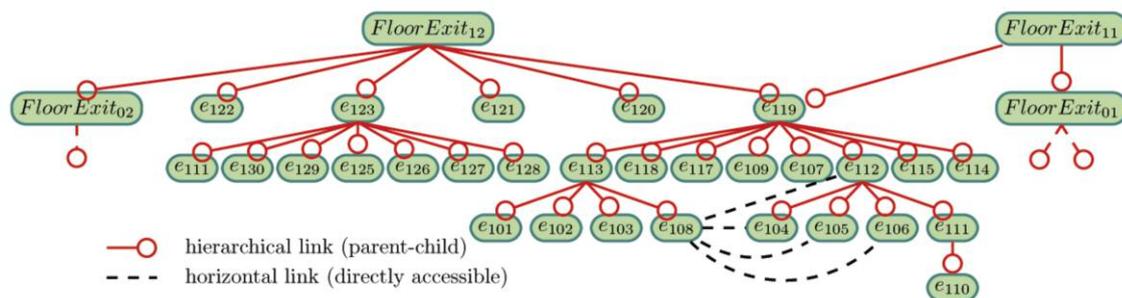


Figure 2: Part of the exit hierarchy derived from the fine-grained graph of the indoor network

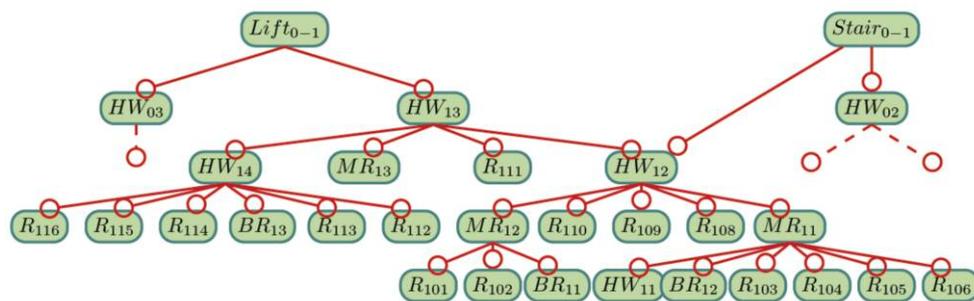


Figure 3: Part of the location hierarchy derived from the fine-grained graph (Figure 1). “HW” means hallway, “MR” meeting room, “BT” bathroom.

amount of data that need to be processed in real-time.

The semantics of a query grammar tied to the indoor data model have also been developed. Such grammar is required to express location-dependent queries. It supports navigation queries and incorporates some other preferences and semantics in the query model. It also supports the hierarchical data model by using the concept of location granules to represent the different levels of abstraction. The use of the location granules concept allows formulation of queries using the location terminology required by the user (eg, vertices at the fine-grained level, rooms, floors, buildings, etc.).

We designed algorithms and a generic architecture usable for the continuous processing of location-dependent queries in indoor environments [2]. Navigation-related queries are processed in accordance with this architecture, and are executed continuously while the request is not explicitly cancelled by the user. Particularly, algorithms for hierarchical path searches and range queries applied to both static and moving objects have been developed [3]. These algorithms take advantage of the hierarchical data model of the indoor environment, and employ an incremental approach in order to effi-

ciently execute continuous location-dependent queries, thus avoiding solving each search problem independently from scratch. Our design and implementation rely on a database extension based on the open source DBMS PostgreSQL. The main parts of the prototype developed are: (i) the hierarchical network-based data model of the indoor environments; (ii) the operators and location-dependent constraints introduced in the query grammar; and (iii) the algorithms to process continuous location-dependent queries over moving objects.

To conclude, a successful integration of indoor knowledge representation and mobile information systems requires the development of hierarchical spatial data models. These models should support continuous processing of location-dependent queries applied to moving objects acting in these environments. This is exactly the scope of the approach that has been developed by our research. Application perspectives are very large, from the development of interactive systems for built environments, to additional professional- and user-oriented services. Not only are the application perspectives promising, but there is also great potential for business as the range of possibilities will surely open many opportunities.

Link:

Naval Academy Research Institute:
<http://irenav.fr>

References:

- [1] I. Afyouni et al: “Spatial Models for Indoor Context-Aware Navigation Systems: A Survey”, *JOSIS*, 4(1), 85-123, 2012
- [2] I. Afyouni et al: “Context-aware modelling of continuous location-dependent queries in indoor environments”, *JAISE*, IOS Press, 5(1):65-88, 2013
- [3] I. Afyouni et al: “Algorithms for continuous location-dependent and context-aware queries in indoor environments”, *ACM SIGSPATIAL 2012*, 329-38.

Please contact:

Imad Afyouni and Cyril Ray and
Christophe Claramunt
Naval Academy Research Institute,
France
E-mail: {imad.afyouni, cyril.ray,
christophe.claramunt}@ecole-navale.fr

Mobile Real Time Applications for Enhancing Public Transport User Experience - The MOVE-ME Project

by João Falcão e Cunha, Teresa Galvão and Jeremy Pitt

Research at the School of Engineering of the University of Porto and Imperial College London has been investigating how smartphones will enable you to enjoy the experience of travelling in a bus.

Pervasive mobile smart devices and sensor networks in public transport vehicles are enabling a new approach for enhancing the experience of public transport customers. The experience may actually start before a journey takes place, by planning a trip, and travellers could still be providing relevant information after reaching their destinations. Transport services and associated information services are closely coupled, and this relationship needs to be better understood. We claim that pervasive mobile services can be used for enhancing user experience in the transport services, and also for enhancing overall public transport services.

Research being conducted in public transport information services at FEUP, the School of Engineering of the University of Porto for the past ten years, and at Imperial College London for the past three years, has been aiming at providing users with real time information, but also finding ways to use the user feedback for benefiting all other transport service stakeholders. Mobile computing is a key enabler for making such a vision possible.

The MOVE-ME service

The MOVE-ME project has developed an infrastructure and a mobile application enabling users to access public transport information in real time. From 2012, this application enables travellers to plan their journeys based on real time or planned data from metro, bus, coach, and train schedules. The infrastructure brings together geo-referenced data from different transport companies, Google map data, and also relevant locations from tourist offices enabling multimodal journey planning. Real time multimodal travel planning can be done in a 60 minute time horizon. When real time data is not available or when the time window is larger than 60 minutes and shorter than three days, travel planning is done based on available published schedules over a three day horizon. With such time horizons the available infra-

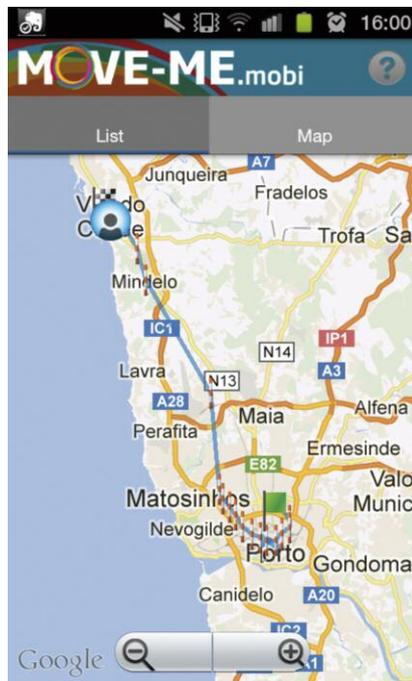


Figure 1: The map shows a multimodal public transport route from Vila do Conde to Porto (in blue) and its stops/stations (red and white poles). Users can follow a trip in real time.

structure has been able to guarantee acceptable response time for a large number of simultaneous users.

The most advanced MOVE-ME service was launched in May 2012 in the Metropolitan Area of Porto, and in other regions of Portugal. Over 25 distinct metro, bus, coach, and train companies share information on their service, and over 20,000 Android and iPhone users are now benefiting from this service.

The MOVE-ME mobile service and infrastructure was developed by OPT (www.opt.pt) with research support from FEUP, in the context of the CIVITAS Elan European project. MOVE-ME won the 2012 CIVITAS European Technical Innovation Award.

Research on mobile advanced traveller information services

More advanced mobile computing research with the objective of

enhancing public transport user experience has been under way at FEUP, in collaboration with Imperial College, London, for the past three years.

Current research is also addressing ways of extending MOVE-ME to enable travellers to pay for their journeys using their mobile devices. The Metropolitan Area of Porto has an open transport network with no barriers. In the future it is possible to envision an entirely pay-as-you-go transport service based on users' mobile devices. As Web enabled smartphones become pervasive, ticket selling machines and validating machines could become redundant and unnecessary.

Mobile devices are also being used to measure users' affective state in real public transport contexts, and how to share and disseminate real time information through social networks for the benefit of other travellers and also for the benefit of the control rooms of transport companies [1]. Preliminary experiments have been conducted in Porto and in London transport services.

A prototype cloud-based mobile service to assess the relationship between affective state and travelling context has been developed and then tested with commuters of the Porto transport network [2]. User's affective state was captured using a simple emotional model of travelling mood, with cognitive pleasure and physical arousal dimensions based on Russell's circumplex of emotion [3]. Travelling context included noise, saturation, smoothness, ambience, speed and reliability. The findings show a strong correlation between mood and context, dependant on the user.

A prototype crowdsourcing application to share information has been developed and tested with commuters of the London transport network. The model borrows key principles from Internet based services. It strives to intensify

win-win relationships between public transport passengers and operators. The structured exchange of information is sustained by a validation mechanism for data reliability, and an incentive mechanism to encourage passenger participation. Passengers benefit from rich real-time data to ease their journeys and improve travel experience, in exchange for their own participation providing and validating information. Operators gain access to rich customer generated data, which in an aggregated format may provide a real-time assessment of customer experience and of local performance across the entire network operation.

It is expected that ubiquitous availability of high quality Web enabled mobile devices and services will improve public transport user experience, both functionally and emotionally. Collaboration with users may also benefit public transport operators, with user's feedback enabling them to enhance service levels.

Better experience leads to increased usage of shared mobility modes, and therefore to more sustainable cities in the future.

Acknowledgments

The authors would like to acknowledge the work of colleagues, PhD students, and collaborators at OPT. Funding from the following projects is also acknowledged: OneStopTransport QREN TICE no. 13843, MOBIPAG QREN TICE 13847, EU CIVITAS Elan, IBM CAS Portugal, INEGI, IDMEC Polo FEUP, UGEI - INESC TEC, PhD grants from FCT, Portugal.

Link: <http://www.move-me.mobi>

References

[1] A. A. Nunes, T. Galvão Dias, J. Falcão e Cunha, J. V. Pitt: "Using social networks for exchanging valuable real time public transport information among travellers", in proc. of the IEEE Conference on Commerce and Enterprise Computing, 365-370, 2011

[2] P. M. Costa, J. V. Pitt, J. Falcão e Cunha, T. Galvão Dias: "Cloud2Bubble: Enhancing quality of experience in mobile cloud computing settings - A framework for system design and development in smart environments", in proc. of the 3rd ACM Workshop on Mobile Cloud Computing and Services (MCS'12), 45-52, 2012

[3] J. A Russell: "Circumplex Model of Affect", journal of personality and social psychology 9, 6 (1980), 1161-1178.

Please contact:

by João Falcão e Cunha and Teresa Galvão
Universidade do Porto, Portugal
E-mail: jfcunha@fe.up.pt,
tgalvao@fe.up.pt

Jeremy Pitt
Imperial College London, UK
E-mail: j.pitt@imperial.ac.uk

U-AirPoll: Mobile Distributed and Collaborative Air Pollution Measurement

by Marino Linaje and Luis Miguel Dominguez-Peinado

Current technology, including smartphones, sensors and the Cloud can revolutionize traditional air pollution measurement. This article summarizes a novel approach focused on cost optimization, collaborative data capture, use of standards and public data distribution.

Generally, pollution measurement projects [1] are location static, and data may be controlled by entities subjugated to economic or politic factors. Consequently, environmental data may be subjective and valid only for a specific location (ie, the area covered by the environmental monitoring station).

U-AirPoll is a small, fully open (hardware and software), web standard-based and autonomous system of air pollution measurement. It can be carried by people or installed in vehicles such as bicycles, motorcycles or buses. To minimize cost, the project relies on the fact that increasing numbers of people are carrying smartphones. Thus, smartphone native sensorial capacities are extended with an external hardware capable to measure air pollution. This hardware base is cheap (around 70€)

when compared to the environmental station's price. Data are captured and sent from the air pollution sensors to the smartphone and uploaded from the smartphone to the Cloud. Thus, data can be collaboratively collected by many nodes (users or vehicles) from the public or private sectors. Data are open published and shared in the cloud to be easily consumed in XML. This approach thus avoids the gaps of traditional pollution measurement projects.

U-AirPoll communications architecture is composed of three parts: the drone, the droid and the Cloud services (Figure 1). In contrast to other solutions, the approach fully exploits the Web of Things paradigm.

The drone is a device that extends the phone sensorial capacities. Two ver-

sions are available: Firstly, the prototype version (Figure 2) including a custom CO2 concentration sensor and an already existing general air quality sensor. Secondly, the ready-to-use version (Figure 2, top-left) that is within an enclosure and avoids unnecessary hardware (eg, grove nest). The platform is extensible with up to 22 environmental sensors. It is built using Flyport by openPicus. Flyport was selected because it is a low cost open hardware with IEEE 802.11b interface and a micro-webserver embedded supporting REST (Representational State Transfer).

Drone hardware could be replaced by Flyport minified and/or washable versions to be embedded into clothes. One aspect of U-AirPoll's versatility is its mobility: a drone can be placed wher-

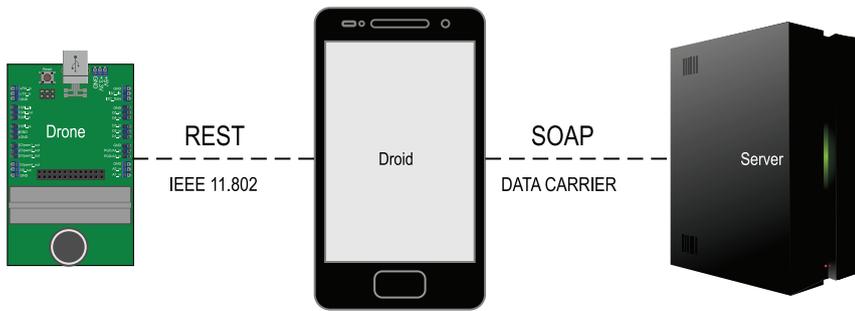


Figure 1: U-AirPoll communications architecture



Figure 2: Drone prototype

ever the user wants within the range of the user's Smartphone Wi-Fi connection (eg, the roof of a bus).

From the droid perspective, the drone is just a service that is used to acquire data. The droid is a mobile application that uses the smartphone geolocation to complete air pollution data from the drone. It is implemented using Apache Cordova multi-device development framework that can run in up to nine different mobile platforms such as Android, iOS or Windows Phone. Apache Cordova was also selected because its applications are coded using standard web development technologies (ie, HTML, CSS and Javascript).

When not connected, the droid tries to locate the drone and checks server connectivity. When the drone is detected, data capturing starts and the droid saves all measurements into a local database on the smartphone, which is cheaper than storing them in the drone (ie, additional storage hardware is not required). In order to upload data, login is required. Once logged in, the droid begins to send data (ie, the pollution data from the Drone augmented with the smartphone geolocation). In par-

allel, the droid requests data from a server in a range of coordinates to show a map in the smartphone. The trend in the Web of Things field is to use REST for the Web services provided. However, to prove that our approach does not fix the Web service technology used, REST and SOAP (Simple Object Access Protocol) are mixed in the final solution.

All the cloud services have been specified using a model driven development environment called WebRatio. All the Java code generated by this tool is open source. The cloud services provide anonymous data pollution concentrations and provide an open access to data. To avoid problems with multiple connections (also known as the c10k problem) a JSP and c10k capable server, such as JBoss, is required.

The main idea behind U-AirPoll is to set the foundations for an open, distributed and mobile air pollution measurement system which allows data to be easily consumed by custom end-user applications and services (eg, a user or organization that wants to create green traffic routes to run or to cycle). Since U-AirPoll is an open hardware project,

other people can add more pollution sensors, design their own enclosures or modify any piece of the project to better suit their needs.

Currently, U-AirPoll supports only 1:1 droid-drone connections. Our research is currently investigating N:N supports (multiple drones and droids shared among users).

This work has been funded by the Spanish Ministry of Science and Innovation (TIN2011-27340) and the European Regional Development Fund (ERDF).

Links:

- <http://sites.google.com/site/U-AirPoll/>
- <http://www.openpicus.com>
- <http://cordova.apache.org>
- <http://www.webratio.com>

Reference:

- [1] D. Vallero: "Fundamentals of air pollution", Academic press, 2007.

Please contact:

Marino Linaje Trigueros
 Universidad de Extremadura, Spain
 Tel: +34 927257258
 E-mail: mclinaje@unex.es

ARGO Sentinel: The Mobile App for Reporting Oil Spillages at Sea

by Massimo Martinelli, Davide Moroni and Ovidio Salvetti

We believe that the contribution of volunteers could play a fundamental role in monitoring and protecting the environment. People at sea sighting pollution caused by oil or hydrocarbon spillages can now immediately report this using a freely downloadable mobile application.

At the Signals and Images Laboratory (SI-LAB) of the Institute of Information Science and Technology of the National

Research Council of Pisa (ISTI-CNR), we have developed a Marine Information System (MIS) for moni-

toring vessel traffic and oil spills within the Mediterranean basin. The MIS collects and integrates geotagged data

related to safety and health issues of the sea from various sources (satellites, optical sensors, electronic noses, autonomous underwater vehicle systems) and provides predictive models to assist the authorities in the management of emergencies at sea.

An integrant part of the MIS is represented by the data that can be collected and shared by volunteers who want to collaborate in monitoring the status of the sea.

For this purpose, we have developed ARGO Sentinel, a free application for smartphones with a Geolocation System (GPS)[1]. Two versions of the app have been implemented: an HTML5 version, running as an Opera widget distributed only to our partners of the ArgoMarine project, and a native Android (v2.2 or higher) publicly distributed because of its level of stability, and downloadable by anyone from Google Play (see Link below).

The app was conceived on the intuition that the contribution of volunteers could play a fundamental role in monitoring and protecting the environment. Using the app, whenever someone at sea sights signs of oil or hydrocarbon pollution they can immediately report this to the SI-LAB in Pisa. The information is recorded in the MIS and complements the data obtained from more traditional sources (eg satellites), improving the quality and coverage of marine monitoring, especially in protected areas. In this way we can build up a detailed map of the status of our seas.

The application – distributed in Italian, English and Greek – sends reports of suspected spills to our Lab, providing a description of the spill and specifying the precise point and severity.

The main screen of ARGO Sentinel shows: a "Message" area where a description of the sighting is entered; a "Red Alert" button to signal a major oil spill with a diameter greater than or equal to 20 meters approx.; a "Yellow Alert" button: to report a mild/moderate amount of oil spill with a diameter less than 20 meters approx.; an "i" button: for information on the use of the app; an "exit" button to close the app (See Figure 1).

By clicking on the red or yellow alarm buttons, an SMS is sent to the CNR

headquarters, which will process the received data.

The app has been installed by about 700 individuals all over the world. By integrating the alerts sent by the volunteers with all the other information sources collected in the MIS, a semi-automatic analysis eliminated the alerts identified as false positive. Even though we have no formal obligations as the project is for research purposes only, we forward



Figure 1: the main screen of the "ARGO Sentinel" mobile app

all the significant information to the General Command of the Italian Coast Guards in Rome, Italy.

Our field tests demonstrate that the use of this new technology could be really important in combating pollution. Our experience suggests that this kind of technology can be applied to many other fields where environmental monitoring and safety is crucial.

An important result is that, during the period of the project's activity, we have acquired a more detailed and immediate knowledge of the conditions of the sea. The use of this application represents a step forward in marine environmental monitoring, because, in addition to the other technologies that are used by the ArgoMarine project, it also adds the contribution of volunteers who can easily communicate the sighting of a

spill. Knowing that deliberate spills can be detected in a timely fashion is in itself a deterrent to malicious actions. Possible developments may allow a more effective intervention by the authorities.

A new version of the ARGO Sentinel app is now under development and will also be released for iOS user.

ARGO Sentinel is has been developed in the framework of the European research project ArgoMarine, ("Automatic Oil spill Recognition and Geopositioning integrated in a Marine Monitoring Network") which aims at traffic and marine pollution monitoring. Coordinated by the Tuscan Archipelago National Park, the partners of ArgoMarine include ISTI-CNR, the National Technical University of Athens, the Nansen Environmental and Remote Sensing Center, the Centro de Investigação Marinha and Ambiental, Universidade do Algarve, the National Maritime Park of Zakynthos, the Joint Research Center and the NATO Undersea Research Center.

Links:

ARGO Sentinel:
<http://tinyurl.com/argosentinel>
 ArgoMarine Project:
<http://www.argomarine.eu>

Reference:

[1] "ARGO Sentinel: the application for reporting oil spillages at sea", http://www.cnr.it/istituti/FocusByN_eng.html?cds=074&nfocus=22, CNR Focus, National Research Council, 2012

Please contact:

Massimo Martinelli
 ISTI-CNR, Italy
 Tel: +39 050 621 28 03
 E-mail: massimo.martinelli@isti.cnr.it

Mobile Devices to Improve Breast Cancer Information Management

by Damià Segrelles, Maite Giménez and Ignacio Blanquer

Intuitive interfaces of mobile devices facilitate the introduction of structured data in breast cancer management, leading to an increase of completeness and accuracy of diagnosis and follow-up evaluations, as well as opening the door for more effective content-based retrieval techniques for clinical decision support.

The exploitation of medical data through distributed computing infrastructures (DCIs) has led to new research lines focused on the extraction of knowledge and the organization of the information to assist radiologists and researchers to retrieve pertinent data from existing medical image repositories. This development not only offers great benefits to training and research, but can also help to improve the clinical management of patients. One significant case is breast cancer screening. Specifically, effective content-based retrieval is based on: (i) the structured representation of the knowledge relative to the images, and (ii) the integration of a means to create and exploit a knowledge database with structured data. However, content-based retrieval is strongly affected by the quality of input data, especially in the medical area. The acquisition of accurate and complete data is a key challenge in clinical healthcare [1].

On the other side, the inherent complexity in the use of DCIs prevents their usage within the wider community. Scientific gateways are being widely used to facilitate the access of scientists to these tools, simplifying the organization of data repositories and the execution of experiments. In the particular case of biomedicine, usability evaluation is widely recognized as critical to the success of any software system with which end users interact [2].

The introduction of mobile computing constitutes a great opportunity for increasing both the quality of introduced data and the content-based retrieval of existing data. Radiology reports are introduced in plain text using voice recognition systems. Coding is limited to conclusions in the best case. Plain



Figure 1: Filling in structured reports for a mammography from a tablet

text-based retrieval is inefficient as documented in the literature due to the ambiguities of terms and the difficulties in dealing with negative statements. Structured reports are playing a key role in the definition of unequivocal pieces of information for comparisons and data mining.

For this purpose, TRENCADIS [3] (Towards a Grid Environment to Process and Share DICOM Objects) is a technology, based on Grid and Cloud computing, created for securely sharing, organizing and searching medical images. Moreover, it provides a toolkit to support the development of applications that are capable of using heterogeneous, distributed computing and storage resources, made available through TRENCADIS services. TRENCADIS has incorporated a mobile-compatible interface to ease the introduction of structured reports. A prototype has been developed in cooperation with the Hospital Universitario Dr. Peset from Valencia for the diagnosis, follow-up and response analysis of breast cancer.

The complete process of breast cancer management has been coded into nine templates that allow the radiologists to

fill in the report faster than using keyboard and mouse without compromising the accuracy and guidance of structured terms. A mobile interface also enables writing the report at the image acquisition point. These components are being integrated in an existing prototype being developed for diagnosis support and research in breast cancer. A usability study will follow the validation of the prototype.

Currently, TRENCADIS software and support is provided to the collaborating centres within research projects and contracts.

An open-source distribution of the toolkit is planned for the first term of 2014.

References

- [1] B. Zheng et al: "Interactive computer-aided diagnosis of breast masses: computerized selection of visually similar image sets from a reference library", *Academic Radiology* 14(8), pp. 917–927, 2007
- [2] C. Maestre et al: "Assessing the Usability of a Science Gateway for Medical Knowledge Bases with TRENCADIS", *J Grid Computing*, 10, pp. 665–688, 2012
- [3] I. Blanquer, V. Hernández, J.E. Meseguer: "Content-Based Organisation of Virtual Repositories of DICOM Objects", *Future Gener Comput System*, 25(6), pp.627–637, 2009.

Please contact:

Ignacio Blanquer
Universitat Politècnica de València,
Spain
E-mail: iblanque@dsic.upv.es

Secure and Privacy-Aware Mobile Identity Management

by Fabio Martinelli

As evidenced in our special theme, mobile devices play an increasingly important role in our everyday lives, not only by enabling us to communicate but also by providing access to a large variety of pervasive services. The global mobile security market [1] is steadily increasing in value (about \$1.6bn in 2012), as companies and organisations seek to secure their smart devices against the dangers of mobile malware. Our software products aim at targeting both this market and that interested in the integration of mobile devices into legacy applications.

We use mobile devices in many roles, eg as employee in relation to our employers, as consumers in relation to commercial service providers, and as citizens in relation to our governments. There are now more mobile devices than desktop computers accessing the Internet. Extrapolating this trend, the number of vulnerabilities affecting these mobile devices and their typologies is increasing with the growing pervasiveness of the services. These devices thus represent an attractive target for attackers.

Mobile devices are not only taking over typical PC activities (social networking, browsing, e-mailing, online shopping) but also offer more sensitive applications in areas such as mobile payment transactions, access to health services, etc. The frequent press coverage of hidden tracking, open profiles, and fraud is one of the factors that has convinced us that security and privacy threats are realistic. The ever-growing sophistication in data mining can easily be misused for commercial or criminal exploitation. A particular concern is that mobile devices are available to a wide range of users (from teenagers to seniors), who are not necessarily experts or educated to these risks. This makes successful attacks on these devices even more dangerous, especially when the devices are then used to access services.

In the framework of the EIT ICT Labs (an initiative of the European Union), several research organizations (CNR, Novay, TU/e, TU Berlin,) and leading European industries (Engineering and SAP) are working together to address the challenge of ensuring the security of mobile devices guaranteeing service access in a privacy preserving way.

The partners in this activity have a significant track record for research and innovation in many mobile security related areas and bring together a significant set of experience ranging from European projects to industry based pilots and experimentations. The field is far from mature, neither in terms of providing solutions nor in terms of fully understanding the threats and the potential on new protection methods. Nonetheless many specific activities have been already performed and are amenable to exploitation. The potential adoption of these technologies –if deployed prop-

erly - is extremely high, while on the other hand the failure to properly secure our mobile devices can be harmful to society at large.

We are thus currently developing a software suite of these mobile devices to secure access to data-sensitive platforms such as financial and e-government services or pervasive ones such as e-health for personal monitoring. The data and protocols used to store and (sometimes even autonomously) use personal identity information must be secure. This entails the protection of the devices themselves, in order to avoid various types of misuse, ranging from the security of communications, to privacy considerations on data disclosed to get access in these communications.

One of our main challenges is the use of security solutions within mobile applications. To address this problem, SAP aims at facilitating the development of secure mobile applications with its SAP Mobile Platform (SMP). SMP offers a unique opportunity for developers to quickly develop business applications interacting with critical systems – typically the company ERP. While SMP is already providing state of the art security, SAP intends to extend the security functionality spectrum and to provide new sensitive services developed within the project.

Besides protecting the private data on the device, we also consider the protection of data used by the services the device connects to. In authenticating to the services, we deploy privacy-friendly biometrics with biometrics templates which are protected, also from 'insider' threats. When using the services, a lot of personal information is shared. We provide tools to analyze the amount of personal data revealed and to validate the privacy-protecting properties of communication protocols used.

Our applications will be tested using the Experience and Living Labs facility offered by EIT ICT. This will allow to test the usability and the social acceptance of the solutions developed. Several business and transfer technology actions are planned in order to make our envisaged mobile centric secure identity managing available on the market place. The activity described above is in collaboration with many groups. As main partners we mention Silvia Boi, Laurent Gomez, Niklas Kirschnick, Jerry Den Hartog, Martijn Oostdijk, Jean-Christophe Pazzaglia and Daniele Sgandurra.

Link

SAP Mobile Platform:

www.sap.com/solutions/tech/mobile.html

Reference

[1] The Mobile Security (mSecurity) Market 2012-2017 March 2012, <https://www.asdreports.com/>

Please contact:

Fabio Martinelli, IIT-CNR, Italy

E-mail: fabio.martinelli@iit.cnr.it

Social Electricity: When Awareness About Electricity Becomes Social

by Andreas Kamilaris, George Taliadoros and Andreas Pitsillides

Although domestic smart metering has been introduced into our lives, it is still not easy for consumers to know how much electricity they are using. Consumers are unable to assess whether their consumption is low, average or high. "Social Electricity" is a Facebook application that aims to make individuals aware of their energy consumption by means of comparisons with the corresponding electrical consumption of their friends, as well as with the total consumption in the area where they live. Effective and realistic social comparisons could raise consumers' awareness of their consumption behaviour, enabling them to take steps to reduce their electricity use and carbon footprint.

Energy conservation is a global issue with huge environmental, social and political implications. Smart metering has been introduced into homes to provide timely energy consumption feedback, helping consumers to more effectively manage their power use. Nevertheless, it is still difficult for people to quantitatively assimilate how much energy they consume. Consumers lack the required metrics to determine whether their total consumption is low, average or high. A promising way to understand the "semantics" of consumed energy is to compare it with the amount consumed by relatives, friends and neighbours. The increasing popularity of online social networking sites (eg Facebook, Twitter), allows the proliferation of social ICT applications targeting energy awareness.

Social Norms for Energy Conservation

People are willing and able to adapt their behaviour to energy-saving lifestyles if given the necessary feedback,

support, and incentives [1]. While detailed feedback can raise people's awareness of their consumption, community influence, which is recognized as an important factor in energy-saving initiatives, has the potential to drive residents towards a more persistent behavioural change [2]. Social norms can motivate people to question their behaviour if they discover it is not "normal" [3].

Social Electricity: Energy Conservation through Social Comparison

"Social Electricity" is a Facebook application that aims to make consumers aware of their electrical consumption, by means of comparisons with the corresponding consumption of their friends, as well as with the total consumption in their street/neighbourhood/village/city/country. By effective and realistic comparisons, consumers may perceive their energy behaviour and take steps to reduce their electricity footprint.

The project has been developed by the Networks Research Laboratory, part of the Computer Science Department, University of Cyprus. It is supported by the Electricity Authority of Cyprus (EAC), which provides data on the domestic consumption in Cyprus every two months.

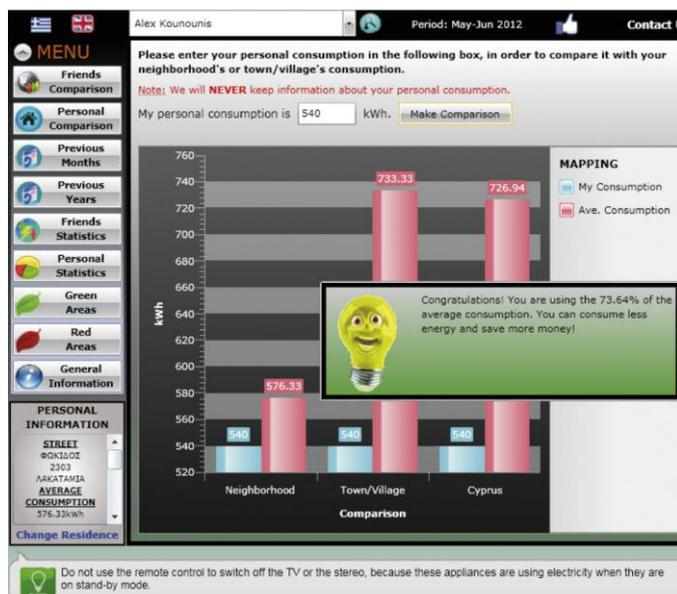
These data are anonymous to protect the privacy of Cypriots, and are aggregated at street level, to further ensure anonymity. Through Facebook, people can compare their consumption with their online friends and view statistics relating to the areas with the lowest/highest energy consumption. Users are also encouraged to add their own consumption, and compare it with that consumed on their street, neighbourhood, village or the whole of Cyprus. Moreover, users may observe the electrical consumption of their street (or their friends' streets) over previous months (or the same month in previous years) and compare it with their present street consumption.

Social Electricity aims to encourage healthy competition, in which users compete with each other to reduce energy consumption. People from the same area are encouraged to cooperate in order to reduce the total aggregated consumption and improve the overall "green ranking" of the area.

This social application is an innovative initiative worldwide and has recently been awarded first prize in the 2nd Green ICT Application Challenge, organized by the International Telecommunication Union (ITU).



Figure 1: A snapshot of Social Electricity application showing a map of Cyprus with a user's friends appearing in green if their consumption is low and red if their consumption is relatively high.



A snapshot of Social Electricity application showing a menu in which the user is encouraged to add his own, personal consumption and compare it with his local neighbourhood, town/village or the whole of Cyprus.

The project only started two years ago, but was officially launched recently and has already amassed over 1,000 users. This is significant, considering that Cyprus has a population of less than a million.

Evaluation

The project was recently evaluated with the use of online questionnaires completed by 190 users of the application. Eighty percent of respondents reported the application to be very useful. Half said they were positively influenced to become more energy-aware. Around 16% realized their energy consumption was high, while 78% said that the application helped them to better perceive and understand their energy behaviour and consumption. Sixty-two percent claimed their consumption was reduced relative to the same period last year. Our measurements in the last 6 months reinforce this claim, since in all these months, the streets of our users consume less energy in comparison to the streets of their postal code (4.3-6.5%) or to the whole of the country (2.5-4.8%).

The primary incentives for using the application are financial and environmental, and secondary incentives are curiosity and social responsibility. The most popular incentive for energy reduction would be a discount on the electricity bills of energy-aware citizens.

Thirty-nine percent of users believe the application will become more useful in coming years, in combination with the massive deployment of domestic smart meters.

Impressively, 64% believe that Social Electricity will help them to reduce their electrical consumption by more than 10% in the long-term.

Conclusion

Social norms may improve people's environmental awareness, while comparison between friends may help them to reduce their electricity footprint. Social Electricity focuses on investigating the effectiveness of these factors on energy awareness and conservation of electricity.

Future work involves enhancing the features of the application, including developing more effective comparisons for electrical consumption between households that share similar physical characteristics (eg size, number of occupants, type of heating), as well as online, real-time competitions for energy savings between friends and neighbours. Accompanying the smart grid is the massive deployment of smart meters for real-time acquisition of household electricity consumption. The concept of Social Electricity is compatible with the future smart grid, since smart meters could enable real-time energy comparisons and promote competition between friends and neighbours, related to energy savings. Evidence indicates that the future Web will be social and pervasive, connected to the physical world. We aim to exploit this trend to promote energy awareness and conservation.

Links:

Social Electricity Application:
<https://apps.facebook.com/socialelectricity/>
Social Electricity Facebook Page:
<https://www.facebook.com/SocialElectricity>
2nd Green ICT Application Challenge:
<http://www.itu.int/ITU-T/climatechange/greenict/201206/index.html>
Networks Research Laboratory:
<http://www.netrl.cs.ucy.ac.cy/>

References:

- [1] A. Kamilaris et al: "Energy Conservation through Social Competitions in Blocks of Flats", in proc. of SMARTGREENS 2012
- [2] R Cialdini: "Influence: science and practice", Allyn and Bacon, 2001
- [3] H. Allcott: "Social norms and energy conservation", Journal of Public Economics, 95(9-10):1082-1095, 2011.

Please contact:

Andreas Kamilaris
 Tel: +357 99 551016
 E-mail: kami@cs.ucy.ac.cy

New Interaction Paradigms in Energy Management

by Paulo Carreira and Alfredo Ferreira

Information technology tools for energy management have, to date, failed to take advantage of the recent advances in human-computer interaction. In fact, considerable expertise is still required to interpret and manage energy data. Other domains, in contrast, including architecture, engineering and construction, have embraced new interaction paradigms aiming at more natural ways to interact with and analyse data. We foresee that integrating natural user interfaces with virtual and augmented reality environments in an energy management system will provide a powerful tool, enabling unskilled users to perform complex energy management activities.

Extracting meaningful information from energy data requires integrating attributes and arranging multiple elements, not only regarding consumption but also building envelope, space function, occupancy history, environmental conditions and equipment status. Traditionally, distinct tools have been used to manage these data, such as CAFM and CAMM tools, as well as energy management (EM) and building control (BCS) systems tools. Since there is no standard integration methodology and framework for energy data, as there is for business data, practitioners must interact with multiple tools to obtain relevant information [1]. Consequently, integration is improvised and largely performed mentally. In a professional setting, the amount of data is overwhelming, hampering meaningful analysis.

In addition, energy data is a type of space related information that is difficult to explore and visualize with traditional tools. Activities such as inspecting the consumption profile of specific types of space, groups of occupants, or specific equipment all have an underlying spatial dimension. Integrated rendering of spatial information is important for visualizing complex phenomena that arise when combining data from multiple sources in order to achieve new insights. Therefore,

most EM tools have to manage spatial information, which can be visualized more effectively when rendered in a graphical representation.

Graphical rendering accomplishes instantaneous identification of the space reality along with the relationships of the elements therein to encourage a fast response. Historically, planimetric CAD drawings and Geographical Information Systems have been used as an effective way to display and manage spatial information related to facilities.

On closely related domains, virtual and augmented reality environments are emerging as a sophisticated and effective way of rendering spatial information [2]. The domains of architecture, engineering and construction (AEC), already have tools that combine these environments with new interaction paradigms and natural user interfaces (NUIs) to improve productivity, such as the tool for HVAC maintenance presented in [3]. However, user interfaces for EM tools are not yet taking advantage of the recent advances in human-computer interaction. In fact, we believe that they are not even close to approaching the spatial dynamism and ease of use offered by virtual and augmented reality environments with NUIs.

Devising a solution of this type presents a twofold research challenge. On the one hand, integration of energy data requires the collection and analysis of large amounts of sensor and meter data. Furthermore, the database workload involves a large number of hybrid queries combining large persisted data with live data posed against an integrated schema - a problem that is not well handled by current database systems. On the other hand, we need to design adequate virtual and augmented reality environments to analyse highly-complex, multi-dimensional and space-related energy data. These challenges are highly interdependent. The more data can be integrated, the richer the interactions can be. Richer interactions, however, require answering even more data integration queries that combine persisted and live data.

Tackling these research challenges calls for a multi-disciplinary approach. At INESC-ID Lisboa, a group of specialists in data management, building automation and energy management have been collaborating with a team of human-computer interaction researchers on a novel approach to EM. We have been designing and experimenting with tools and proto-



Figure 1: Screenshot of the VRE prototype application displaying energy data and rendering status of devices at INESC-ID Taguspark building.

types that augment a virtual facility with energy consumption information along with spacial characteristics, location and status of equipment, while providing simple ways to control them. Unlike previous work that has largely relied on CAD or VRML for scene generation, this work takes advantage of recent game engine technologies for fast and real-time rendering of feature-rich representations of the facility, along with spacial information, equipment conditions and device status.

One of the above-mentioned tools has been rolled out commercially. Our vision is to create a platform for data integration and visualization that enables easy and efficient ways for unskilled users to explore energy data. We believe that a better, integrated ICT solution for handling energy data may translate to faster and more effective energy management through more accurate energy usage diagnostics and just-in-time corrective action that will eventually translate to a more rational usage of energy.

References:

- [1] J. Granderson, M. A. Piette, G. Ghatikar, P. Price: "Building Energy Information Systems: State of the Technology and User Case Studies", technical report, LBNL-2899E, Environmental Energy Technologies Division (2009)
- [2] P. Fite-Georgel: "Is there a reality in Industrial Augmented Reality?", in proc. of ISMAR'11, IEEE Computer Society, Washington, DC, USA, 201-210 (2011)
- [3] Viet Toan Phan and Seung Yeon Choo: "Using AR for real-time crosscheck of ventilator ducts at worksite", in proc. of VRCAI'10, ACM, New York, 293-298, 2010.

Link:

<http://www.inesc-id.pt>

Please contact:

Paulo Carreira, Alfredo Ferreira
 INESC-ID Lisboa and IST/Technical University of Lisbon, Portugal
 Tel: +351 214 233 287, +351 214 233 512
 E-mail: paulo.carreira@ist.utl.pt, alfredo.ferreira@ist.utl.pt

VMC: A Tool for the Analysis of Variability in Software Product Lines

by Maurice ter Beek, Stefania Gnesi and Franco Mazzanti

Researchers from the Formal Methods and Tools group of ISTI-CNR have developed a tool for the computer-aided verification of behavioural variability in product families.

During the last decades, we have witnessed a paradigm shift from mass production to mass customisation in an attempt to serve as many individual customer's needs as possible. A typical example is the production of mobile phones. Software Product Line Engineering (SPLE) translates this paradigm into a software engineering approach aimed at developing, in a cost effective way, a variety of software-intensive products that share an overall reference model, ie that together form a product family. Usually, commonality and variability are defined in terms of features, and managing variability is about identifying variation points in a common family design and deciding which combinations of features are to be considered valid products.

Feature models are widely used for variability management in (S)PLE. A feature model provides a compact representation of all the products of a product family in terms of their features. Figure 1 shows an example feature model inspired by the mobile phone industry, adapted from [1]. The features are represented as nodes in a tree, with different relationships between them. We see that all mobile phones must run software to support calls and to display information on either a basic, a colour or a high resolution screen. The software may include support for a GPS and for media devices such as a camera, an MP3 player or both. Finally, mobile phones including software for a camera must also include software to support a high resolution screen, whereas software for a GPS cannot run on a basic screen.

There is a large body of literature on the computer-aided analysis of feature models to extract valid products and to detect anomalies, ie, undesirable properties such as superfluous or - worse - contradictory variability information (for instance so-called false optional and dead features). These analyses however do not take into account any behavioural variability as only the presence of the software implementing the features is considered, not their temporal ordering. Since software products are often large and complex, and many are used in safety-critical applications in the avionics, railways, or automotive industries, it is unrealistic and practically infeasible to specify - let alone verify - the behaviour of each product individually. Note that our mobile phone example would already require behavioural models for 14 largely identical products.

For this reason, we recently launched a research effort to 1) develop a formal modelling and analysis framework capable of dealing with behavioural variability and 2) pro-

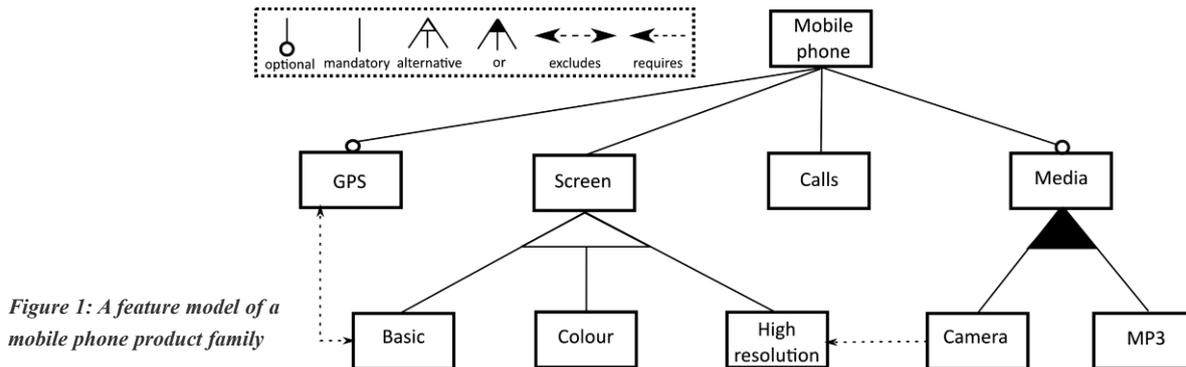


Figure 1: A feature model of a mobile phone product family

vide tools to support this framework with automated verification. This has resulted in the use of Modal Transition Systems (MTSs) as a formal method for describing the combined operational behaviour of an entire product family. An MTS is a Labelled Transition System (LTS) that distinguishes between optional (may) and mandatory (must) transitions. Since MTSs cannot model all variability, we enrich them with an additional set of variability constraints.

We have also developed a tool for the computer-aided verification of behavioural variability in product families, namely VMC (Variability Model Checker). VMC takes as input a high-level description of an MTS together with a set of textual constraints. In our mobile phone example the latter would include “GPS EXC Basic” and “Camera REQ High resolution”.

VMC allows the user to interactively explore an MTS of a product family; model check properties (branching-time temporal logic formulae) over an MTS; visualise the (interactive) explanations of a verification result; generate one, some, or all of the family's valid products (represented as LTSs); browse and explore these; model check whether or not products (one, some, or all) satisfy certain properties; and, finally, help the user to understand why a certain valid product does or does not satisfy specific verified properties

by allowing such a product to be inspected individually. Figure 2 shows VMC's capability to model check a temporal logic formula over all valid products of a family.

The core of VMC consists of a command-line version of the model checker and of a product generation procedure. These programs are stand-alone executables written in Ada that can easily be compiled for the Windows, Linux, Solaris, and Mac OS X platforms. These core executables are wrapped with a set of CGI scripts handled by a web server, making it easy to build an html-oriented GUI as well as to integrate graph drawing tools. The development of VMC is ongoing, but a prototypical version is publicly usable online (see the link below) while its executables are available upon request.

Currently, VMC is not targeted at very large families. Its main limitation, however, is the generation of the model from its input language, while its on-the-fly verification engine and advanced explanation techniques are those of the highly optimized family of on-the-fly model checkers developed at ISTI-CNR over the last few decades for verifying formulae in action and state-based branching-time temporal logics derived from the CTL family of logics. The on-the-fly nature of VMC means that in general it is not necessary to generate and explore the whole state space. This feature improves its performance and allows it to deal with infinite-state systems.

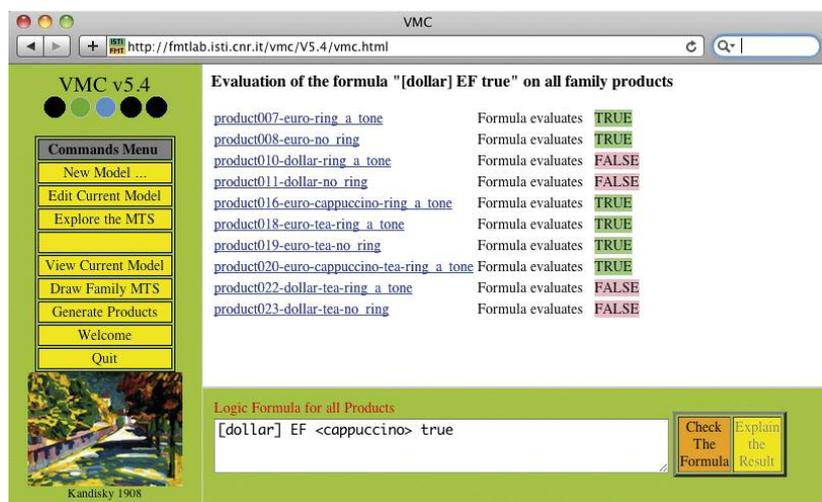


Figure 2: Verification result of a property for all products of a family of coffee machines

Link:

VMC: <http://fmt.isti.cnr.it/vmc/>

References:

[1] D. Benavides, S. Segura, and A. Ruiz-Cortés: “Automated analysis of feature models 20 years later: A literature review.” Information Systems 35(6) pp. 615-636, 2010. <http://dx.doi.org/10.1016/j.is.2010.01.001>

Please contact:

Maurice ter Beek
ISTI-CNR, Italy
E-mail: maurice.terbeek@isti.cnr.it

Wood Variety Recognition on Mobile Devices

by Pavel Vácha and Michal Haindl

Each type of wood has its own specific physical, aesthetic and economic properties; thus correct identification of wood species is required in numerous practical applications, from construction industry, manufacturing, furniture design, and restoration to pricing evaluation of wooden items. Fast, reliable, and practical recognition of wood species is therefore important, having potential impacts in a range of areas, including: the intended application, construction safety, and detecting illegal logging of endangered species. A mobile Android application that can automatically recognize wood species from a phone photo has been developed at the Institute of Information Theory and Automation AS CR in Prague..

The traditional method of identifying wood species involves manual browsing through digital wooden veneer catalogues and making a subjective judgement. This is labour intensive, and concentration problems can lead to errors. Additionally, gradual changes and changing shades due to variable light conditions are difficult for humans to detect. As an alternative, we aimed to develop an application to identify wood

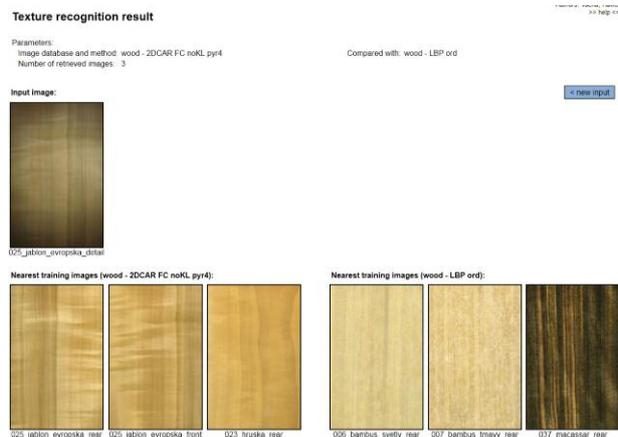


Figure 1: Apple wood sample taken by a smartphone camera and the three closest query results using Markovian and LBP textural features.

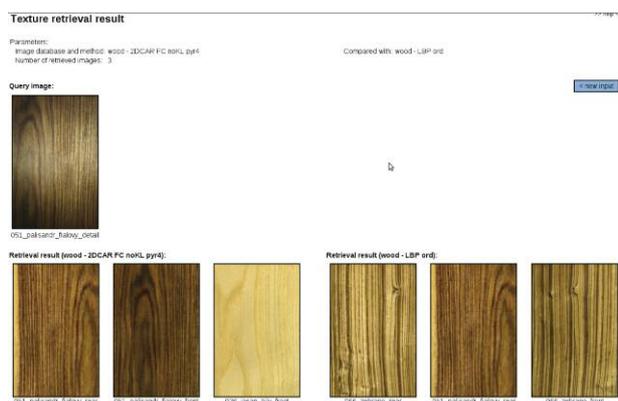


Figure 2: Comparison of results for Palisander wood smartphone retrieval.

species using a smartphone camera, which returns the resulting species name and a corresponding high quality database wood specimen image. This computer-aided wood identification system retrieves a wood template from a digital wood database, selecting that which most closely resembles the query sample.

A wooden surface is captured by a smartphone camera with the developed Android application, and the image is transmitted to the server side which computes the advanced multi-spectral Markovian textural features [1] and finds the most similar wood species from its database. Markovian features are not only very efficient, compacting the representation of visual properties of wood, but are also invariant to illumination colour and robust to illumination heterogeneity and illumination direction, therefore the retrieval result is not influenced by illumination properties. The recognized wood species together with its high quality database pattern is sent back to the user so he or she can verify the classifier's result. The challenging part of the method is to compare poor quality smartphone images taken under variable illumination and resolution conditions with high quality high resolution matte wooden textures stored in the wood database.

The performance of our application was verified on the wood database, which contains about sixty different European as well as exotic wood species, each with two sample images. Our Markovian textural features (2D CAR) [1] compared favourably with alternative Local Binary Patterns (LBP) [2] and Gabor textural features [3]. The test images were photographed from approximately the same distance and with camera oriented along tree rings to reduce variation in scale and rotation. These simple non-restrictive conditions were enough to raise the performance of the features above their rotation invariant versions. Our correct recognition rate was 66 %; LBP features 22 %, while the Gabor features completely failed due to non-homogeneous smartphone flash illumination. The wood recognition rate significantly improves (ie, Markovian features 81 %, LBP 55 %, Gabor features 48 %) if the input samples are measured with a high-end consumer digital camera. Although smartphone cameras have sufficient resolution (up to 10 mega pixels), their poor quality lenses and aggressive JPEG compression result in overall inferior image quality and thus a more demanding recognition task. Figures 1 and 2 illustrate the system's typical performance applied to European and exotic wood samples.

References:

- [1] P. Vácha, M. Haindl, T. Suk: "Colour and rotation invariant textural features based on Markov random fields", Pattern Recognition Letters, vol.32, no.6, pp.771-779, 2011
- [2] T. Ojala, M. Pietikäinen, T. Mäenpää: "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns", IEEE Trans. Pattern Anal. Mach. Intell. 24(7), 971-987 (2002)
- [3] B. S. Manjunath and W. Y. Ma: "Texture Features for Browsing and Retrieval of Image Data", IEEE Trans. on Pattern Anal. and Mach. Intelligence, 18(8), 1996.

Please contact:

Pavel Vácha, CRCIM (UTIA), Czech Republic
Tel: +420 2 6605 2365, E-mail: vacha@utia.cas.cz

The New SHIELD Architectural Framework

by Mariana Esposito, Andrea Fiaschetti, Francesco Flammini

New SHIELD (nSHIELD) is a European research project co-funded by the Artemis Joint Undertaking (Sub-programme SP6) focused on the research of innovative solutions for security, privacy, dependability (SPD) in the context of embedded systems (ES), including those employed in railway security applications.

The nSHIELD project aims at addressing SPD issues as “built in” rather than as “add-on” functionalities, by adopting an innovative holistic approach. We perceive this strategy as being the first step towards SPD certification for future ES.

The leading ideas at the basis of this research are: i) to enrich the state-of-the-art with new SPD solutions and ii) to enable the composability of these (new or already existing) solutions. This will be achieved in two steps. First, starting from current SPD solutions, the project will develop new technologies and consolidate those already explored in pSHIELD (a SHIELD pilot project) in a solid basement that will become the reference milestone for a new generation of “SPD-ready” ES. Second, these technologies will be then enhanced with the “composability” functionality that is being studied and formalized.

In a nutshell, composability is the possibility of dynamically activating one or more SPD functionalities in order to achieve a desired SPD level. This is possible with the implementation of the following enabling mechanisms and technologies:

- Semantic description of security domain and system components, in order to have a machine-understandable language to drive the automatic composition.
- SPD Metrics, in order to quantify the security needs and the achieved security level over heterogeneous environments
- Security Agent, the engine is in charge of continuously monitoring the environment to look for new components or new security needs
- Policies and control algorithms to provide a solution for the “composition problem”, ie how to put together the available SPD technologies in order to achieve the security target.

nSHIELD will approach SPD at 4 different levels: node, network, middleware and overlay (see Figure 1). For each level, the state of the art in SPD of individual technologies and solutions (ranging from hardware and communication technologies to cryptography, middleware, smart SPD applications, etc.) is expected to be significantly improved and inte-

grated into the so-called SHIELD architectural framework, which will represent the breakthrough result of the project.

The nSHIELD consortium comprises 6 manufacturers and system integrators, 7 universities, 9 SMEs and 2 Industrial R&D organizations, mostly members of ARTEMISIA, the European community of experts in the Embedded Security domain. The project is led by an industrial partnership (70% of the effort) although the universities and research centres involved in the project have an important role (30%) in contributing the needed innovation.

The main objective of the project is to conceive and design an innovative, modular, composable, expandable and high-dependable architectural framework. nSHIELD will achieve the desired SPD level in the context of integrated and inter-operating heterogeneous services, applications, systems and devices, and will develop concrete solutions capable of achieving this objective in specific application scenarios

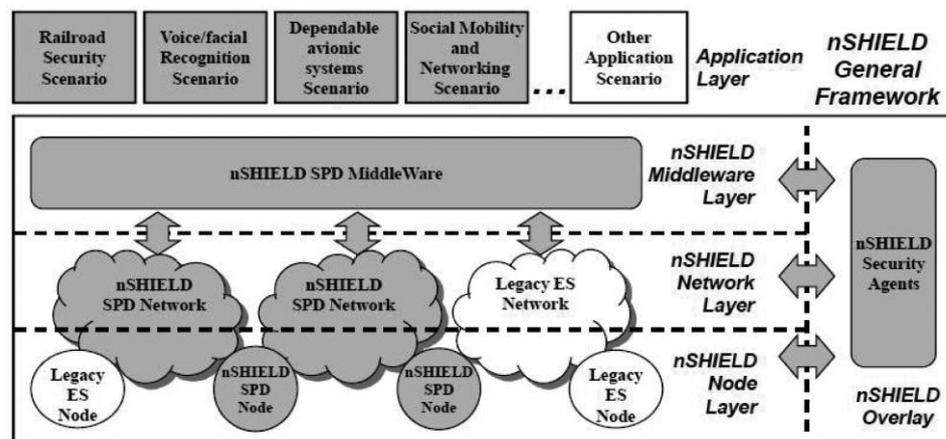


Figure 1: The nSHIELD framework

with minimum engineering effort. Four scenarios have been carefully selected in order to cover a wide and significant range of expected industrial needs.

One of these scenarios addresses dependable surveillance systems for rail-based transit security, but the aim is to extend applicability also to safety-critical (the so-called “vital”) subsystems in railway signalling, control and supervision. In these contexts, the composability of the SHIELD architectural framework will have great impact on the system design costs and time to market of new products and solutions. At the same time, the integrated use of SPD metrics in the framework will impact on the development cycles of SPD in ES because the qualification, (re-)certification and (re-)validation process of a SHIELD framework instance will be faster, easier and widely accepted.

Link: <http://www.newshield.eu/>

Please contact:

Mariana Esposito, University of Naples, Italy

E-mail: mariana.esposito@unina.it

Andrea Fiaschetti, University of Rome, Italy

E-mail: fiaschetti@dis.uniroma1.it

Francesco Flammini, Ansaldo-STTS, Italy

E-mail: francesco.flammini@ansaldo-sts.com

A Biofeedback System for Self Empowerment and Improved Quality of Life

by Johanna Mercurio

We have designed "Affective Health", a biofeedback system that supports personal empowerment rather than setting tasks for our bodies. Rather than reducing our bodies to objects that can be kept in perfect shape, our approach looks at what it means to be human.

A wave of wellness and health systems focusing on biofeedback are gaining popularity in the commercial world through apps that mirror factors such as one's pulse, sweat levels or energy consumption. Commercial systems such as the Nike Fuel-band, Polar's pulse meter and Affectiva's bracelet are allowing users to interact with their arousal reactions or



Researcher Johanna Mercurio with the Affective Health system.

energy consumption in real-time, as part of their everyday life. Most existing biofeedback systems aim at relieving pain, creating interesting experiences, or making us healthier.

Our interest in bio-sensor-based interaction began a decade ago when we were exploring affective interaction design [1]. In affective interaction design, meaning, dialogue, and emotions are created through the interaction itself – they are not given entities to be recognized, identified and treated by the system.

Affective Health is a system consisting of a bio-sensor wristband, developed by Philips Research, which measures movement (tri-axial accelerometer) and arousal level (Galvanic Skin Response). The bio-sensor data is transferred, via Bluetooth, to a mobile phone in real-time and it logs the data on the phone. A mobile application visualizes the move-

ments' trough shapes while arousal is fluently represented by a colour scheme. The measurements can indicate one's current bodily state relative to the history of previous states. It can portray situations that are stressful and engaging as well as peaceful moments.

The system is designed to encourage the users to understand and interpret the feedback it gives them. The system is built on the notion of interactional empowerment [2] with the design allowing users to form their own, personal interpretation with the bio-sensor data and construct the meaning of the picture painted by the system over time. In this way the system empowers users to identify, reflect on, and find patterns in, their behaviour through a value-free design stance. It is through the interaction over time that the system starts making sense, mirroring bodily reactions or the users' activities back to them. An interactional perspective on design will not aim to detect, and inform the user of, a singular account of the "right" or "true" interpretation of the user's behaviour, but rather make experiences available for reflection and interaction. By carefully crafting the feedback from the system, timing its response, and making the design mirror the users' visceral, emotional reactions, the system pulls the users into what we term an "affective loop experience" [3]. In this interaction, users actively take control over their own somatic reactions, learn how to master their own bodies and acquire a better understanding of themselves.

Our long-term trials of the Affective Health system show that people need more social solutions. People want to engage socially, exercise together, and chat. This raises the question: To what context would Affective Health be best suited? We have developed initial design concepts to understand which data users would like to share, for instance contextual data from social platforms or location or calendar information. It is also important to consider how the data could be stored in the cloud in a safe and secure way. The next step is to further explore how bio-data could be shared socially in real-time or past time, remotely or directly.

The Affective Health system is developed by SICS and funded in collaboration with EIT ICT Labs. The start-up company for the system, Mirrormirror, is in the initial stages of commercialization.

Link: <http://www.sics.se/ah>

References:

- [1] P. Fagerberg, A. Ståhl, and K. Höök: "Designing gestures for affective input: an analysis of shape, effort and valence", in *proc. of MUM 2003, Norrköping, Sweden, 2003*
- [2] K. Höök, A. Ståhl, P. Sundström, and J. Laakso: "Interactional empowerment, in *proc. of CHI'08, ACM Press, 647-656, 2008*
- [3] A. Ståhl, P. Sundström, and K. Höök: "A foundation for emotional expressivity", in *proc. of Designing for User eXperience 2005, AIGA: American Institute of Graphic Arts, 33.*

Please contact:

Johanna Mercurio, SICS Swedish ICT
E-mail: johmerc@sics.se

Software and Hardware-Intensive Activities for Supporting Creative Learning

by Michail Giannakos and Letizia Jaccheri

Our Toys project is investigating how teenage girls participate in the creative production of digital artifacts. Furthermore, the project introduces them to digital and computer literacy. We capture the students' attention with character development using physical artifacts, and use computer programming to bring the character to life.

Digital artifacts that enable people to exchange, create, and distribute information have, in the past couple of decades, profoundly reshaped the way we work and live. The creative production of digital artifacts in learning activities has been linked to teaching new computer literacy skills [1]. Beyond desktops, ubiquitous technologies not only allow a more

appealing objects where children start to work without being prompted by adults.

The children attending the workshops were instructed and assisted by the programming artist, the leader of the ReMida centre, a senior researcher, the project manager and two HCI researchers. The children completed and published six interactive works and eleven installations (see example, Figure 1) based on the software/hardware and the recycled materials within the centre. At the end of the project, students' interactive projects were presented in an official exhibition in the university (<http://www.itovation.org>).

Children engaged in programming languages (ie, Scratch) and programmable hardware platforms (ie, Arduino), which enable them to engage in the world of creativity with digital enriched artifacts, like robots and interactive installations.

The early results of our initiative [2] showed that introducing young people (in particular, girls) to creative software and hardware intensive activities is a good way to attract students to the field of Computer Science by increasing their interest and engagement. In particular, we found out that our approach [2]: raises awareness of technology, makes the

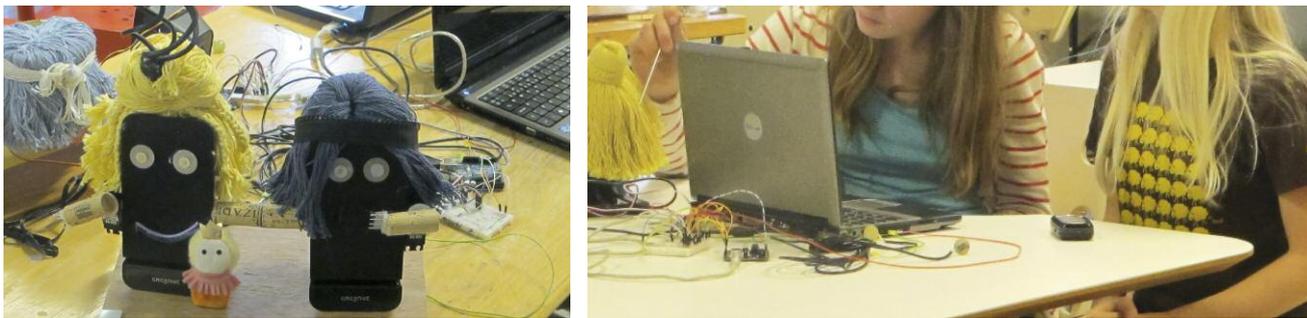


Figure 1. Example of interactive installations (left); giving life on the installations through programming (right).

active and physical engagement, but also provide the opportunity for novel and creative interactions. Physical and digital enhancement also provides the facility to convey experiences that are not possible in the physical world, for example, turning our thoughts into reality (ie, through a 3D printer). In turn, this physical and digital enhancement can provide opportunities to encourage or even enhance further exploration, discovery, reflection, and collaboration.

In our approach, we designed and developed a workshop program called “Our Toys”. Our Toys is based on open source software and hardware and consists of tutorials on open source tools, artifact-enriched creative sessions and students' demonstrations/presentations.

To date, 66 children have participated in our workshop program. The workshops took place at creative centres of University and ReMida centre. ReMida houses a variety of materials for use in creative and educational projects. The centre is a cooperation between the municipality, the education project, Reggio Children, the municipal waste company (recycling) and the local business community. Students worked according to Reggio Emilia education principles. The main idea is that the child takes the initiative in creative actions. ReMida centres are creative places with many

experience more intense, invites children to explore boundaries, and increases collaboration and the exchange of views and ideas.

Our Toys project is partly supported by the ERCIM “Alain Bensoussan” Fellowship Programme and NTNU, in collaboration with ReMida centre and Trondheim Electronic Art Centre.

References:

[1] L. Buechley, M. Eisenberg, J. Catchen, A. Crockett: “The LilyPad Arduino: Using Computational Textiles to Investigate Engagement, Aesthetics, and Diversity in Computer Science Education” in proc. of CHI '08, ACM Press, 423-432

[2] M.N. Giannakos, L. Jaccheri and R. Proto: “Teaching Computer Science to Young Children through Creativity: Benefits and Perspectives”, in proc. of CSERC '13, ACM Press.

Please contact:

Letizia Jaccheri, Michail Giannakos (ERCIM Fellow)

NTNU, Norway

E-mail: letizia.jaccheri@idi.ntnu.no, michailg@idi.ntnu.no

W3C at the Mobile World Congress 2013

by Marie-Claire Forgue

The World Wide Web Consortium (W3C) was present at the Mobile World Congress 2013 25-28 February 2013 to demonstrate the impact of the Open Web Platform and HTML5-based Web apps, proposed as the best development solutions addressing reduction in costs and complexity. Nowadays, HTML5-based Web applications are enhanced with new capabilities such as responsive design, WebRTC (Web Real-Time Communication), near field communication (NFC), Web payments, and full integration of video.

In 2005, W3C launched its Mobile Web Initiative, and, since then, has been present at almost every Mobile World Congress, one of the biggest events



W3C booth at the Mobile World Congress

focused on mobile devices and networks. Back in 2005, any claim that the Web had a big role to play on mobile devices was met with incredulous smiles and funny stares at best.

Fast forward to the 2013 edition of the congress, and while W3C certainly had the largest HTML5 logo on display at its booth, the real message came from the spread of these logos on many, many other stands. Not only has the Web gained visibility on mobiles, with several systems now providing HTML5-based development platforms for their systems (Windows 8, Blackberry, Tizen), but also, one of the highest profile announcements at the congress was around FirefoxOS, Mozilla's mobile operating system entirely based on Web technologies.

In 2013, W3C showcased demonstrations of HTML5 and other open W3C

Web technologies at its booth in the "App Planet" Hall of the new congress venue, Fira Gran Via, in Barcelona. W3C presented two Web apps featuring the creation of an online photo gallery for a variety of devices. In addition, W3C Members such as Adesis, Ayuntamiento de Zaragoza, Intel, Joshfire, Mozilla, Nokia, Tapquo, and partners such as the EU project webinos, demonstrated their unique Open Web Platform applications.

The W3C booth received much attention from developers, industry leaders, the media, analysts and many others. With this large and broad audience, W3C staff and W3C members were pleased to discuss how the Web is transforming all industries, including mobile, television, advertising, games, publishing, automotive, and health care. Specifically, a wide range of stakeholders are now looking at accelerating the adoption of Web technologies in the automotive industry, leading to the

recent creation of the W3C Automotive and Web Business Group. Also, the February 2013 W3C electronic books workshop reported technical discussions focused on Open Web Platform technologies currently used in eBooks and the need for improvements of these technologies for future digital publications.

ERICIM is the European host of W3C.

Links:

W3C@MWC'13:

<https://www.w3.org/2013/MWC/>

webinos project: <http://webinos.org/>

W3C eBooks workshop report: <http://www.w3.org/2012/08/electronic-books/rapportebok.html>

W3C Automotive and Open Web Platform Business Group:

<http://www.w3.org/community/autowebplatform>

PROMISE Winter School 2013 on Bridging between Information Retrieval and Databases

by Nicola Ferro

The main mission of the PROMISE EU FP7 network of excellence is to advance the evaluation and benchmarking of multimedia and multilingual information access systems. Together with the ELIAS research network on information access system evaluation, funded by the European Science Foundation, PROMISE organized a winter school on "Bridging between Information Retrieval and Databases" as a week long event in Bressanone, Italy, from 4-8 February 2013.

The aim of the school was to give participants a grounding in the core topics that constitute the multidisciplinary area of information access to unstructured, semi-structured, and structured information. The idea of the school stemmed from the observation that, nowadays, databases are increasingly using techniques that have traditionally been typical of information retrieval and, viceversa, the use of database-oriented techniques is becoming more common in information retrieval.

17 experts from academia and industry were invited to speak on a large variety of topics ranging from introductory talks on databases, information retrieval, experimental evaluation, metrics and statistics to advanced topics such as semantic search, database keyword search, semi-structured search, and performance evaluation. Focused lectures were devoted to bridging gaps between information retrieval and databases and to the management and sharing of research data via evaluation infrastructures. Hot topics concerned evaluation with respect to usefulness, crowdsourcing, evaluation on social media, and moving from evaluation to applications.

52 participants from 16 countries attended the courses (17% MsC students, 63% PhD students, 10% post-



Participants of the PROMISE Winter School

docs, 10% academic) with a background mostly on databases (32%), information retrieval (40%), both (15%), and natural language processing (9%). 15 scholarships (supported by ELIAS) were made available to assist participation. The multidisciplinary of the participants and lectures stimulated lively exchanges of ideas and many questions. Most of the speakers stayed for the entire week, enriching the discussions. Interestingly enough, the school turned out to be a brainstorming and discussion opportunity also for the lecturers, since they had the occasion of meeting colleagues from different fields with different perspectives, on a common ground of shared topics and issues.

To favour discussion and knowledge sharing, participants were asked to bring a poster describing their own research activities and plans. The three best posters were awarded a small prize and the winners were invited to contribute a short paper on their activities for the volume on the school lectures, currently under preparation.

The school presentation slides can be found at the link below. Information on all PROMISE publications, including the Proceedings of the PROMISE 2012 winter school on “Information Retrieval meets Information Visualization” can be accessed on the main PROMISE web site.

Links:

PROMISE: <http://www.promise-noe.eu/events/winter-school-2013/>

ELIAS: <http://www.elias-network.eu/>

Please contact:

Nicola Ferro

University of Padua, Italy

E-mail: ferro@dei.unipd.it

ATLANTIS PREMIUM PROCEEDINGS

Maximum visibility for your Conference and Workshop Proceedings

- **Open Access:** all papers are available in Open Access and included in the major search engines, thus offering maximum visibility;
- **Professional:** our publishing services include ISBN, DOIs, dedicated site, search engine and indexing services;
- **Fast:** production time is 4 weeks only;
- **Cheap:** only 30-50 euro per article.



For more details and examples, please check:
<http://www.atlantis-press.com/proceedings/>

Call for Contributions

ERCIM/EWICS Workshop on Dependable Embedded and Cyberphysical Systems at SAFECOMP 2013

Toulouse, France, 24 September
2013

The ERCIM Working Group on Dependable Embedded software-intensive Systems will again organize a workshop at the SAFECOMP conference jointly with EWICS TC7, the European Workshop on Industrial Computer Systems, TC7, Safety, Reliability and Security.

The first day is dedicated to workshops and tutorials. The organisers would like to invite ERCIM members to contribute to and participate in our workshop which can be attended independent from SAFECOMP. No extra costs will arise for participants since the additional cost is covered by the dissemination budgets of the organizers.

Embedded systems are everywhere – comfort, health, services, safety and security of people depend increasingly on them. In combination and close interaction with the real-world environment and humans, they become so-called “Cyber-physical Systems”, acting independently, co-operative or as “systems-of-systems”. The impact on society as a whole is tremendous. Smart (embedded) systems are regarded as the most important business driver for European industry. They are a targeted research area for European Research Programmes in Framework 7, in the

ARTEMIS Joint Undertaking, and in several dedicated Programmes and European Technology Platforms (ARTEMIS, EPoSS) or the future JTI Electronics in Horizon 2020. Their application is not only in the traditional areas of aerospace, railways, automotive, or process industry and manufacturing, but also in robotics and services of all kind, in home appliances (smart environments, smart homes, ambient assisted living) and health care.

Sessions are planned on the following topics:

- dependable and resilient embedded systems
- autonomous systems and robotics
- systems-of-systems.

These sessions will cover aspects from design, development, verification and validation, certification, maintenance, standardization and education & training. In contrast to the SAFECOMP conference mainstream, the workshop will include reports on on-going work aiming at discussions and experience exchange. Reports on European or national research projects as well as industrial experience reports are welcome.

Workshop proceedings will be published by LAAS-CNRS, and distributed to participants during the workshop. The workshop papers will also be published online on the HAL/Arxiv open publication site.

Important dates:

- 21 May 2013: Full paper submission deadline
- 15 June 2013: Notification of authors
- 28 June 2013: Final camera-ready paper due

The international programme committee is composed of selected EWICS and ERCIM members, led by the workshop organizers (see contacts below).

More information:

<http://conf.laas.fr/SAFECOMP2013/>

Please contact:

Erwin Schoitsch, AIT Austrian Institute of Technology, Austria
E-mail: erwin.schoitsch@ait.ac.at

Amund Skavhaug

NTNU, Norway

E-mail: amund.skavhaug@ntnu.no

Call for Abstracts

NETTAB 2013 - Workshop on Semantic, Social and Mobile Applications for Bioinformatics and Biomedical Laboratories

Venice, Italy 16-18 October 2013

NETTAB Workshops are a series of international meetings on “Network Tools and Applications in Biology” introducing the most innovative ICT tools as they are being applied to the biomedical domain. In the past, workshops have been focused on themes like multi-agent systems, scientific workflows, Web Services, Semantic Web, collaborative research, and biological wikis.

ICTs have permeated society with newer forms of social participation. In biology, we already rely on many social tools and applications, eg for distributed annotations, Wiki knowledge bases, documentation and productivity. Internet is increasingly accessed with mobile devices. Health and lifestyle mobile applications are widely used and the rapid adoption of mobile solutions in medicine and healthcare is already a reality, but we cannot say the same for life sciences. Semantic methodologies and technologies are instead well established in “-omic” projects and the bioinformatics community was an early adopter of Semantic Web technologies.

In the NETTAB 2013 workshop, mobile, social and semantic solutions for bioinformatics and laboratory informatics problems will be explored. A savvy combination of these technologies could greatly enhance the research outcome of life scientists and markedly simplify the workflows in biomedical laboratories.

Deadlines

- 5 July: Oral communications submission
- 31 July: Posters submission.

More information:

<http://www.nettab.org/2013/>

Call for Papers

Eighteenth International ERCIM Workshop on Formal Methods for Industrial Critical Systems

Madrid, 23-24 September 2013

The aim of the FMICS workshop series, is to provide a forum for researchers who are interested in the development and application of formal methods in industry. In particular, FMICS brings together scientists and engineers that are active in the area of formal methods and interested in exchanging their experiences in the industrial usage of these methods.

Topics of interest include:

- Design, specification, code generation and testing based on formal methods. Methods, techniques and tools to support automated analysis, certification, debugging, learning, optimization and transformation of complex, distributed, real-time systems and embedded systems.
- Verification and validation methods that address shortcomings of existing methods with respect to their industrial applicability (eg, scalability and usability issues).
- Tools for the development of formal design descriptions.
- Case studies and experience reports on industrial applications of formal methods, focusing on lessons learned or identification of new research directions.
- Impact of the adoption of formal methods on the development process and associated costs.
- Application of formal methods in standardization and industrial forums.

Important Dates

- 3 May: Paper submission
- 24 June: Notification
- 12 July: Final version due
- 23-24 September: Workshop

The workshop is organised by the ERCIM Working Group “Formal Methods for Industrial Critical Systems

More information:

<http://lvl.info.ucl.ac.be/Fmics2013/Fmics2013>

CWI researcher Floor Sietsma Youngest PhD in the Netherlands

Floor Sietsma (20) of CWI in Amsterdam is the youngest PhD in the recent academic history of the Netherlands. Recently elected “Nerd of the Year” by readers of the Dutch magazine *Quest*, she was 20 years old when she defended her thesis “Logics of Communication and Knowledge” on 13 December 2012 at the University of Amsterdam.

Sietsma’s research concerns knowledge transfer in communication and makes use of epistemic logic, a logic of knowledge. She applies this to analyze shared knowledge in complex e-mail exchanges with visible (cc) and hidden (bcc) receivers. She also develops a complete logic of manipulative communication, for instance in a game of Liar’s Dice, where players mutually agree to lie and deceive.

Sietsma started her PhD research at CWI in 2010 at age 17. She held a Bachelor’s degree in computer science and a Master’s degree in logic, and was at the age of 12 the youngest student ever at a Dutch university. For her research at CWI she was awarded a special, personal grant by the Netherlands Organisation for Scientific Research (NWO), because of her unique talents. 2.5 years later, she already had gathered enough material to defend her thesis. Sietsma also studies Pedagogical Sciences to be able to help gifted children.



Photo: CWI

Warsaw Center of Mathematics and Computer Science Established

Warsaw Center of Mathematics and Computer Science (WCMCS) is a joint project of two scientific institutions: the Faculty of Mathematics, Informatics and Mechanics of the University of Warsaw (MIMUW), and the Institute of Mathematics of the Polish Academy of Sciences (IMPAN). The Center is built on the longstanding cooperation, in both teaching and research, between the two. The Center was designated as the “Leading National Research Center” (Krajowy Naukowy Ośrodek Wiodący, KNOW) by the Polish Ministry of Science and Higher Education in July 2012. The Center organizes scientific seminars, schools and conferences, and also participates in several research grants.

WCMCS offers a wide range of cooperation projects addressed to students, PhD students, young and senior researchers. The current projects include student internships, PhD internships, PhD fellowships, post-doctoral positions, as well as research positions for the leaders of research projects. For more detailed information about the research areas, the current projects, and the application procedures, please visit the WCMCS homepage (<http://wcmcs.edu.pl>).

Advertisement

OpenConf

**The world’s most user-friendly
peer-review, abstract, and
conference management software**

Simple.
OpenConf seamlessly automates submission, reviewer assignment, discussion and selection, allowing for multi-file uploads, program creation and much more

Reliable.
Celebrating 10 years since its first public release, OpenConf delivers proven reliability and constant innovation, backed by the industry’s best customer service and support

International.
OpenConf has powered thousands of events in more than 100 countries, and provides authors and reviewers access in their native language with 12 included translations

Affordable.
Simplify your next event for less with an exclusive ERCIM reader discount:

OpenConf.org/ERCIM

WHAT'S ON MY
CARPET
TRIANGLES
OR SQUARES?

What mathematician in Warsaw developed the fractal construction named after him?
Find out at wcmcs.edu.pl

Warsaw Center
of Mathematics
and Computer Science

» phd program in mathematics and computer science
» post-doctoral positions
» research positions

MIMUW
KNOW



AARIT
Austrian Association for Research in IT
c/o Österreichische Computer Gesellschaft
Wollzeile 1-3, A-1010 Wien, Austria
<http://www.aarit.at/>



CNR
Consiglio Nazionale delle Ricerche, ISTI-CNR
Area della Ricerca CNR di Pisa,
Via G. Moruzzi 1, 56124 Pisa, Italy
<http://www.isti.cnr.it/>



Crcim
Czech Research Consortium
for Informatics and Mathematics
FI MU, Botanicka 68a, CZ-602 00 Brno, Czech Republic
<http://www.utia.cas.cz/CRCIM/home.html>



CWI
Centrum Wiskunde & Informatica
Science Park 123,
NL-1098 XG Amsterdam, The Netherlands
<http://www.cwi.nl/>



Fonds National de la Recherche Luxembourg
Fonds National de la Recherche
6, rue Antoine de Saint-Expéry, B.P. 1777
L-1017 Luxembourg-Kirchberg
<http://www.fnrl.lu/>



FWO
Egmontstraat 5
B-1000 Brussels, Belgium
<http://www.fwo.be/>

F.R.S.-FNRS
rue d'Egmont 5
B-1000 Brussels, Belgium
<http://www.fnrs.be/>



FORTH
Foundation for Research and Technology – Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
<http://www.ics.forth.gr/>



Fraunhofer IUK-TECHNOLOGIE
Fraunhofer ICT Group
Anna-Louisa-Karsch-Str. 2
10178 Berlin, Germany
<http://www.iuk.fraunhofer.de/>



Inria
Institut National de Recherche en Informatique
et en Automatique
B.P. 105, F-78153 Le Chesnay, France
<http://www.inria.fr/>



NTNU
Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and Electrical Engineering, N 7491 Trondheim, Norway
<http://www.ntnu.no/>



I.S.I. - Industrial Systems Institute
Patras Science Park building
Platani, PATRAS, Greece, 265 04
<http://www.isi.gr>



PEG
Portuguese ERCIM Grouping
c/o INESC Porto, Campus da FEUP,
Rua Dr. Roberto Frias, n° 378,
4200-465 Porto, Portugal



Science & Technology Facilities Council
Rutherford Appleton Laboratory
Chilton, Didcot, Oxfordshire OX11 0QX, United Kingdom
<http://www.scitech.ac.uk/>



SparCIM
Spanish Research Consortium for Informatics and Mathematics
D3301, Facultad de Informática, Universidad Politécnica de Madrid
28660 Boadilla del Monte, Madrid, Spain,
<http://www.sparcim.es/>



SICS Swedish ICT
Box 1263,
SE-164 29 Kista, Sweden
<http://www.sics.se/>



SIRA
Swiss Association for Research in Information Technology
c/o Professor Abraham Bernstein, Department of Informatics,
University of Zurich, Binzmühlestrasse 14, CH-8050 Zürich
<http://www.sarit.ch>



MTA SZTAKI
Magyar Tudományos Akadémia
Számítástechnikai és Automatizálási Kutató Intézet
P.O. Box 63, H-1518 Budapest, Hungary
<http://www.sztaki.hu/>



University of Cyprus
P.O. Box 20537
1678 Nicosia, Cyprus
<http://www.cs.ucy.ac.cy/>



University of Southampton
University Road
Southampton SO17 1BJ, United Kingdom
<http://www.southampton.ac.uk/>



University of Warsaw
Faculty of Mathematics, Informatics and Mechanics
Banacha 2, 02-097 Warsaw, Poland
<http://www.mimuw.edu.pl/>



University of Wrocław
Institute of Computer Science
Joliot-Curie 15, 50-383 Wrocław, Poland
<http://www.ii.uni.wroc.pl/>



VTT
Technical Research Centre of Finland
PO Box 1000
FIN-02044 VTT, Finland
<http://www.vtt.fi/>

ERCIM offers fellowships

- in Informatics and Applied Mathematics
- for PhD holders from all over the world
- in leading European research institutes

Deadline for applications:
30 April 2013

fellowship.ercim.eu

