Number 98 July 2014

ERCIM NEWS

www.ercim.eu

Special theme:

Smart Cities

Also in this issue:

Keynote:
Smart Cities
by Eberhard van der Laan, Mayor of Amsterdam

Research and Innovation:
How to Detect Suspect Behaviour at Sea

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 6,500 printed copies and is also available online.

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

Contributions
Contributions should be submitted to the local editor of your country

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

Next issue
October 2014, Special theme: Software Quality

Smart Cities

Smart Cities are a bit like football: Every city has a team working on a "Smart City" and wants to be the Smartest City in the world, and at the start of every season every supporter thinks his or her team will be the global champion. Various ranking systems exist, comparing cities on indicators ranging from energy consumption per capita to life expectancy; from Wifi coverage to crime-rates. In other words: Smart Cities are about everything and therefore about nothing. To be frank: I do not really believe all this Smart City marketing. I do believe that innovation and technology gives us the opportunity to improve the quality of life of the citizens and make our cities more competitive.

As a mayor of a European capital city, which is ranked reasonably high in most European rankings and comparisons, I will try to paint a picture of what a Smart City means to me and what it means for our city. As well as the benefits of Smart Cities, I also see a few risks that sometimes go unrecognized. But first a bit of history, because I strongly believe that the DNA of the city will never change and every development in Amsterdam is based on its (long and prosperous) history.

One could argue that Amsterdam was the first Smart City in the world; not with the start of our Smart City program in 2009, but as early as the end of the 16th century when Amsterdam began to grow in importance as a trading city. One of the key reasons that Amsterdam became one of the richest cities of that time, with over 50% of all the sea-vessels in the world departing and arriving at its harbour, was the availability of data on trade and cargo. In a physical square of 400 by 400 meters all the information on the cargo, destination and ownership of all these vessels was gathered. This enabled tradesman to trade cargo that hadn't even been offloaded from the vessel yet and to compare one another's products. The rich data enabled the start of the first stock exchange and the opportunity for everybody to invest in trade, generating investments never seen before. Financial newspapers shared the information with everybody who wanted it. Essentially, it was a data driven economy with Open Data avant la lettre.

The sharing of information provided a whole new dimension: "normal" citizens of Amsterdam could take part in the city's welfare and economic growth. This is why Amsterdam is not a city of bishops and kings but of citizens and merchants, without castles and cathedrals. In Dutch art from the 17th century ordinary people were painted; Vermeer's Milkmaid now displayed in the national museum is a great example.

Back to the now: what is the DNA of Amsterdam and how does it apply to our approach? In this context I would state that it is about: openness, entrepreneurship, collaboration and inclusion.

Smart Cities are about openness: open data, open infrastructures and open innovation are at the very heart of the development of our Smart City. Like the data on cargo and the access to the canals, every "Amsterdammer" should have

Photo: S. van der Torren

Eberhard van der Laan, Mayor of Amsterdam

access to our infrastructure. The sharing of information drives innovation because individuals and business can build on jointly gathered knowledge.

The second subject is entrepreneurship. New technologies will lead to new business opportunities, new business models and new types of companies; companies that are based on communities, sharing, or new products and services. This includes companies that have business models that will lead to reduced, rather than increased, energy consumption or car miles - companies such as Car2Go, Spaces or energy service companies.

The third subject, "Collaboration", is of course important on a European scale, for example within Eurocities or the Smart Cities and European programmes, but it is particularly important that strong collaborations occur between companies, knowledge institutes, government(s) and citizens on a local level. Everybody is a stakeholder in the development of the city and all stakeholders are equally important, which means there are hundreds of thousands of stakeholders.

This brings me to the last point: inclusivity - a city for all "Amsterdammers" and its visitors. All "Amsterdammers" should have the opportunity to be part of the city we are living in.

I also see many new challenges - challenges relating to privacy, the ability for everybody to benefit from these new technologies, and the rapidly changing business which will lead to a large shift in labour. I expect companies, governments and scientists to take responsibility for these challenges and find solutions to address the downside of these smart developments.

Having said this I would like to conclude on a simple note: A city is as great as the people living within it. This will still hold true in a smart city: it is all about the people who live, work and play there. As a Mayor I will make sure that our citizens are always ranked first. This will of course make us the Johan Cruyff of the world's Smart Cities.

Eberhard van der Laan

RESEARCH AND INNOVATION

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

EVENTS, IN BRIEF

ERCIM “Alain Bensoussan” Fellowship Programme

ERCIM offers fellowships for PhD holders from all over the world.

Topics cover most disciplines in Computer Science, Information Technology, and Applied Mathematics.

Fellowships are of 12-month duration, spent in one ERCIM member institute.

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 the 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 (a proof will be requested).

Application deadlines

30 April and 30 September

More information and application form:
http://fellowship.ercim.eu/

ERCIM 25 Years Celebration

The 25th ERCIM anniversary and the ERCIM fall meetings will be held at the CNR Campus in Pisa on 23-24 October 2014.

On the occasion of ERCIM’s 25th anniversary, a special session and panel discussion will be held on 23 October in the afternoon in the auditorium of the CNR Campus. Speakers and representatives from research, industry, the European Commission, and the ERCIM community will be invited to present their views on research and future developments in information and communication science and technology.

For more information, please contact:
Adriana Lazzaroni, CNR, Italy
E-mail: adriana.lazzaroni@iit.cnr.it

W3C and the Web of Things

W3C held a workshop on the Web of Things in Berlin on 25-26 June 2014, hosted by Siemens. The Web of Things is expected to have broad and sweeping economic and societal impact. Open standards will be critical to enabling exponential growth of the kind we experienced with the early days of the Web.

Until now, the focus has been on the devices and communications protocols. There is a growing awareness that the business opportunities will be centered on the associated services, moreover, the current situation is one of fragmentation with products being developed in isolation due to a plethora of IoT protocols and a lack of a shared approach to services.

The workshop examined the opportunities for open Web standards for service platforms in the network edge (e.g. home gateways) and the cloud, along with the challenges for security, privacy and the integration with the Web of data. A workshop report is now in preparation. This workshop was supported by the EU project Compose .

Links:
Compose project:
http://www.compose-project.eu/
WoT workshop:
http://www.w3.org/2014/02/wot/

COMPOSE - Converging the Internet of Services with the Internet of Things

COMPOSE (Collaborative Open Market to Place Objects at your Service) will create an ecosystem for unleashing the power of the Internet of Things (IoT) via an easy transformation to an Internet of Services (IoS). The technology developed by the project will enable the creation of base services, composite services, and applications stemming from and operating on smart objects.

COMPOSE aims at enabling new services that can seamlessly integrate real and virtual worlds through the convergence of the Internet of Services (IoS) with the Internet of Things. COMPOSE will achieve this through the provisioning of an open and scalable marketplace infrastructure, in which smart objects are associated to services that can be combined, managed, and integrated in a standardised way to easily and quickly build innovative applications. The project will develop novel approaches for virtualising smart objects into services and for managing their interactions. This will include solutions for managing knowledge derivation, secure and privacy- preserving data aggregation and distribution, dynamic service composition, advertising, discovering, provisioning, and monitoring.

COMPOSE is expected to give birth to a new business ecosystem, building on the convergence of the IoS with the IoT and the Internet of Content (IoC). The COMPOSE marketplace will allow SMEs and innovators to introduce new IoT-enabled services and applications to the market in a short time and with limited upfront investment. At the same time, major ICT players, particularly cloud service providers and telecommunications companies, will be able to reposition themselves within a new IoT- enabled value chain.

Technical Approach

The vision of the COMPOSE project is to advance the state of the art by integrating the IoT and the IoC with the IoS through an open marketplace, in which data from Internet-connected objects can be easily published, shared, and integrated into services and applications. The marketplace will provide all the necessary technological enablers covering both delivery and management aspects of objects, services, and their integration:

- Object virtualization to enabling the creation of standardized service objects
- Interaction virtualization - abstract heterogeneity while offering several interaction paradigms
- Knowledge aggregation to create information from data
- Discovery and advertisement of semantically-enriched objects and services
- Data Management to handle massive amounts and diversity of data/metadata
- Ad hoc creation, composition, and maintenance of service objects and services
- Security, heterogeneity, scalability, and resiliency, incorporated throughout the layers

Expected Impact

COMPOSE strives for a strong impact on a developing market by lowering barriers to develop, select, combine, and use IoT-based standardized value added services. This will be achieved by providing a complete ecosystem, and having it adopted by enterprises, SMEs, government-related bodies, and individual developers and end-users.

We hope that opening the door to this realm for smaller entities will lead to higher innovation. COMPOSE expects to aid by fostering a developers' community and advocating an open source/interfaces

Use-Cases

COMPOSE design, development, and validation will be based on innovative use cases highlighting different aspects of the platform:

- Smart Shopping Spaces: this use cases will pilot COMPOSE in shopping environments, focusing on the dynamic composition and delivery of services starting from products available in shops.
- Smart City (Barcelona): Ample amount and diversity of sensors are deployed at a Barcelona district under the supervision of a COMPOSE partner. Along with Barcelona's OpenData, COMPOSE intends to showcase life in a smart city by creating a group of city services for the citizens.
- Smart Territory (Trentino): With the collaboration of regional network providers, the tourism board, and meteorological data providers, COMPOSE will explore innovative services for tourists. This pilot aims to enhance the tourist experience by exploiting COMPOSE technologies for the creation of personalized, social- and environmentally-aware (web and mobile) tourism services and territory monitoring services that leverage the regional networking and environmental infrastructures

Survey on promising IoT Application Domains

The project is currently collecting useful insights about existing and future exploitation opportunities in the IoT field for the expected project outcomes. The questionnaire is targeted at industry and technology experts and asks to identify, on the basis of their experience, the most promising IoT application domains and then specify them in terms of: (i) existing market drivers and inhibitors and (ii) success stories and associated business models. Eventually, with respect to the outlined context, experts are kindly asked to provide their insights about the most promising features of the COMPOSE project. The questionnaire is available on the project web site.

W3C is a partner of the project through ERCIM EEIG. A task of W3C is to facilitate standardization and exploitation of the activities. Fraunhofer FOKUS the project comprises 12 partners from industry and academia as well as standardization bodies, including the ERCIM member Fraunhofer FOKUS.

Link: http://www.compose-project.eu

Please contact:
Philippe Hoschka, W3C
E-mail ph@w3.org

Introduction to the Special Theme

Future Cities and Smart Technologies: A Landscape of Ambition and Caution

by Theo Tryfonas and Ioannis Askoxylakis

The old Chinese curse of "may you live in interesting times" has never been more relevant than in the field of urban development. For some time now, the global urban population has exceeded the global rural population. Cities and city regions have therefore emerged as the only growth models capable of meeting the increased demands and strains facing global supply systems, which are seriously affected by population growth, climate change, globalization and international security issues. Given these constraints, urban development is becoming a tough challenge.

Smart technologies may play a role in addressing these issues. From sensor networks to wireless connectivity and big data analytics to the Internet of Things, our planners, architects, infrastructure engineers and facilities managers have the ability to harness the power of information generated within the boundaries of cities and the built environment. Transport planners can better understand the demand for services and passenger behaviour; similarly, energy providers can tailor their supply to the real needs of their customers. End users can benefit from personalized and timely service provision that takes into account the individual's location and situation - from way finding to micropayments to finding a good place to dine. There is enormous potential to use these technologies to establish more sustainable behaviour: for example, with energy use monitoring informed by smart meters, personalized transport assisted by e-ticketing and mobile technology, wireless infrastructure, to name just a few possibilities.

This capability, however, comes at a significant cost. Real time monitoring and bulk data collection call for high performance computational infrastructure, huge storage capacity, wide reaching connectivity, digital skills and a positive governance attitude. Furthermore, this technology creates the potential for mass surveillance operations by the state as well as global corporations, triggering grave concerns for individuals' privacy. This is a promising yet challenging field that has a significant role to play in creating a more sustainable urban future. Technologies deployed and tested across multiple cities provide insights into a future that we may wish to see, but also give us forewarnings about the challenges that lie ahead.

This special issue of ERCIM News captures the state of art of European research and provides an insight into smart systems and technologies that are already, or may become in the future, part of our daily urban lives. Although it is difficult to provide an exact definition for it so early in its development, these technological advances contribute fundamentally to a forward-thinking concept that has been termed 'Smart City', i.e., an urban environment that capitalizes on contemporary information management capabilities, including ICT infrastructure and applications, to improve on the delivery of public services, transport, health, sustainability, the economy and the overall wellbeing of its residents. A recent report from the Department of Business, Innovation and Skills (BIS) in the UK estimated the emerging global market of Smart Cities to worth a minimum of $400 billion. As an emerging domain, however, there are still a number of uncertainties to be addressed.

With a wide selection of articles from novel networking technologies to digital entrepreneurial processes, it is clear from this special issue that no single technology or vendor can deliver the promises of the Smart City. No amount of sensors or data will offer a solution without purposeful assembly of the right components, including technology, skills and governance. This is an important perspective which highlights the need for European policy, system architecture and systems integration to be informed by the challenges of understanding human behaviour, fostering creativity by developing the human capital and ensuring compliance with the data protection framework through balancing the need for security with respect for individual rights. What interesting times, indeed!

Please contact:
Ioannis Askoxylakis
ICS-FORTH, Greece
E-mail: asko@ics.forth.gr

Theo Tryfonas,
Faculty of Engineering, University of Bristol
E-mail: theo.tryfonas@bristol.ac.uk

Urban Future Outline - A Roadmap on Research for Livable Cities

by Martina Ziefle, Christoph Schneider, Dirk Vallée, Armin Schnettler, Karl-Heinz Krempels and Matthias Jarke

Urban Future Outline (UFO), is a project funded by the Excellence Initiative of the German states and federal governments at RWTH Aachen University (2013-2015). It aims at developing a holistic approach, in which energy needs, mobility demands, ecological and climatological requirements and human demands are addressed, against the background of socially responsible technology development in urban areas.

One of the key challenges for contemporary research is to develop and maintain livable, sustainable and resilient cities. For the first time in history, the majority of the world's population live in urban spaces. The complexity of urban areas necessitates interdisciplinary approaches to be taken in research and development. Goods and services have to be available at short notice. Mobility must be affordable and reliable. Technical infrastructure needs to be easily accessible. Citizens require appealing living quarters embedded in open spaces with green and blue infrastructure. Often, these demands are contradictory, thus imposing considerable challenges on urban planners.

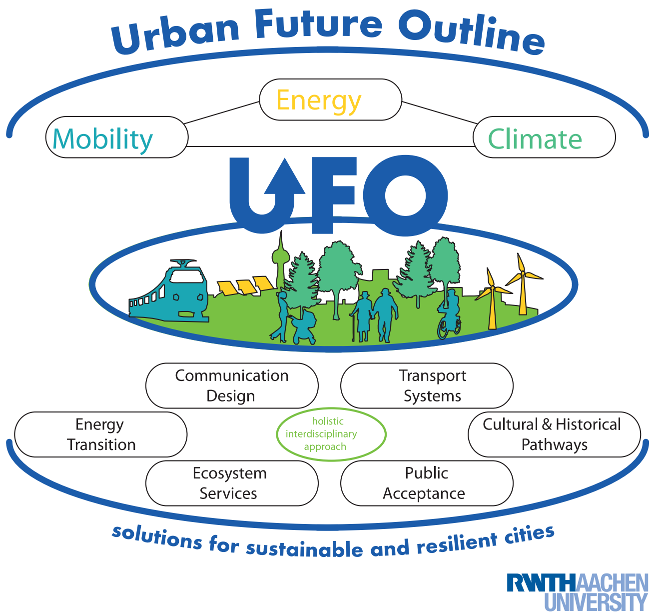

Figure 1: The overall approach and the topics addressed within UFO

UFO is organized in three subprojects: Future Mobility refers to sustainable infrastructure; Future Energy tackles questions regarding energy supply and the transition to renewable energy production; and the fields of biodiversity conservation and climate change adaptation are the targets of Future Ecology.

Future Mobility

Mobility is an essential prerequisite for people participating in social and economic lifes. In light of the peak oil debate, the EU guidelines for air pollution control and environmental noise and demographic developments, mobility concepts must change drastically. Transport systems need to be highly adaptive to periods of peak usage (i.e., rush hours) at multiple scales (e.g., day-by-day or seasonally). They must meet a range of community needs including accessibility, comfort, safety, sustainability, affordability and environmental justice. Mobility options must be intermodal, flexible and designed as "door-to-door" mobility chains. The predominantly technology-centered planning of infrastructural mobility concepts, without integrating citizens into the decision making processes, is no longer viable.

This sub-project combines city and mobility planning with the technical possibilities of information and communication, the cognitive and communicative requirements of urban populations, discursive practices, media strategies and gender aspects. The goal is to develop a multifactorial information and communication strategy that involves the social, individual, communicative and cognitive needs of the citizens, as well as technical and urban planning related aspects of feasibility.

Future Energy

This sub-project focuses on energy turnaround and the interrelation between the technical, economical, IT, ecological and social perspectives. Its aim is to develop a holistic model and a methodology for the implementation of sustainable, robust energy systems that systematically include social factors (i.e., user perceptions of energy systems) into the identification, planning and realization of energy scenarios. Within the context of environmental and technical conditions, acceptance-relevant factors (perceived as the benefits and drawbacks) from different regional contexts are collected, evaluated and modeled to determine their relationships. The inclusion of social knowledge is accomplished via three data sources: the empirical modelling of cognitive-affective attitudes, the analysis of opinion-forming processes on the Internet and an assessment of environment-related aspects. Potential decision trade-offs are identified using conjoint analyses.

The results are integrated into the development of technical, economical and IT transformation processes. Based on this modelling, an understanding on the technical, economical and social perspectives on energy transition are derived. These results are useful for politicians and decision makers and make an important contribution in developing transparent and adequate communication policies.

Future Ecosystem

Planning, implementing and sustaining livable urban spaces that pay equal attention to humans, technology and nature is a task that has not been solved holistically to date. Situations where combined stressors are acting pose specific challenges for urban development because the negative effects of superimposed stresses such as heat, noise and air pollution cannot be easily discerned when taking a single stressor perspective. This deficiency is all the more sensitive at certain times, for example summer, when thermic differences are more extreme and higher levels of air pollution are expected due to climate change.

This sub-project aims to (1) consider combined stress situations in a measurement and model chain, (2) derive combined stress indices and (3) make these accessible in a virtual environment (aixCAVE). Apart from the thermic component, air quality, acoustic perception and user perception are all considered. Which mix of combined stresses for urban residents inhabiting differently designed open spaces will be investigated, as will future scenarios. As opposed to previous approaches to modelling the effects of thermic and actinic stresses, (4) stresses will be diversified by user profiles, gender, and age.

Overall, the research design in UFO aims to make use of the most recent developments in technology, information sciences and natural sciences. At the same time, it sought to develop a feasible solution to integrating the social, economic and cultural needs associated with urban development [1]. An integral part of the UFO project, therefore, is the multi-disciplinary approach [2] [3]. The consortium consists of nine scholars from RWTH Aachen University:

- Christoph Schneider - Physical Geography
- Martina Ziefle - Communication Science
- Dirk Vallée - Urban and Transport Planning
- Matthias Jarke - Information Systems
- Armin Schnettler - High Voltage Technology
- Eva-Maria Jakobs - Textlinguistics
- Carmen Leicht-Scholten - Gender and Diversity
- Janina Fels - Technical Acoustics
- Andreas Schäffer - Environmental Biology
- Peter Russell - Computer Aided Architectural Design
- Thomas Niehr - German Linguistics
- Thorsten Kuhlen - Virtual Reality

References:

[1] C. Schneider, B. Achilles, H. Merbitz: "Urbanity and Urbanization: An Interdisciplinary Review Combining Cultural and Physical Approaches", Land 2014, 3(1), 105-130

[2] S. Krengel, T. Falke, A. Schnettler: "Optimization Model for the Energy Supply in City Quarters", in proc. of CIRED 2013, Stockholm

[3] S. Himmel, M Ziefle, K. Arning,: "From Living Space to Urban Quarter: Acceptance of ICT Monitoring Solutions in an Ageing Society", in M. Kuroso (ed.), Human-Computer Interaction, Users and Contexts of Use. Springer LNCS 8006, pp. 49-58.

Please contact:

Matthias Jarke,
Fraunhofer-FIT, Germany
E-mail:
matthias.jarke@fit.fraunhofer.de

Martina Ziefle, Christoph Schneider
RWTH Aachen, Germany
E-mail: ziefle@comm.rwth-aachen.de,
christoph.schneider@geo.rwth-aachen.de

Urban Civics - Democratizing Urban Data for Healthy Smart Cities

by Sara Hachem, Valérie Issarny, Alexey Pozdnukhov and Rajiv Bhatia

The Urban Civics project brings together a multi-disciplinary, trans-atlantic team of experts working towards a common goal that citizens and governments collaborate in achieving participatory democracy in healthy cities.

Technologies that bring together the emerging visions for smart, connected and resource-efficient cities, and the revolutions in social media and mobile networking, offer unprecedented opportunities for participatory democracy, health and sustainability. Early proponents of Smart Cities have advanced the idea of hyper-connected city information systems as a means of promoting operational efficiency and natural resources conservation. However, the idea of citizen participation in the design and control of these systems has been missing. In democratic political systems, optimization in cities also means being responsive to the observations, attitudes and demands of their citizens.

The widespread adoption of social networking and mobile communication, as well as the Internet-of-Things (IoT), citizen-science and open data, generate new information resources for understanding the needs and preferences of citizens from government services. Information generated via citizen science and crowd sensing sources, are particularly good at reflecting the priorities and perceptions of citizens. Inclusive and participatory, "smarter" cities can integrate and leverage these diverse information resources to improve responsiveness to citizens and improve the urban quality of life. For example, environmental conflicts such as air and noise pollution can be better managed. Still, supporting participatory information gathering at urban scales, and integrating that sensed data into information that is relevant to urban governance and moreover, actionable, remains a major research issue. The Urban Civics project tackles this challenge by bringing together a multi-disciplinary team of experts, so that citizens and governments may work together toward healthy cities.

In a nutshell, Urban Civics is developing, from the design to prototype implementation, a middleware solution that tackles the following research question (Figure 1): 'How can urban data sources be leveraged and comprehensively integrated to accurately capture nuisance issues facing our cities?'. We have specifically concentrated on addressing the following associated research challenges:

- How can we develop a distributed system architecture that leverages the richness of the urban sensors of the new digital era (featuring IoT, open data, citizen science, social networking and mobile computing) at large scales? Our approach lies in leveraging relevant probabilistic protocols to overcome the urban scale challenges together with building upon semantic Web technologies to aggregate physical and social data sources [1].

- How can we ensure citizen participation? Our approach lies in studying dedicated machine-learning algorithms (e.g., see [2]), to mine leadership and thereby prompt citizen participation, as well as optimizing data collection strategies through incentive-driven, pro-active citizen engagement and a more informed approach to crowd-sensing.

- How can we assimilate the rich urban data to develop significant city models? We address this challenge by building on the assimilation of observations that deal with time-varying distributions, and take into account both model and observational errors to produce the most accurate map of the monitored phenomenon (e.g., see [3]).

The Urban Civics middleware solution applies to a variety of environmental issues that our cities are currently facing. Indeed, the physical and social sensing approach implemented by Urban Civics remains similar across use-cases dedicated to environment monitoring. Specifically, the approach requires plugging dedicated sensors to provide quantitative data and applications to enable end-users to provide qualitative data, both of which are aggregated through data modelling & assimilation functionalities. The mathematical details vary depending on the use-case specific phenomenon. As an illustration, we discuss noise pollution and air quality below:

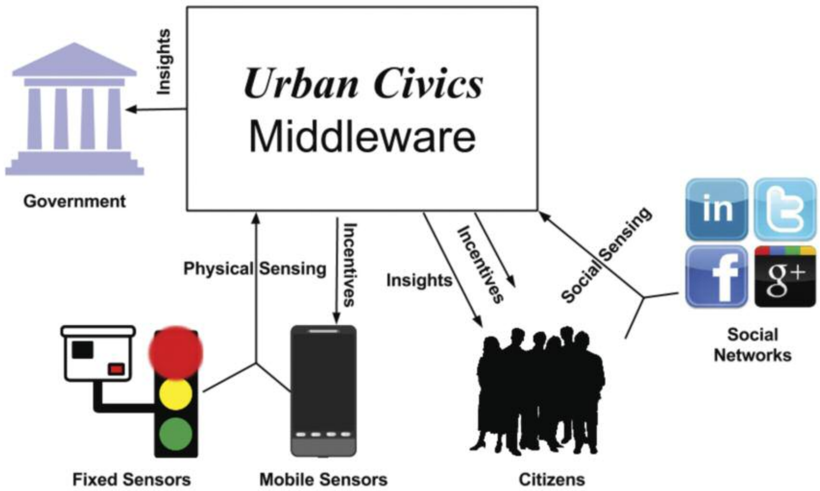

Figure 1: The role of Urban Civics in urban democracy

- Noise pollution, i.e., the undesired sounds that harm one's well being, come from several sources including traffic, neighbors and construction works. Noise can be measured using a microphone, hosted on sound/noise level meters which can be deployed throughout a city. However, whilst highly accurate, static sensors alone cannot account for the fine-grained spatial variations that occur.. Consequently, smartphone microphones can be exploited as a complementary data source; end-users describe noise levels through a dedicated application plugged in with Urban Civics. Additional information can also be automatically extracted from social networks, e.g., if users are attending a public event at the measurement time. Once the various inputs are provided, the data is analyzed and proper city models (of noise pollution) are generated.

- Air quality monitoring, in particular the tracking of nitrogen dioxide (which is mostly emitted from vehicles), can be performed by deploying static sensors throughout a city. Further, since smartphones do not yet host sensors that can monitor air pollution, end-users can themselves be equipped with such sensors and can provide their own qualitative input to describe air quality at their locations. Data assimilation for air quality is already applied in Paris and leveraged in several cities in France. However, urban-scale air quality monitoring is still challenged by the required deployment of sensors, which is made much more tractable using Urban Civics and its diverse urban sensors.

Urban Civics is under active development, as part of larger research initiatives oriented toward Smart Cities that are being launched by both University of California, Berkeley, and Inria. The initiatives will promote collaboration between the two organisations. In parallel, we are working closely with the city administration to acquire insights on the broad applicability of Urban Civics.

Link:
CityLab@Inria Project Lab on Smart Cities: http://citylab.Inria.fr

References:

[1] S. Hachem, A. Pathak, V. Issarny: "Service-oriented middleware for large-scale mobile participatory sensing", Pervasive and Mobile Computing, 2014
[2] F. Kling, A. Pozdnoukhov: "When a city tells a story: urban topic analysis", in proc. of the 20th International Conference on Advances in Geographic Information Systems, 2012
[3] A. Tilloy et al.: "BLUE-based NO2 data assimilation at urban scale", Journal of Geophysical Research, 118(4), 2013.

Please contact:

Valerie Issarny,
Inria@Silicon Valley, France, USA
E-mail: Valerie.Issarny@inria.fr

AppCivist - A Service-oriented Software Platform for Social Activism

by James Holston, Rodrigo Ochigame, and Animesh Pathak

Anthropologists and ICT researchers from France and USA join forces to enable social activists to leverage the tools of the digital revolution.

The increased adoption of mobile devices and social networking mobilizes citizens to monitor their living environment actively. Ultimately this engagement is likely to prompt governments to take action. Therefore, the study of socially-grounded software systems that are dedicated to making cities both smarter and more inclusive holds great potential for cross-domain research among social scientists and computer scientists.

In this context, mobile social networking appears as a significant tool to support social activism within cities. However, the development of information and communication technologies (ICT) for social activism is still in its infancy, ranging from the use of email for internal communication to the Web and social media for mass communication. In terms of integrated platforms, there have been a few initial offerings such as Google Apps for Non Profits. However, Google Apps is a generic platform for organizations and requires significant customization to meet the needs of social activists. At the same time, initiatives like the Social Apps Lab at CITRIS at the University of California, Berkeley, clearly show the potential of mobile devices to generate an active, critical and direct engagement of citizens in social issues, which may increase the impact of social activism.

The AppCivist Vision

AppCivist is a software system being developed by the Social Apps Lab in partnership with Inria, France. It provides a platform for democratic assembly and collective action that lets users make their own applications, called Assemblies, with modular components. An underlying premise of the project is that the development of ICTs for social activism crucially depends on the provision of adequate software tools so that activists can easily assemble applications matching their needs. From a software engineering perspective, architectural principles like service-orientation, coupled with the latest advances in distributed computing, appear as relevant base building blocks. From the social science perspective, social activism relies on a number of key functionalities that may be leveraged and re-used effectively across applications. Such functionalities include the following, among others:

- Proposing, deliberating, and voting allow citizens to have meaningful discussions based on available information, leading to action. While simple plurality-based voting systems exist, they fail to enable the nuanced discussion, filtering, consideration, and consensus-building that is fundamental to both social activism and democratic citizenship.
- Collaborative mapping helps people crowd-source the creation of custom maps that display phenomena of their interest. The Social Apps Lab has already successfully deployed this functionality as part of the Dengue Torpedo project in the favelas of Rio de Janeiro, Brazil.
- Data visualization enables non-specialist users to gather and convert information quickly into easily-consumable graphics and reports for wide dissemination.
- Mobilizing citizens encourages people to specify what needs to be done and to get involved. AppCivist helps users raise issues, propose solutions, and evaluate proposals. As activism depends on follow-through, it also tracks the progress of activities, allows citizens to collaborate, and rewards those who participate.
- Linkage with media and administration enables activists to report issues to those in power and to track their resolution.

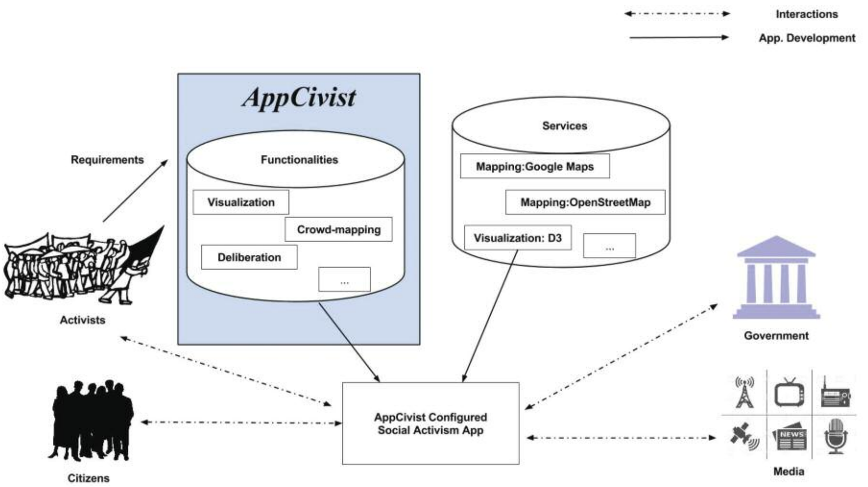

Figure 1: AppCivist core platform and associated services for social activism.

These functionalities encourage proposal-making that is more directly democratic and enable people to make better arguments for their proposals. Making better arguments is key to advancing democratic deliberation. At the same time, the field of social activism brings unique challenges to service-oriented computing in particular and to ICT in general, including the following:

- Focus on Usability. Since the users of AppCivist are not expected to be ICT experts or "digital natives", it is imperative that the use of the platform's modular functionalities be intuitive. The same applies to the administration of plug-and-play and interdependent individual services.
- Cost-sensitivity. The platform should accommodate the vastly different budgets of activist organizations. Consequently, it should carefully

consider the use of costly services such as SMS gateways and incorporate ways to define and manage budgets.
- Privacy. Since activist organizations and their members may have significant security issues, AppCivist must provide strong guarantees regarding the privacy of the data and identity of users.
- Accountability. Nevertheless, certain activities (such as proposing and deliberating) may require that people be accountable for their inputs. Therefore, AppCivist should also provide strong guarantees for non-repudiation, where needed.
- Inclusiveness. We do not want the use of AppCivist to lead to exclusion - something that can happen, for example, if the platform restricts access to messaging to those with smartphones. Consequently, the platform should enable heterogeneity in different dimensions of communication, such as language, technology, and so forth.

Next Steps: Enabling Rapid Integration of ICT in Social Activism

In view of the above, the objective of our research is to study, from design to prototype implementation, a service-oriented architecture and supporting platform aimed at facilitating the assembly of applications by social activists. Jointly performed at Inria and the Social Apps Lab at the University of California, Berkeley, this research will also involve close consultation with activists.

As illustrated in Figure 1, the development of the core platform and associated services for social activism will include algorithms and protocols for filtering and voting, data visualization, assembly of modular components, real-time multi-platform communication between activists and among the general public, mobile-based crowd-sourcing/sensing, and dynamic task-allocation and management among participants. Further, the interplay between mobile and cloud computing entails assessing both cost and social effectiveness.

Links:
Social Apps Lab at CITRIS: http://socialappslab.com/
CityLab@Inria Project Lab on Smart Cities: http://citylab.inria.fr/

References:

[1] J Holston: "Insurgent Citizenship: Disjunctions of Democracy and Modernity in Brazil", Princeton University Press, 2008.
[2] A. Toninelli, A. Pathak, V. Issarny: "Yarta: A Middleware for Managing Mobile Social Ecosystems", GPC 2011.

Please contact:

James Holston
University of California, Berkeley, USA
E-mail: jholston@berkeley.edu

Animesh Pathak
Inria, France
E-mail: Animesh.Pathak@inria.fr

CityLab@Inria - A Lab on Smart Cities fostering Environmental and Social Sustainability

by Valérie Issarny

The Inria Project Lab CityLab@Inria which is currently under creation, studies information and communications technology (ICT) solutions that promote social and environmental sustainability and facilitate the transition to Smart Cities. The Lab places a strong emphasis on multi-disciplinary research, integrating relevant scientific and technology studies from sensing up to analytics and advanced applications. The idea is that the research environment will mirror the predicted Smart City Systems of Systems. A central concern of the Lab is running experiments so that we are able to investigate proposed approaches in real-life settings.

Our Motivation

The world is in the midst of an immense population shift, as people move from rural to urban areas. This has led governments, businesses and communities to rely on technologies, in particular, information and communication technologies (ICT), to overcome the challenges posed by this rapid urbanization. As a result, various academic, industry and city-led ICT initiatives have been launched in recent years with a view to building "smart urban infrastructures". These provide detailed information about the functioning of a city to its citizens and businesses, thereby enabling them to better understand their infrastructures and resources and thus, improve the management of them.

While environmental and economical sustainability have been on the ICT research agenda for some time, the equally important social sustainability has been overlooked in the context of Smart Cities. Indeed, cities are first and foremost places for people, and thus building cohesive, inclusive and flourishing communities should be at the forefront of our research agenda. Without the right social infrastructure in place, problems such as isolation, mental health problems, anti-social behaviours and crime are more likely to arise, pushing communities into decline.

Research Themes and Challenges

The objective of the CityLab@Inria is to study ICT-based Smart City systems from supporting "sensing" systems up to advanced data analytics and new services for the citizens that promote social and environmental sustainability.

Specifically, CityLab@Inria brings together Inria project teams in networking (FUN and URBANET), distributed software systems (ARLES-MiMove and MYRIADS), data management (DICE, OAK and SMIS) and data analytics (CLIME and WILLOW) to investigate the following research questions:
- How can urban-scale sensing that needs to combine both physical and

social sensing be effectively sustained while accounting for the requirements associated with the target network? These include scalability, energy-efficiency and privacy preservation. Sensing the 'city pulse' brings challenges for the supporting data management which must scale-up, as well as integrate highly heterogeneous data of various qualities. The literature is rich with papers addressing these concerns individually. However, they are seldom tackled together, especially while simultaneously considering the urban scale. Our approach to overcome these challenges lies in the study of scalable protocols from the networking up to the middleware layers, together with advanced techniques for privacy enhancement and semantic-aware data management.

- How can the data be aggregated so that the evolution of a city can be not only understood, but also anticipated and perhaps even influenced? Data analytics is at the core of Smart Cities, making big data available to us through sensing. Based on the open data trend, this can become very useful in providing knowledge on the cities. It is a very active area of research. However, numerous open problems remain regarding how large-scale data is analyzed and the uncertainty associated with urban-scale, crowd-sourced data collection must also be overcome. Our contribution in this area leverages advanced research results on data assimilation and machine learning.
- While city-scale sensing and data analytics are two complementary aspects of Smart City systems, they are also inter-related as one may adequately inform the design of the other. Therefore, it is essential to design crosscutting architectures for Smart City systems based on the comprehensive integration of the custom data sensing and analytics that we will investigate.
- Finally, the Smart City vision will only come true if it is accompanied by concrete urban services that do make our (future) cities sustainable and agile. A number of application areas have been suggested and these include smart energy, smart health and smart transportation. However, we are still lacking disruptive services that will contribute to making our cities better places to live while also addressing the central challenge of growth. One important question is how the use of city-scale sensing can impact city governance, particularly its social dimension? Our research will be guided by the study of new urban services which will be undertaken in close collaboration with external partners (especially city representatives) as well as researchers from the social science field.

While the scientific focus of CityLab@Inria is broad, the Lab's research leverages relevant effort within Inria project-teams that is further revisited as well as integrated to meet the challenges of smart cities.

An International Lab

CityLab@Inria research builds on the collaborative effort of the international research community, especially the Inria@SiliconValley program. Indeed, a key characteristic of the CityLab@Inria Lab is its international dimension which began with the Paris-San Francisco cooperation agreement toward smarter cities . This agreement, signed on March 20, 2013, is dedicated to developing smarter cities and includes support for targeted research programs among which is the Joint Inria-CITRIS CityLabs Program.

Links:
CityLab@Inria Project Lab on Smart Cities: http://citylab.Inria.fr
Inria@Silicon Valley program: https://project.inria.fr/siliconvalley/

Please contact:
Valérie Issarny, Inria@Silicon Valley,
E-mail: Valerie.Issarny@inria.fr

A Framework for Improving the Multi-Device User Experience in Smart Cities

by Luca Frosini and Fabio Paternò

The current evolution of pervasive technologies in our cities means that traditional access modalities are not always suitable for a particular service. Consequently novel solutions are needed. Freely moving citizens in urban environments would often like to be able to better exploit available devices to access services. In such cases, a multi-device user interface in which interactive components can be dynamically distributed over devices would be extremely useful: unfortunately, this is not currently supported by available development toolkits. Here, we propose a solution that extends existing Service Front Ends for Smart Cities and supports multi-device interaction by exploiting personal devices and public displays.

Our life is becoming a multi-device user experience. In the last decade, a wide variety of interactive devices have penetrated the mass market and people are spending increasing amounts of time using them. There are two modes of usage, sequential or simultaneous. The sequential user moves from one device to another at different times to accomplish a task whilst the simultaneous user interacts with more than one device at the same time, for either related or unrelated activities [1]. Recent technological advances have meant that large display screens are now available on the mass market and these greatly enhance usage. Given their low cost, such devices are being installed in great numbers in many public places (e.g., train stations, airports, hospitals, public offices, museums, universities, shop centers, bars and restaurants). Thus, deployment is occurring in both outdoor and indoor environments, providing various types of content (e.g., informative, entertainment or advertisements) [2]. Unfortunately, insufficient attention

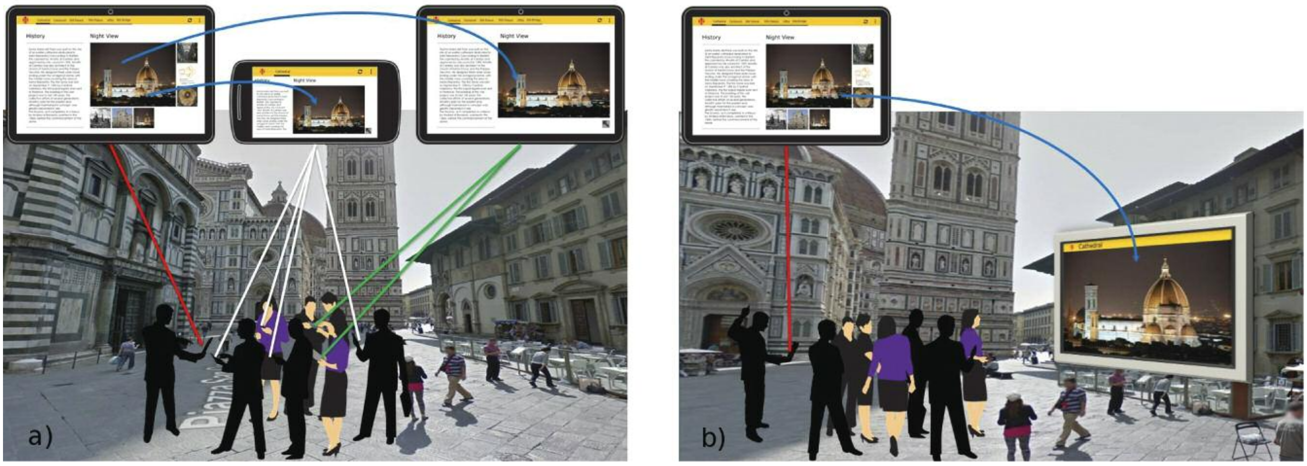

Figure 1: An example of a supported multi-device user interface, where a) shows the application providing content in the tourist version and b)shows how the application can be used to exploit public displays.

has been paid to how information is provided accessed and through such devices, thus diminishing their potential effectiveness.

People would like to be able to better exploit available technology when using multiple devices to interact with their applications. For example, users may wish to dynamically move components of their interactive applications across different devices with various interaction resources [3]. However, managing information across devices is challenging. Unfortunately, the development of multi-device user interfaces (UIs) is limited by the interaction development toolkits currently available. These toolkits are still designed under the assumption that they are supporting the development of UIs for single devices, rather than providing support for multi-device access. Consequently, the main issues with multi-device UIs are their poor adaptation to their usage context, the lack of coordination among tasks performed through different devices and inadequate support for seamless cross-device task performance.

To address this gap we have designed and developed a novel framework which is capable of:
- providing developers with an API that can be exploited by both Web and Android applications in order to obtain application UIs that can be more easily distributed dynamically and/or migrated in multi-device and multi-user environments,
- dynamically creating sessions between groups of users/devices with a distributed UI. The elements of the UI can be distributed according to specific device(s), group(s) of devices, specific user(s) or groups of users classified by roles, and
- avoiding the need for a fixed server to manage the distribution, which can be useful when connectivity is limited.

This framework is publicly available at http://giove.isti.cnr.it/tools/MUDUIDME/home. It provides developers with a small set of commands that allow them to easily support dynamic situations in which flexible UI distribution and good performance are requested, with limited impact on the code.

Various applications have already been developed using this framework. One application was designed to improve the user experience during a museum visit. A single user application which was able to exploit mobile devices, in conjunction with a public display in an indoor environment, was implemented. The museum has some large public displays, which allow visitors to access multimedia content. When users are near the large screens, they can use their smartphone to select content of interest. For example, visitors can select and display some high-resolution images of artworks that cannot be viewed from close up (e.g., because of security or art preservation issues). Another application supports city guides who are accompanying groups, with either tablets or smartphones. The application shows information supporting the mobile visit and allows the visitor to select what content is to be shown in the version of the tourist role (Figure 1a) or exploit the public displays they encounter (Figure 1 b). Yet another application is a set of games for multiple users who can participate with the support of any device, either mobile or stationary.

The framework we propose can be used to easily obtain interactive multi-device applications for different domains and in different contexts in Smart Cities including single or multi-user applications, indoor or outdoor environments and mobile and stationary devices.

Links:
HIIS Lab: http://giove.isti.cnr.it
Framework available at:
http://giove.isti.cnr.it/tools/MUDUIDME/home
IUSDM Project:
http://giove.isti.cnr.it/IUDSM

References:
[1] Google: “The new multi-screen world: Understanding cross-platform consumer behavior”, Technical report, August 2012, http://www.google.com/think/research-studies/the-new-multi-screen-world-study.html
[2] J. Müller, F. Alt, A. Schmidt, D. Michelis: “Requirements and Design Space for Interactive Public Displays”, ACM Multimedia, 1285–1294 (2010)
[3] S. Hosio et al.: “Supporting distributed private and public user interfaces in urban environments”, in proc. of HotMobile 2010, 25-30.

Please contact:
Luca Frosini, Fabio Paternò
CNR-ISTI, Italy.
E-mail: {luca.frosini, fabio.paterno}@isti.cnr.it

Moving Towards Interoperable Internet-of-Things Deployments in Smart Cities

by Gregor Schiele, John Soldatos and Nathalie Mitton

The objectives of the European FP7 VITAL project are to overcome Internet-of-Things (IoT) silos in Smart Cities and enable the deployment and integration of Internet-Connected-Objects (ICO) independently of the underlying IoT architecture. In achieving these goals, the project assists in the transition towards smart, secure and cost-effective cities.

Using the Internet-of-Things (IoT) to deploy Internet-Connected-Objects (ICOs) in Smart Cities raise several challenges, especially in terms of interoperability. Drawing on semantic interoperability solutions developed for IoT systems, the VITAL project aimed to bridge the numerous silos which exist in IoT deployments in Smart Cities. This was achieved by repurposing and reusing sensors and data streams across multiple applications, without compromising citizens' security and privacy. This approach promises of increase Return-On-Investment (ROI) figures which are associated with the typically costly development of Smart City infrastructures by expanding the number and scope of potential applications.

IoT technology enables a large number of physical and virtual ICOs to be coordinated so as to provide human-centric services in a variety of sectors, especially in Smart Cities [1]. Many research, pilot and commercial applications have been developed thus far, ranging from those which use a RFID/Wireless Sensor Network to those which involve large numbers of different devices and ICOs. Collectively, these applications have a significant impact on both business and society. Thus, the IoT paradigm is relevant to the development and implementation of sustainable urban development policies. However, a number of IoT systems have been created in parallel which is leading to the creation of IoT application silos within modern Smart Cities (Figure 1). These silos often reflect the organizational structures of local governments. For example, the separation of law enforcement, transportation and public works into three separate departments results in different IoT deployments and the creation of three associated technical silos.

Figure 1: IoT application silos within modern Smart Cities

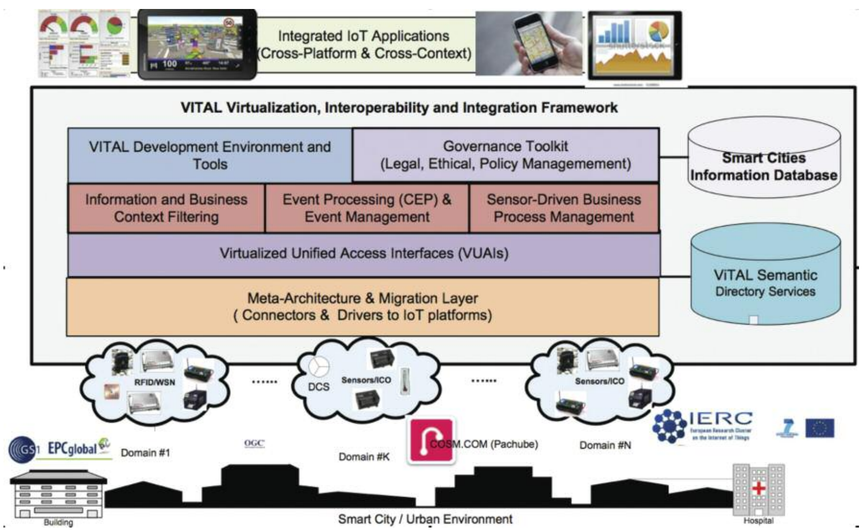

Figure 2: VITAL architecture

Semantic Interoperability Solutions for the IoT

The interoperability of IoT systems extends beyond technical or syntactic interoperability, towards semantic interoperability which mainly explores the use of a common ontology for describing resources across disjoint IoT systems [1]. Recently, several research initiatives have extended the ontology of the W3C SSN incubator group, which aims to overcome the limitations of pre-existing XML-based formats [2] and the fragmentation of sensor ontologies into specific domains or applications [3]. This ontology describes sensors, observations and related concepts, but not domain concepts (e.g., for Smart Cities).

An example of a system which provides a common semantic layer is the OpenIoT project (see link below). This system enables humans, devices and services to announce and annotate different (virtual) sensors/devices as W3C SSN compliant sensors. OpenIoT offers an open source cloud-based IoT platform which includes components like sensor middleware, cloud data storage and a scheduler (see link below).

VITAL: Semantic Smart City Interoperability

The VITAL project builds on these approaches and extends two main aspects to ensure the semantic interoperability of evolving Smart City IoT applications and projects. It uses the SSN ontology to model data and the OpenIoT as a common data management component. The first extension provides a much richer data model for Smart City applications, including city-wide information (e.g., demographics and stakeholder details) as well as city-specific Key Performance Indicators. The second extension provides interoperable access, not only for data coming from different IoT systems (e.g., OpenIoT) but also for services provided by these systems (e.g., discovery, monitoring and complex event processing). This allows for higher level services to be created which can then be integrated into a single federated service view.

VITAL will be tested in two Smart Cities, London and Istanbul, who are represented in the project consortium by the London Borough of Camden and the Istanbul Metropolitan Municipality, respectively.

Conclusion

VITAL will enable applications and service providers to integrate services and data streams stemming from multiple IoT ecosystems, architectures and middleware infrastructures. This will allow for existing sensors and IoT systems to be reused and repurposed, increasing the ROI on Smart City infrastructures. We hope this will reduce the costs associated with developing new Smart City applications for city authorities and the open developer community and bring time efficiencies, leading to a new wave of applications in cities across Europe.

VITAL is a three year joint European project, started in September 2013 by a consortium of ten partners from Ireland, France, Greece, Italy, Spain, the UK and Turkey.

Links:
http://www.vital-iot.com
http://openiot.eu/
http://github.com/OpenIotOrg/openiot

References:
[1] M. Presser, P. Barnaghi, M. Eurich, C. Villalonga: "The SENSEI project: Integrating the physical world with the digital world of the network of the future", IEEE Communications Magazine, pp. 1-4, 2009
[2] A. Sheth, C. Henson, S. Sahoo: "Semantic Sensor Web", IEEE Internet Computing, 12 (4), 2008
[3] M. Compton, P. Barnaghi, L. Bermudez, R. G. Castro, O. Corcho, S. Cox, et. al.: "The SSN Ontology of the Semantic Sensor Networks Incubator Group", Journal of Web Semantics: Science, Services and Agents on the World Wide Web, ISSN 1570-8268, Elsevier, 2012.

Please contact:
Gregor Schiele, NUIG, Ireland
E-mail: gregor.schiele@deri.org

John Soldatos
Athens Information Technology, Greece
E-mail: jsol@ait.gr

Nathalie Mitton, Inria, France
E-mail: nathalie.mitton@inria.fr

Realizing Smart City Scenarios with the ALMANAC and DIMMER Platforms

by Mark Vinkovits, Marco Jahn and René Reiners

Information and communication technology (ICT) is becoming a key factor to develop green and sustainable applications within Smart City scenarios. Effective management of resources, gathering and interpreting data as well as ecological considerations are prerequisites for realizing the vision of smart cities. The two European FP7 projects ALMANAC and DIMMER address these issues by providing a generic, flexible and quickly customizable platform for application development.

A Smart City service Platform (SCP) for application design and implementation needs to collect, aggregate, and analyse real-time or near real-time data from, e.g., appliances, sensors and actuators or smart meters that have been deployed to implement processes running over a pervasive data communication network. The platform must allow decision support, and implement an intelligent control of devices through an M2M network as well as the management of local installations. The FP7 project ALMANAC combines real-life validations with simulation techniques leading to large-scale hybrid simulations that allow to test the entire system and simulate critical loads conditions within large scale deployments.

ALMANAC focuses on the development of integrated cross-application information systems. A service delivery platform with corresponding technology solutions integrates Internet of Things (IoT) edge networks with communication access for green and sustainable applications. The developed services contain management of public and private network access - containing Telco M2M infrastructure, common abstraction of resources, handling of gathering latest and historical values of heterogeneous information sources, and reuse of data models provided by individual applications. The ALMANAC architecture has been designed to emphasize these services, which are then demonstrated through three example applications: Waste management, water management, and citizen engagement.

The four main services for the SCP that were indicated as initial starting points cover interoperability over devices, such that applications can make use of any protocol over a uniform web-service based interface. Service virtuality makes it possible to overcome physical network boundaries and to simplify application design. Composing rules and data caches help to define thresholds and detect trends when concrete values are not needed or their retrieval takes too long for the current use case. Privacy policies of individual providers enforce applications to only access data and functionalities for which they have explicit access rights.

Another aspect of smart city scenarios is covered by smart grids where "prosumers" that are both, energy producers and consumers at the same time must be integrated into the grid infrastructure. This flexibility pose new challenges to the ICT infrastructure for planning and managing smart cities. The FP7 European Project DIMMER (District Information Modeling and Management for Energy Reduction) - addresses this emerging aspect that plays a major role in the field of renewable and independently generated energy. DIMMER leverages real-time data from heterogeneous data sources and subsystems such as smart meters and wireless sensor networks to model the environmental characteristics and energy consumption of districts. At the district level this information is used to optimize district heating and cooling and the energy grid, acting as an enabler for the visions of smart cities and the smart grid. The project involves relevant stakeholders as central players in the research and development process, such as distribution system operators (DSOs), energy service companies (ESCOs), residents, building managers, and city planners. End user applications are developed targeting the needs of the different stakeholders and aiming at improving the mutual relationship between them, taking into account the fact that the concept of a smart city does not only have a technological dimension but must also consider people and institutions [1].

Through advanced visualization and simulation, DSOs will be able to better understand the usage of their networks and take informed decisions e.g. to improve peak demand management and reduction, energy and cost-analysis, tariff planning and evaluation, failure identification and maintenance. Applications for residents and building managers will help to create bi-directional feedback channels between energy providers, consumers and prosumers allowing more effective control of the energy distribution network. In conclusion, the expected results are a consistent reduction in both energy consumption and CO2 emissions by enabling efficient energy distribution policies, according to the real characteristics of district buildings and inhabitants as well as efficient utilization and maintenance of the energy distribution network.

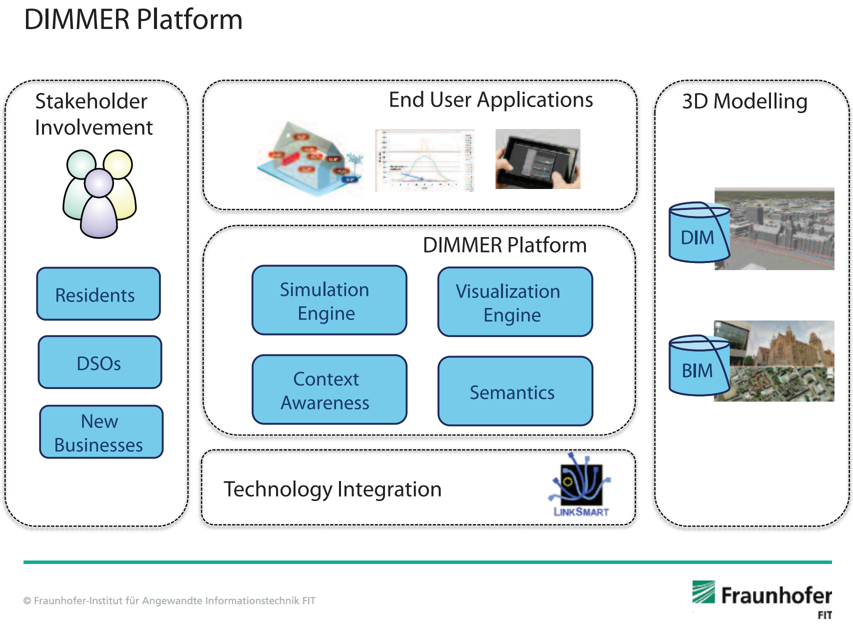

Figure 1: The three pillars of the DIMMER Project: stakeholders, the technology platform and the BIM/DIM.

Figure 1 shows the described main pillar of the project, namely stakeholder involvement, technology platform and building/district information modelling. In order to validate the platform, both public and private buildings included in mixed urban districts are considered in two different cities, i.e., Turin (IT) and Manchester (UK).

Fraunhofer FIT is responsible for the design and development of the software platforms' architecture, device integration as well as leading the requirements engineering process. FIT will bring into the project its extensive knowledge and experience in middleware development. It will employ and further develop the LinkSmart Middleware to allow integration of heterogeneous technologies and systems in the smart city districts [2].

LinkSmart, itself was established as an EU research project and is continuously reused and further developed in follow-up research.

Links:
http://www.fit.fraunhofer.de/en/fb/ucc/projects/almanac.html
http://www.fit.fraunhofer.de/en/fb/ucc/projects/dimmer.html
http://www.almanac-project.eu/news.php
http://dimmer.polito.it

References:
[1] T. Nam, T. A. Pardo: "Conceptualizing smart city with dimensions of technology, people, and institutions", in proc. of dg.o '11, ACM, New York, NY, USA, 282-291, 2011, DOI=10.1145/2037556.2037602 http://doi.acm.org/10.1145/2037556.2037602
[2] M. Eisenhauer, P. Rosengren, P. Antolin: "Hydra: A development platform for integrating wireless devices and sensors into ambient intelligence systems", in The Internet of Things (pp. 367-373). Springer New York, 2010.

Please contact:
Mark Vinkovits, Marco Jahn, René Reiners
Fraunhofer FIT, Germany
E-mail:
mark.vinkovits@fit.fraunhofer.de
marco.jahn@fit.fraunhofer.de
rene.reiners@fit.fraunhofer.de

Internet of Things Applications for Neighbourhood Embedded Devices

by Joan Fons, Daniel Gaston, Christophe Joubert and Miguel Montesinos

Smart Cities enable the physical and virtual worlds to merge, based on Internet Of Things (IoT), data and services. Through Web of Things (WoT), IoT is realized by using the existing web architecture as a platform, thus making smart things directly accessible as web services on the Internet. This approach simplifies object and application deployment, commissioning, maintenance, operation and service composition within both city and building infrastructures. In this work, we focus on cooperative objects to develop an open smart neighbourhood.

The European project Web of Objects WoO (ITEA2 - 10028, January 2012-December 2014) aims to develop a service infrastructure for IoT business applications that enables multi-tenancy, interoperability, and service composition, thus facilitating the building of applications on top of embedded devices. The infrastructure inherits many features from the Web, including bookmarking, caching, linking, searching and securing, and enables mashups for dynamic and ad-hoc composite applications involving embedded devices. Figure 1 describes the WoO architecture based on IoT-A standard.

One of the technical challenges we faced on the project was creating social objects with web capabilities. Such features enable to have complex user-created domain oriented applications that can be used both in indoor and outdoor contexts.

At home, several sensors, actuators, and appliances are available (through home automation technologies). In this project, all these resources and capabilities were mapped as services, and provided by means of either Web Services Technologies or REST Technologies [1]. In addition, we used a mobile device platform that provides access to different sensors and objects (GPS locations, NFC reader, vibrator, accelerometer, etc.) to enable the interaction between home objects (NFC door, home appliances, and other sensors managed by a WSAN [2]) and mobile device sensors (NFC reader/writer, GPS locations, etc.). Table 1 presents an overview of the different smart objects that were deployed on the WoO platform.

In order to semantically define the smart objects, we used and extended the SSN ontology (for sensors and actuators) proposed in [3] and the OWL-S (for services) to define a WoO ontology. SSN ontology was also extended to support the definition of the "actuator" concept. By using those semantically described smart objects, it is possible to facilitate the integration between smart objects inside the WoO architecture.

On top of the WoO architecture, we provided a service layer implemented as REST services. For instance, a door lock web accessible service was designed. This door lock can be discovered by means of an NFC Tag and the NFC interaction is performed through

Mobile device object	*Makes use of multiple sensors (GPS, NFC, Accelerometer, etc.) to measure physical quantities. It is used to identify the user in other systems and to access the user's profile. It supports the user with GNSS (Global Navigation Satellite System) capability.*
NFC Door object	*Operates an NFC-tagged electromechanical door lock using a mobile device (equipped with an NFC sensor).*
User Profile object	*Stores users' preferences to query about them.*
GNSS Location object	*Represents the user's location.*
Location Rules object	*Represents the rules concerning location sensor-based services in order to trigger the corresponding requests to the home automation platform.*
Request object	*Stores the URI and Payload corresponding to the request performed by the mobile device.*
Home Automation Platform object	*Handles the mobile device requests, and also monitors and manages the home automation appliances.*

Table 1: Social Objects in Open Smart Neighbourhood services

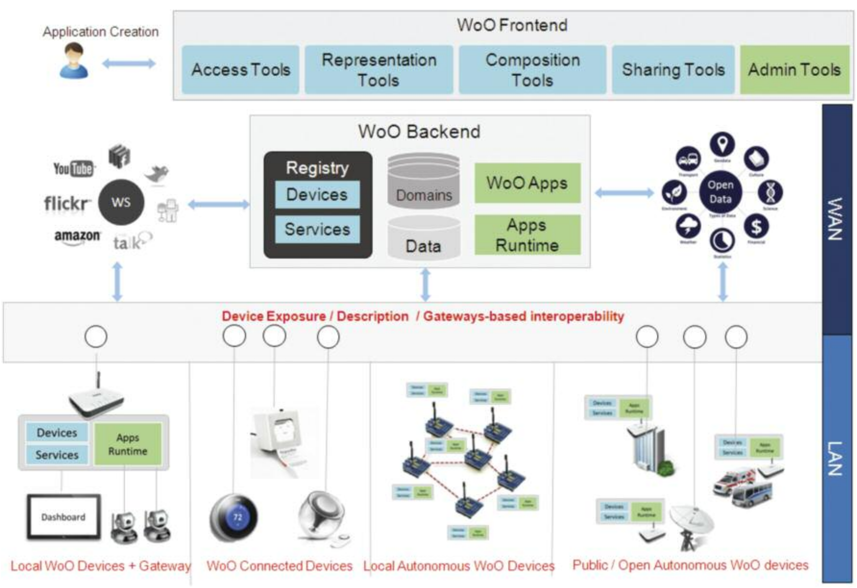

Figure 1: WoO architecture

an open-source mobile agent (Gloo Android application) with profile-management capabilities.

The main challenges that were faced in this context were: smart object discovery (NFC tags), universal smart object accessibility, mobile interface and sensor access, user profiling, and ad-hoc "on-the-fly" interaction between smart objects (heterogeneity). We developed several scenarios, such as the Location Sensor-Based Services outside Home (see Figure 2). The mobile agent running on the device continuously receives the location coordinates from its GNSS sensor and compares those locations with a user defined proximity area. Whenever the user enters this proximity area, the smartphone automatically triggers a series of requests (based on user-profile and preferences) to control smart objects (for example switching on the heating to a certain temperature).

The results to date are very promising: we integrated the previous smart objects with many partners in order to compose an Open Smart Neighbourhood ecosystem on top of the Web of Objects platform. This allowed us to deploy and test some recent research results in projects related with the Internet of Things, Machine to Machine communication (M2M) and Ambient Intelligence systems development. We have acquired insight into the use of smart WoO objects such as mobile devices, user profiles, location rules, requests, doorlocks and smart home objects.

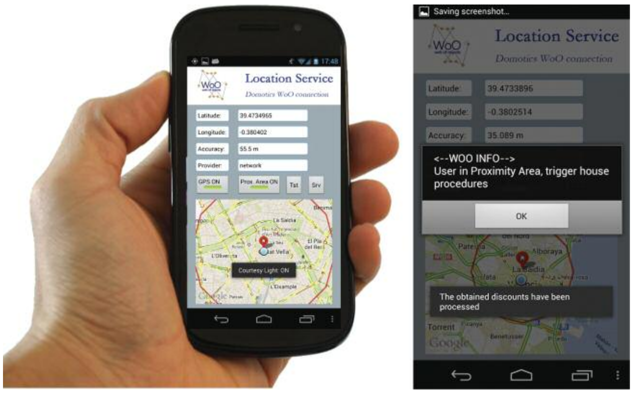

Figure 2: GNSS location-based automatic house entities requests

In our research, we collaborated with several SMEs, Universities and Research Centres, such as DEIMOS, VISUAL TOOLS, TELESPAZIO, ETIC, UPC, UPM, UPV (Spain), Thales, Odonata, Sogeti, UPEM, CEA, IMT (France), Concordia University (Canada), KAIST, KT, ETRI, Kwangwoon and Hankuk University, Innopia, Miksistem (Korea) and Smartec, Nma, University of Cairo (Egypt).

This research is also part of a horizontal task force with other ICT Future Internet projects - such as FIWARE, BUTLER and SOFIA - that deals with building new innovative applications and services for every-day working and living environments. Our work is partially supported by the Spanish MEC INNCORPORA-PTQ 2011, MiTYC TSI-020400-2011-29, and FEDER programs.

Links:
http://www.web-of-objects.com
http://www.youtube.com/user/webofobjectsproject

References:
[1] F. Belqasmi, R. Glitho, C. Fu: "RESTful web services for service provisioning in next-generation networks: a survey", Communications Magazine, IEEE , vol.49, no.12, pp.66-73, 2011
[2] D. Niyato, L. Xiao, L, P. Wang: "Machine-to-machine communications for home energy management system in smart grid", IEEE Communications Magazine, pp. 53-39, vol 49, 2011
[3] L. Ni, Y. Zhu, and M. Jian: "Semantic Sensor Net: An Extensible Framework", Springer, 2005.

Please contact:
Christophe Joubert
Prodevelop, Spain
E-mail: cjoubert@prodevelop.es

Internet of Things: A Challenge for Software Engineering

by Charles Consel and Milan Kabac

The Internet of Things (IoT) has become a reality with the emergence of Smart Cities, populated with large amounts of smart objects which are used to deliver a range of citizen services (e.g., security, well being, etc.) The IoT paradigm relies on the pervasive presence of smart objects or "things", which raises a number of new challenges in the software engineering domain.

The Object's World project

There are an abundance of research and industry initiatives that have been undertaken with the aim of promoting the emergence of IoT [1]. In line with this goal, the Object's World project brings together stakeholders from different domains to build and support the emergence of an IoT sector in France and beyond. The project is lead by SIGFOX, the world's first cellular network operator dedicated to low-bandwidth wireless objects. The cooperation between industry and research partners (e.g., sensor manufacturers, computer science and electrical engineering research labs) is of uttermost importance in overcoming technological barriers. This issue is currently hindering the development of an IoT sector. The main objectives of this project are the development of:

- expertise in the low-bandwidth network sector,
- low-cost transmitter/receiver chips,
- low-energy autonomous sensors, and
- software frameworks which cover the entire lifecycle of IoT applications.

Network infrastructures which support huge numbers of objects open up a range of opportunities for innovative services. Critically, these new opportunities rely on the ability to address the software engineering challenges of this new sector. We promote an approach that revolves around software frameworks. In areas such as mobile and web development, this approach has already been shown to facilitate software development by abstracting over implementation details and guiding the programmer.

A design-driven development approach

To guide and support the development of applications which orchestrate networked objects, our research group introduced a design-driven software development approach which draws on principles and techniques taken from the programming language domain.

In particular, we developed DiaSuite [2], a tool-based methodology which guides the developer through the entire lifecycle of an orchestrating application (Figure 1). DiaSuite offers a design language, providing high-level, declarative constructs that are dedicated to describing the application's architecture, along with the smart objects it orchestrates. The methodology relies on a compiler that generates support in the form of a Java programming framework, customized with respect to a given application design. By providing the developer with a programming framework, our approach ensures conformity between the design and implementation. The generated programming framework provides the developer with an abstract class per component declaration. The application logic is implemented by subclassing each abstract class and programming its abstract methods. To further ease the development of orchestrating applications, DiaSuite relies on Eclipse to guide developers during the implementation phase by introducing placeholders that need to be provided with code. Finally, our approach provides developers with a back-end to address the deployment and execution of orchestrating applications.

Orchestrating smart objects at a large scale

The development of orchestrating applications which are responsible for large numbers of smart objects raises a number of challenges. We have addressed these by introducing a new design language.

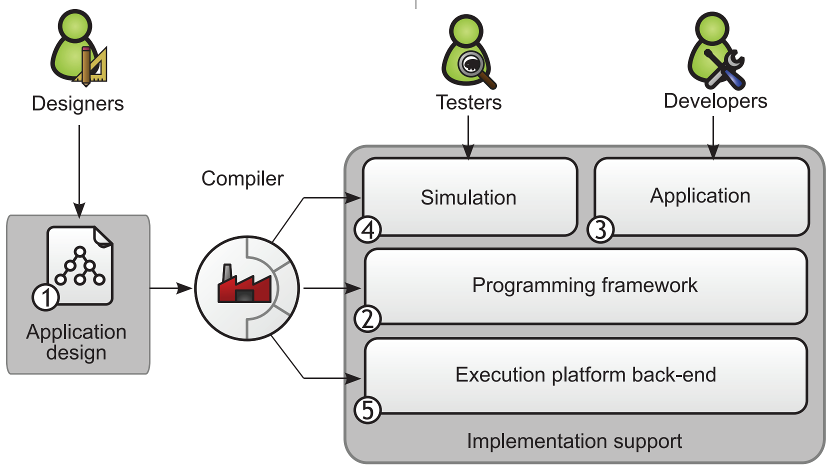

Figure 1: The DiaSuite tool-based methodology

Service discovery

Standard service discovery at the individual object level does not address the needs of applications orchestrating large numbers of smart objects. Instead, a high-level approach which provides constructs to specifying sub-sets of interest is needed. Our approach allows developers to introduce application-specific concepts (e.g., regrouping parking spaces into lots or districts) at the design time and then these can be used to express discovery operations. Following our design-driven development approach, these concepts are used to generate code to support and guide the programming phase.

Data gathering

Applications need to acquire data from a large number of objects through a variety of delivery models. For instance, air pollution sensors across a city may only push data to the relevant applications when pollution levels exceed tolerated levels. Tracking sensors, however, might determine the location of vehicles and send the acquired measurements to applications periodically (e.g., 10 min. intervals). Data delivery models need to be introduced at design time since they have a direct impact on the application's program structure. In doing so, the delivery models used by an application can be checked against sensor features early in the development process.

Data processing

Data that is generated from hundreds of thousands of objects and accumulated over a period of time calls for efficient processing strategies to ensure the required performance is attained. Our approach allows for an efficient implementation of the data processing stage by providing the developer with a framework based on the MapReduce [3] programming model which is intended for the processing of large data sets.

Future work

We envisage to enrich our design-driven methodology with support for simulation of infrastructures of smart objects. To achieve this, we will leverage design-time declarations to generate application-specific simulation support, while keeping the application code unchanged.

Links:

http://www.sigfox.com/en/
http://www.telecom-design.com/en/
http://phoenix.inria.fr/

References:

[1] D. Miorandi et al.: "Internet of things: Vision, applications and research challenges," Ad Hoc Networks, vol. 10, no. 7, 2012

[2] D. Cassou, et al., "Towards A Tool-Based Development Methodology for Pervasive Computing Applications," IEEE Transactions on Software Engineering, 2011

[3] J. Dean and S. Ghemawat, "MapReduce: simplified data processing on large clusters," Communications of the ACM, 2008

Please contact:

Charles Consel
University of Bordeaux / Inria Bordeaux - Sud-Ouest,
E-mail: charles.consel@inria.fr

Milan Kabac
Inria Bordeaux - Sud-Ouest,
E-mail: milan.kabac@inria.fr

Semantic Management of Moving Objects in Smart Cities

by Sergio Ilarri, Dragan Stojanovic and Cyril Ray

Smart Cities depend on information regarding moving objects (e.g., people, vehicles, assets, etc.) being processed. We propose a framework to enable a fully-fledged semantic management of moving objects that can be efficiently and flexibly exploited in Smart Cities.

The different types of moving objects in a city, such as people, vehicles and assets, that are equipped with mobile devices (e.g., smartphones, wearable devices, smart sensors, etc.) have increasing computing, communication and sensing capabilities. Thus, they will play a key role in Smart Cities. The efficient management of the information generated, provided and used by these objects (including information on their locations, trajectories, features, behaviors, activities and environmental contexts) facilitates the improved understanding and analysis of how a city performs. In addition, this type of information assists in the goal of providing citizens with contextual and adapted services in a range of areas including traffic management, urban dynamics analysis, ambient assisted living, emergency management, and mobile health. Whilst significant effort has been invested in modelling and managing moving objects, and some progress has been made regarding the representation of semantics associated with them, further efforts are needed to achieve a fully-fledged semantic management approach for moving objects that can be efficiently and flexibly exploited.

By integrating methods and techniques developed in different fields, including moving object databases, spatio-temporal data mining, and the Semantic Web, it is possible to enhance the way moving object information is managed (with respect to the modelling, querying, processing and analysis components). For example, recent proposals claim that linking Location-Based

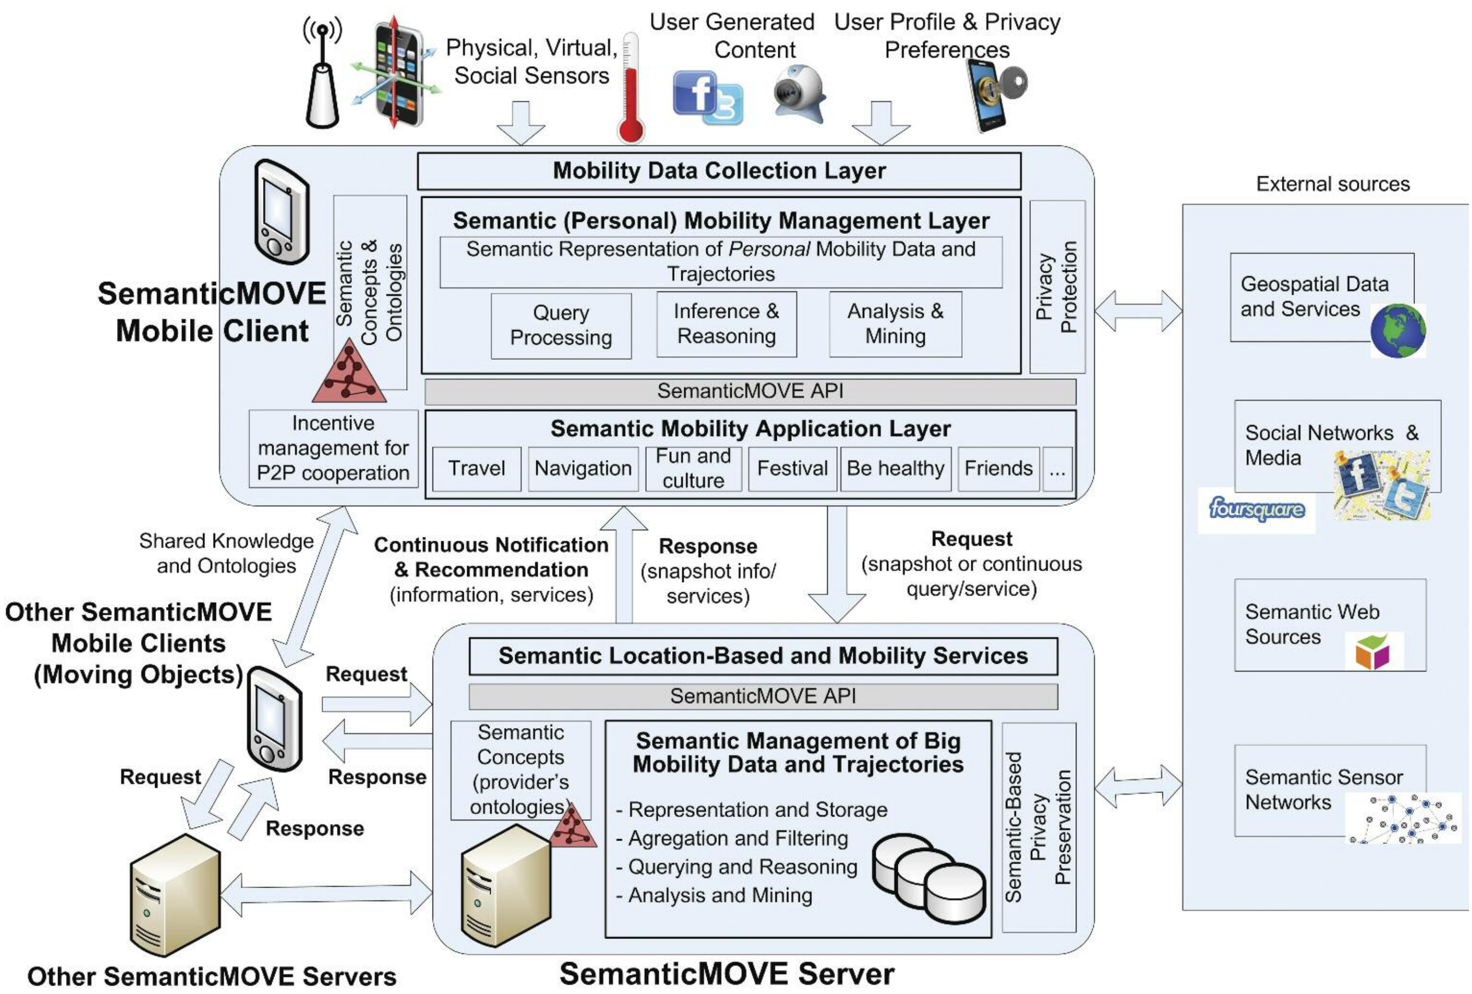

Figure 1: The architecture of the SemanticMOVE framework

Services (LBSs) and semantics can provide interesting benefits [1], such as flexible querying, management of semantic locations and trajectories, interoperability among different LBSs and providers, protection of personal location information, and reasoning in complex and dynamic contexts. Our idea is to develop and provide LBSs that understand the user's requests and interactions, implicitly based on the semantics of mobility and contextual information, and know how to behave and adapt in dynamic and unexpected situations.

Incorporating semantics in such an Internet of (Moving) Things provides valuable semantic information and services to mobile users, thus supporting enhanced mobility in dynamic environments such as those found in Smart Cities. However, there are numerous research challenges that need to be adressed to make this idea a reality. The first lies in the collaborative collection of the data required (e.g., the data measured by sensors on the moving objects). The data collection can be achieved through participatory/ collaborative sensing but this is subject to difficulties related to the correlation and analysis of those data, especially when this analysis is performed in a distributed way on mobile devices. Another challenge relates to the semantic representation of moving objects, where a unified approach that takes all the mobility aspects of moving objects (i.e., their trajectories, contexts, activities, goals, the physiological status of the user, environment, and the accessible services) into account, at different levels of granularity, is still missing. Thirdly, while interesting work has progressed on context-awareness, there is still a need for efficient integrated approaches for query processing in mobile environments, reasoning, and semantic searches, as well as an appropriate abstraction layer that enables the exploitation of the available functionalities. The fourth challenge relates to the inference of higher-level semantics (such as group mobility behaviours) from a large dataset of individual semantic mobility data and trajectories. Finally, we have to mention the problem of privacy protection. This is a particularly critical issue in Smart City environments, as the basic movement data is enhanced with rich semantics of the participating moving objects and their trajectories. Thus, it should be preserved according to the required privacy preferences of the participating users.

Prompted by the COST Action IC0903 on knowledge discovery from moving objects, we have collaboratively designed a generic and scalable distributed framework (SemanticMOVE), whose realization will enable the comprehensive management of the semantics of moving objects. Achieving this will leverage increased sensing, processing, interaction and communication capabilities in mobile devices in a scalable and effective way (Figure 1). As opposed to other related work, we envisioned a quite generic architecture which supports a fully distributed and interoperable scenario for the management of semantics of moving objects. Thus, each moving object can collect, store and analyze its own semantic mobility data and trajectories, and reason over them locally. It can share and exchange semantic mobility data and semantic concepts/knowledge (e.g., regarding the location, personal and social status, vehicle conditions, activities, behaviours, environment and traffic conditions, air pollution levels, etc.), with other moving objects/users/services in the vicinity, over ad-hoc wireless networks. It can also access and share information and knowledge through geo-social networks, social media services, and geospatial information services. Our framework includes fixed servers as an additional element of an ecosystem where the distributed and ad-hoc cooperation among moving objects, encouraged by some incentive mechanisms, plays a key role.

We believe that the development of such an approach will facilitate the deployment of interesting semantic mobility applications for Smart Cities. Whilst there are still some significant research issues that require further investigation, the initial results represent promising steps towards the intended direction [2, 3].

Links:
MOVE: Knowledge Discovery from Moving Objects:
http://dmsl.cs.ucy.ac.cy/projects/cost/index.html
SemanticMOVE:
http://webdiis.unizar.es/~silarri/SemanticMOVE/

References:

[1] S. Ilarri et al.: "Semantics in Location-Based Services - Guest Editors' Introduction for Special Issue", IEEE Internet Computing, ISSN 1089-7801, 15(6): 10-14, IEEE Computer Society, 2011
[2] B. Predic ,D. Stojanovic: "Localized Processing and Analysis of Accelerometer Data in Detecting Traffic Events and Driver Behaviour", Journal of Universal Computer Science, ISSN 0948-695x, 18(9): 1152-1176, 2012
[3] I. Afyouni, C. Ray, C. Claramunt: "Spatial Models for Context-Aware Indoor Navigation Systems: A Survey", Journal of Spatial Information Science, ISSN, 4(1): 85-123, 2012.

Please contact:

Sergio Ilarri
University of Zaragoza, Spain
E-mail: silarri@unizar.es

Dragan Stojanovic
University of Nis, Serbia
E-mail:
dragan.stojanovic@elfak.ni.ac.rs

Cyril Ray
Institut de Recherche de l'École Navale (IRENav), France
E-mail: cyril.ray@ecole-navale.fr

Flexible Access to Services in Smart Cities: Let SHERLOCK Advise Modern Citizens

by Roberto Yus, Eduardo Mena, Sergio Ilarri and Arantza Illarramendi

Citizens can access a variety of computing services to get information, but it is often difficult to know which service will offer the best information. Researchers in the SHERLOCK (System for Heterogeneous mobilE Requests by Leveraging Ontological and Contextual Knowledge) project, from the University of Zaragoza and the Basque Country University, address this by providing mobile users with interesting Location-Based Services (LBSs).

As Smart Cities become a reality, citizens are starting to be overwhelmed with the amount of data they receive from different sources. This is partly caused by the sheer number of apps they can download to obtain information: most apps are designed for specific scenarios and goals and are embedded with implicit knowledge about the application context. They also receive data almost continuously (e.g., information about pollution, traffic, parking spots, lighting, etc.), which makes it difficult to distinguish which information is valuable. In this scenario, the use of semantic techniques becomes particularly relevant: a system that uses semantic information can help users select the most appropriate and relevant apps to the user's interests and transform the raw received data into smart data that represents actionable information.

Mobile devices (e.g., smartphones and tablets) have become a fundamental part of our everyday lives. According to a report by BI (Business Intelligence), one in five people worldwide owns a smartphone and one in 17 owns a tablet. These devices not only consume information but also create huge volumes of data (e.g., geo-tagged images, videos, text, etc.). In addition, sensors are ubiquitous and a common feature on smart and wearable devices (e.g., activity trackers, smart glasses and watches) as well as city-based sensors (e.g., sensors that measure air pollution, noise and traffic levels). These city-based sensors are useful in the context of Smart Cities. New management approaches are required for all this information to ensure users do not become overwhelmed. The semantic management of data in wireless environments can help users in a variety of ways including assisting with the process of determining what information a user really needs to finding the most appropriate information from a range of different sources and presenting it in an integrated way. For example, imagine tourists who arrive on an evening flight and need to reach their city hotel. At first, it would be useful to infer what kind of information the tourists might need, for example, with regards to transport information they might need to know the different options (e.g., buses, metros, taxis or car rental options), traffic conditions and perhaps even where available parking spaces are located. This information comes from a variety of heterogeneous sources (e.g., websites, city-based sensors, other users and vehicles) and it is difficult to reconcile and present these data as they are needed.

In developing the SHERLOCK [1] system, we sought to provide mobile users with interesting Location-Based Services (LBSs). Begun in 2011, this collaborative project brings together the Distributed Information Systems (SID) group at the University of Zaragoza (Spain) and the Interoperable Database Group (BDI) group at the Basque Country University (Spain). Since we started, we have developed an Android prototype [2] that can be downloaded from the website of the project (see Figure 2 for a screenshot of the app).

Our main focus is to develop a general and flexible system that is able to respond to the information needs of mobile users. We met a number of challenges in this project, but there were three main problems. Firstly, the system has to be able to understand the information needs of the user. For this task, SHERLOCK leverages the user’s context (e.g., location, activity, etc.) that can be extracted for a range of sources, for example, from sensors in his/her mobile device. A personal agent guides the user in the process of selecting an interesting LBS. This agent uses context information along with

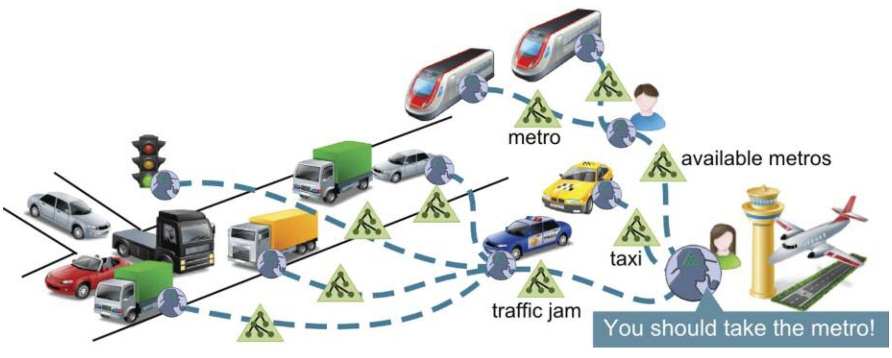

Figure 1: SHERLOCK finding appropriate transport in a city

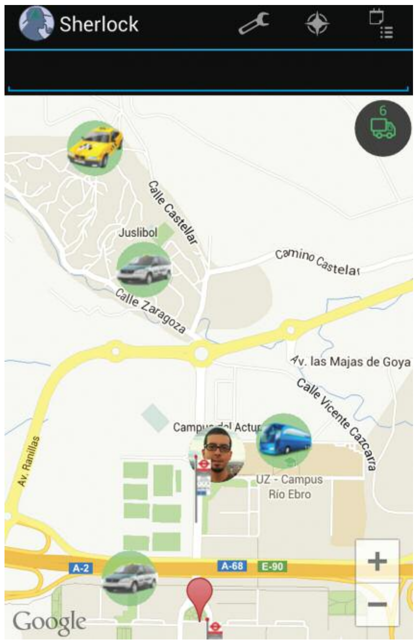

Figure 2: SHERLOCK app for the transports scenario

information about the services, encoded in ontologies (i.e., a formal representation of knowledge) and a semantic reasoner (i.e., a software to infer logical consequences), to deduce which LBSs would be useful for the user. Secondly, the system has to be able to find the information the user needs. For this task, SHERLOCK deploys a network of mobile agents (i.e., software that is able to move from one device to another autonomously) to search all available information among the distributed data sources and bring back the results to his/her device. Finally, the system has to be able to present the results to the user and keep this information updated for as long as the user requires.

We believe that a system such as SHERLOCK is very useful in the context of Smart Cities, as it can support citizens and visitors in achieving their tasks more efficiently and effectively. Any developed city services could easily be made available to SHERLOCK, simply by defining it in the form of an ontology.

Links:
SHERLOCK:
http://sid.cps.unizar.es/SHERLOCK
Smartphone and tablet penetration in 2013: http://www.businessinsider.com/smartphone-and-tablet-penetration-2013-10

References:
[1] R. Yus, et al.: "SHERLOCK: Semantic Management of Location-Based Services in Wireless Environments", Pervasive and Mobile Computing, 2013.
[2] R. Yus, et al.: "SHERLOCK: A System for Location-Based Services in Wireless Environments Using Semantics", WWW 2013

Please contact:
Roberto Yus
University of Zaragoza, Spain
Tel: +34 976 76 26 50
E-mail: ryus@unizar.es

Quantifying the Benefits of Taxi Trips in New York through Shareability Networks

by Paolo Santi, Giovanni Resta and Carlo Ratti

Shareability networks demonstrate that more than 95% of taxi trips taken in New York can be shared with minimal passenger discomfort.

The increasing pervasiveness of digitalized information has unleashed unprecedented opportunities for understanding aspects of human behaviours and social lives, including individual mobility. An enormous amount of digital traces are now obtainable from a range of sources (e.g., cell phone records, taxi GPS traces, etc.) which allows human mobility to be analyzed to an extent that would have been inconceivable several years ago [1]. Although this raises privacy concerns, this "big data" era offers unique opportunities to improve understanding around human mobility needs and, hence, improve transportation system efficiencies.

The goal of this project, performed in cooperation with a team from the MIT Senseable City Lab (Michael Szell, Stan Sobolevsky and coordinated by Carlo Ratti), is to quantify the benefits of taxi sharing in New York City (NYC). Our analysis was based on a dataset which captured all the taxi trips taken in NYC in 2011 (over 150 million trips). For each trip, the dataset captures the pick-up time and location and drop-off time and location.

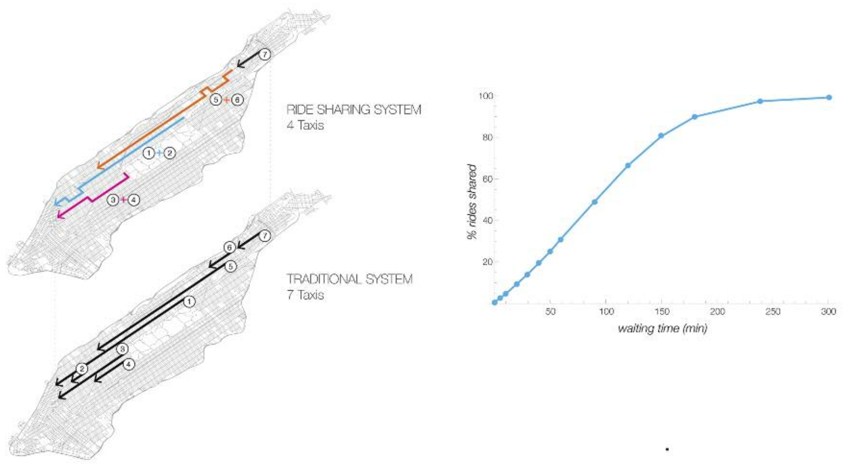

Figure 1: The shareability of taxi rides taken in New York City is constrained by the length of time customers are willing to be delayed. As the delay time lengthens, trip sharing opportunities increase. Results indicated that a delay of 5 min meant more than 95% of the taxi rides could be shared.

In this project, we posed the fundamental question, "How many taxi trips can be shared in NYC?". To answer this question, the intrinsic trade-off between shareability opportunities and passenger discomfort must be considered: the longer a passenger is willing to wait for a shared trip, the higher the sharing opportunities. This tradeoff is made explicit by the novel notion of a shareability network in which we have defined the model sharing opportunities: each network node represent a separate trip and links between two nodes represents a sharing opportunity between those trips. The criterion used to determine whether two trips can be shared is based on spatial and temporal constraints. For two trips, T_1 and T_2, a sharing opportunity only exists if a route connects the respective pick-up and drop-off points such that both passenger groups can be picked up and delivered to their destinations with a delay of no more than Δ,where Δ explicitly models the tradeoff between passenger discomfort and shareability. A higher value of Δ results in a denser shareability network and corresponds to more opportunities for trip sharing.

Using a shareability network allows an optimal solution to be found to an otherwise computationally intractable problem. Shareability networks impose a structure to an otherwise unstructured, immense search space by constraining the number of trips that can be shared (up to k, where k is a user-defined parameter, set to 2 or 3 in this study) and considering only static trip sharing. This term means that once two or more trips are combined into a shared trip, the combined trip is served by a single taxi that cannot be rerouted for further sharing. Once the search space has been reduced and structured, an optimal trip sharing figure can be computed in approximately 0.1 seconds by running a computationally efficient maximum matching algorithm on the shareability network (10,000 nodes and 100,000 links) [2] using a standard Linux workstation. Thus, our proposed methodology is suitable for real-time implementation.

A notable result of our analysis is that constraining the search space to reduce computational complexity does not impair trip sharing opportunities. The percentage of shareable trips is shown to increase with the delay parameter and we found it was as high as 95% if a delay in the order of five minutes was permitted. Thus, the vast majority of taxi trips taken in NYC can be shared with minimal passenger discomfort.

We also investigated the effects of the sharing penetration rate which accounts for the true fraction of passengers that want to use a shared taxi service. We are extending this analysis to other cities (Singapore, Vienna, San Francisco, etc.) to investigate whether similar sharing opportunities arise in other urban contexts.

References:

[1] M. Gonzales, C. Hidalgo, A.L. Barabasi, “Understanding Individual Human Mobility Patterns”, Nature, 2008

[2] Z. Galil, “Efficient Algorithms for Finding Maximum Matching in Graphs”, ACM Comp. Surv. 1986.

Please contact:

Paolo Santi

Istituto di Informatica e Telematica and MIT Senseable City Lab

E-mail: paolo.santi@iit.cnr.it; psanti@mit.edu

Integrated Electric Vehicles Sharing and Pooling Mobility Solutions for Smart Cities

by Marie-Laure Watrinet, Gérald Arnould, Hedi Ayed and Djamel Khadraoui

Integrating public transport systems with individual car-and-ride sharing concepts is considered as an attractive, convenient and emissions reducing mobility concept in the frame of Smart Cities. The pooling of mobility services is considered to be an important enabler of the Smart Cities’ concept, especially with regards to achieving flexibility and integration with existing transport modes (mostly public transport). Even if the levels of user acceptance towards new smart sustainability concepts are still challenging their uptake, it is important to address ICT-related challenges to ensure adaptive solutions are found to reduce complexities. This is especially relevant in the case of electro-mobility related systems.

Tudor, via its mobility projects, has developed a concept for sharing electric vehicles (EVs) and cross-company optimization solutions. This concept aims to increase sustainable resource productivity by sharing EVs across various companies or organizations. This approach is likely to increase the overall usage of each vehicle but reduce costs per kilometre and users. During the work-day the vehicles are used for professional activities, but for the rest of the time, they become a collective resource that can be placed in carpooling mode.

Such a combined usage presents significant algorithmic complexity and solving this so that an e-fleet can be used for public and professional purposes is challenging, The actual concept was evaluated via simulations of different scenarios using MATSIM, A multi-agent approach was considered, where each agent represented a traveler. A previous national project, Moebius [1,2], provided the simulation data. The main objectives were to validate the concept of combined car-sharing and car-pooling from the perspective of resource optimization and find strategic charging locations [3].

The EV sharing service concept presents users with a hop-on, hop-off system that has demand responsive fleet management (which includes predictions of what locations are going to have the highest user needs). The system’s resources (e.g., EVs, e-bikes and associated infrastructures) are mutualized along with an information system, the service and, in future, public transport services.

The goal of the mutualisation concept is to maximize the use of vehicles across the day (both work and recreational times) but decrease the residual cost of the EVs. This can be achieved with an optimization algorithm [1] that can minimize the required number of vehicles at all times, based on planned usage (determined using statistics and simulation tools), real-time demand, and third party companies that can manage the fleets (by zones) with new business models. EV sharing can also be associated with other transport means such as car-pooling and public transport.

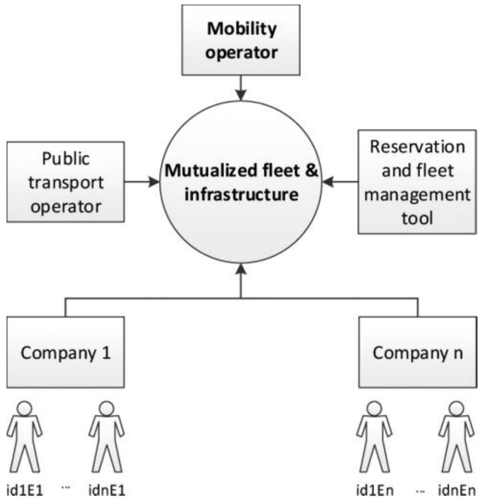

Figure 1: Zac-eMovin mutualized mobility service infrastructure

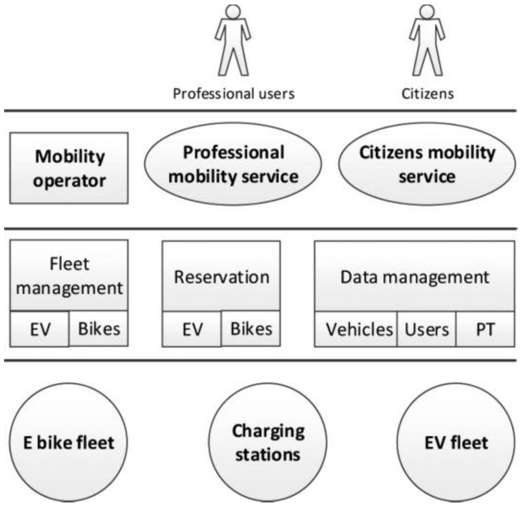

Figure 2: Nordstad-eMovin mutualized mobility services

The experimentations

Two real experimentations (via two European Regional Development Fund (ERDF) projects called ZAC-eMovin and Nordstad-eMovin) have been deployed under the actual concept. The first project is on the professional use of shared EV fleets and infrastructure between companies. The second focuses on the citizen oriented integration of several types of e-mobility services. In both projects, a coherent approach is considered, including the use of public transportation systems and the identification and integration of the required IT services. Each project provide recommendations on the best way to integrate e-mobility services, based on in-depth analyses of relevant systems and user behaviours.

In the ZAC-eMovin project, employees of the three project partners are using the EV in their company's fleet for day-to-day professional travel as well as travels between home and work (and visa versa). The goal of the project is to mutualize the resources, EV and charging infrastructure in an activity zone between the companies.

The fleet management service relies on a resource allocation algorithm which take the limited range of EVs and their charging periods into account, following which it optimizes the EV. The service proposes public transport trips instead of EV leases where possible, and allows for resources to be shared across the participants which guaranteeing safe data exchange and privacy.

The Nordstad-eMovin project is considering the shared usage of EVs and e-bike in Nordstad (Luxembourg) from the citizen's perspective. This project is also considering the integration of public transport systems. Citizens can use any of the three services offered by the system in an integrated way: EV short-time renting, e-bike renting or private vehicle charging. E-bikes are preferred for short-distance travel or when there are heavy traffic conditions. EVs can be used to go shopping, or for longer-distance travel, depending on the requirements of the user. A typical usage could cover the last mile of travel from the train station to work or home. Resources will be mutualized to allow for both the professional and private use of the services.

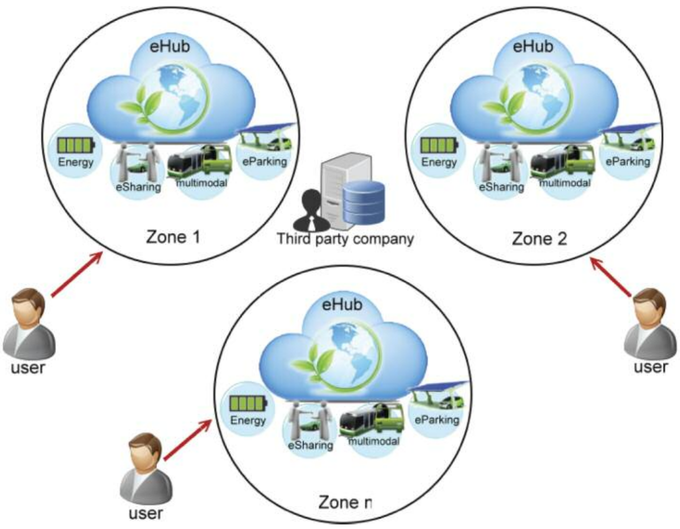

Figure 3: Service mutualisation in three activity zones

Links:
http://www.nordstad-emovin.lu
http://www.zac-emovin.lu

References:

[1] H. Ayed, D.Khadraoui: "Local Search Algorithm for the Daily Car Pooling Problem", International Conference on Metaheuristics and Nature Inspired Computing 2014
[2] H. Ayed, B. Gateau, D. Khadraoui: "Simulating daily mobility in Luxembourg using multi-agenda based system", 28th European Conference on Modelling and Simulation 2014.
[3] G. Arnould, et al.: "Modelling and simulating a dynamic carpooling system for improving citizens mobility", International Journal on Space-Based and Situated Computing, Vol 1, 2013.

Please contact:

Marie-Laure Watrinet, Gérald Arnould, Hedi Ayed, Djamel Khadraoui
CRP Henri Tudor, Luxembourg
E-mail: marie-laure.watrinet@tudor.lu, gerald.arnould@tudor.lu
hedi.ayed@tudor.lu,
djamel.khadraoui@tudor.lu

A Carpooling Recommendation System in the Smartphone Age

by Marcell Fehér and Bertalan Forstner

Worldwide, the one billion cars in use emit enormous amounts of carbon dioxide. Potentially these emissions could be drastically reduced if people shared their daily commutes. The idea of carpooling has been around for a long time and already, millions of gallons of gas have been saved, but the widespread adoption of this strategy is yet to come. Here, we introduce a carpooling service which is easy to use and aims to minimize user effort, thus encouraging the switch to this smart, environmentally friendly transportation option.

In recent history, cars have become our primary choice of transport. However, this high number of motor vehicles have become a significant contributor to climate change, emitting approximately five billion tons of carbon every year. The transportation sector currently accounts for approximately 15% of overall greenhouse gas emissions, a figure which grew by 45% between 1999 and 2007. While getting drivers to swap the comforts of a car for public transport seems like an unattainable goal, encouraging them to carpool is a more appealing prospect and with an estimated 3.75 empty seats per car, per trip [1], the efficiency potential is significant. Changing the mindset of car commuters might not be the task of computer science, but making carpooling services more accessible definitely is. In the era of smartphones, with location tracking abilities and always-on connectivity, providing the public with an easy-to-use recommendation system is no longer impossible.

According to our vision here at the Budapest University of Technology and Economics, the ideal computer-aided carpooling system should require very little effort from the users. In the last two years, research engineers at the Department of Automation and Applied Informatics developed the underlying algorithms such a system would use, and we describe that system below.

It depends on a light-weight application runs constantly on the smartphones of both the drivers and the passengers, tracking their locations and learning their transport usage habits. This data is not uploaded to any server, but rather movement patterns are extracted on the phone itself using advanced, mobile-optimized algorithms. The recurring patterns, identified as routines, consist of a day reference (i.e., what days of the week they happen on), the approximate time window (e.g., between 7:30 and 8:20) and the actual route typically taken by the user.

When a user wants to either volunteer as a driver or seek a ride, they can securely upload their routines to a central server, where the matching process takes place. Here, the routines of the drivers and passengers are compared to find good matches which are then suggested to the user. For example, the usual departure and destination locations of Passenger X are indicated by blue markers in Figure 1. After cross-comparing his departure and destination preferences with the morning routines of drivers in the same area, the system found two potential cars Passenger X may be able to join (displayed by the red and green tracks). Note that the full routes are intentionally not shown since they considered sensitive, private information. When Passenger X receives his suggested options, he can clearly see the proposed pick-up and drop-off locations, as well as approximate timings. He can also call the driver using the phone icon displayed next to their name.

Figure 1: An example of the proposed carpooling system where Passenger X's usual departure and destination locations are indicated with the blue markers. Following a comparison process with drivers in the same location, two possible carpooling suggestions are identified (indicated by the red and green tracks).

This system assumes that arrival times are strictly held on morning routes since people are typically going to work at this time. Conversely, with afternoon routes, departure times are more important since users don't want to wait long for their rides after work. Therefore, when a carpooling search is performed for the morning, the system considers the distance between the drop-off and destination points and (using a five kilometer per hour walking speed) calculates the possibility of being able to walk to the destination and arrive before the strict arrival time.

While these carpooling recommendations a very useful feature by themselves, the large set of learned routines enable a much broader family of services. For example, in case of a traffic jam or road works, the system is able to identify their location(s) by identifying the lower car speeds of drivers who usually take that route. Other regular route users could then be automatically notified of such a disruption, thus enabling them to leave at a different time or take a different route. Our team is currently working on supporting several usages of this geo-social information.

Links:
http://amorg.aut.bme.hu/
http://www.internationaltransportforum.org/2010/

Reference:
[1] Mayinger, F. (Ed.). Mobility and Traffic in the 21st Century. Springer.

Please contact:
Marcell Feher
Budapest University of Technology and Economics, Hungary
E-mail: Marcell.Feher@aut.bme.hu

A Smart Parking Campus: An Example of Integrating Different Parking Sensing Solutions into a Single Scalable System

by Enrique Moguel, Miguel Ángel Preciado and Juan Carlos Preciado

Smart Parking is based on a software system that links hardware and software, with support from Augmented Reality technologies, to provide an enhanced solution in the query of information regarding parking spaces. Smart Parking responds to the need of people with disabilities who have to know the availability of adapted parking spaces.

Under the promise of improving the management and efficiency of cities, the concept of Smart Cities has gained relevance over recent years. Local councils have promoted individual Smart City practices through investments in different sensor systems, infrastructures, measurements and modernisation. An enormous amount of data is now available, collected from sensors with different architectures and specifications. The challenge in turning this into relevant information for citizens. This project aimed to research and develop a single framework for the design of an integral Smart City strategy based on a single plug&play scalable system and diverse data sources.

The Quercus Software Engineering Group, University of Extremadura in collaboration with the Vodafone Spain Foundation designed (2013) and launched (2014) the first version of a Smart City living lab at the Polytechnic School of Cáceres, thus transitioning it into a Smart Campus. A campus environment features the majority of behaviours that occur across the broader city environment.

One of the main problems that citizens and municipalities face is the mobility and parking. This is the main motivation for the Smart Parking project. The goal of this project was to design a low-cost single system that allowed for different data sources to connect in a plug&play mode, being scalable and mainly based on software engineering techniques. This system would 1) offer drivers real-time assistance to find available parking spaces near them, 2) know the occupation ratios for taking decisions about parking space vehicle occupancy and 3) provide useful information for disabled people who typically require more time to find free parking spaces.

In this context, sensors were installed in the parking spaces of the Cáceres campus, using three different data source technologies. We aim to compare reliability, costs, performance, and other factors between the different technologies [1] in an empirical way to determine which is best.

Hardware architecture

The architecture of the system is based on a three-layer model, Producers-Server-Consumers, in which sensor devices are responsible for producing (PRODUCERS) and sending data to the server [2]; the server (SERVER) gives intelligence to the system using algorithms that transforms the data into relevant information; and the users (CONSUMERS) that query this information in real-time.

The layer of the data producers (PRODUCERS) is formed by three types of sensors which are outlined in Table 1.

The data produced by the sensors are sent to the server, where the occupancy status of each parking spaces is stored. This layer is responsible for providing intelligence to the system, processing all data captured and transforming them into relevant information. Then the users can access the services offered by the system from any device with internet access. They can check the real-time occupancy status of parking spaces and view the history of occupation. It is also possible to know if a space is under- or over-utlized. The Smart Parking architecture is shown in Figure 1.

Software architecture

The software layer has been developed using web engineering technologies (WebRatio 7.1/HTML 5) and provide access to any user from any device with internet access. The main view of the system is shown in Figure 2.

Types of data producers

Sensor	Operation	Scope	Asking time	Connection mode
Ultrasound	Emits an ultrasound and measures the time it takes to return	Up to 8 meters	40 seconds	RF or WiFi
Infrared	Emits infrared radiation and measures the time it takes to return	Up to 6 meters	40 seconds	RF or WiFi
Electromagnetic	Creates a magnetic field and check if it is interfered by a metallic object	Up to 3 meters	30 seconds	WiFi

Table 1: The three types of data producers

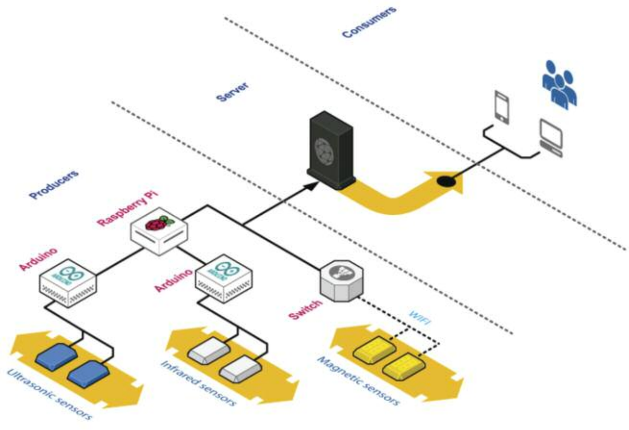

Figure 1: The hardware architecture of the Smart Parking system

Figure 2: Real-time view of the occupancy status for each parking space in the campus parking lot.

Figure 3: Distribution of devices by type and status of parking spaces.

Augmented reality layer

In innovative addition has been development of an augmented reality system [3] that provides the user with context to around where the parking spaces are placed. This differential technology solution can help any user find the correct position of a parking space and its occupancy status.

Conclusions and future lines

Developing a single plug&play scalable multi-layer system has enabled us to conceptualize the design of a Smart City (parking dimension). This development process was founded on our extensive R&D of systems engineering but innovation opportunities that came from lessons learnt also played a key role. These lessons were from experiences in a wide range of fields including big data, cloud computing, business intelligence, business analytics, data visualization, open data and packed services.

This experience has allowed us to manage Smart City concepts at the laboratory level and investigate different software/hardware alternatives. A system based on software engineering offers a scalable and plug&play design and at the sensing level we can develop a low-cost product.

We are also managing cameras (available in Figure 3) which sense parking places via software. In investigating this avenue, we are seeking to maximise the opportunities (i.e., cheaper infrastructure, more control about cars) but solve the associated problems (i.e., more server processing).

Acknowledgement:

Work funded by Spanish Contract MIGRARIA - TIN2011-27340 at Ministerio de Ciencia e Innovacion and Gobierno de Extremadura (GR-10129) and European Regional Development Fund (ERDF)

Figure 4: The augmented reality layer which provides users additional spatial context so that they can find the correct parking space.

Results obtained in the sensors comparison

Sensor	Development	Reliability	Cost	Scalability (installation)
Ultrasound	Medium	High	Low (12 $ each place)	High
Infrared	Medium	Medium	Low (15 $ each place)	High
Electromagnetic	High	High	High (400 $ each place)	Medium

Table 2: Results obtained in the sensors comparison.

Link: http://uex.be/SmartSpaces

References:

[1] G. Revathi, V.R.S.Dhulipala: "Smart parking systems and sensors: A survey", in proc. of Computing, Communication and Applications (ICCCA), Prag 2012, p.1-5

[2] U. Männi: "Smart sensing and time of arrival based location detection in Parking Management services", in proc. of Electronics Conference (BEC) 2010, p. 213-214.

[3] P. Pulli: "Mobile Multimodal Virtual Reality for Personal Teleservices", ERCIM News No.31, October 1997.

Please contact:

Enrique Moguel, University of Extremadura,, Spain.
E-mail: enrique@unex.es

Miguel Ángel, Preciado, Homeria Open Solutions, Spain.
E-mail: mapreciado@homeria.com

Juan Carlos, PreciadoUniversity of Extremadura, Spain.
E-mail: jcpreciado@unex.es

Stochastic Travel Planning for Unreliable Public Transportation Systems

by Tim Nonner, Adi Botea, Marco Laumanns

As part of an IBM First-Of-A-Kind project, the IBM research labs in Dublin and Zurich joined forces to investigate planning approaches for unreliable and highly stochastic public transportation systems (PTS).

When commuting via public transport, a common problem is finding the right direction. Typically getting to the correct destination requires several bus or train changes. Many public transport providers, as well as independent companies, offer online planning tools to simplify this task. The use of these tools usually involves the user entering the origin and destination locations for a given journey along with nominating several preferences (e.g., maximum number of interchanges). The results then provide the user with a sequence of directions, nominating the different buses and trains that must be taken. Ideally, multiple options are suggested, each with its own advantages and disadvantages. For example, Option 1 might present a route with a short travel time, but requiring multiple changes, Option 2 a slow but direct route, and Option 3 may be the cheapest option but it includes an extended wait time at an intermediate stop. For any given option, missing one step in the sequence can result in the route requiring re-planning to address the different conditions now faced by the user.

A traveller who is familiar with a PTS might blend aspects from the variety of solutions presented to develop their own unique solution. This offers the user the advantage of being able to adjust their journey dynamically in response to their own shifting needs. For example, if the user misses a connecting bus, he/she may be aware of an alternative bus which also suits their current needs. This dynamism is especially necessary if the PTS is unreliable and features "rule of thumb" scheduling or only publishes vehicle frequencies as opposed to actual times (e.g., three buses per hour vs. three buses at 11:20am, 11:40am and 12:00pm).

Challenge and Policy-based Planning

Turning this intuitive planning approach based on personal experiences into a stochastic planning tool presents an ongoing challenge which the authors have addressed using a number of different approaches. Instead of providing a simple sequence of options to follow, these approaches all suggest a tree of options, called a policy. The policy outlined above (Figure 1) includes four stops A, B, C and D: at each stop the traveller can pick one of the provided alternatives. For example, at Stop A the user can select Bus 1 or Metro 1. Choices can be made according to a first-come-first-serve rule which is easy to execute (even in an offline mode) or to a more detailed timing regime which might be supported by a smartphone application. We have developed several algorithms to compute such policies.

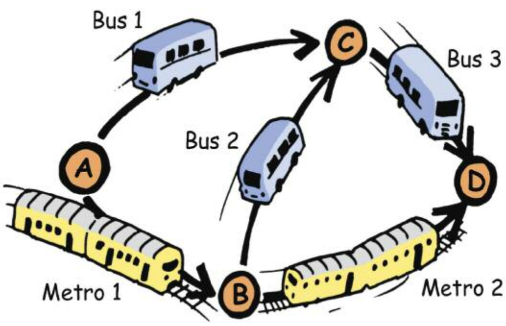

Figure 1: An example policy for travelling from Stop A to Stop D.

Drawing on the heuristic search used in the Artificial Intelligence area, the first algorithm uses a forward planning perspective (i.e., the user is at the beginning of their journey) [1] [2] and attempts to factor in all the possible events and options to find a policy that is flexible within each situation. The second algorithm, inspired by the classical shortest path algorithms, takes a backward planning perspective using dynamic programming techniques [3]. Thus, this algorithm iteratively builds stochastic plans starting from the destination. Both the algorithms we developed are fast enough to allow for real-time application.

There is, of course, the need to build an abstraction of reality in both algorithms. For example, the stochastic behaviour of buses must be simplified to make their movements computationally tractable. Therefore, to make a final check and comparison of the algorithms proposed solutions, we used a more detailed simulation that aims to more closely approximating the real world. For example, tests undertaken on public transport timetable data showed that waiting times can be significantly decreased in many cases by providing more alternatives at each stop: in fact our results showed a 25% decrease in waiting time in 20% of the scenarios considered. Therefore, mainstreaming such policy-based planning approaches in every-day life would significantly improve the user experience of many PTS.

Link:
http://resweb.watson.ibm.com/researcher/view_group_subpage.php?id=4130

References:

[1] A. Botea, E. Nikolova, M. Berlingerio: "Multi-modal journey planning in the presence of uncertainty", in proc. of ICAPS'13, AAAI, 2013.
[2] B. Chen, et al.: "Uncertainty in urban mobility: Predicting waiting times for shared bicycles and parking lots" in proc. of ITSC'13, IEEE, 2013.
[3] T. Nonner: "Polynomial-time approximation schemes for shortest path with alternatives" in proc. of ESA'12, Springer, 2012.

Please contact:

Tim Nonner, Marco Laumanns
IBM Research - Zurich, Switzerland
E-mail: tno@zurich.ibm.com, mlm@zurich.ibm.com

Adi Botea
IBM Research - Dublin, Ireland
E-mail: adibotea@ie.ibm.com

A Quantitative Approach to the Design and Analysis of Collective Adaptive Systems for Smart Cities

by Maurice ter Beek, Luca Bortolussi, Vincenzo Ciancia, Stefania Gnesi, Jane Hillston, Diego Latella and Mieke Massink

Figure 1: public bike-sharing system (CicloPi) in Pisa

It's smart to be fair. Researchers from the Formal Methods and Tools group of ISTI-CNR are working on scalable analysis techniques to support smart applications for the efficient and equitable sharing of resources in the cities of our future. The research is being carried out under the European FET-Proactive project, QUANTICOL.

The Smart City concept is on the research agenda of many European Union (EU) and other international institutions and think-tanks. As urban populations grow, innovative information and communication technology (ICT) initiatives are seen by many as one of the key factors that will allow modern cities to reach or maintain a good and sustainable quality of life for their inhabitants, allowing for the timely and equitable distribution of resources.

These ICT-based systems are based on decentralised and distributed designs, comprised of many autonomous and interacting entities, known as collective adaptive systems (CAS). CAS are required to adapt their services seamlessly to the changing needs of their users, who also form an integral part of the system. They typically consist of a large number of spatially distributed, heterogeneous entities with decentralised control and varying degrees of complex autonomous behaviour. This requires the development of novel scalable analysis techniques to investigate their dynamic behaviour and support the design and operational management of a wide range of such systems. In the QUANTICOL project [1], three principal case studies drive the development of a design and analysis framework for CAS: two smart urban transportation systems (smart bus systems and bike-sharing schemes) and smart grid applications. We present some of the QUANTICOL research performed at ISTI-CNR.

In the first year of the project, we developed several scalable analysis techniques that exploit mean field and fluid flow techniques, in combination with logic-based model-checking, to support the investigation and prediction of dynamic resource usage. Mean field techniques were originally developed in the field of statistical physics to cope with the analysis of very large scale systems composed of interacting objects such as molecules in a gas. The possibly non-linear behaviour of such systems is conveniently modelled by a deterministic approximation, i.e., the limit for an infinite number of agents, given as the solution of a set of differential equations (in the continuous case) or difference equations (in the discrete case). Their combined application with model-checking techniques provides a way to verify properties of individual entities in the context of a large system on which they depend, but also properties of the global system or combined local and global properties. An example is the study of the potential effects of user-incentives on maintaining a satisfactory distribution of bikes and empty parking slots over time. The extension of these techniques to address spatial aspects, including spatial model-checking, is a major objective of the project [2].

A further objective of QUANTICOL is to study the relationships between (representations of) small populations and a compact (family) representation of a large population 'built' from these smaller populations, by indicating the commonalities and variabilities of single entities in their overall environment. As an initial step in this direction, we performed variability analyses on a bike-sharing product line, considering its behaviour to exhibit variability, not only in the kind of features involved but also in the timing and probability characteristics of these features.

In this context, ISTI-CNR initiated a collaboration with "PisaMo S.p.A. azienda per la mobilità pisana", an in-house public mobility company in the Municipality of Pisa, that had recently introduced a public bike-sharing system (CicloPi) in Pisa. This led to an initial feature model of a family of bike-sharing systems, annotated with attributes and global quantitative constraints aiming to minimize the total cost of a chosen configuration while simultaneously aiming to maximize customer satisfaction and capacity (of docking stations).

We have studied the specification and analysis of the possible behaviour of a family of bike-sharing systems in terms of the capacity of their docking stations in a value-passing modal process algebra, considering a dynamic redistribution scheme as an optional feature. Future work includes studying a further parametric extension of the value-passing modelling and verification environment as well as the addition of a quantitative dimension to the behavioural model.

QUANTICOL will run until March 2017 and is coordinated by Jane Hillston from the University of Edinburgh (UK). Other partners are EPFL (Switzerland), IMT Lucca (Italy), University of Southampton (UK), LMU (Germany), INRIA (France) and ISTI-CNR (Italy). We thank Marco Bertini from PisaMo S.p.A. for generously sharing his knowledge on bike-sharing systems with us.

Links:
QUANTICOL: http://www.quanticol.eu/
PisaMo: http://www.pisamo.it

References:
[1] L. Bortolussi et al.: "A Quantitative Approach to the Design and Analysis of Collective Adaptive Systems, FoCAS'13, Taormina, Sicily, Italy, 2013.
[2] J. Hillston: "Challenges for Quantitative Analysis of Collective Adaptive Systems", TGC 2013, Buenos Aires, 2013, Springer LNCS, Vol. 8358, pp.14-21, 2014, DOI: 10.1007/978-3-319-05119-2_1.

Please contact:
Mieke Massink, ISTI-CNR, Italy
E-mail: mieke.massink@isti.cnr.it

Query-Driven Smart Grid City Management

by Stamatis Karnouskos

Cyber-physical systems underpin modern Smart Grids and Smart Cities. New approaches are required to enable efficient, secure and decentralized use of the huge amount of generated data in value-added applications. The SmartKYE project (www.smartkye.eu) is investigating such directions for query-driven integration of Smart Grid City Systems.

We witness a revolution that capitalizes on the prevalence of networked (embedded) devices ranging from simple ones, such as sensors, to complex Cyber-Physical Systems (CPS), in order to empower a new generation of innovative anytime-anywhere services. The CPS revolution enables system interconnection and fine-grained information acquisition at unprecedented scales. The analytics of CPS data pave the path towards informed real-time decision making in complex infrastructures. For Smart Grid Cities this implies innovative solutions [1] in energy management, sophisticated demand-response (DR) and demand-side management (DSM), optimal resource usage, reliability and security, new business opportunities etc. The Smart Grid and its services [2] are seen as an integral part of the Smart City of the future. Several ongoing efforts [1] strive towards capitalizing on the hyper-connected information infrastructure and the collaboration among the things, services, and systems expected to exist in the Smart City.

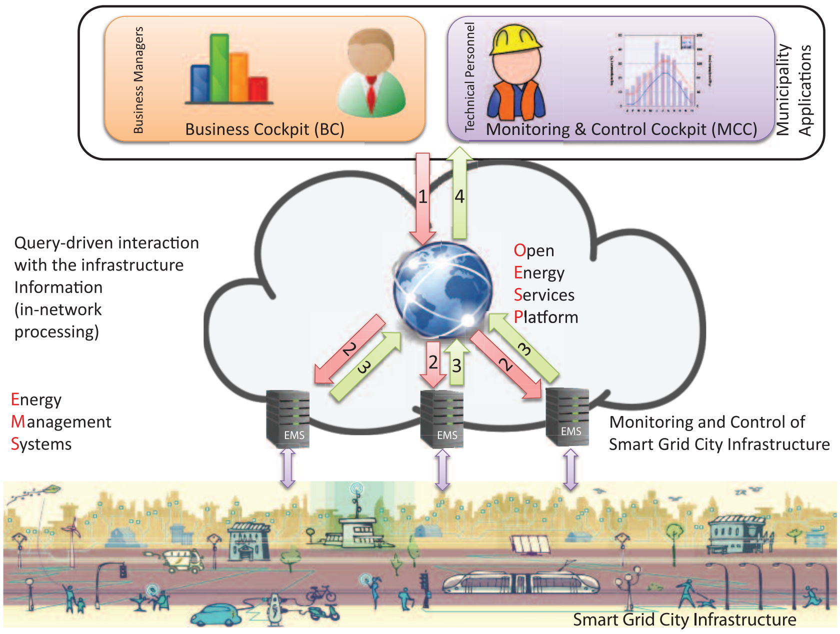

Figure 1: The SmartKYE concept of query-driven integration and management

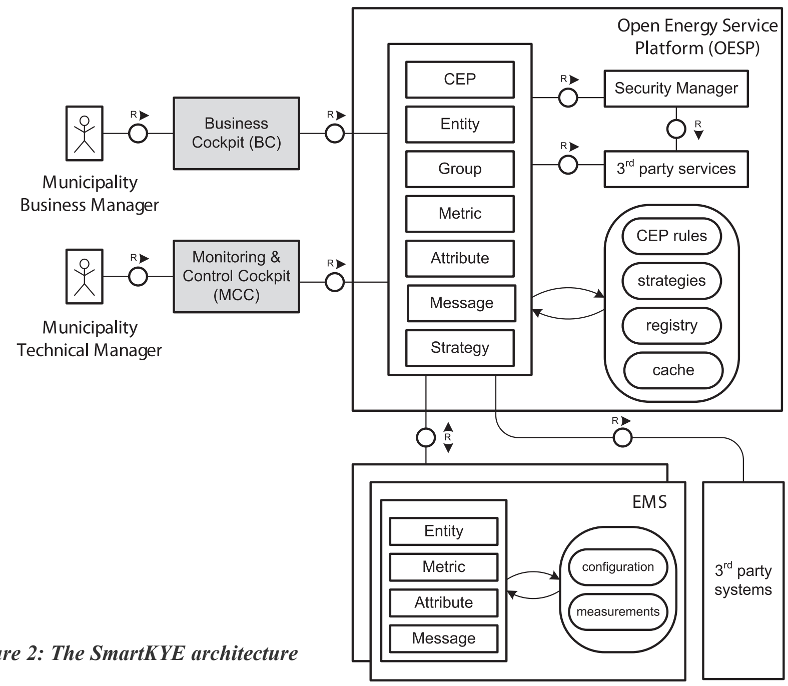

Figure 2: The SmartKYE architecture

The SmartKYE project project goes beyond existing efforts towards a realistic approach for future Smart Grid Cities. Most research projects [1] that deal with aspects of energy management are still focusing on core infrastructure issues that will enable them to measure energy production and consumption and potentially manage it [3]. Some go a step further towards building analytics and value-added services based on the data acquired. Recent advances in cloud computing and high performance in-memory databases tuned for delivering timely analytics and various web tools/apps [3] empower such efforts.

One key point that may be problematic, however, is information ownership and management. Most of the approaches developed today assume usually centralized unconditional access to the data [2]. This is a valid assumption within the scope of an organization, such as an energy provider, that controls the whole value chain. However, in large-scale infrastructures multiple such stakeholders will emerge with portions of the data, and many are likely to be reluctant to share their data unconditionally since it represents a key asset and competitive advantage of an enterprise. Hence, business models need to be extended to include interactions among stakeholders that will deliver value-added services without revealing all of the data, i.e., that will reveal relevant details without handing over an enterprise's core advantages.

The SmartKYE project capitalizes exactly on that vision. Contrary to currently overwhelmingly centralized information gathering approaches, SmartKYE considers a query-driven interaction with the vast number of stakeholders and their systems (as illustrated in Figure 1). A loosely coupled infrastructure enables stakeholders to join or leave the SmartKYE system, and the core infrastructure data may still be owned by the respective systems. To what extent and to what resolution such data is communicated to other stakeholders, depends on the queries issued and the willingness or contract-based negotiated actions. By not owning the data, and having an infrastructure that intelligently handles the communication between consumers and producers of information, scalability can be achieved, as well as coevolution at both ends. As depicted in Figure 1, several stakeholders monitor and control parts of the smart grid city infrastructure. As Figure 2 illustrates, the acquired information is gathered via various Energy Management Systems (EMS) which also host control capabilities for actuating the underlying infrastructure. Each EMS has its own local view and can operate autonomously from the rest. In the cloud, the Open Energy Services Platform (OESP) acts as the glue among the disparate systems and the information consumers who, in our example case, are portrayed at the top as the municipal information systems. The OESP hosts several advanced functionalities including: the capability of accepting queries from the end-user applications, disaggregating them to see which EMS need to be addressed, contacting the EMS to acquire the necessary information, aggregating the answers, and delivering the information as requested. The latter is done with additional indicators on the quality of data delivered. In addition, OESP offers services via which management actions can be undertaken on the infrastructure. Finally the end-users interact via cockpits with the system. In SmartKYE, as a proof of concept, two cockpits are addressed with different contexts:

- the Business Cockpit (BC) targets the municipal administrators and decision makers who can get high-level information and can be assisted in their decision processes; and
- the Monitoring and Control Cockpit (MCC) which targets the municipality technical personnel and infrastructure technical managers, by providing them with fine-grained technical information as well as control (management) capabilities.

References:

[1] V. Giordano, A. Meletiou, C. F. Covrig, et al: "Smart Grid projects in Europe: Lessons learned and current developments 2012 update," Joint Research Center of the European Commission, JRC79219, 2013

[2] S. Karnouskos, D. llic, P. Goncalves Da Silva: "Assessment of an Enterprise Energy Service Platform in a Smart Grid City Pilot", in IEEE 11th International Conference on Industrial Informatics (INDIN), Bochum, Germany, 2013

[3] S. Karnouskos, P. Goncalves Da Silva, D. Ilic: "Developing a Web Application for Monitoring and Management of Smart Grid Neighborhoods", in IEEE 11th International Conference on Industrial Informatics (INDIN), Bochum, Germany, 2013.

Link:
SmartKYE Project Web Site:
http://www.smartkye.eu

Please contact:
Stamatis Karnouskos
SAP, Germany
E-mail: stamatis.karnouskos@sap.com

'U-Sense', A Cooperative Sensing System for Monitoring Air Quality in Urban Areas

by Giuseppe Anastasi, Paolo Bruschi and Francesco Marcelloni

Air quality has a serious impact on public health, the environment and, ultimately the economy of European countries. In this article we present U-Sense, a cooperative sensing system for real-time and fine-grained air quality monitoring in urban areas. U-Sense allows for monitoring to occur in places where people spend the majority of their day-to-day lives.

Air quality continues to be a serious issue for public health, the environment and ultimately, the economy of European countries. Poor air quality results in ill health and premature deaths and damages ecosystems, crops and buildings. Urban areas, where the majority of European's live, are most seriously affected. In recent years, Europe has significantly reduced the emissions of several air pollutants such as sulphur dioxide (SO_2), carbon monoxide (CO), benzene (C_6H_6) and lead (Pb). However, particulate matter (PM), ozone (O_3), nitrogen dioxide (NO2) and some organic compounds still represent a serious threat. The report, "Air Quality in Europe", published in October 2013 by the European Environment Agency (EEA) [1] describes the effects of air pollution on health, ecosystems and the climate. It also provides an overview of policies and measures introduced in Europe to improve air quality and minimize the impacts of air pollution impacts.

Real-time, accurate monitoring of air pollution levels in urban areas plays a key role in enabling appropriate and timely public health decisions to be made and thus, is of paramount importance. Currently, air quality is typically monitored through large and expensive sensing stations, installed at strategic locations (few) such as intersections.

These stations do allow for accurate monitoring but their limited spatial coverage means this information is restricted to specific areas. In addition, the sensing stations are managed by public authorities which means that the pollution data they gather are often not publicly available. Citizens, however, are typically interested in knowing the air quality conditions in places which are relevant to their daily lives such as at home, school, work, or public species.

With this need in mind, we have developed U-Sense, a cooperative sensing system which allows for real-time, fine-grained air quality monitoring in urban areas (Figure 1). U-Sense relies on low-cost sensor nodes, equipped with appropriate gas sensors, which can be privately installed by citizens. The sensor nodes are powered by batteries which allow for flexible deployment and easy relocation. Users can share their measurements using social networking which enables cooperating sensing.

Data transfer can occur in a number of different ways. For example, sensor nodes can be connected directly to the database through a home-based WiFi router. Alternatively, data can be first transferred to an intermediate relay node and then on to the database. This relay node can be the user's smartphone or, more commonly, any mobile agent (e.g., a mobile relay node mounted on a bus or taxi). The sensor nodes can only transfer their locally acquired data when the relay node comes nearby. Hence, an opportunistic communication protocol is used for data transfer. Data accumulated on the database can be used for providing a number of community services through a variety of different user interfaces (e.g., the Web, smartphone or tablet). The range of services envisaged include access to air quality data measured by individual sensor nodes at any given location, the visualization of pollution maps and search facilities for less polluted paths. The web-based user interface for U-Sense is shown in Figure 2.

The U-Sense project is still in progress. We have completed the software implementation component and we have deployed a number of sensor nodes in across our city (Pisa, Italy). Libelium Waspmote sensor nodes have been used for this testing phase which are equipped with a sensor board that allows for the following air quality parameters to be measured: CO (carbon monoxide), CO_2 (carbon dioxide), NO_2 (nitrogen dioxide), O_3 (ozone), VOC (volatile organic compound), temperature and humidity.

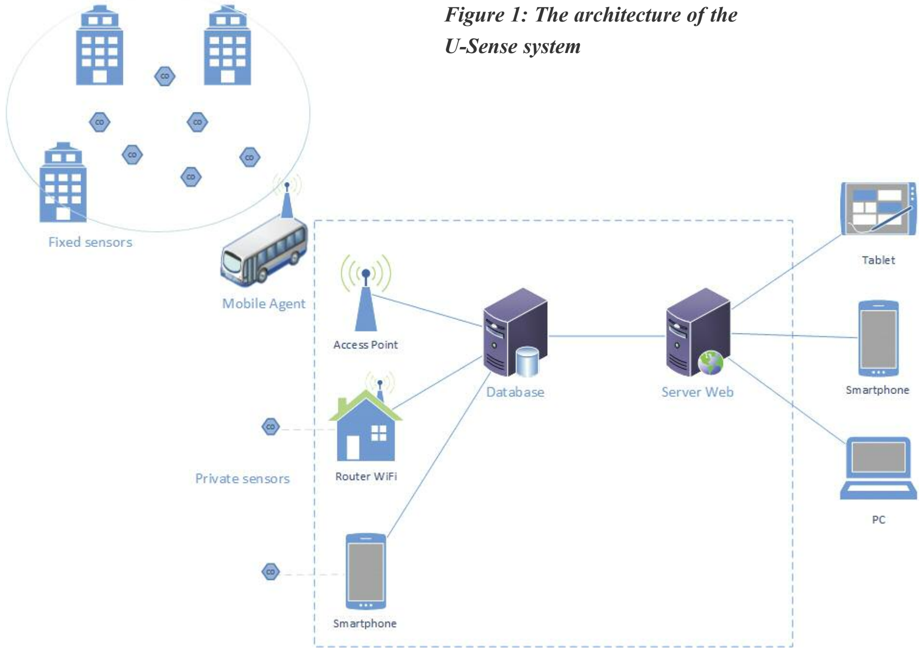

Figure 1: The architecture of the U-Sense system

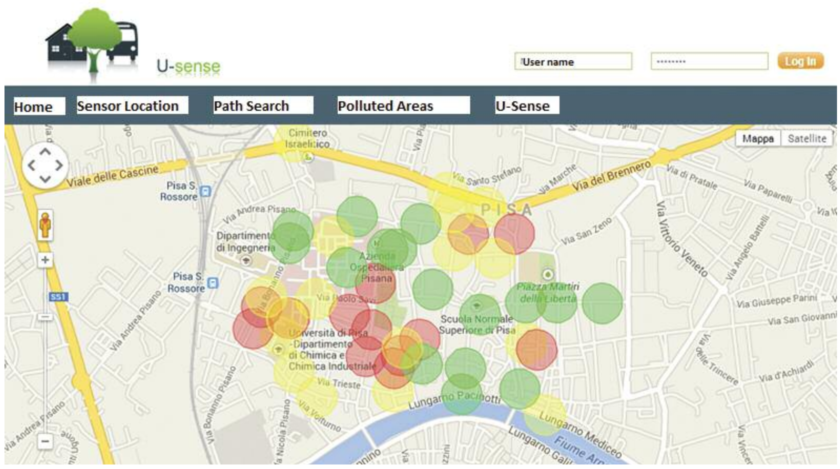

Figure 2: Web-based user interface for the U-Sense system.

This research activity is being undertaken as part of the "SMARTY" Project, funded by the Regional Government of Tuscany using European funds. The project aims to develop innovative services for sustainable transport and mobility in smart cities. Along with the University of Pisa, the University of Florence and a number of small and medium enterprises are contributing to this project.

Reference:

[1] European Environment Agency, "Air quality in Europe - 2013 Report", October 2013. Available at: http://www.eea.europa.eu/publications/air-quality-in-europe-2013.

Link:

http://www.iet.unipi.it/~anastasi/

Please contact:

Giuseppe Anastasi
Dept. of Information Engineering, University of Pisa,
Tel: +39 050 2217 559
E-mail: giuseppe.anastasi@unipi.it

Monitoring and Controlling Energy-positive Public Lighting: The E+grid System

by Balázs Csanád Csáji, Borbála Háy, András Kovács, Gianfranco Pedone, Tibor Révész, and József Váncza

The concepts of smart cities and self-sustaining renewable energy systems are revolutionizing the world of public lighting. This paper presents the architecture of a novel, adaptive and energy-positive outdoor lighting system, as well as the IT solutions that control and monitor the operation of the whole system.

This research was motivated by the potential opportunities to use renewable, solar energy in public lighting services via the appropriate combination of LED luminaries, energy generation and storage, as well as sensor technologies with novel data processing, communication and control methods. To this end, an industry-academy consortium formed by GE Hungary, the Budapest University of Technology and Economics, and two institutes of the Hungarian Academy of Sciences (MFA and SZTAKI) developed a so-called energy-positive community microgrid (E+grid). E+grid balances energy demand against production and guarantees the required level of street lighting even at times of moderate-duration power outages. In this article, we outline the key aspects of the information, communication and control features of the E+grid system.

System architecture

The E+grid system reduces the energy consumption of street lighting by using LED luminaries that regulate their lighting levels according to the surrounding environmental conditions, thus providing just the required level of lighting at all times. This is achieved by mounting motion sensors and smart controllers to each light pole, which are then connected by wireless communication. A positive yearly energy balance is achieved by photovoltaic (PV) energy generation, whereas protection against power outages is guaranteed by battery storage. The system has a bi-directional grid connection, which enables trading with electricity and can be used to exploit variable energy tariffs. Energy flows are monitored by smart meters. A local weather station collects weather data and hosts a twilight switch which allows for the lighting periods to respond to the current environmental conditions. The overall system (Figure 1) is monitored and controlled by a central computer (CC): black connectors indicate power flow while red lines show information flow (smart meters are not depicted).

The central computer

The CC of E+grid, an innovative cloud-deployed software application, is one of the key orchestrators of the system, enabling the adaptive lighting behaviour. Its core functions include:

- which are delivered by a web application and the extensive use of JavaScript. The graphical user interface (GUI) (Figure 2) provides friendly access to, among others, the geographic information system (GIS) of the installed luminaries, the management of the energetic components and the visual analysis and export of the collected data; and
- which enable the outer platform components to communicate with the CC. These are provided through a dedicated layer which is responsible both for handling the connections and validating the semantics. Each component is assisted by a proper data synchronization process, whereas the control of the outdoor lighting system is delegated to a specific lighting scheduler.

Controlling the energy flow

One of the main roles of the CC is controlling the energy flow, not only for minimizing the total energy cost, but also for ensuring the robustness against power outages through the use of battery storage. To achieve this, energy production and consumption must be forecasted. This is carried out by fitting time series models to the available data [1]. An assessment of several different models showed that good predictions of energy production can be obtained using nonlinear autoregressive exogenous (NARX) models. These models use wavelet-type nonlinearities, where the exogenous components come from a clear-sky model. Efficient energy consumption forecasts can also be achieved using Box-Jenkins type models, where the exogenous inputs come from averaged historical data. The controller itself uses a model-predictive approach

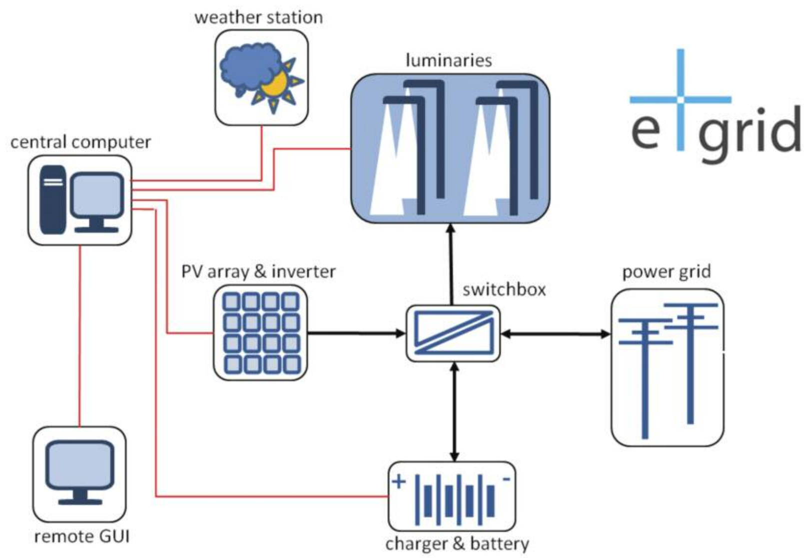

Figure 1: Architecture of the E+grid system

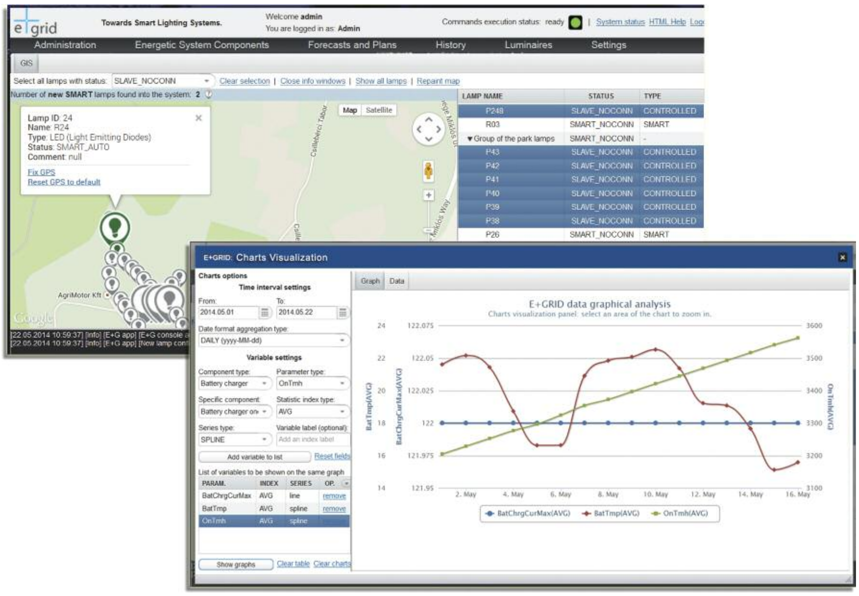

Figure 2: A graphical user interface (GUI) of the central computer

and works on a rolling-horizon where each time step involves solving a linear optimization problem [1].

Prototype system

The E+grid system design has been validated through simulation experiments which confirmed that the system can achieve a slightly positive yearly energy balance. A physical prototype, containing 130 luminaries and 152 m2 of PV panels is currently being deployed at the campus of the MFA Institute of the Hungarian Academy of Sciences. This prototype adopts three different battery technologies and four types of PV panels in order to evaluate their long-term performance under real-life environmental conditions.

Acknowledgement

This project has been supported by grants from the National Development Agency, Hungary, under contract numbers KTIA KMR 12-1-2012-0031 and NFÜ ED-13-2-2013-0002. B. Cs. Csáji acknowledges the support of the János Bolyai Research Fellowship No. BO/00683/12/6.

Reference:

[1] B. Cs. Csáji, A. Kovács, J. Váncza: "Prediction and robust control of energy flow in renewable energy systems", in proc. of the 19th IFAC World Congress, to appear, 2014.

Please contact:

András Kovács
Fraunhofer PMI, SZTAKI, Hungarian Academy of Sciences, Hungary
Tel: + 36 1 279 6299
E-mail: andras.kovacs@sztaki.mta.hu

Demand-Side Management in Smart Micro-Grids: An Optimization Perspective

by Talbi El-Ghazali

With the smart grid revolution, the energy consumption of houses will play a significant role in the energy system. Indeed, home users are responsible for a significant portion of the world's energy needs but market prices for this energy remain inelastic (i.e., energy demands do not follow energy prices). Thus, the performance of the whole energy generation and distribution system can be improved by optimizing the management of household energy use. There are a number of challenges associated with this goal including cost, the environmental considerations, user comfort and the presence of multiple decision makers (e.g., the end users and energy operators).

The smart micro-grid is an integrated system which supports distributed energy resources and multiple electrical loads but operates as a single power grid [1]. It is a smaller version of the traditional electrical grid or the smart digitalized grid which working independently of, or interconnected with, a larger existing power network. Smaller scale grids can deliver a wide range of improvements including greater reliability, fewer line losses, better fail recoveries, increased energy efficiencies, carbon emission reductions, reduced demands on the transmission infrastructure, cost reductions. Finally, it introduces the possibility of using alternative energy sources because more localized sources of power generation can be relied on. The global interest in reducing fuel consumption in favor of renewable energies has been a catalyst for the growth in the use of smart micro-grids. They are an ideal way to integrate renewable resources at the community level and allow for customer participation in the electricity enterprise. They enable electricity to be locally generated, distributed and electricity flows to consumers to be regulated.

A smart building system has its own micro-grid and some decentralized resources (e.g., a wind generator, CHP generator, boiler, thermal storage and electrical storage), to provide the basic electricity needs (Fig.1). It may also feature a grid connection which allows it to obtain electricity during peak hours or sell electricity (back to the grid) when surplus electricity is generated. Electric micro-grids are also regarded as a security mechanism as they ensure homes have access to energy during events such as weather-related black-outs or military attacks. All of these benefits are stimulating an increased demand worldwide for micro-grids which are being applied in a variety of contexts including campus and residential environments, military sectors and commercial and industrial markets. The global market for micro-grids topped $5billion in 2011, according to the 2013 Microgrid Market Report and is it seems likely this figure will reach $27 billion by 2022.

Optimization perspective

With the Smart Grid revolution, the energy consumption of households will play a significant role in the energy system. Indeed, home users are responsible for a significant portion of the world's energy needs but market prices for this energy remain inelastic (i.e., energy demands do not follow energy prices). Thus, the performance of the whole energy generation and distribution system can be improved by optimizing the management of household energy use. Smart Grid technologies have been deployed in a variety of contexts including smart sensors on distribution lines, smart meters in homes, advanced metering, transmission efficiencies, smart switches, enhanced demand response capabilities, distribution automation, electric vehicle integration and integrated communication systems. These pave the way for the creation of smart self-managing energy grids within a time varying pricing scheme (i.e., where prices may vary according to time of day). The Smart Homes of the future will include automation systems which may simultaneously provide lower energy consumption costs and a comfortable and secure living environment for end users. Residential users are expected to play a key role in improving network efficiency,through the adoption of intelligent mechanisms that manage energy demand. These days, most electricity consumers act as price takers with flat rates and no acknowledgement of the differences that exist in electricity prices. Therefore, they don't have any incentives to adjust their electricity consumption patterns. However, consumers can optimally adjust their energy consumptions by participating in a house management program which seeks to minimize electricity bills. To achieve this end goal, consumers must respond to shifting electricity prices and redesign their consumption patters to match with periods of lower electricity prices. The user can define the earliest and the latest starting times that a service can be processed, with a utility value identified for each time slot which indicates his or her comfort expectations. Thus, the challenge for the household energy management is to calculate when the optimal starting point occurs across the different household appliances in order to minimize energy costs for the user whilst maximize their satisfaction. In addition to reducing the energy bills of single users, joint management of multiple users can also be advantageous. If energy consumption tasks occurring across multiple homes can be scheduled, based on the suite of users' requirements, there is potential for reducing 'peak demand' periods. As all electricity grids require a minimum energy generation capacity so as to prevent failures and power cuts, overall grid costs might be reduced if peak demands can also be reduced. Peak shaving, as this phenomena is known, reduces the burden on the central grid and any upgrade expenses which may be incurred to ensure the grid can fulfill increasing energy demands. Finally, it also offers savings to both the utility companies and customers by reducing the need for generation capacity and energy amounts the utility companies must purchase on the open market during peak demand periods to prevent failures and power cuts. Thus, reducing the global contractual peak would bring multiple advantages, not just to the electric grid but also to residential users and energy retailers. In the DOLPHIN project, these challenges have been formulated as multi-objective optimization problems.

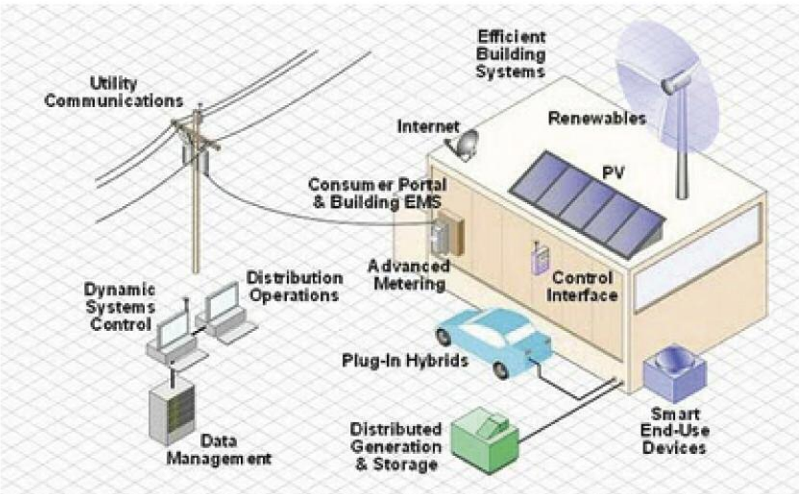

Figure 1: Demand-side management in a smart home.

Energy prices have to be defined not only to retrieve their production costs but also to understand consumer behavior. Consumers choose their services or energy providers with the aim of minimizing their disutility values. A failure to recognize this may lead to significant revenue losses. To ensure this hierarchical and multi-objective decision process is captured, where the leader (energy provider) explicitly takes the reaction of a follower (consumer) into account in their decision process, the energy pricing and management problems can be formulated as bi-level multi-objective problems [2].

References:

[1] J. Ekanayake et al: "Smart grid: Technology and applications", Wiley 2012.

[2] E. Alekseeva, L. Brotcorne, E-G. Talbi: "A new exact method finding a feasible solution for the bi-level energy pricing problem with two objectives at the lower level", in META'2014, Marrakech Morocco, October 2014.

Please contact:
Talbi El-Ghazali
University of Lille, Inria, CNRS, France
E-mail: el-ghazali.talbi@inria.fr

Cyber Physical Systems give Life to the Internet of Energy

by Giampaolo Fiorentino and Antonello Corsi

A new overlay network which runs on top of the electricity grid, puts together old and new renewable energy sources, giving life to the Internet of Energy. A coordinated and controlled Internet of Things will optimize the Distribution Grid Control and the Demand Side Management (DSM) based on business, comfort and occupancy dynamic, introducing the autonomous intelligent Commercial Prosumer Hub.

The introduction of new kinds of energy mixes to the electricity grid is a challenging environmental task for present and future generations as they fight the pollution and global warming issues associated with urbanization. Individual appliances and whole buildings that continuously incorporate local intelligence which originates from the new technologies of Internet Of Things (IoT) are the new infrastructure that this integration will be based on. Smart Electricity Grids are becoming more intensively integrated with tertiary building energy management systems and distributed energy generators such as wind and solar. Optimal energy management strategies require the ability to control and predict energy consumption incorporating all types of DER (distributed energy resources), at both local and global scales.

The INERTIA project

In the context of the EU FP7- ICT project INERTIA, a consortium of 10 complementary partners from six different European Countries aims at addressing the important challenges that future energy systems will face by building a new modelling methodology that integrates physical components and cyber technologies, thus creating Cyber Physical Energy Systems. The INERTIA project works on the definition of a new data management infrastructure to allow electricity production and consumption to be measured, reported and controlled, delivering an efficient Distribution Grid overlay control/management infrastructure (Figure 1). This new infrastructure will maximize the response capacity of the vast, small-commercial prosumer base (e.g., tertiary buildings, offices, etc.), presenting incentives and delivering benefits through their automated active participation in the energy market.

INERTIA extends the DSM strategies by incorporating various types of DER going beyond simple consumer loads, and treats both local generation and consumption under a single unified framework. The concept of DER flexibility will also be extended to incorporate various local context parameters, which affect DER operation and performance, as well as potential capacity of DER to provide flexible services to the Distribution Grid. This model includes the impact of communication networks and further, cyber components as well as the relevant information of the physical system.

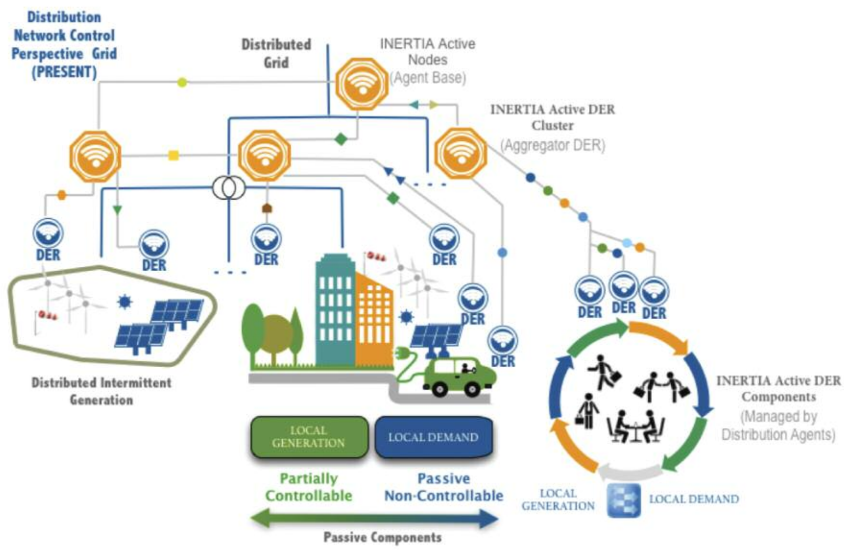

Figure 1 - The concept behind the INERTIA project. The overlay network will have the task of balancing distributed energy resources (micro-generation, RES, energy storage systems and demand) in the grid through the intelligent mining of energy flexibility, as a balancing power between local generation and local demand.

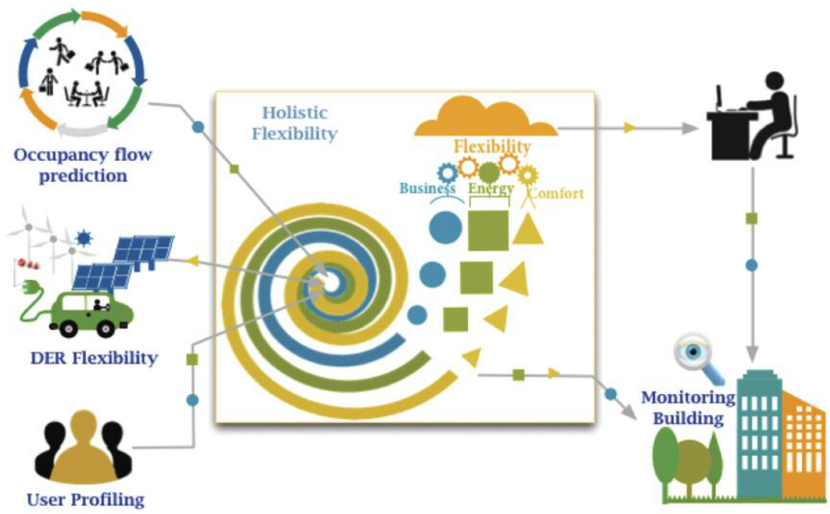

Figure 2: Holistic flexibility engine - This component is able to integrate and provide forecast energy consumption providing a real time tool for integrating medium voltage points of the net in the new smart grid energy.

The INERTIA framework

The old approach [1, 2] was to try to align consumption by asking consumers to reduce their power usage rather than increasing the power generation facilities. Under this approach, consumers that cooperated might receive incentive payments from the power company. In comparison, the INERTIA project claims to solve this issue using IoT technologies to expand DR services based on the analysis of users' historical interactions with the lighting, HVAC and other appliances controls. This has been realized calculating energy consumption, flexibility and forecast demand using a holistic flexibility model, which calculates several key performance indicators which can be categorized in Business, Energy, Comfort and Flexibility. Furthermore, real-time employee activities including arrival and departure times and meetings attendance (Figure 2) are also simulated.

Besides, INERTIA takes into account how demand response can change the end user comfort using a Bayesian formalism. This formalism is applied to infer the probability that any environmental situation should be considered by the user as environmentally uncomfortable, defining the upper and lower limits of the energy flexibility that can be offered to the whole energy system.

So, the holistic flexibility model is able to safeguard the electricity grid, facing a trade-off between energy provided by the local generation facility and consumers' consumption. In this respect, the holistic flexibility model generates the suitable energy demand strategy deploying DSM operations, which takes into account the final users comfort. The comfort is evaluated also by the facility manager selected to enroll a specific DR program

Conclusion

New IoT technologies are essential for implementing electrical power systems in Smart Cities and reducing environmental impacts and social costs. At the moment, appliances can self-adjust to consume less power or even turn themselves off. Using the data collected from all the devices, algorithms can calculate the exact energy demand so to reduce the need for standby generation. Over time, generation can be automatically shifted according to predicted increases and decreases in demand. Morover, a framework that receives input from relevant sources such as weather information, heating, cooling, lighting equipment and usage scenarios (occupancy rates) can be implemented to compute energy consumption and identify available flexibilities from DER. This could be conducted in a non-invasive way that would respect the comfort of the end user.

Acknowledgements

This work was partially supported by the EU project "INERTIA - Integrating Active, Flexible and Responsive Tertiary Prosumers into a Smart Distribution Grid" (FP7-ICT-2011-8, Grant Agreement No. 318216).

Link: http://www.inertia-project.eu/

References:

[1] A Di Giorgio, L Pimpinella: "An event driven Smart Home Controller enabling consumer economic saving and automated Demand Side Management"

[2] P Palensky, D Dietrich: - Industrial Informatics, IEEE …, 2011 - Demand side management: Demand response, intelligent energy systems, and smart loads

Please contact:

Giampaolo Fiorentino, Antonello Corsi

R&D Lab - Engineering Ingegneria Informatica S.p.A, Italy

E-mail: {giampaolo.fiorentino, antonello.corsi}@eng.it

When Smart Cities meet Big Data

by Vincenzo Gulisano, Magnus Almgren and Marina Papatriantafilou

Sharing information is a key enabler in the transition of a city becoming smart. Information, generated by the ICT backbone of a city, and maintained by distinct public and private entities, comes with processing challenges that must be addressed in order to increase citizens' quality of life and make their cities sustainable. In CRISALIS and SysSec, we investigate such challenges from a security perspective in order to protect and enhance smart cities' sensitive infrastructures.

The possibilities enabled by Information and Communication Technologies (ICTs) are driving the evolution and transition of cities to Smart Cities. The ultimate goal is to increase the awareness of citizens', companies' and authorities' and improve their quality of life while also making it sustainable. A considerable number of research directions embrace Smart Cities: users' privacy protection [1], detection of malicious actions and misuses and users' awareness through social media. More research efforts are dedicated to specific features of a Smart City. As an example, the energy forecast techniques used to predict consumption and allow the usage of alternative energy resources (e.g., solar or wind power) to be scheduled. What all these research fields have in common is their dependency on the (possibly sensitive) data produced by the devices forming the Internet of Things (IoT) of a city. The possibilities enabled by Smart Cities demand for novel data processing paradigms to form the expertise of public and private companies. Based on our experience with both academic and industrial partners, in this article we discuss some of the challenges associated with data processing in Smart Cities.

Scalable and online access to the data

In a Smart City, millions of messages will be exchanged on a daily basis by hundreds of thousands of devices (e.g., mobile phones, electrical meters, weather stations, etc.). For example, more than 1.2 million messages are exchanged on a daily basis within an AMI infrastructure (owned by one of our industrial partners) that covers a metropolitan area with roughly 600,000 inhabitants [2]. The information generated by such devices could be matched and joined to enhance the management of Smart Cities. For example, energy or

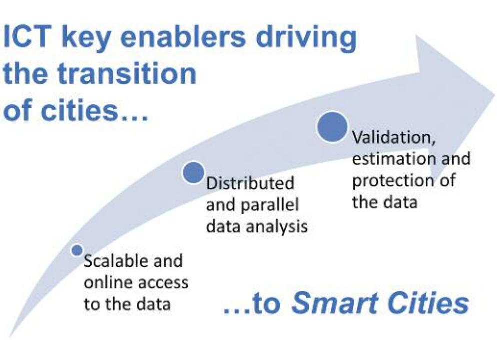

Figure 1: Key enablers

water losses caused by faulty devices could be reduced by matching the consumption measured by users' meters with the one measured by other utilities' systems. To this end, on-the-fly processing of data becomes all the more important while traditional store-then-process approaches in which each company retrieves its data and stores it in order to access it sometime in the future might be no longer appropriate.

Think in a distributed and parallel fashion

Smart Cities will be composed of several independent networks (even within the same stakeholder). Hence, no centralized application will embrace the information carried by the messages exchanged by the devices. At the same time, the huge volume of information shared by ICT devices will make parallel processing a necessity [3]. To this end, pushing the analysis closer to the sources of information would be a natural way of analyzing the messages exchanged by them and leverage the information they carry. Challenging aspects in this context will be imposed by the constrained resources of such devices.

Validate, estimate and protect the data

Cheap, resource-constrained devices are largely employed to build the networks that will form the IoT of a Smart City. Unfortunately, the data measured and reported by such devices (e.g., energy consumption readings) is usually noisy and lossy. Reasons of this are not limited uniquely to the devices themselves (e.g., faulty or badly calibrated devices, lossy or overloaded communication channels) but can also be caused by (possibly malicious) citizens. As an example, incorrect consumption readings could be manipulated by malicious users aiming to lower their bills. To this end, validation schemes, estimation schemes and security countermeasures must be adopted in order to ensure that who leverages the information is not mislead by incorrect, partial or malicious data.

The shift from cities to Smart Cities depends on the efficiency with which information is shared among citizens and private and public companies. This information brings challenges, and, following the big data revolution, novel processing schemes must be adopted to enable the possibilities that exist of this domain. All the possibilities enabled by smart cities, like improved quality of life or energy efficiency, shall build on top of efficient data processing and users' privacy protection schemes.

CRISALIS may be contacted at contact@crisalis-project.eu. SysSec may be contacted at the corresponding contact@syssec-project.eu, followed in twitter (twitter:syssecproject) and Facebook (http://www.facebook.com/SysSec).

References:

[1] V. Tudor, M. Almgre, M. Papatriantafilou: "Analysis of the Impact of Data Granularity on Privacy for the Smart Grid", in proc. of WPES'13

[2] Z. Fu et al.: "Managing your Trees: Insights from a Metropolitan-Scale Low-Power Wireless Network", in proc. of CCSES'14, held in conjunction with the INFOCOM 2014

[3] V. Gulisano, M. Almgren, M. Papatriantafilou: "METIS: a Two-Tier Intrusion Detection System for Advanced Metering Infrastructures", in ACM e-Energy, 2014.

Please contact:

Vincenzo Gulisano
Chalmers, Sweden
E-mail: vinmas@chalmers.se

cItyAM: Managing Big Urban Data for Analyzing and Modelling Cities

by Alessandro Bozzon, Claudia Hauff and Geert-Jan Houben

As we race towards fully connected living environments at full speed, urban stakeholders and decision makers are demanding analytic solutions that are able to provide actionable insights about citizens and their interactions with these environments. The Web Information Systems group at Delft University ofTechnology responded to this challenge by developing cItyAM, an integrated platform that supports the analysis and modelling of urban data. This data is received from a wide range of sources including social media, (social) sensors, open government data and a multiple knowledge repositories.

An open and extendable urban analytics platform, cItyAM is designed to cope with big urban data, providing stakeholder and decision makers with a set of tools to 1) integrate and analyze various urban data types and 2) model, engage and interact with people in urban spaces.

In view of the trend toward making cities fully connected and rising urban population densities, cItyAM provide the wherewithal to exploit big urban data, thus allowing urban phenomena to be mapped (e.g., the mobility, environmental, sustainability and social aspects associated with urban lifestyles).

The cItyAM platform has been developed as part of a research initiative started by the Web Information Systems group at the Delft University of Technology in 2012. It was driven and inspired by several European- and national-level initiatives including the City Data Fusion EIT ICT Lab project (conducted in collaboration with Politecnico di Milano, CNR Pisa, SIEMENS, and Telecom Italia), the Amsterdam Institute for Advanced Metropolitan Solutions (AMS) initiative (which includes WAGENINGEN UR, MIT, TNO, and other societal and industrial partners), the SHINE project, the ImReal EU FP7 project, and the COMMIT SEALINCMedia project (conducted in collaboration with the VU Amsterdam, CWI, the Rijksmuseum, and others), as well as the experience gained from the spinoff activities around Twitcident. The goal of the cItyAM platform is to progress urban data management, by advancing our ability to “feel the pulse” of cities: it provides methodologies and systems for data fusion, analysis and visualization.

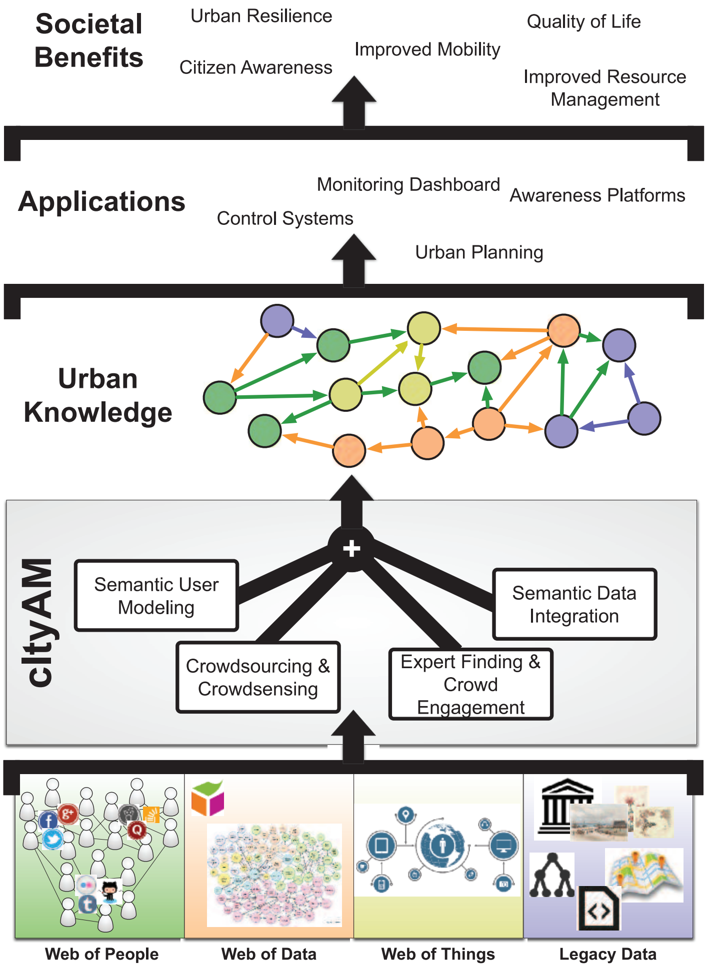

Figure 1: The role of the cItyAM platform

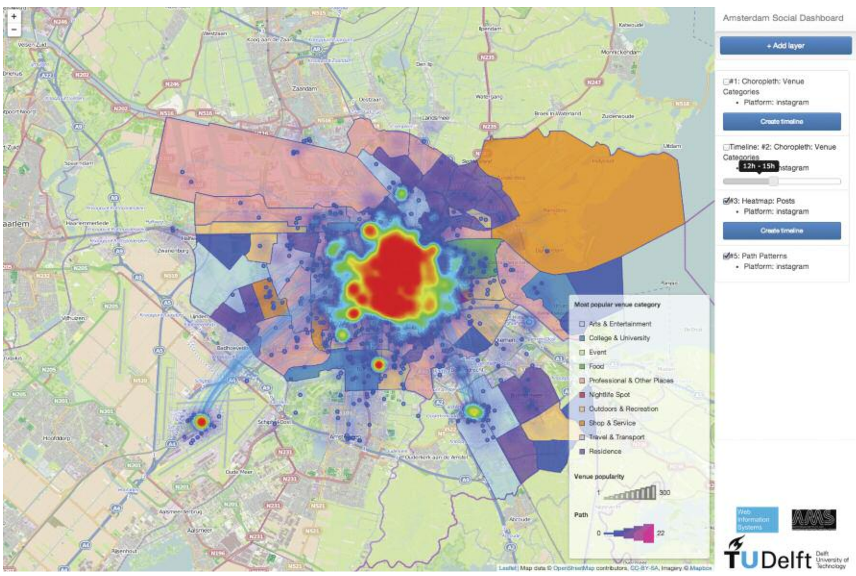

Figure 2: The Amsterdam Social Dashboard, developed on top of the cItyAM platform for the Amsterdam Institute for Advanced Metropolitan Solutions (AMS) opening day.

The cItyAM platform capitalizes on the experience of the Web Information System group and employs four main pillars:

1. Semantic user modelling techniques, drawn from the experiences gained during the ImReal project [1], and in this project tailored to the needs and features of urban citizen modelling. It allows for the semantically enriched analysis of conversations and activities which take place on social media channels (e.g., Twitter, Instagram and Foursquare).
2. Crowdsourcing and crowdsensing techniques for data collection, refinement and verification, integration and interpretation and processing. City-associated data can be patchy, noisy and disparate. Creating city related knowledge is driven by an extendable set of social sensing tools, like the CrowKnow system. Automatic techniques for data cleansing and linking often fail to achieve the levels of accuracy required for the effective fusion of urban data. Crowd-sourcing this cleansing and linking function is emerging as an effective means to solve the task when automatic methods fail.
3. Expert finding and crowd engagement strategies, developed within the COMMIT SEALINCMedia project [2] which actively drive the routing and engagement of citizens upon request from stakeholders. An example of this is when a crowdsensing campaign is initiated. The cItyAM incorporates incentive and engagement techniques that are specifically designed for urban environments and life styles.
4. Using the semantic data integration techniques that were also developed in the context of the ImReal project,

cItyAM can interface with a broad range of data sources. The RDFGears framework [3] allows for the integration, publication and retrieval of data coming from physical sensors, crowds and existing knowledge repositories such as DBpedia and Freebase.

Selected components of the cItyAM platform have been deployed in the last two years to monitor and analyze city-scale events like the Milano Design Week, the Lucca Comics & Games, the Dutch King's Day and others. In the future, we plan to deploy cItyAM as a continuous service operating in multiple cities around the world. In the context of Amsterdam, for instance, we plan to integrate with urban data sources managed by the municipality to support studies on urban mobility and social city life. In Rotterdam, the plan is to create social sensing solutions for water management, with a specific focus on emergency management for flooding events. Finally, we plan to provide a 24/7 social recommendation service covering the EXPO 2015 event in Milan.

Links:
http://citydatafusion.org
http://www.ams-amsterdam.com/
http://direct.tudelft.nl/shine-117.html
http://www.imreal-project.eu
http://www.commit-nl.nl
http://twitcident.com
http://www.wis.ewi.tudelft.nl/CroKnow
http://www.wis.ewi.tudelft.nl/research-lines/rdfgears

References:
[1] Qi Gao, et al.: "A Comparative Study of Users' Microblogging Behavior on Sina Weibo and Twitter", UMAP 2012
[2] J. Oosterman, et al.: "Crowd vs. experts: nichesourcing for knowledge intensive tasks in cultural heritage", WWW (Companion Volume) 2014
[3] M. Grabowski, J. Hidders, J. Sroka: "Representing MapReduce, Optimisations in the Nested Relational Calculus", BNCOD 2013

Please contact:
Alessandro Bozzon, Claudia Hauff, Geert-Jan Houben
TU Delft, The Netherlands
E-mail: a.bozzon@tudelft.nl, c.hauff@tudelft.nl, g.j.p.m.houben@tudelft.nl

Urban-Scale Quantitative Visual Analysis

by Josef Sivic and Alexei A. Efros

Urban-scale quantitative visual analysis opens up new ways Smart Cities can be visualized, modelled, planned and simulated by taking into account large-scale dynamic visual inputs from a range of visual sensors.

Map-based street-level imagery, such as Google Maps with Street View, provides a comprehensive visual record of many cities around the world. For example, the visual appearance of Paris has been captured in almost 100,000 publically available Street View images. We estimate there are approximately 60 million Street View images for France alone, covering all the major cities. In the near future, additional visual sensors are likely to become more wide-spread, for example, cameras are being built into most newly manufactured cars. Another increasing trend is the ability for individuals to continuously capture their daily visual experiences using wearable mobile devices such as the Google Glass. Collectively, all this data provides large-scale, comprehensive and dynamically updated visual records of urban environments.

Automatic analysis of urban visual data

The next opportunity lies in developing automatic tools that allow for the large-scale quantitative analysis of the available urban visual records. Imagine a scenario where we could provide quantitative answers to questions like:

- What are the typical architectural elements that characterize the visual style of a city (e.g., window or balcony type)? (Figure 1a)
- What is their geo-spatial distribution? (Figure 1b)?
- How does the visual style of an area evolve over time?
- What are the boundaries between visually coherent areas in a city?

These examples just touch on the range of interesting questions that can be posed regarding the visual style of a city. Other types of questions concern the distribution of people and their activities For example, how does the number of people and their activities at a particular place evolve during a day, the seasons or years? Or perhaps you might want to know the make-up of activities on a given street: the presence of tourists sightseeing, locals shopping, the elderly walking their dogs, or children playing. This type of data can also be used to respond to significant urban issues, for example, what are the major causes of bicycle accidents?

New applications

To progress the way we can respond to these types of questions would open-up new ways Smart Cities can be visualized, modelled, planned and simulated by taking large-scale dynamic visual inputs from a range of visual sensors into account. Some examples of how this data might be applied include:

- the real-time quantitative mapping and visualization of existing urban spaces [1] to support architects and decision makers (Figure 1),
- the ability to predict and model the evolution of cities [3] (e.g., land-use policies and the way they impact on the visual appearances of different neighbourhoods),
- obtaining detailed dynamic semantic city-scale 3D reconstructions and using them to simulate different environmental scenarios, e.g., levels of noise, energy consumption or illumination, and
- the analysis of human activities, e.g., evaluating the future success of a restaurant or the need of to introduce new traffic security measures.

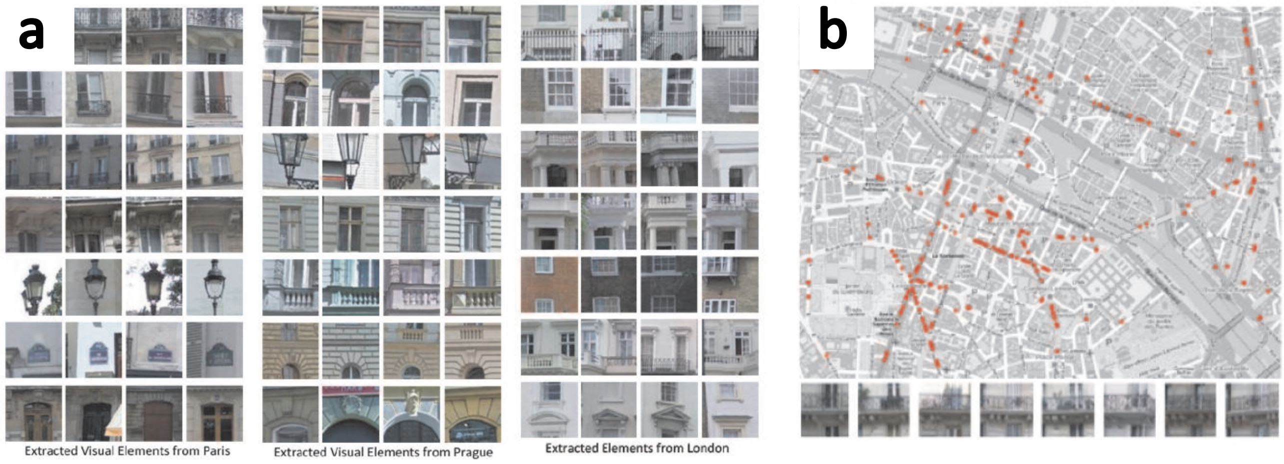

Figure 1: Quantitative visual analysis of urban environments from street view imagery [1].
1a: Examples of architectural visual elements characteristic of Paris, Prague and London, identified through the analysis of thousands of Street View images. 1b: An example of a geographic pattern (shown as red dots on the map of Paris) of one visual element, balconies with cast-iron railings, showing their concentration along the main boulevards. This type of automatic quantitative visual analysis has a potentially significant role in urban planning applications.

The challenge of urban-scale visual analysis

Impressive demonstrations of the analysis of large-scale data sets have already started to appear in other scientific disciplines. In natural language processing, an analysis of more than 5 million books published between 1800 and 2000 revealed interesting linguistic, sociological and cultural patterns [2]. In the visual domain, however, a similar large-scale analysis has yet to be demonstrated. As visual data and computational resources are becoming more widely available, the key scientific challenge now lies in developing powerful models which can competently meet the spatio-temporal, widely distributed and dynamic characteristics of this visual data. For example, while the vocabulary and grammar of written text are well defined, there is no accepted visual vocabulary and grammar that captures the subtle but important visual differences in architectural styles, or the different visual appearances of human activities on city streets.

Example: quantitative analysis of architectural style

In this first phase of investigation, we considered quantitative visual analysis of architecture style [1]. Using the large repository of geo-tagged imagery, we sought to find a way of automatically identifying which visual elements, e.g., windows, balconies and street signs define a given geo-spatial area (in this case Paris). This is a tremendously difficult task as the differences between the distinguishing features of different places can be very subtle. We were also faced with a difficult search problem: given all the possible patches in all the possible images, which patches are both geographically informative and occur frequently? To address these issues, we proposed a discriminative clustering approach which took into account the weak geographic supervision. We show that geographically representative image elements can be discovered automatically from Google Street View imagery in a discriminative manner. We applied the algorithm on image datasets from 12 cities (Paris, London, Prague, Barcelona, Milan, New York, Boston, Philadelphia, San Francisco, San Paulo, Mexico City and Tokyo), with each dataset featuring approximately 10,000 images. An example of the results was discussed above (and illustrated in Figure 1). This example demonstrates that these learnt elements are visually interpretable and perceptually geo-informative. We further demonstrate that the identification of these elements can support a variety of urban-scale quantitative visual analysis tasks, such as mapping architectural correspondences and influences within and across cities, or finding representative elements at different geo-spatial scales [1].

Link:
CityLab@Inria Project Lab on Smart Cities: http://citylab.inria.fr

References:
[1] C. Doersch, S. Singh, A. Gupta, J. Sivic, and A. Efros. What makes Paris look like Paris? ACM Transactions on Graphics (SIGGRAPH 2012)
[2] J.B. Michel et al. Quantitative analysis of culture using millions of digitized books. Science, 331(6014):176–182, 2011
[3] C.A. Vanegas et al. Modelling the appearance and behaviour of urban spaces. Computer Graphics Forum, 29(1):25-42, 2010.

Please contact:
Josef Sivic
Inria and Ecole Normale Supérieure, France
E-mail: Josef.Sivic@ens.fr

Alexei A. Efros
UC Berkeley, USA
E-mail: efros@eecs.berkeley.edu

Mobile Augmented Reality Applications for Smart Cities

by Mathieu Razafimahazo, Nabil Layaïda, Pierre Genevès and Thibaud Michel

The TyReX research team at the Inria Grenoble - Rhône-Alpes Research Centre, France is working on the design of a high-accuracy localization system. This key technology allows for a smooth integration of vision-based techniques in augmented reality applications [1].

Since 2007, the widespread use of mobile devices (e.g., phones, tablets) and emergence of a new generation of devices which are equipped with more precise sensors and cameras has placed augmented and virtual-reality technologies within reach. Such applications are very resource-demanding and critically interdependent on the quality of user-location information and environment tracking (e.g., buildings, monuments). Several mobile services based on augmented-reality technologies have already been developed. These can aid navigation, offer context-aware access to a wealth of geo-tagged content related to cultural heritage, provide information on city events and assist with urban planning issues. However, these services still lack enhanced and robust tracking techniques that are sufficiently effective enough to make the applications reliable and widely deployed.

Today, one of the main challenges in mobile augmented reality applications design is understanding how our perception of reality can be profitably augmented and how this augmentation can be made to fit seamlessly with the user's interaction with the world. The European VENTURI project which began in 2011 is aiming to develop the first generation of ubiquitous augmented reality (AR) tools that meet real user needs and fit within the context they operate in. To illustrate the potential impacts such technologies might have in our everyday lives, a demonstration will be conducted in front of a large audience in late 2014. This demonstration will be a cultural-heritage visit of Grenoble (France) presented on smartphones (Figure 1). The application features an augmented tour of Grenoble's antique neighbourhood as far as the 19th century fort (476m elevation). Along the route, the application behaves like an audio AR guide, delivering walking directions and historical information. It also includes a set of immersive virtual history galleries. Overall, it exposes the user to Grenoble's past and present by combining 3D-reconstructed monuments and old geo-positioned pictures. For example, the city's old roman fortifications are overlaid on a live camera view using visual-recognition techniques. Finally, to further enhance user immersion, the application provides different 3D soundscapes synchronized with the visual rendering.

This demonstration addresses a wide spectrum of the technological challenges a project faces including audio-visual scene analysis to understand the user's context, the collection, creation, fusion and delivery of AR content, 3D audio rendering, mobile human-machine interactions and finally the provision of a high-accuracy localization system. To this end, the VENTURI project consortium brings together different ERCIM members such as Inria, Fraunhofer, FBK, mobile device manufacturers SONY and ST Microeletronic and software companies like Metaio GmbH and EDIAM Sistemas. The common goal of all these partners is to design a hardware and software platform dedicated to such applications.

Of these technologies, the high-accuracy localization is central to allowing for smooth integration with vision-based techniques. The TyReX research team, located in the Inria Grenoble - Rhône-Alpes Research Centre, France has been working since 2010 on the design of such a localization system, commonly known as Pedestrian Dead-Reckoning (PDR). This system is based on the traditional micro-electro-mechanical systems (MEMS) used in smartphones. The on-board accelerometer and compass are used to provide an estimation of the user's relative position

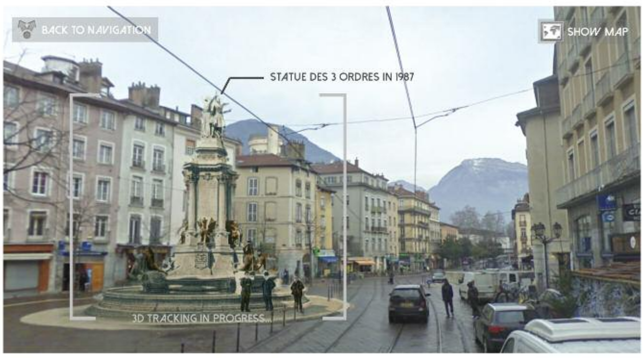

Figure 1: Examples of user interfaces seen as part of the cultural-heritage visitor application developed for Grenoble which uses augmented reality tools.

and orientation in 'urban canyon' environments (defined as areas with poor GPS-signal strength such as indoors). Combined with geographical data, this component can also be used to identify user-activity patterns (e.g., walking, running or being in an elevator). The interpretation of sensor values, coupled with different walking models, allows us to ensure continuity in determing the user's location when they are both indoors and outdoors. However, dead recknoning is subject to cumulative errors that are driven by multiple factors (e.g., sensor drift, missed steps or poor estimation of stride length). These errors can be reduced by fusing various external signals from sources such as GPS and Wi-Fi or by relying on the analysis of a user's trajectories with the help of a structured map of the environment.

PDR has played a key role in the VENTURI application since the delivery of its AR content is mainly based on the user's estimated location. PDR and the other technologies mentioned in this article aim to provide available information in a 'user-centric' way as opposed to a 'device-centric'. The 2014 demonstration in Grenoble, which combines advanced AR, indoor and outdoor PDR and 3D vision tracking will be this project 's final deliverable.

Links:
TYREX team: http://tyrex.inria.fr
VENTURI project: http://venturi.fbk.eu

Reference:
[1] J. Lemordant, T. Michel, M. Razafinahazo: "Mixed Reality Browsers and Pedestrian Navigation in Augmented Cities" in proc. of the Graphical Web Conference, Oct 2013, San Francisco, USA, http://hal.inria.fr/hal-00872721/en

Please contact:
Nabil Layaïda
Inria Rhône-Alpes, France
E-mail: nabil.layaida@inria.fr

Monitoring People's Behaviour using Video Analysis and Trajectory Clustering

by Francois Bremond, Vania Bogorny, Luis Patino, Serhan Cosar, Guido Pusiol and Giuseppe Donatiello

Activity discovery within transportation systems (for example subways and roads) or home care monitoring based on cognitive vision and data-mining technologies, are the core activities of a project at Inria.

It is well known that video cameras provide one of the richest and most promising sources of information about people's movements. New technologies which combine video understanding and data-mining can analyse people's behaviour in an efficient way by extracting their trajectories and identifying the main movement flows within a scene equipped with video cameras. For instance, we are designing an activity recognition framework which can monitor people's behaviour in an unsupervised manner. For each observed person, the framework extracts a set of space-time trajectory features describing his/her global position within the monitored scene and the motion of his/her body parts. Based on trajectory clustering, human information is gathered in a new feature that we call Perceptual Feature Chunks (PFC). The set of PFC is used to automatically learn the particular regions of a given scene where important activities occur. We call this set of learned scene regions the topology. Based on a k-means algorithm, a clustering procedure over the PFCs is performed in order to construct three topology layers, organized from coarsest to finest. Using topologies and PFCs, we are able to break the video into a set of small events or primitive events (PE), each of which has a semantic meaning. The sequences of PE and the three layers of topology, are used to construct a hierarchical model with three granularity levels of activity.

The proposed approach has been experimented in collaboration with Nice Hospital within the FP7 European project DEM@CARE collecting datasets to supervise patients suffering from dementia. These datasets contain older adults performing everyday activities in a hospital room, equipped with a monocular and RGBD cameras (resolution: 640x480 pixels). The activities considered include "watching TV", "preparing tea", "answering phone", "reading newspaper/magazine", "watering plant", "organizing the prescribed drugs", "writing a check at the desk" and "checking bus routes on a map". The mono-camera dataset consists of 41 videos and the RGBD dataset contains 27 videos. For each person, the video lasts approximately 15 minutes. For the monocular camera, person detection is performed using an extension of the Gaussian Mixture Model algorithm for background subtraction. For the RGBD camera, we used a person detection algorithm that detects people's heads and shoulders. Trajectories of people in the scene are obtained using a multi-feature algorithm that uses features such as 2D-size, 3D-displacement, colour histogram, dominant colour and covariance descriptors.

Figure 1: Extracted trajectories of a person at home during the preparation and eating of a meal (left)and people in a metro station while buying tickets (right).

Experimentations on both datasets show that the framework achieves a high rate of True Positives and a low rate of False Negatives. In total, 99% of the performed activities (in real-life) are recognized by the framework [1]. Furthermore, the duration of the recognized activities is matched with more than 80% accuracy to the ground truthed activities, which means that not only can the system count the amount of activity instances but also pretty accurately detect their duration's. Although a few activities were missed because of failure to detect finer motions, the experimental results show that this framework is a successful system that can be used to automatically discover, learn and recognize Activities of Daily Livings (ADLs). In addition, it can be observed that the framework is useful in medical applications for supporting the early diagnosis of Alzheimer or dementia in older adults. The framework can successfully distinguish people suffering from Alzheimer from those with Mild Cognitive Disorder or Normal Controls.

This framework can also be used in many other fields, such as the video surveillance of metros and roads. We have applied a variant of this framework in two scenarios within the FP7 European project VANAHEIM. The first scenario relates to the monitoring of activities in different locations (e.g., the entrance concourse) of Paris and Turin metro stations. The results obtained show which zones have the most intense activities (called Heat Map). In the scene under observation, rare or uncommon behaviours such as "jumping above the barrier", "fainting" and "loitering" or frequent behaviours such as "buying tickets" were identified. In the second scenario, a road lane reserved for buses was surveyed. Again, we were able to learn the topology of the scene and reveal which were normal activities (i.e., the passage of a bus into the zone) and abnormal activities (i.e., the passage of other vehicles into the reserved lane). When the results were groundtruthed, a high level of recall and an acceptable degree of precision were obtained [2, 3].

A variety of domains can benefit from smart video analysis. For instance, in the near future, more than 2 billion people will be over 65 years old, and video analysis has the potential to help aging adults in their daily life through the use of smart home environments and to help providing doctors with activity observations that can be used to detect possible anomalies for disease prevention. Movement analysis in crowded areas, such as metro stations, can detect and alert anomalous behaviours, identifying simple issues such as a broken ticket machine to more complex events like a robbery or terrorism act. In supermarkets, the analysis of customer movements can provide information on how to enhance the shopping experience. However, this analysis still has many remaining challenges, such as the accuracy of activity recognition within huge amounts of noisy metadata.

References:

[1]. G. Pusiol, F. Bremond and M. Thonnat: "Unsupervised Discovery, Modeling and Analysis of long term Activities, in proc. of ICVS 2011, Sophia Antipolis, France, 2011.

[2] L. Patino, F. Bremond, M. Thonnat: "Online learning of activities from video", in proc. of AVSS 12, Beijing 2012.

[3] L. Patino, H. Benhadda, F. Bremond: "Data Mining in a Video Database", in "Intelligent Video Surveillance Systems", Wiley, Online Library, 2013, DOI:10.1002/9781118577851.ch14

Please contact:

François Bremond
Inria, France
E-mail: Francois.Bremond@inria.fr

Trusted Cells: Ensuring Privacy for the Citizens of Smart Cities

by Nicolas Anciaux, Philippe Bonnet, Luc Bouganim and Philippe Pucheral

The Smart City concept is founded on the collection, sharing and analysis of data that is either about citizens or produced by them, with the view to enhancing efficiencies and the social sustainability of cities. The current Web model which is fully centralised is not appropriate for managing such data as it raises potential privacy abuse and misuse issues. In the Trusted Cell project, we propose the addition of a personal dimension to the Web model: each citizen would possess their own personal data server which would provide tangible privacy and security guarantees and help individuals to share and disseminate their data properly.

In smart cities, the convergence of mobile communication, sensors and online social networks technologies has lead to an exponential increase in the creation and consumption of data which can be linked to individuals (i.e., personal data). These data are considered by the World Economic Forum as "the new oil", creating an unprecedented potential for new applications and business. However, there are privacy concerns linked to the Web model currently in use and its underlying businesses. Frequent breaches of users' privacy are thwarting enthusiasm for the use of this data.

In smart cities, privacy is a fundamental value, upheld by social sustainability. The personal data given back to citizens can potentially describe all of the activities undertaken in day-to-day lives. To exploit the data, a new sharing and usage model which is more user-centric and prioritises privacy must be invented. Our vision [1] is to launch a sea change in the way data and applications are managed, enabling users to exercise control over how their data is used by introducing trusted cells in the architecture. Trusted cells are units of hardware and software, owned by citizens, which are able to perform data

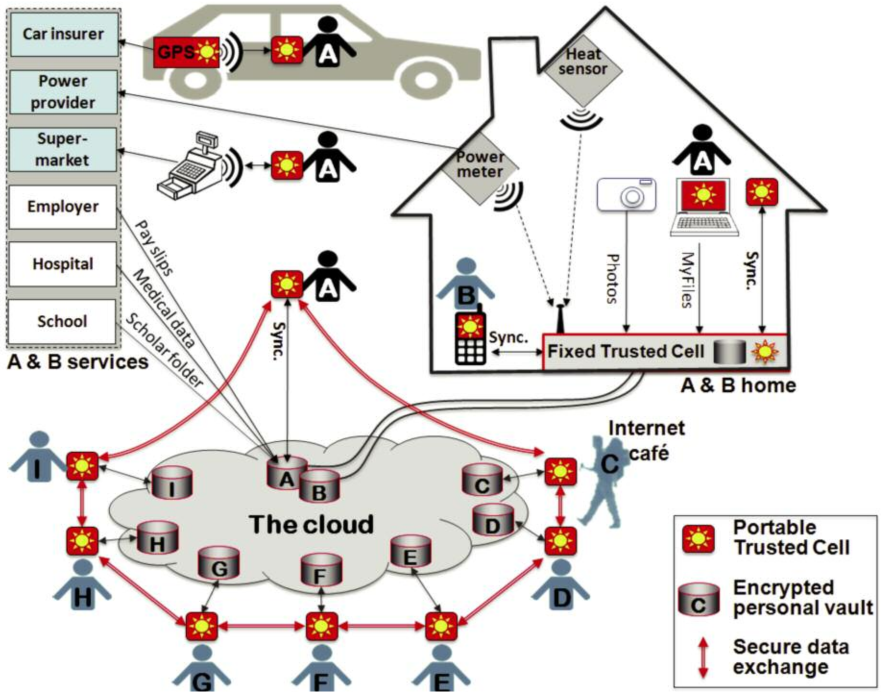

Figure 1: Personal data management in Smart Cities using trusted cells.

Alice (A) and Bob (B), citizens of a Smart City, are equipped with trusted cells, which acquire data from several sources. For example, their home is equipped with smart meters generating hundreds of measurements per hour. Companies provide them with energy management services based on smart meter data disaggregation (i.e., machine learning algorithms detecting appliance signatures and deriving their use in time). These data are managed into trusted cells which act as trusted gateways that enforce a usage control model, in connection with an untrusted cloud infrastructure.

management tasks in a privacy-preserving manner. Our goal is (i) to provide data management and access control techniques embedded into trusted cells which limits the personal data exposed to the services and applications to a minimum subset; and (ii) to rely on trusted cells to implement a usage control model which protects the personal data used by applications against unexpected disclosures once they have left the security sphere of a trusted cell.

Our approach is based on the vision that privacy-preserving trusted cells, i.e., trusted gateways deployed at home (fixed) or in users' hands (portable), can enforce a usage control model in connection with an untrusted cloud infrastructure. From a technical point of view, the challenges associated with this approach are:
- Developing trusted cells which provide strong security guarantees: in a Smart City, many devices can be considered as potential trusted cells including microcontroller based tokens (e.g., SIM cards or sensors) or mobilephones endowed with an ARM Trustzone processor. Preliminary proposals start addressing the support of relational operations like selection, projection and join within secure microcontrollers [2].
- Designing usage control models, as well as the mechanisms needed to enforce them. We are currently exploring a solution which relies on two building blocks that enforce access and usage control models when the data is transmitted outside the trusted cell to perform a computation. We designed and implemented a set of rich operations within the trusted cell that can be combined to reduce the computations to be performed outside the trusted cells. Then we investigated solutions based on the definition of sandboxed computation containers, that are able to connect to the trusted cell and extract the raw data whilst guaranteeing that unexpected data disclosures will not occur. We expect that these efforts will give rise to a new paradigm for the design of privacy preserving, user-centric applications.
- Global and anonymous computations. To compute results at a population-level, a large number of trusted cells may be required. The aim is to only reveal the results, but organize the computation so that the raw data and their owner’s identities remain hidden. Large sets of trusted cells with limited resources and potentially, low connectivity, may be involved. The feasibility of this approach has been illustrated in previous work [3] which addressed the problem of Privacy-Preserving Data Publication in this context. Computing regular queries, e.g., aggregations to discover overcrowded roads or conduct public surveys, is still an open issue.

Link:
CityLab@Inria Project Lab on Smart Cities: http://citylab.inria.fr

References:

[1] N. Anciaux et al: “Trusted Cells: A Sea Change for Personal Data Services”, in CIDR, 2013.
[2] N. Anciaux et al: “MILo-DB: a personal, secure and portable database machine”, in Distributed and Parallel Databases 32(1), 2014.
[3] T. Allard, B. Nguyen, P. Pucheral: “METAP: revisiting Privacy-Preserving Data Publishing using secure devices”, in Distributed and Parallel Databases 32(2), 2014.

Please contact:

Nicolas Anciaux
Inria and University of Versailles, France
E-mail: Nicolas.Anciaux@inria.fr

Smart City Operation Center: A Platform to Optimize Urban Service Rendering

by Filippos Gouidis, Theodore Patkos and Giorgos Flouris

The delivery of Public Services to citizens is becoming a challenging task, as urban populations increase and the related networks are becoming more complex and entangled. In this article, we present a platform designed for the optimal coordination of a city's public agencies and emergency response teams. Such coordination is particularly relevant in emergency situations when the infrastructures are pushed to their limits.

In a typical modern city, many different agencies are activated, in order to provide citizens with the services that underpin the seamless operation of the urban centres, such as the fire brigade, police department and health care services providers. The management of the actions necessary for this objective is in principle carried out by each organization independently, since a continuous coordination is considered too costly and arduous. A notable exception is the handling of emergency incidents due to their critical nature.

Over the last few years, many Smart City projects have been launched with the purpose of augmenting the capabilities of the relevant agencies in carrying out their tasks. For example, in Amsterdam, the Smart Traffic Management Project is used to administer traffic flow and the European Union (EU) is preparing to launch the new emerging trend of Connected Cars which aims at substantially improving ambulance response times. However, the approach underlying these applications is specialized rather than holistic, as is also the case with the various public agencies involved.

Nonetheless, constantly maintaining a coordinated and synergistic plan can be significantly optimized for a number of reasons. First, this approach translates to a substantial cost reduction in terms of human resources and infrastructure. In addition, the activities of two or more servicing organizations are often interdependent: properly regulated traffic congestion by the traffic police can facilitate enormously ambulances to reach their destination on time. Hence, a unified approach could better serve the interests of the stakeholders and ultimately the citizens of a city. Furthermore, the experience and expertise gained by exercising this type of administration, might prove to be crucial in the event of a large-scale emergency incident, such as an earthquake or extreme weather event.

The ISL laboratory of FORTH-ICS is developing the Smart City Operation Center (SCOC), a platform for the central coordination of the various urban services (Figure 1 shows the main window of the platform). The system receives requests from different locations in the city for servicing tasks and organizes the routes of the different service providing units (e.g. police trucks or ambulances). This is accomplished in such a way that the aforementioned demands are satisfied in the least amount of time possible. Figure 2 presents a toy example of agents cooper-

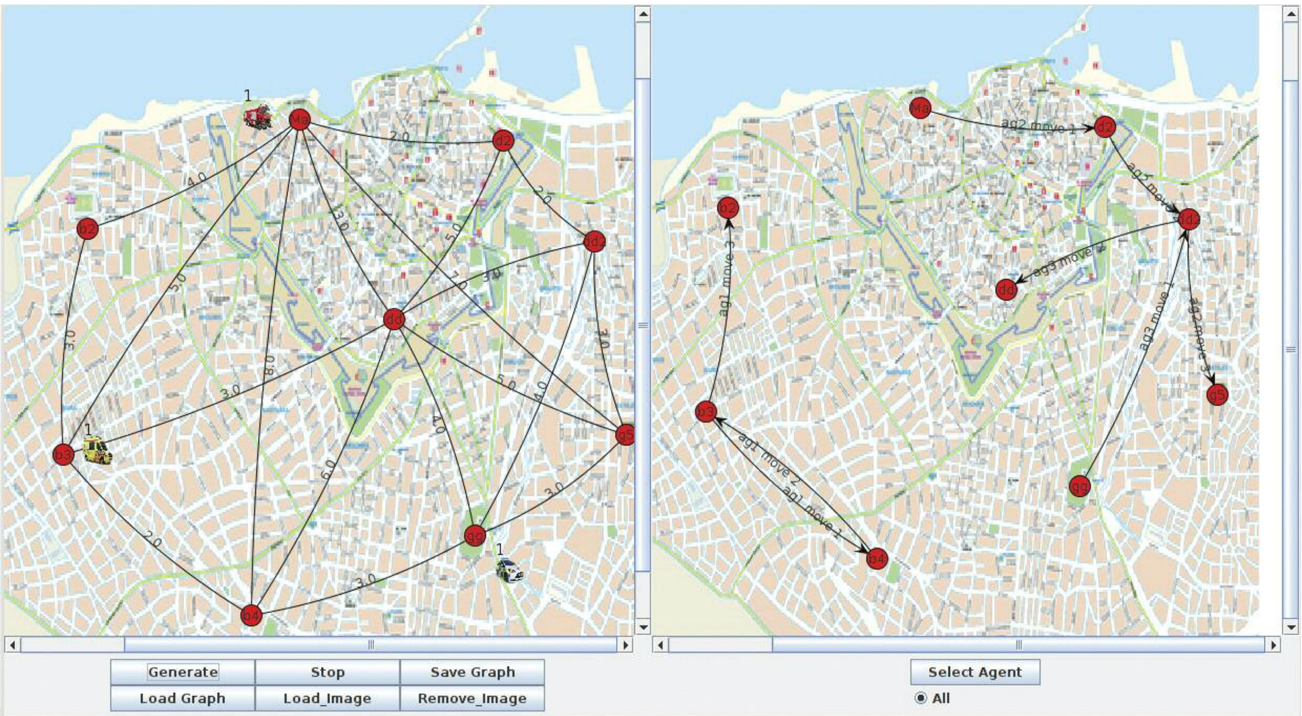

Figure 1: The main window of the SMOC application. In the left sub-window, the current service requests are displayed and in the right, the optimal agent allocation required to fulfil them.

ating, in order to minimize the time of service. The development of the application begun in December 2013 and in its current version it can support the coordination of small scale emergency scenarios. The plan for the next stage of the development includes the extension of the platform so that it can cope with larger scale emergency events, such as those that occur in mid-sized cities. The long-term objectives are to address the daily administrative requirements associated with all servicing organizations operating in a medium to large city.

The efficient operation of SCOC relies on the optimal coordination, cooperation and communication between a large number of different agents, each of which has certain capacities and, often, oriented towards separate goals. For this reason, we opted for a synthesis of formal methods adopted from the field of Artificial Intelligence (AI) and Knowledge Representation and Reasoning. Therefore, this work can be regarded as an amalgam of features stemming from solutions of problems like the multiple Traveling Salesperson Problem, graph traversal, planning with temporal and causal constraints and temporal reasoning. The efficient combination of these techniques, however, is far from trivial. The interdependencies between the agents, the partial observability of the environment, the non-deterministic nature of certain actions, the unpredictability of many factors, the appropriate management of vast amounts of information data and other issues related to imperfect communication are some of the challenging complications that have to be considered.

On the technical side, our implementation utilizes two eminent AI paradigms: Action Languages (ALs) and Answer Set Programming (ASP). The former is a family of logical formalisms that is suitable for representing actions and reasoning about their effects. Due to their efficiency in dealing with properties of dynamically-chaning worlds, ALs hold a pivotal role in providing solutions, like model checking and planning, for a multitude of dynamic domains. In our case, they formally represent the dynamic aspects of problem and allow for the formal validation and verification of the solutions generated by the system.

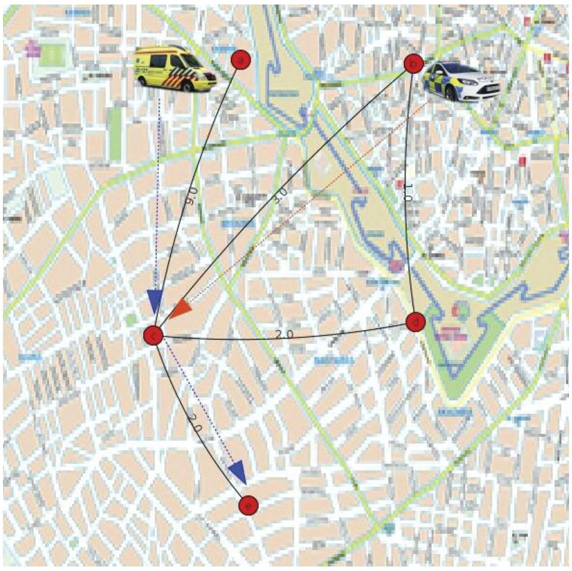

Figure 2: An example of a planned coordination generated by the system. In order for the ambulance to reach its destination in the least possible time (node e), a police car is dispatched in an intermediate location (node c).

ASP, on the other hand, is a logic-based paradigm that is well suited for performing commonsense reasoning, for modelling an agent's belief sets and for handling defeasible inferences, preferences and priorities. Some of ASP's main advantages include its efficiency in dealing with incomplete information and the usage of a very expressive, yet, non-complex modelling language. In addition, ASP solvers are rapidly improving, rendering this declarative paradigm to be considered a prominent candidate for addressing large-scale problems. For instance, recently ASP was applied with great success to the highly demanding task of team building in the Gioia-Tauro seaport [1]. However, despite their merits, each of the paradigms has its own drawbacks and one of the biggest challenge is to combine them in a way that harnesses the best of both worlds without assimilating the corresponding pitfalls.

In conclusion, we believe that this work offers practical benefits and will contribute to the more profound comprehension of the theoretical foundations behind related challenging problems. Moreover, it could serve as a guide for the development of other Smart City applications revolving around the emerging field of Multi-Agent Systems.

Link:
http://www.ics.forth.gr/isl/

Reference:
[1] Grasso, Giovanni, et al. "An ASP-based system for team-building in the Gioia-Tauro seaport." Practical Aspects of Declarative Languages. Springer Berlin Heidelberg, 2010. 40-42

Please contact:
Filippos Gouidis, Theodore Patkos and Giorgos Flouris
FORTH-ICS, Greece
Tel:+30-2810-391632
E-mail: gouidis@ics.forth.gr; patkos@ics.forth.gr; fgeo@ics.forth.gr

Building Smarter Cities through ICT-driven Co-Innovation

by Christophe Ponsard, Robert Viseur and Jean-Christophe Deprez

The Smart City concept is actively being discussed in many places across Europe. However, turning this ideal vision into a practical roadmap is challenging because it requires effective coordination and cooperation between the range of organizations involved in the city operations. In this article, we discuss the use of co-innovation techniques in a number of Belgian cities which gather people with complementary skills for collaborative projects, thus enabling a range of potential projects to be considered. In these projects, information and communication technologies (ICT) would play a key role.

Many European cities are currently aiming to become Smart Cities, with more than 1000 having already joined the European Smart Cities Stakeholder Platform. However, many face problems in defining a 'high-level inclusive vision' that is not just driven by a few decision key makers. Using such a top-down process is risky, not only because of the risk that it might address the wrong priorities, but because it doesn't accurately reflect the complex web of interactions between multiple people and organizations that cities are built on. A key factor of success is to actively involve stakeholders and empower them as actor in the process of change. Over the last decade, a number of new concepts have emerged to support this preferred collaborative approach including co-creation, user innovation and open (source) innovation [1]. They involve specific instruments such as Living Labs [2], multidisciplinary incubators, co-working spaces and creative hubs.

It is widely recognized that ICT plays a key role in supporting the development of Smart Cities. Automating and optimizing the management of transportation, energy and water requires the efficient monitoring and processing of data, making use of concepts such as sensor networks, Internet-of-Things (IoT), big data and open data. It is essential that ICT experts are involved in the co-creation process of a Smart City.

Here, we share some lessons learnt from two co-innovation stories that feature a strong ICT component, currently taking place under the eGov Wallonia initiative with which our research center is deeply involved. The co-innovation process is organized at the regional level because multiple cities share similar concerns and many have common public transportation systems, power management systems and water distribution networks.

Story 1: How open public transportation data can be enabled

The Walloon public bus transport system is operated by the public TEC company which has its own website offering a few services (e.g., timetables) but leaves a number of needs unanswered. These include questions around mobile access and multi-modal schedules. This gap led to some independent actors developing their own mobile apps using web-scrapping techniques. The TEC responded aggressively to this move, a common reaction in these situations. The company put pressure on the web application developer, accusing developers of data theft. Subsequently, the company realized that the transport data belonged to the public and learnt about the open data movement. They then joined the eGov hackathon initiative, providing open access to their data at a hackathon specifically dedicated to enhancing mobility organization in the co-working space of Wallonia's capital city, Namur. This move resulted in a "TEC Real Time" project which is currently exploring how best to exploit (anonymous) data from registered users and the bus GPS information. This scenario illustrates how important it is to first establish trust and then explore new possibilities together, many of which are made possibly by ICT.

Story 2: How communities can be mobilized efficiently to improve the accessibility of public infrastructures

This issue is particularly important for people with reduced mobility. Changing public infrastructure is a long-term process, however, it is possible to improve information flows regarding accessibility issues, thus making short-term adaptations possible. Associations all over Europe are currently fighting for this. In Wallonia, CETIC is helping the local association, CAWAB, to use ICT tools to achieve a balance between mobilizing communities (using many non-expert volunteers) and relying on just a few experts (who are unable to cope with massive workloads) [3]. This resulted in the deployment of two community websites. The first, Access-i, is mainly dedicated to informing people about accessibility with a possible return channel. The second, CENA, is a Moodle-based open collaborative platform for accessibility experts, enabling new comers to learn from the established experts.

The stories showcased in this article illustrate the smart use of tools like hackathons, co-working spaces and community platforms can yield benefits which make a concrete contribution towards developing Smart Cities. Our next step is to optimally use the Living Labs that are currently being deployed (or will be deployed are as part of the next FEDER program 2014-2020).

Links:
http://eu-smartcities.eu
http://hackathonegovwallonia.net
http://www.access-i.be
http://cena.accessible-it.org

References:
[1] A. Kambil, G.B. Friesen, A. Sundaram: "Co-creation: A new source of value", Outlook Magazine, 1999
[2] B. Bergvall-Kareborn, M. Host, A. Stahlbrost: "Concept design with a living lab approach", in proc. of 42nd Hawaii International IEEE Conference on System Sciences, 2009
[3] C. Ponsard, V. Snoeck: "Unlocking Physical World Accessibility through ICT: a SWOT Analysis", ICCHP'14, 2014.

Please contact:
Christophe Ponsard
CETIC, Belgium
E-mail: christophe.ponsard@cetic.be

Supporting the Design Process of Networked Control Systems

by Alexander Hanzlik and Erwin Kristen

The increasing complexity of electronic control systems requires new methods for design, analysis and test of such systems. The DTFSim Data Time Flow Simulator developed at the Austrian Institute of Technology is a discrete-event simulation environment for model-based design, analysis and test of networked control systems for automotive applications.

Automotive Networked Control Systems

An automotive control system consists of a set of devices that control one or more functions of the vehicle, such as an electronic brake. Modern automotive control systems are networked control systems (NCS). A typical NCS contains the following components:

- Sensors to acquire information from the physical environment
- Controllers to provide decision and commands
- Actuators to perform the control commands
- Communication network for information interchange between sensors, controllers and actuators

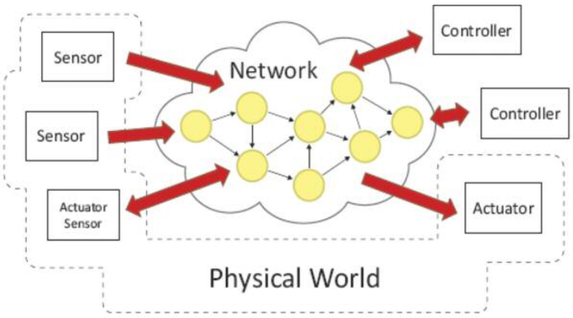

Figure 1: Networked control systems containing sensors, controllers, actuators and a communication network.
Source: http://www.hycon2.eu/

In an NCS, control and feedback signals are exchanged among the different system components in the form of information packages through the communication network.

Principle of Operation

The DTFSim is based on the idea that a NCS can be built by repetitive use of the following parts:

Component part: Components are typically a piece of hardware, such as an ECU, a sensor or an actuator. Each component executes one or more functions. Components are spatially separated and connected via direct links (wire) or a communication network.

Network part: The network part is responsible for communication between components that are not directly linked. This part comprises the network architecture and the network protocol, such as an Ethernet bus.

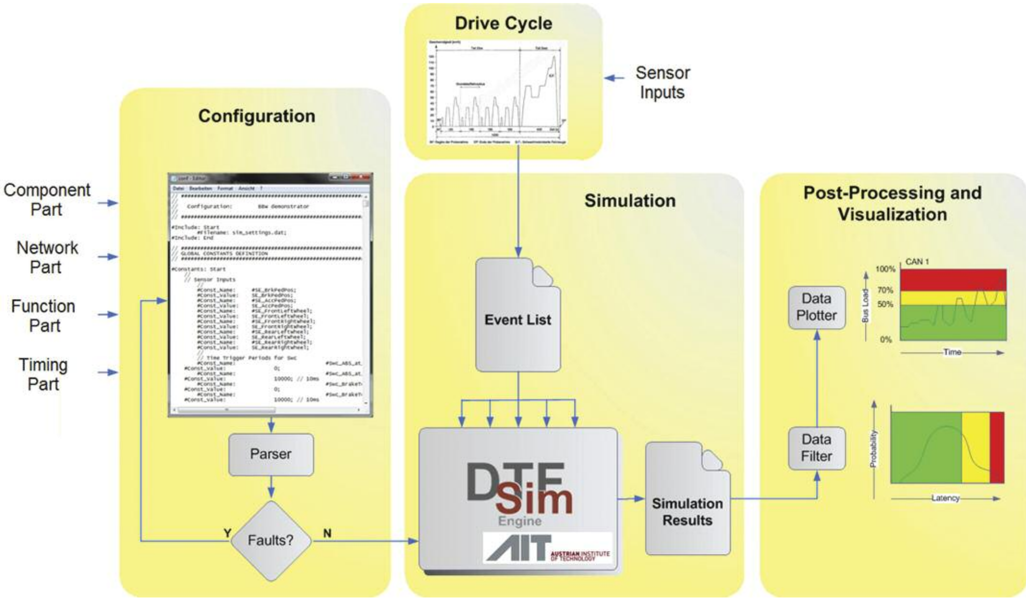

Figure 2: The DTFSim Data Time Flow Simulator workflow comprising the steps configuration, drive cycle, simulation, and postprocessing and visualization.

Function part: The function part contains the implementation of the component functions, for instance, a brake torque calculation algorithm.

Timing part: The timing part contains the timing properties of the system, such as the worst-case execution times of component functions.

The system architecture, consisting of the component part and the network part, is built using elements from a modular assembly system provided by the DTFSim. These elements are grouped together to from more complex structures, such as an ECU or an Ethernet bus. Typical elements are sensors, processors and actuators. Component functions (function part) are added to the different elements to be able to perform specific control tasks. Finally, each element has a propagation delay (timing part).

DTFSim models consist of “event chains”, where an event chain is a directed path from a sensor to an actuator, with an arbitrary number of elements in between. The simulation aims to determine essential system characteristics, such as the bus load (the load of the communication network) and the control signal latencies (the propagation times of control signals from sensors to actuators). This is achieved by stimulating the system at the model inputs (sensors) and by observing the event and data flow along event chains over time to the actuators.

DTFSim Workflow

A typical workflow comprises the following steps:

- Configuration: The first step is the generation of the system configuration, which consists of the system architecture (component and network part) and the system behaviour (function and timing part).
- Drive Cycle: A drive cycle consists of a set of input events for each sensor of the system, e.g., the different positions of the vehicle brake pedal during simulation. The sum of all input events is an event list used for stimulation of the system.
- Simulation: Based on the configuration and the drive cycle, the simulation is executed until all events have been processed. Each event that occurred during simulation is stored in the simulation results log file.
- Post-Processing and Visualization: The last step is the post-processing of the simulation results. In this step, data of interest is extracted from the simulation results log file. This data is then used to produce visualizations of system characteristics of interest, such as the bus load over simulation time or the propagation time distributions of dedicated control signals.

Current Applications and Future Work

The DTFSim has been developed in the course of the EU ARTEMIS project POLLUX, which was related to the design of electronic control systems for the next generation of electric vehicles. Currently, the DTFSim is deployed in two EU ARTEMIS projects: In MBAT (Combined Model-based Analysis and Testing of Embedded Systems), the DTFSim is deployed in an automotive used case where the performance of an electronic braking system is analyzed by means of a simulation model. In CRYSTAL (Critical System Engineering Acceleration), the DTFSim is also deployed in an automotive use case where it is used for timing analysis of an automatic speed limitation application.
Future work relates to the integration of the DTFSim into an automated verification and validation process for embedded control systems.

Links:
http://www.mbat-artemis.eu/home/
http://www.crystal-artemis.eu/

Reference:
[1] A. Hanzlik, E. Kristen: “A Methodology for Design, Validation and Performance Analysis of Vehicle Electronic Control Systems”, in proc. of AMAA 2013, Berlin

Please contact:
Alexander Hanzlik, Erwin Kristen, AIT / AARIT, Austria
E-mail: alexander.hanzlik.fl@ait.ac.at,
erwin.kristen.fl@ait.ac.at

How to Detect Suspect Behaviour at Sea

by Anders Holst

A new model to detect and visualize ships behaving in a strange or suspicious manner can help authorities to prevent or mitigate accidents and detect illegal activities at sea.

The three-year SADV project, which finished in 2013, has been a collaboration between SICS Swedish ICT, Saab AB, the Swedish Coast Guard, the Swedish Customs Service, the Swedish Armed Forces and the Swedish Space Corporation. The aim of the project has been to support the maritime surveillance operators by providing them with an increased situation awareness. This is done using automatic tools which detect vessels bahaving suspiciously or uncharacteristically. Increasing volumes of maritime traffic, combined with threats from criminal activities and environmental issues, just to name a few, there is a high demand for tools that can support the real-time analysis of the huge amounts of maritime situation data that are produced.

Three approaches for anomaly detection

There are essentially three different approaches to detect anomalies and the framework developed in this study supports all three:

1. statistical anomaly detection, where a statistical model is built on normal situations and situations that are very unlikely to come from that model are considered anomalies;
2. rule based anomaly detection, where rules are designed to detect situations of interest; and
3. model based (or simulator based) anomaly detection, where real observations are compared to simulated results (i.e., what would have happened in a normal situation). A difference indicates an anomaly.

While it is easily seen that approaches two and three require extensive expert involvement to design the rules and simulator, approach one would appear to require no expert involvement. However, this approach still requires extensive domain knowledge, first to determine what kind of anomalies would be interesting to detect and secondly, to ensure that the model considers the features that are relevant for detecting those anomalies. Typically, the data requires some processing as these features are not directly represented in the raw data.

Uncharacteristic behaviours considered

A number of uncharacteristic vessel behaviours have been considered in this study. They are:

- unusual speeds and directions (e.g., going the wrong direction in a route),
- unusual movement patterns (e.g., too many stops or turns at-sea),
- making a rendezvous at sea.
- unusual route choices (as compared with other ships), and
- data inconsistencies (e.g., a ship changing identity or using the same identity as another ship).

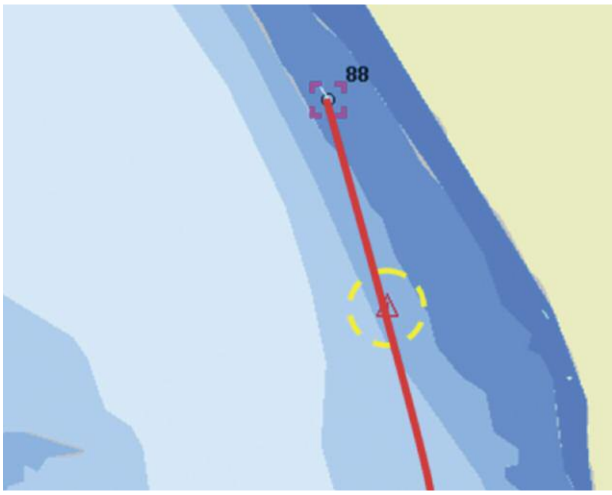

Figure 1: The risk of grounding was detected (yellow ring) three minutes before it actually occurred (purple square).

Figure 2: A small fishing boat (red) repeatedly stopped in the middle of the shipping route. It was almost hit by a large freight ship (purple).

Each of these behaviours requires different anomaly detectors which consider different features calculated from different parts of the data, often on different time scales. Some are suitably handled with statistical methods, others with rule based approaches and yet others with a combination. Most statistical anomaly detection approaches used in this study were based on the Incremental Stream Clustering framework [1] developed by SICS, and most of the rule based detection is based in the Situation Detector platform [2] developed by Saab.

Results

The methods developed in this study were evaluated on real maritime data collected from the Baltic sea. The success of this approach in detecting interesting situations is illustrated by the following examples:

- A risk of running ashore - Saab has combined a rule based detection of ships approaching shallow waters with a statistical anomaly detector for ships that deviate from the usual routes in the area. The detector detected actual groundings up to three minutes prior to such an event occurring (see Figure 1).
- Dangerous behaviour - the SICS's movement pattern behaviour analysis detected a small fishing boat that was

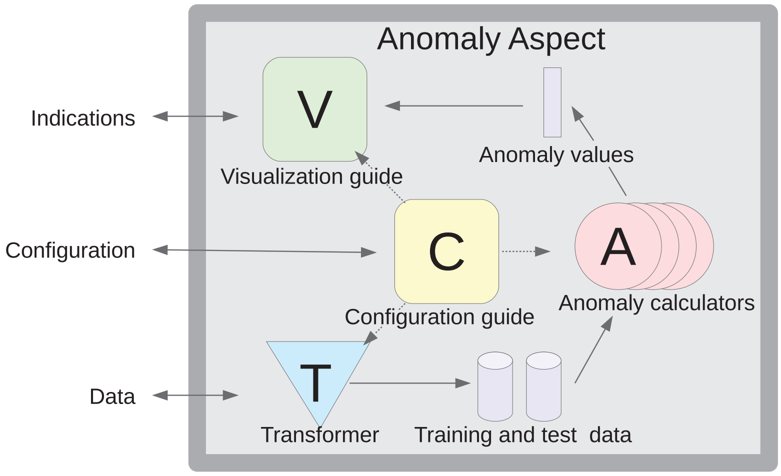

Figure 3: The architecture of the anomaly detection module

repeatedly stopping in the middle of a shipping route and moving irregularly in the wrong direction (apparently fishing). It is almost hit by a large freight ship that managed to veer at the last minute (Figure 2).

- Several ships using the same identity - the SICS's inconsistency detector identified several ships using the wrong identity number. These are mostly genuine mistakes; however, there are examples of ships in other parts of the world intentionally changing their identity numbers to circumvent trading embargoes.
- Meetings between ships at sea - most meetings, unless they are between certain ship types (i.e., coast guard, rescue vessels, etc.), may potentially indicate illegal activities such as smuggling or possibly trawling in prohibited areas.

How the model works

The task of combining all three different types of anomaly detection in one framework was not trivial, but we have managed to design a general architecture that fits with several existing maritime surveillance platforms and is also able to handle several anomaly detectors of different kinds. The architecture of the anomaly detection module is schematically shown in Figure 3.

The key to this architecture is identifying information flows that are common across all the surveillance platforms and for all kinds of anomaly detectors. All detectors need situation data from the platform, and they need to provide anomaly indications back, based on that data. They also need to support user configuration and inspection.

The next step is to identify similarities between the different kinds of anomaly detectors. They all need to extract the relevant features from the provided data, use their respective methods to determine whether those features represent an anomaly, and then finally transform that judgement into an "indication" which is complete with details of what, when and why.

If the indications from one anomaly detector can be used as an input feature to another anomaly detector, it is possible to combine the rule based and statistical detectors in a more advanced ways, as opposed to just having them run in parallel in the same system [3]; statistical detections can be referred to in a rule, or the number of times that a rule matches can be considered in a statistical detector.

Conclusion

The SADV project has significantly extended previous methods used for anomaly detection within the maritime surveillance sphere and shown that these methods can detect interesting and relevant situations.

Link:
http://www.sics.se/projects/sadv-statistical-anomaly-detection-and-visualization-for-maritime-domain-awareness

References:
[1] A. Holst and J. Ekman: "Incremental stream clustering for anomaly detection and classification", in Eleventh Scandinavian Conference on Artificial Intelligence, SCAI 2011
[2] J. Edlund, M. Grönkvist, A. Lingvall and E. Sviestins: "Rule-based situation assessment for sea surveillance", in proc. of SPIE Defence and Security Symposium 2006
[3] A. Holst et al.: "A Joint Statistical and Symbolic Anomaly Detection System: Increasing performance in maritime surveillance", in FUSION 2012.

Please contact:
Anders Holst
SICS Swedish ICT
Tel: +46 (0)8 633 1593
E-mail: aho@sics.se

A Breakthrough for Balanced Graph Partitioning

by Fatemeh Rahimian

Researchers at SICS have invented a new way to partition graphs into a given number of clusters, such that the related information belongs to the same cluster and the number of connections between clusters are minimized. Graph partitioning, a well known NP-Complete problem, has been already thoroughly investigated for small and medium sized graphs, but this distributed algorithm works on very large graphs.

Our world is increasingly connected, and this is changing our lives in ways that we still do not fully understand. New tools and technologies have given us the unprecedented potential to make sense of the huge inter-connected datasets that exist in various fields of science, from social networks to biological networks. Relations can be mathematically described as “graphs”; for example, a network of friends on Facebook can be described by a graph, which shows a user’s closest friends and who they are related to. The increasing size of datasets means that graph sizes are also increasing: this means that they must be partitioned into smaller clusters so they can be more easily managed on multiple machines in a distributed fashion.

The new solution for this is a distributed heuristic based algorithm that can efficiently partition big graphs into a given number of clusters of equal size or any given size. The key point in this approach is that no global knowledge is required, and simultaneous access to the entire graph is not assumed.

The algorithm used to solve the problem is very intuitive. It starts by randomly assigning vertices to partitions. Over the course of the algorithm, vertices from different partitions iteratively negotiate with each other and exchange their partition assignments, if that exchange results in a better partitioning [1]. When this iterative optimization converges, neighboring vertices of the graph will mostly end up in the same partition, while there will be very few edges that cross different partitions (See Figure 1). The algorithm is inherently a local search optimization enhanced using a simulated annealing technique. A variant of this work has been specifically designed for graphs with power-law degree distribution, where vertex-cut partitioning has proved more efficient than edge-cut partitioning [2]. Both solutions are kept as simple as possible, so that the algorithm can be easily adopted by various real world applications.

With the recent advances in graph database technologies, it is anticipated that leading companies in this field will soon move to make use of graph partitioning for scaling up. Successful partitioning is an important ingredient in efficient big data analysis and that is why this solution has attracted a lot of attention in the community.

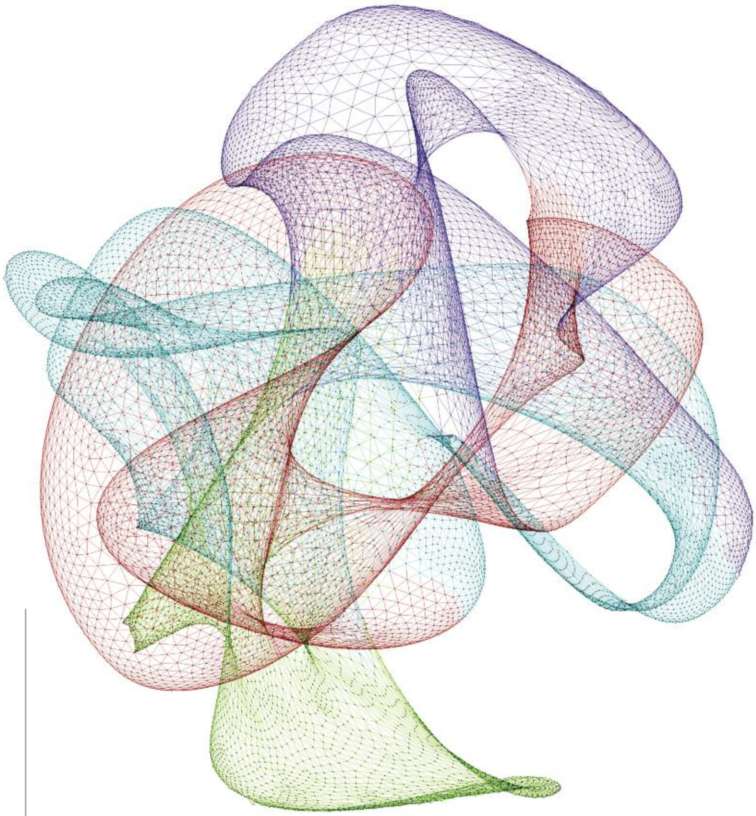

Figure 1: 4elt sparse graph, partitioned into 4 balanced components that are indicated by colors. Very few links are on the border of regions with different colors. For her achievement [1] Fatemeh Rahimian received the Best Paper Award at the 7th IEEE International Conference on Self-Adaptive and Self-Organizing Systems in Philadelphia in September 2013. Moreover, in May 2014 she obtained a doctoral degree that was partly based on this work.

Link:
http://e2e-clouds.org/

References:

[1] Rahimian, Fatemeh, Amir H. Payberah, Sarunas Girdzijauskas, Mark Jelasity, and Seif Haridi. "Ja-be-ja: A distributed algorithm for balanced graph partitioning." In the 7th IEEE Conference on Self-Adaptive and Self-Organizing Systems (SASO), 2013.

[2] Rahimian, Fatemeh, Amir H. Payberah, Sarunas Girdzijauskas, and Seif Haridi. "Distributed vertex-cut partitioning." In the 14th IFIP International conference on Distributed Applications and Interoperable Systems (DAIS), 2014.

Please contact:

Fatemeh Rahimian
SICS Swedish ICT
E-mail: fatemeh@sics.se

Call for Participation

FMICS'14 - 19th International Workshop on Formal Methods for Industrial Critical Systems

Florence, Italy ,11-12 September 2014

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 (but are not limited to):
- 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 (e.g., 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.

FMICS'14 is co-located with QEST 2014, SAFECOMP 2014, EPEW 2014, and FORMATS 2014 in the frame of FLORENCE 2014 - a scientific week on computer safety, reliability, performance, and quantitative analysis.

More information:
http://fmics2014.unifi.it/
http://www.florence2014.org/

Call for Participation

7th International Conference of the ERCIM Working Group on Computational and Methodological Statistics

Pisa, Italy, 6-8 December 2014.

The 7th International Conference on Computational and Methodological Statistics (ERCIM 2014) is organized by the ERCIM Working Group on Computational and Methodological Statistics (CMStatistics) and the University of Pisa. The conference is held jointly with the 8th International Conference on Computational and Financial Econometrics (CFE 2014). The conference has a high reputation of quality presentations. The last editions of the joint conference CFE-ERCIM gathered over 1200 participants.

Aims and Scope

All topics within the Aims and Scope of the ERCIM Working Group CMStatistics will be considered for oral and poster presentation.

Topics include, but are not limited to: robust methods, statistical algorithms and software, high-dimensional data analysis, statistics for imprecise data, extreme value modeling, quantile regression and semiparametric methods, model validation, functional data analysis, Bayesian methods, optimization heuristics in estimation and modelling, computational econometrics, quantitative finance, statistical signal extraction and filtering, small area estimation, latent variable and structural equation models, mixture models, matrix computations in statistics, time series modeling and computation, optimal design algorithms and computational statistics for clinical research.

Sessions

The joint conference CFE-ERCIM will have five plenary zessions. Moreover, ERCIM 2014 will have four "special invited sessions" and a significant number of "organized invited sessions" on key topics, and "contributed sessions" that will run in parallel during the three days of the conference.

Tutorials

Tutorials will be given on Friday 5 December 2014. The number of participants to the tutorials is limited and restricted only to those who attend the conference.

Publications

The journal Computational Statistics & Data Analysis will publish selected papers in special peer-reviewed or regular issues.

More information:
http://www.cmstatistics.org/ERCIM2014/

Call for Participation

SAFECOIM 2014 - 33rd International Conference on Computer Safety, Reliability and Security

Florence 10-12 September 2014

SAFECOMP is an annual event covering the state-of-the-art, experience and new trends in the areas of safety, security and reliability of critical computer applications. SAFECOMP provides ample opportunity to exchange insights and experience on emerging methods, approaches and practical solutions. It is a one-stream conference without parallel sessions, allowing easy networking.

SAFECOMP 2014 will be part of FLORENCE 2014, a one-week scientific event with conferences and workshops in the areas of formal and quantitative analysis of systems, performance engineering, computer safety, and industrial critical applications.

More information:
http://www.safecomp.org
http://www.ewics.org
http://www.florence2014.org

Call for Papers

International Workshop on Computational Intelligence for Multimedia Understanding

Paris, 1-2 November 2014

Multimedia understanding is an important part of many intelligent applications in our social life, be it in our households, or in commercial, industrial, service, and scientific environments.

Analyzing raw data to provide them with semantics is essential to exploit their full potential and help us managing our everyday tasks. Our purpose is to provide an international forum to present and discuss current trends and future directions in computational intelligence for multimedia understanding. The workshop, organized by the ERCIM MUSCLE Working Group, also aims at fostering the creation of a permanent network of scientists and practitioners for an easy and immediate access to people, data and ideas.

Scope and Topics

Topics include, but are not limited to:

- multimedia labeling, semantic annotation
- multimodal and cross-modal data analysis, clustering
- big data, distributed digital libraries
- activity and object detection and recognition
- text and speech recognition
- multimodal indexing and searching in very large data-bases
- visual data processing techniques
- medical image segmentation and model-based representation
- microscopic image feature extraction and classification

Important dates

- Paper Submission: 25 July, 2014
- Notification: 12 September, 2014
- Final Regular Paper Submission: 26 September, 2014

More information:
http://iwcim.isep.fr

Call for Participation

SERENE'14 - 6th International Workshop on Software Engineering for Resilient Systems

Budapest, 15-16 October 2014

Unprecedented level of complexity of modern software and software-based systems makes it difficult to ensure their resilience - an ability of the system to persistently deliver its services in a dependable way even when facing changes, unforeseen failures and intrusions. Yet we are observing the increasingly pervasive use of software in evolvable and critical systems like transportation, health care, manufacturing, and IT infrastructures. This trend urges the research community to develop powerful methods for assuring resilience of software-intensive systems.

These challenges have also appeared in the scope of the current Horizon 2020 calls that aim at developing tools and methods for incorporating resilience into evolving software systems; and also in calls related to specific application areas like advanced cloud infrastructures and services, smart objects, and robotics.

The SERENE 2014 workshop provides a forum for researchers and practitioners to exchange ideas on advances in all areas relevant to software engineering for resilient systems, including, but not limited to:

- design of resilient systems
- requirements engineering & re-engineering for resilience
- frameworks, patterns and software architectures for resilience
- Engineering of self-healing autonomic systems;
- design of trustworthy and intrusion-safe systems
- resilience at run-time (mechanisms, reasoning and adaptation)
- verification, validation and evaluation of resilience
- modelling and model based analysis of resilience properties;
- formal and semi-formal techniques for verification and validation
- experimental evaluations of resilient systems
- quantitative approaches to ensuring resilience
- resilience prediction
- case studies & applications
- empirical studies in the domain of resilient systems
- cloud computing and resilient service provisioning
- resilient cyber-physical systems and infrastructures
- global aspects of resilience engineering: education, training and cooperation.

The workshop is organised by the ERCIM Working Group SERENE.

The proceedings of SERENE 2014 will be published as a volume in Springer Lecture Notes in Computer Science (LNCS).

Autumn School

An autumn school will be held right before the SERENE workshop. The autumn school will explore the resiliency of cyber physical systems.

More information:
http://serene.disim.univaq.it/

International Innovation Award for Martin Kersten

Martin Kersten, Research Fellow at CWI, is awarded the 2014 SIGMOD Edgar F. Codd Innovations Award, the most prestigious prize for researchers who have made innovative and significant contributions to database systems and databases. Kersten received the award on 26 June at the annual ACM SIGMOD/PODS Conference. The Awards Committee recognized the influential contributions of Kersten to the big data problem, his outstanding achievements in scientific research in advanced database architectures and his pioneering work in the realization of MonetDB.

According to Kersten, massive accumulation of data requires fundamental research and changes in database management systems. In MonetDB he has pioneered the columns-store technology since 1993. Nowadays it is the most widespread open-source column-store database management system, which is worldwide used in more than 130 countries with over 300,000 downloads. Since 2011, column-store technology as pioneered in MonetDB has found its way into the product offerings of all major commercial database vendors in relational database systems.

More information:
http://www.cwi.nl/news/2014/international-innovation-award-big-data-research-martin-kersten

European Project on Precision Farming

Increasingly, remote sensing data is made publicly available by organizations like ESA and NASA. Based on these data, and in combination with other sources, new valuable applications can be created. In the European research project Linked Open Earth Observation Data for Precision Farming (LEO) researchers from CWI and the University of Athens join their forces with industry partners to develop an application for precision farming. The new application is based on remote sensing and geospatial data. With precision farming - advanced agriculture using GPS, satellite observations and tractors with on-board computers - the farming process is performed as accurately and efficiently as possible. This is achieved by combining data from earth observations with other geospatial sources such as cadastral data, data on soil quality, vegetation and protected areas. This enables farmers to find the optimal trade-off in maximizing yield while minimizing fertilizers and pesticides. At CWI, the researchers will combine various data sources, transfer them into RDF format and publish them as Linked Open Data.

More information: http://www.linkedeodata.eu

Roberto Scopigno receives the Eurographics "Distinguished Career Award" 2014

Roberto Scopigno is the recipient of the Eurographics "Distinguished Career Award" 2014, an award given every other year to a professional in computer graphics who has made outstanding technical contributions to the field and has shaped computer graphics in Europe.

Roberto Scopigno's work has had a profound impact on the field of visual computing. Of particular importance has been his work on surface simplification, LOD and multiresolution representations for surfaces and volumes. His papers on these topics have been cited widely and stimulated considerable follow-up work. The BDAM algorithm and its 3D extension, Tetra-Puzzles, grew over ingenious novel data structures for hierarchical seamless space subdivision and have inspired new multiresolution applications for the inspection of terrain models and gigantic meshes. He is a world-wide leader in the development of novel algorithms and techniques for Cultural Heritage and for the acquisition, preservation and visualization of digital copies of physical artifacts inherited from the past.

Roberto Scopigno is a recognized scientific leader. He has created a very successful research group at CNR Pisa, and several of his former students have themselves become well-known productive researchers. He has demonstrated a strong leadership in Computer Graphics research in Europe, and he has been a major actor in shaping Computer Graphics in Italy. Roberto has significantly advanced the field through his work and energy, and by setting a personal example.

More information:
The complete award citation is available at:
https://www.eg.org/index.php/component/content/article/57-awards/distinguished-career-award-recipients/340-scopigno-2014

https://www.eg.org/index.php/awards/career-award

ERCIM
European Research Consortium for Informatics and Mathematics

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

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

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

Portuguese ERCIM Grouping
c/o INESC Porto, Campus da FEUP,
Rua Dr. Roberto Frias, nº 378,
4200-465 Porto, Portugal

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

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

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

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/

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

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

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

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

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/

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

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

University of Geneva
Centre Universitaire d'Informatique
Battelle Bat. A, 7 rte de Drize, CH-1227 Carouge
http://cui.unige.ch

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

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

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

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

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

Universty of Wroclaw
Institute of Computer Science
Joliot-Curie 15, 50–383 Wroclaw, Poland
http://www.ii.uni.wroc.pl/

I.S.I. - Industrial Systems Institute
Patras Science Park building
Platani, Patras, Greece, GR-26504
http://www.isi.gr/

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

Subscribe to ERCIM News and order back copies at http://ercim-news.ercim.eu/