Number 94, July 2013

ERCIM NEWS

www.ercim.eu

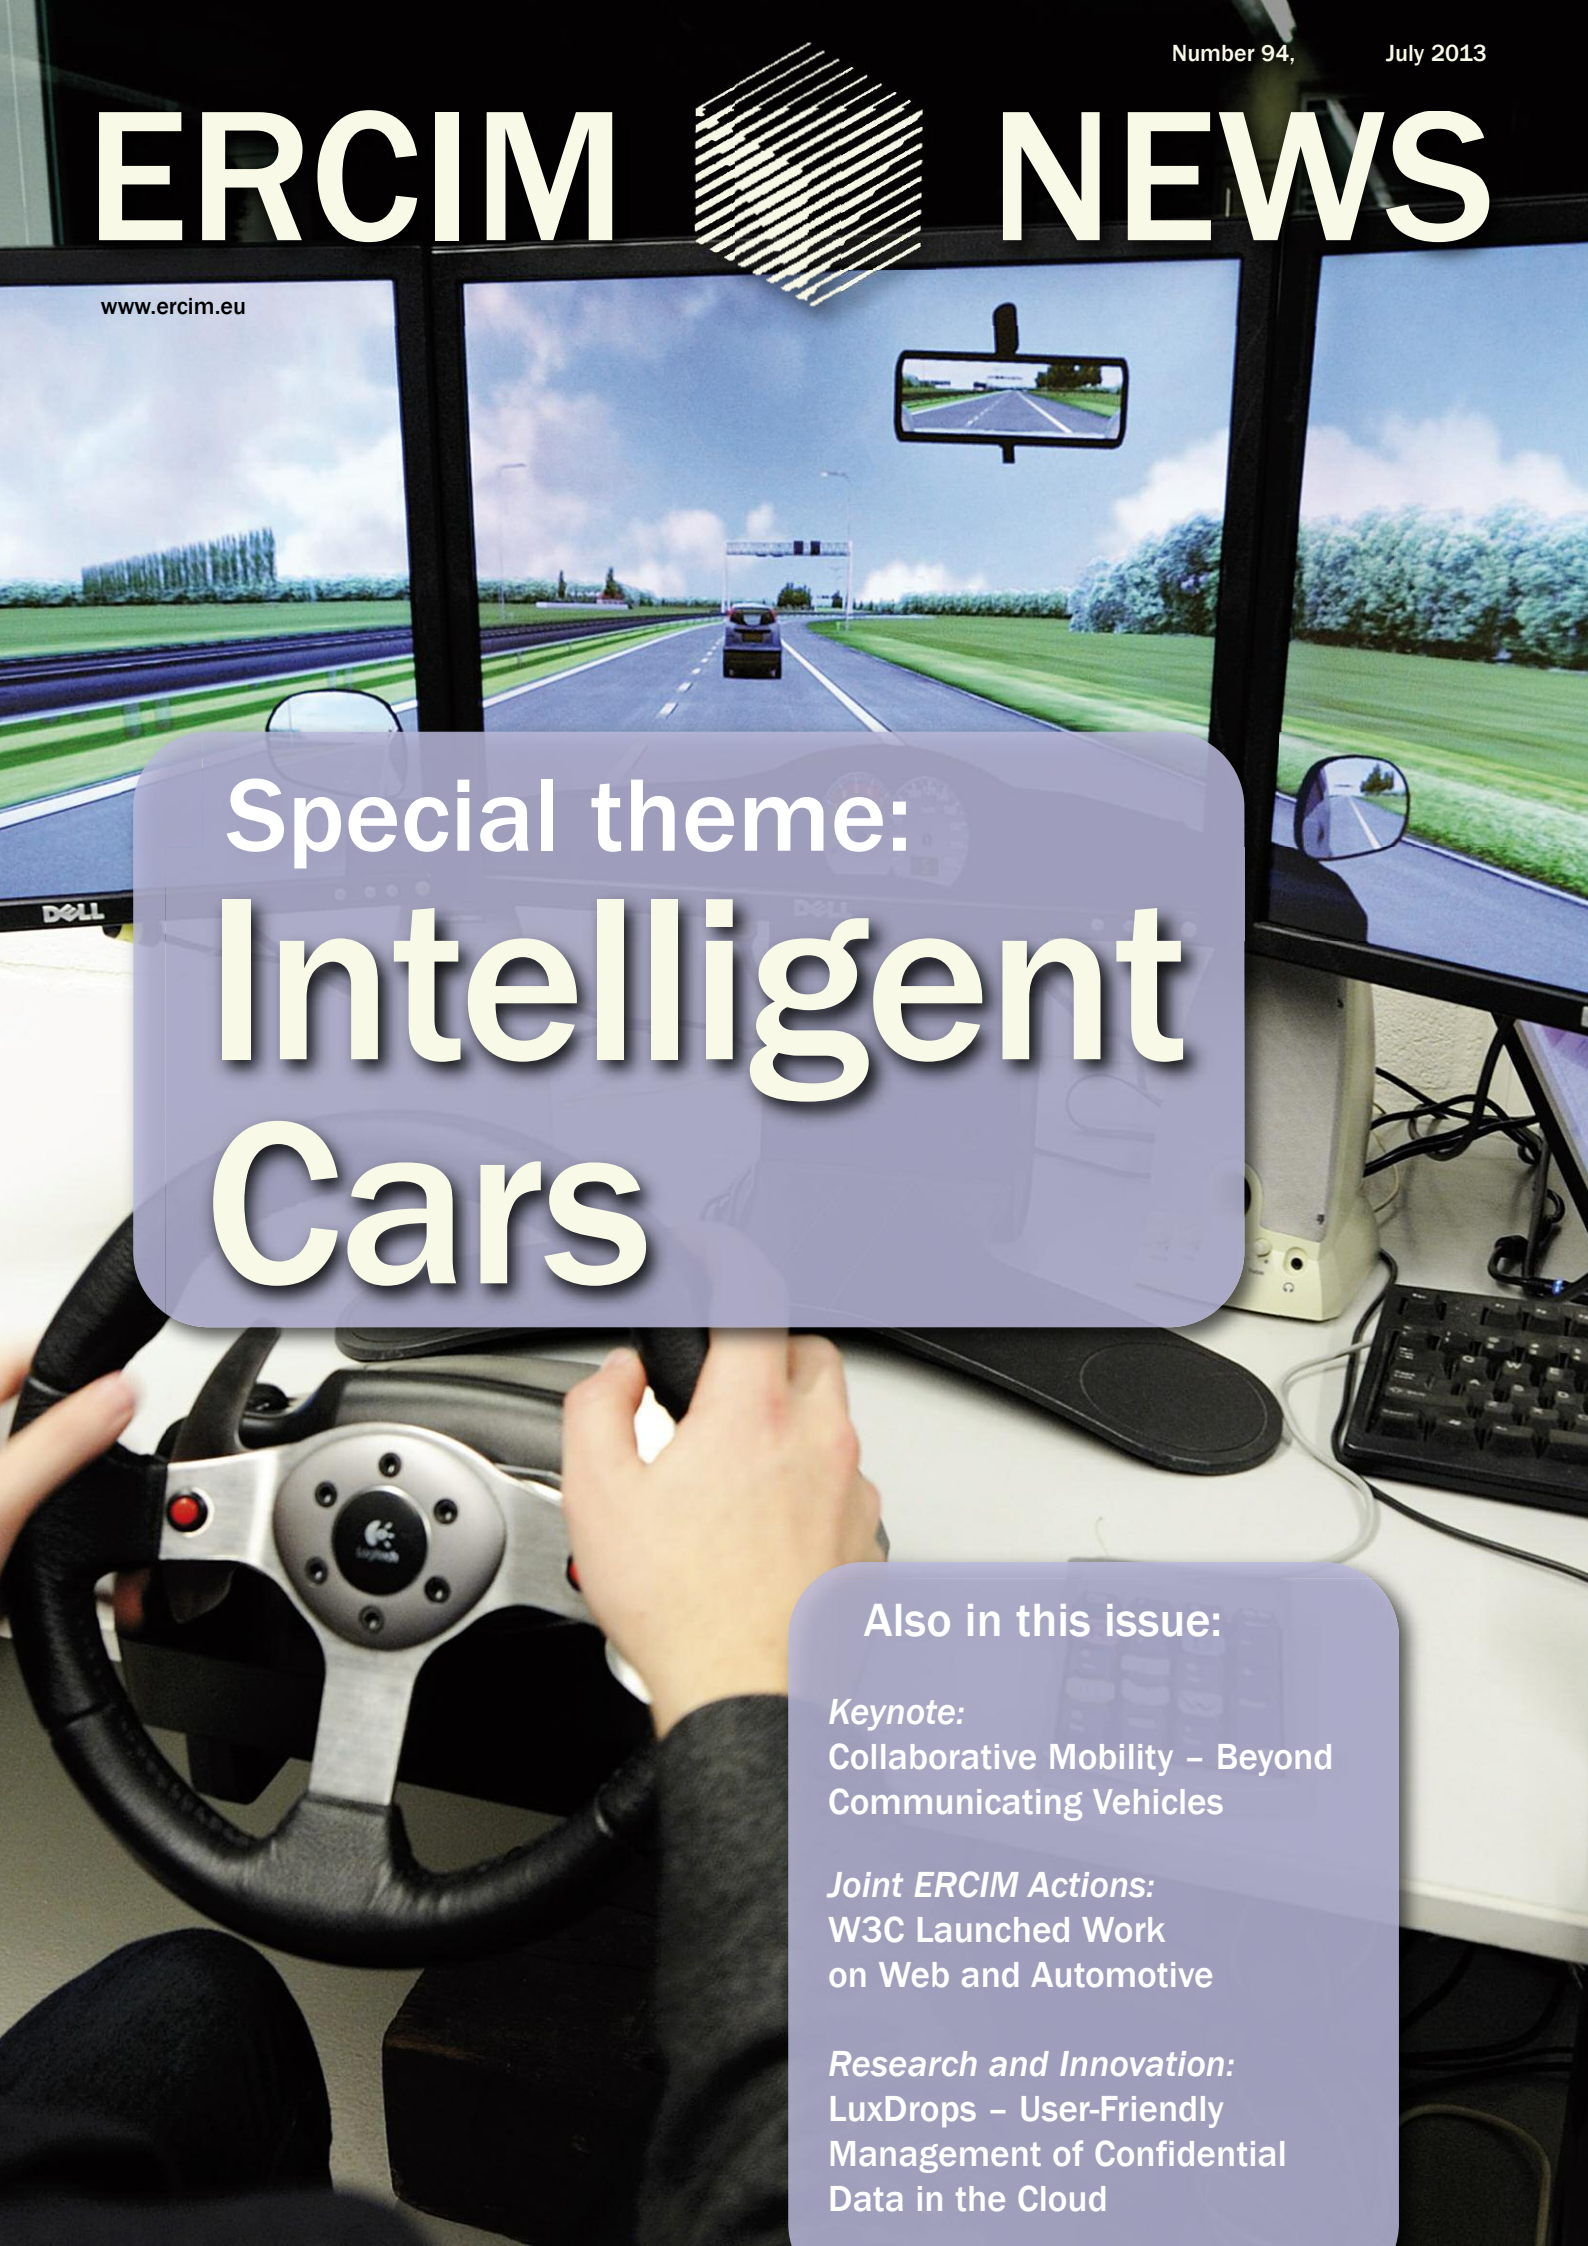

Special theme:

Intelligent Cars

Also in this issue:

Keynote:
Collaborative Mobility – Beyond Communicating Vehicles

Joint ERCIM Actions:
W3C Launched Work on Web and Automotive

Research and Innovation:
LuxDrops – User-Friendly Management of Confidential Data in the Cloud

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of about 8,500 copies.

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

KEYNOTE

4 Collaborative Mobility – Beyond Communicating Vehicles
by Ilja Radusch

JOINT ERCIM ACTIONS

5 W3C Launched Work on Web and Automotive
by Philipp Hoschka

SPECIAL THEME

The special theme section "Intelligent Cars" has been coordinated by Jaroslav Machan, ŠKODA AUTO a.s., Czech Republic and Christian Laugier, Inria, France

Introduction to the Special Theme

6 Intelligent Vehicles as an Integral Part of Intelligent Transport Systems
by Jaroslav Machan and Christian Laugier

8 ABV – A Low Speed Automation Project to Study the Technical Feasibility of Fully Automated Driving
by Evangeline Pollard, Fawzi Nashashibi and Paulo Resende

9 Cooperative Systems for Car Safety Improvement
by Zdeněk Lokaj, Tomas Zelinka and Miroslav Šrotýř

11 Location Assurance and Privacy in GNSS Navigation
by Xihui Chen, Carlo Harpes, Gabriele Lenzini, Sjouke Mauw and Jun Pang

12 Paving the Way for Apps in Vehicles
by Avenir Kobetski and Jakob Axelsson

13 Mobile Opportunistic Traffic Offloading for Map-Based ADAS Applications: the MOTO Project Approach
by Raffaele Bruno, Andrea Passarella, Leandro D'Orazio and Filippo Visintainer

14 Autonomous Traffic Warning System with Car-to-X Communication
by Josef Jiru

16 Understanding V2X Communication Dynamics Through Complex Network Science
by Nicholas Loulloudes, George Pallis and Marios D. Dikaiakos

17 The VESPA Project: Driving Advances in Data Management for Vehicular Networks
by Thierry Delot and Sergio Ilarri

19 Intelligent Vehicles: Complex Software-Based Systems
by Javier Ibanez-Guzman and Christian Laugier

20 **Modularity Analysis of Automotive Control Software**
by Yanja Dajsuren, Mark G.J. van den Brand and Alexander Serebrenik

22 **Electro Vehicle Car-Sharing System – A Simulation Model for Big Cities**
by Dmitry Rozhdestvenskiy and Petr Bouchner

23 **The Car that Looks Like Me: Similarity Cues can Increase Trust in the Self-Driving Cars of the Future**
by Frank Verberne, Jaap Ham and Cees Midden

25 **Digital Road Authority for Coordination between In-Car Navigation Systems and Traffic Control Centres**
by Rob van der Mei, Maaike Bots and Frank Ottenhof

RESEARCH AND INNOVATION

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

26 **Imperia: The 21st Century Eco-Friendly Sports Car**
by François Henry

28 **An Innovative Active Suspension System for Autonomous Vehicles: A Safe and Comfortable Ride and Good Handling**
by Katerina Hyniova

29 **Wire-Speed Regular-Expression Scanning at 20 Gbit/s and Beyond**
by Jan van Lunteren and Christoph Hagleitner

31 **DRMSim: A Routing-Model Simulator for Large-Scale Networks**
by Aurélien Lancin and Dimitri Papadimitriou

32 **LuxDrops – User-Friendly Management of Confidential Data in the Cloud**
by Pascal Bauler, Daniel Heinesch, Antoine Schlechter and Fernand Feltz

33 **InGeoCloudS: A Cloud-Based Platform for Sharing Geodata Across Europe**
by Claudio Lucchese, Raffaele Perego, Salvatore Trani, Makis Atzemoglou, Benoit Baurens and Dimitris Kotzinos

34 **Electromagnetic Fields for Neuron Communications**
by Rié Komuro and Ilangko Balasingham

36 **A Patented Position Determination Solution**
by Per Kreuger

36 **SoFWIReD – Southampton Fraunhofer Web and Internet Research & Development**
by Sepideh Chakaveh

38 **Preserving Linked Data**
by Carlo Meghini, Anna Molino, Francesca Borri and Giulio Galesi

EVENTS, BOOKS, IN BRIEF

Event Reports

39 **MediaEval 2012 Evaluation Campaign**
by Gareth J. F. Jones and Martha Larson

40 **TransForm School on Research Directions in Distributed Computing**
by Panagiota Fatourou

Announcements

40 **AITA 2013 – 12th International Workshop on Advanced Infrared Technology and Applications**

40 **W3C Workshop on Publishing using the Open Web Platform**

41 **Second Plenary Meeting of the Research Data Alliance**

41 **Collaboration meets Interactive Surfaces: Walls, Tables, Tablets and Phones**

41 **5th Joint Virtual Reality Conference (JVRC 2013)**

Books

42 **Formal Methods for Industrial Critical Systems: A Survey of Applications**

42 **Set-theoretic Fault-tolerant Control in Multisensor Systems**

42 **Digital Stereoscopy – Scene to Screen 3D Production Workflow**

Jobs

42 **Research positions at The Warsaw Center of Mathematics and Computer Science**

In Brief

43 **European Laboratory on Big Data Analytics and Social Mining**

43 **Michael Wilson Obituary**

43 **National ICT Awards 2013 for Python and SIG**

43 **Cryptographer Ronald Cramer appointed Fellow of IACR**

Ilja Radusch, head of department "Automotive Services and Communication Technologies" at Fraunhofer FOKUS and head of the Daimler Center for Automotive IT Innovations.

Collaborative Mobility – Beyond Communicating Vehicles

European efforts for more intelligent transport systems have a long history with the first programmes focusing on Road Transport Informatics (RTI) dating back to the 1980's. The corresponding Intelligent Vehicles Safety Systems started even earlier in the late 1950's by progressing from passive safety systems towards truly pro-active safety functions we envision today. The state of the art in road and vehicle safety is today represented by the concept of cooperative driving. Current cooperative driving is based on vehicles communicating with each other – widely regarded as vehicle-2-vehicle communication – and with the infrastructure – accordingly named vehicle-2-infrastructure communication. This development has been driven mainly by safety needs to provide travellers sufficient information early enough to be able to respond to dynamic traffic situations. Additionally, vehicles anonymously announcing their general position and velocity are set to be far more effective for assessing current traffic conditions and improving general traffic efficiency. Thus, vehicle-2-x communication technology (the "x" stands for vehicles and infrastructure) extends the range of vehicle sensors to yet unseen distances allowing vehicles to "see" around corners and warn drivers of upcoming dangers, making sure she or he can act in due time to avoid an accident or at least mitigate its impact.

Paving the road for communication till 2015

Recent years and research projects have been dedicated to test and assess vehicle-2-x communication under real world conditions. Europe's largest field operational trial was successfully concluded in Germany with over 500 drivers travelling 1.6 million kilometres in one of 120 vehicles. By 2015 this technology will be included in new vehicles saving up to 11 billion euros in Germany alone due to accidents avoided and travel times reduced. These estimated savings, however, rely on a rapid adoption of vehicle-2-x communication not only by vehicle manufacturers but also road and infrastructure operators. Therefore, in a first step to full deployment, Dutch, German, and Austrian authorities agreed to establish an intelligent corridor ranging from Rotterdam in the Netherlands via Frankfurt/Main in Germany to Vienna in Austria where construction areas are announced wirelessly to the equipped vehicles. This corridor will be the first installation of cooperative traffic where vehicles and road infrastructure work together.

The road ahead

However, the possibilities of cooperative traffic do not stop there. Just like the Internet, vehicle-2-x communication technology is providing a link between travellers and not just their vehicles. Therefore, the next task is to join those travellers (which include drivers and any type of road users) and the infrastructure operators in a collaborative network to solve various travel needs all the way from eco-friendly parking to short-term decisions on trip planning. Collaboration is the key concept of a future mobility approach, which extends the cooperative concept of the first generation systems and applications – such as travel time optimizing navigation systems – by including the human user in a highly integrated cooperative, interactive, and participatory network. In this collaborative concept, it is not only the systems and vehicles that communicate, but all actors (systems and humans) are engaged in a continuous bi-directional, dynamic exchange of information allowing for pro-active traffic system management which encourages active participation and interaction of road users.

Mobility research so far has been focused on safety systems and applications relieving drivers from the most exhausting tasks. However, further improvements require a behavioural change in all road users and operators in the direction of a collaborative pro-active mobility management that steers the network of users towards an optimum of network level benefits (as opposed to simply maximizing individual benefits). Acceptance of such a collaborative mobility concept – that ties together interaction with participation and will achieve progress beyond simple cooperation – requires continuous coaching of travellers and drivers not only while driving but also within the realm of other co-modality options. The objective of this coaching is to support the traveller pro-actively with adequate hints tailored to his or her current situation.

From static to elastic infrastructures

The overall effect of collaborative mobility will be the necessary balancing of the needs of all road users and road operators alike, thereby extending towards all citizens. Future cities will be built on benefits that will accrue in making the transition from the static concept of mobility arising from the needs of individual road users only, to a community-aware and adaptive concept of mobility using reliable real-time data, capturing the needs and intentions of all travellers. Social awareness will regard actions encouraging road users to follow collaborative strategies that benefit all road users as a group. Elastic infrastructures refer to the ability of the road operator to react to the needs of all road users without having to strain the environment by building new roads.

Ilja Radusch

W3C Launched Work on Web and Automotive

by Philipp Hoschka

Supported by the "Webinos" project, W3C has created a new Automotive and Web Platform Business Group to accelerate the adoption of Web technologies in the automotive industry. The group convenes developers, automotive manufacturers and suppliers, browser vendors, operators and others to discuss how to enhance driving, safety, and passenger entertainment with the Open Web Platform. The group will first focus on defining a Vehicle Data API that will create new opportunities for automotive services via the Web.

W3C first explored the impact of the Open Web Platform on the automotive industry at the November 2012 Web and Automotive Workshop. Participants discussed how location-based services, enhanced safety, entertainment, and integration of social networking will benefit drivers and passengers. In addition, they looked at business drivers for Web technology adoption such as the ability to attract customers with convenient and innovative services, maintain ongoing customer relations, address regulatory requirements, manage mobile payments, and lower development costs.

The more than 40 organizations that participated in that conversation expressed broad consensus that HTML5 is a compelling platform for providing these services and lowering the risk of fragmented solutions. The new Automotive and Web Platform Business Group picks up where the Workshop left off.

The initial plans of the new Business Group are to:

- Create specifications, starting with a Vehicle Data API Specification.
- Create conformance tests to cover new specifications that get defined.
- Provide use cases and other reports to identify additional needed standards work and to drive successful automotive web deployment.

The W3C Automotive and Web Platform Business Group held its first face to face meeting in Barcelona on 22 April this year.

A lot of ideas were discussed during this meeting between BMW, VW, PSA, Visteon, Continental, Intel, KDDI, LG, Magneti Marelli, QNX, Ford, Strategy Analytics, Genivi and W3C with agreement on the steps needed to prepare the next face to face meeting in Tokyo which took place on 29 May.

In particular, a first round of proposals for a Web Vehicle API was discussed, with proposals from Intel, QNX, Genivi (LG) and Webinos. The group decided to first focus on an API that would provide read access to vehicle data. Next steps include:

- Creating an overview of the superset of the proposed Vehicle APIs to look for overlaps and gaps.
- Considering whether the mandatory dataset provided in OBDII could serve as a starting point for a common dataset to be shared by different OEMs via a Vehicle API.
- Creating a document that maps current W3C work onto the requirements of the automotive sector (e.g. geolocation, packaging, etc.).

W3C invites all organizations and individuals interested in bringing the Web and automotive industries together to join the Automotive and Web Platform Business Group.

W3C Business Groups give W3C Members and non-Members a vendor-neutral forum for collaboration on the development of the Web in the near-term. Business Groups do not create W3C standards, but they do develop reports, use cases, requirements, and other forms of input to the W3C standardization process.

ERCIM is the European host of W3C.

Links:
http://www.w3.org/community/autowebplatform
http://www.w3.org/2012/08/web-and-automotive/

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

Introduction to the Speical Theme

Intelligent Vehicles as an Integral Part of Intelligent Transport Systems

by Jaroslav Machan and Christian Laugier

Socio-economic context

The recent global economic growth and the related overall increase in volumes of transportation, along with stronger public demand for mobility, are leading causes of road infrastructure congestion and of a large number of accidents and deaths (about one fatality per every 100 million km of driving and a social cost in France of 23 billion Euros in 2011). These factors have lead to a heavy rise in both energy consumption and also environmental and social problems.

Reaction to this situation cannot be limited to the classic countermeasures, which typically consist in a push to expand road transport infrastructure, but must be based on a search for suitable solutions with an ever stronger emphasis on innovation. Intelligent Transport Systems (ITS) is a research field that aims at integrating all types of transportation and tries to optimise energy efficiency.

An integral part of Intelligent Transport Systems are Intelligent Vehicles, such as the "intelligent car" which constitutes the main subject of this special issue.

ITS & ICST

Intelligent Transport Systems are advanced applications which, while not intelligent in their own right, have been created with the objective of providing intelligent services to various types of transportation and traffic control [1]. The aim is to increase the amount of information available to users in order to ensure a safer, better coordinated and more efficient - thus "intelligent" - utilisation of transportation networks and vehicles.

ITS strategies have an impact on how communication and information technologies are used for strengthening or optimising the performance, efficiency, safety, economy, environmental impact and comfort of transportation. All of these parameters are based on how individual systems communicate with each other, and on faster and more precise methods of relaying information.

Intelligent Transport Systems bring together telecommunications, electronics and information technologies with transport engineering in an effort to plan, design, operate, maintain and control transport systems. Using information and communication technologies in the field of road transport and interfacing them with other types of transportation will significantly contribute to decreasing the impact of road transport on the environment while improving its efficiency and, in particular, its energy consumption. It will also contribute to increasing road safety by considering outside threats, such as the transportation of dangerous goods. It will lead to improvements in the mobility of both goods and people, and will foster a healthy internal market, increased competitiveness and a higher rate of employment.

The current progress made possible by applying information and communication technologies to other areas of transport should also influence the development of road transport, with the primary objective of achieving a higher degree of integration between the road and other modalities of transportation.
Figure 1 shows how various elements in transport are connected together in an Intelligent Transport System.

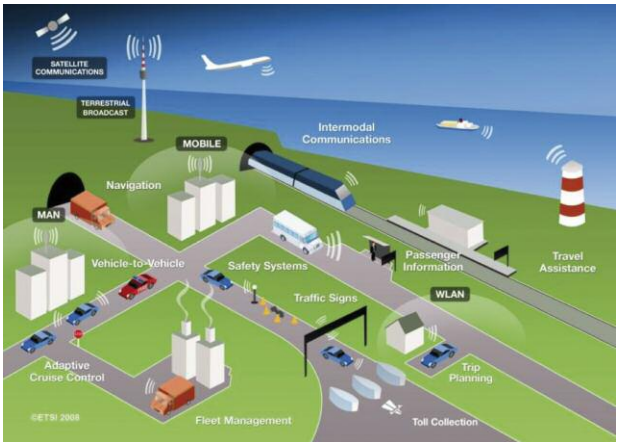

Figure 1: All newly developed vehicles should be compatible with ITS guidelines and standards. Source: ETSI.

Standardisation

All newly developed vehicles should be compatible with ITS guidelines and standards. Two existing technical committees, the European CEN TC278 Road Transport and Traffic Telematics committee and the global ISO TC204 Intelligent Transport Systems committee, generate dozens of standards which have a deep multinational impact on the whole field of transport telematics due to their integration at a European and a global level. It is essential that all countries have thorough knowledge of these standards and capability to apply them. Unfortunately, the number of standards is very high, with, in 2010, approximately 300 standards in place at a varying state of completion. These standards are usually highly complex, their structure is based on IT standards, such as UML, and they often run up to hundreds of pages in length. The only viable solution is to develop correct methods to apply them in practice.

Safety issues

Ensuring a maximum level of safety in the field of road transport is a top priority. ITS has been investigating for many years the feasibility of implementing eCall automatic emergency call systems, to monitor shipments of dangerous goods, to weigh transport vehicles while underway, and to monitor the conditions of live animals under transit. The biggest obstacle for the future development of ITS lies in the non-existence of regulatory measures ratified at a European-wide level.

In combination with the intelligent car, ITS are addressing and reducing risks in many of the major phases surrounding a vehicle collision. Figure 2 shows a timeline of individual collision phases consisting of:

- The information exchange phase. In this phase, it is possible to warn the vehicle of a potentially dangerous situation.
- The safety systems function phase. In this phase, it is possible to adjust the parameters of anti-collision systems according to information from ITS.
- The passive protection function phase. ITS is practically not utilised in this phase.

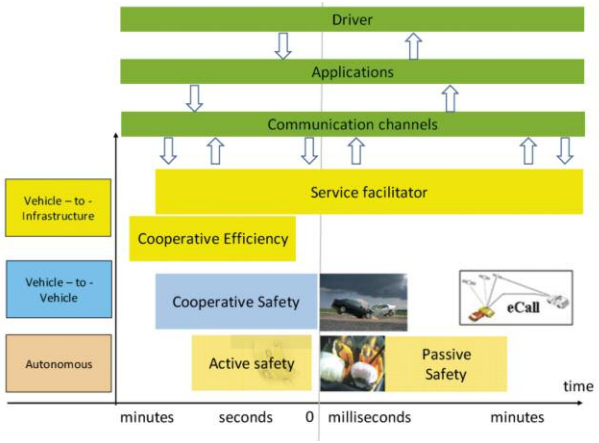

Figure 2: Integrated vehicle safety intelligent system (integration of intelligent vehicles into ITS) [3]. Source: EUCAR.

- The post-collision phase. In this phase, emergency units are notified with the help of ITS and eCall, and vehicles in the surrounding area are warned of potential danger.

The development and implementation of strong and reliable ITS is essential in order to increase traffic safety.

Current and Future Developments

Further ITS development will also entail a change in approaches to the design of motor vehicles, which will need to acquire a capability of dual-way communication. Vehicles will have to be equipped with the necessary sensors to be able to relay information not only about themselves, but also about their surroundings. Figure 3 defines the proportion of the integration of individual ITS elements into vehicles and infrastructure.

Authorities in individual countries are also aware of the need to introduce globally compatible ITS networks. A strong emphasis is being placed on the so-called 'global', but currently tripartite (EU-US-Japan), concept of communication network architecture [2]. Attention is focused not only on standardising ITS, but also on E-Mobility, where it is expected that a high level of interconnection will be achieved.

The following areas have been given top priority from the technical perspective: Driver Distraction; cooperative systems

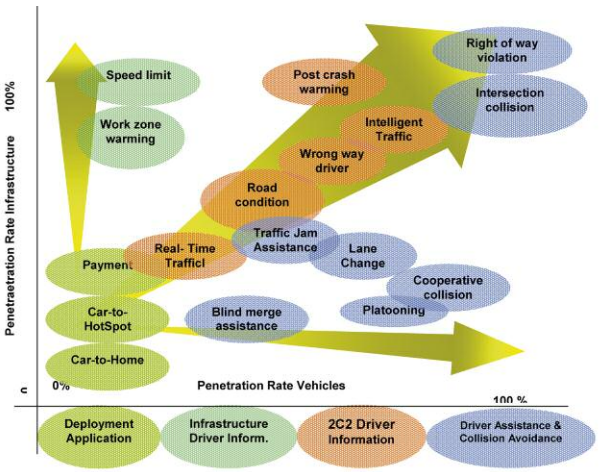

Figure 3: Integration of individual ITS elements into vehicles and infrastructure [3].

(Vehicle-to-Vehicle, Vehicle-to-Infrastructure); safety improving systems such as Collision Warning and Intersection Warning; and more generally ADAS (Advanced Driving Assistance Systems) technologies and recent promising technologies for Fully Autonomous Driving (see for instance the Google Car project). A description of the state-of-the-art of these technologies is presented in [5], and some emerging approaches for Risk Assessment and Collision Warning are presented in [6] . Efforts are being made to find solutions in hardware (HW) and software (SW) compatibility (support for joint HW and SW platforms), that will improve robustness and safety of the system as a whole.

From the perspective of satellite navigation systems, emphasis is being placed on the capability of system interoperability, namely between the Galileo, GPS and GLONASS systems developed by the EU, U.S. and Russia, respectively.

From this discussion it is clear that the issue of intelligent vehicles and their integration into Intelligent Transport Systems is crucial for the development of future transportation, which is why this special edition of our magazine is dedicated to this topic. The articles in these special theme sections provide a panorama of European research in the field. They address several relevant and active subtopics: advanced driving assistance systems and autonomous driving, software engineering methodologies and best practices, vehicle-to-vehicle and vehicle-to-infrastructure communications, data management and privacy, traffic monitoring and control systems and human-machine interaction. This diversity of topics and challenges does not only illustrate the great complexity of the task, but also the progress we are making towards providing viable solutions for it.

Acknowledgement: We would like to thank Dr. Dizan Vasquez for his support during the review process and the finalization of this special theme.

References:

[1] European Parliament and Council Directive 2010/40/EU of 7 July 2010 on ITS in the field of road transport
[2] EUCAR: Sustainability of Road Transportation Research Priorities & Research Targets by 2013, contribution to FP7 Specific and Work Programs
[3] Přibyl et al.: "Transport Terminology Dictionary (Slovník dopravní terminologie)", ČVUT Publishing, 2010, ISBN 978-80-01-04654-8
[4] Strategic Research Agenda "Vehicles for Sustainable Mobility", 2/2013, ISBN 978-80-260-3952-5.
[5] A. Eskandarian (ed.): "Handbook of Intelligent Vehicles, section 10 "Fully Autonomous Driving", edited by C. Laugier, Springer, New York, March 2011, ISBN: 978-0-85729-084-7 (Print) 978-0-85729-085-4 (Online).
[6] C. Laugier et al: "Probabilistic Analysis of Dynamic Scenes and Collision Risk Assessment to improve Driving Safety", in Intelligent Transportation Systems Journal 3.4, Nov. 2011.

Please contact:

Jaroslav Machan, ŠKODA AUTO a.s., Czech Republic
E-mail: Jaroslav.Machan@skoda-auto.cz

Christian Laugier, Inria, France
E-mail: Christian.Laugier@inria.fr

ABV – A Low Speed Automation Project to Study the Technical Feasibility of Fully Automated Driving

by Evangeline Pollard, Fawzi Nashashibi and Paulo Resende

The purpose of the ABV project was to demonstrate the technical feasibility of fully automated driving at speeds below 50 km/h in urban and suburban areas with adequate infrastructure quality (no intersections, known road geometry and lane markings available). Researchers of Inria were in charge of the automation of an electrified Citröen C1 Ev'ie.

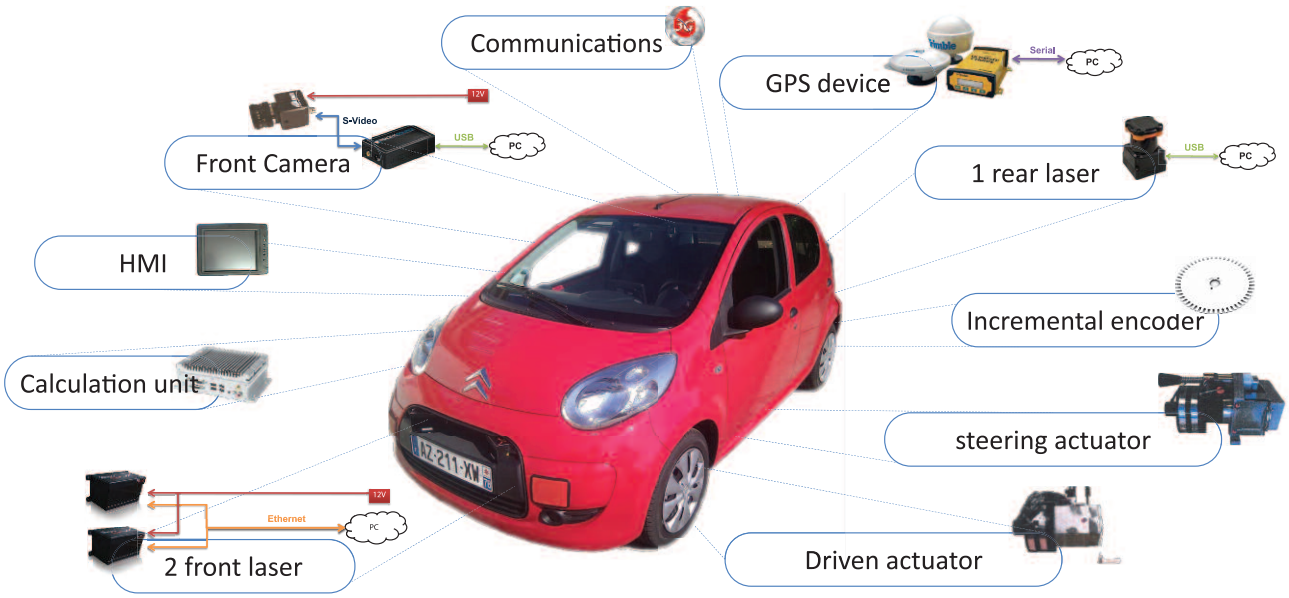

Figure 1: Vehicle equipment.

The Intelligent Transportation Systems (ITS) community has been focusing on vehicle automation since 1987, when the European Commission funded the 800 million Euros EC EUREKA Prometheus Project on autonomous vehicles. In this project, vehicles such as the Dickmann's Mercedes-Benz were able to perform vision-based autonomous navigation on empty roads. Since this first European initiative, much work has been conducted in this area. Initiatives including the Cybercar (90's), the PATH initiative (1997), the Darpa Grand Challenge (2004), the VisLab intercontinental autonomous challenge (2010) and Google car (2011) have contributed to this big challenge.

However, although more aspects of driving are being automated, full automation is still impossible in Europe for legal reasons. In order to create a legal framework for road automation and driving sharing, many projects, including the European project HAVEit, have tried to demonstrate its feasibility. The French Automatisation Basse Vitesse (ABV) project, standing for Low Speed Automation, differs from HAVEit by focusing on congested and heavy traffic in urban and suburban roads at speeds below 50 km/h and adding the fully automated driving capability to the automation spectrum.

By automatically following congested traffic, the ABV system relieves the human driver from performing monotonous tasks such as holding the brake pedal or rather risky maneuvers like changing lanes or simply keeping a safe distance from the vehicle in front. During fully automated driving inside the application zone, the human driver is still responsible for the vehicle, and is involved in the task of driving (e.g. activating a blinker to overtake) but with much less engagement. When reaching the end of the application zone the driver is required to take over the control of the vehicle. If the driver fails to do so, the vehicle will automatically stop.

The automation algorithms, following the classical scheme (1) perception (2) planning (3) control have been integrated into an electrified Citröen C1 Ev'ie and tested on several experiments on the Satory tracks (Versailles, France).

The Citröen C1 was equipped with sensors (odometers, lasers and frontal camera) in order to detect obstacles to avoid and determine which lane to follow, on-board computer to process data and make decisions, acceleration/brake and steering actuators (PARAVAN) to control the vehicle, and communication devices.

Perception issues are mainly divided into three axes. First the ego-localization allows the vehicle to know its position and orientation in the map (in the application zone). Position measurements coming from the GPS devices are fused with the velocity measurements coming from the odometers by using an Extended Kalman Filter (EKF). Obstacle detection and tracking is made with frontal and rear laser-scanners. Laser data are segmented using a Recursive Line Fitting algorithm to

extract objects which are tracked with a Nearest Neighbour (NN) approach. Lane detection is made using the frontal camera. Images are transformed into bird's-eye view images by Inverse Perspective Mapping, then a Canny detector is used to detect edges, which are filtered using a priori rules. Finally, extracted lines are filtered using a particle filtering.

Lane and obstacle detection are both used to plan the trajectory. A Lane Keeping System (LKS) allows the vehicle to stay in the middle of the current lane knowing the marking position. An Adaptive Cruise Control (ACC) is used to keep the vehicle at a safe distance from the vehicle in front. Relative velocity of the front obstacle is calculated by numerical derivation of the relative distance. An Obstacle Avoidance System (OAS) launches secured overtaking if activated by the driver. Finally, emergency braking can be activated by the driver or by the presence of an obstacle on the lane.

Control is achieved by automatically calculating the steering wheel's angle and the velocity of the vehicle. Lateral control consists of following the reference line provided by the lane detection algorithm. For longitudinal control, a PID (proportional-integral-derivative controller) is used to maintain the velocity at the minimum velocity from among the following: The driver's desired velocity; the legal speed limit; the speed limitation due to road geometry; and the front obstacle velocity.

Some work has been done on interactions between the system and the driver. A specific human-machine interface in the vehicle has been designed as a touch screen to let the driver choose the preferred speed and to initiate lane changes by activating the blinkers, if possible. A situation awareness component monitors the current state of the lanes (number of lanes, lane marking type, and current lane index), obstacles, ego vehicle state and current maneuver.

Numerous experiments have been conducted on the Satory tracks under several conditions. The vehicle has demonstrated its automation ability, as well as its driving sharing ability with the driver. Long term perspectives tend to ensure automation outside the application zone. Work is still in progress on the interaction between the driver and the car. There is no doubt that the future of automation lies in intelligent interaction, capable of assessing the driver's, as well as the system's, state.

Links:
https://www.youtube.com/watch?v=xsV8P7X4gvw
http://youtu.be/DoM06ho7JC0
http://www.projet-abv.fr/en/

References:
[1] H. Li and F. Nashashibi: "Robust real-time lane detection based on lane mark segment features and general a priori knowledge", IEEE ICRB, 2011.
[2] P. Resende, et al.: "Low Speed Automation: technical feasibility of the driving sharing in urban area", IEEE ITSC, 2013

Please contact:
Evangeline Pollard, Fawzi Nashashibi and Paulo Resende
Inria, France
E-mail: evangeline.pollard@inria.fr

Cooperative Systems for Car Safety Improvement

by Zdeněk Lokaj, Tomas Zelinka and Miroslav Šrotýř

Cooperative systems are playing an increasingly important role in the automotive industry, particularly in relation to improvement of safety. Telecommunication solutions are crucial in cooperative systems, owing to the need to deliver data to appropriate places at appropriate times. This paper describes state-of-the-art of cooperative systems and describes one real application which is under development.

Many research and development projects are in progress around the world with the goal of developing new cooperative applications and testing them in real environments. Generally the approach of R&D in cooperative systems focuses on the following applications:

- Mitigation of risk of accidents at high speed – Electronic emergency brake lights, front collision warning systems, blind spot warnings / warnings during lane changing, overtaking warnings, left turn assistant
- Mitigation of risk of accidents at low speed – Crossroads motion assistant
- Safety awareness – Remote diagnostics (battery monitoring), stationary vehicle (or pedestrian) warnings, road conditions warnings
- Emergency applications – eCall
- Eco-Green Mobility / Vehicle Speed – Electronic payments (e.g. Electronic Fee Collections – EFC), traffic data collection, dedicated eco-lane, green wave, traffic signal timing for optimal stopping of the vehicle
- Applications for driver convenience – Electric-vehicle charging stations guides, eCommerce, internet access, multimedia files download, video telephony.

Telecommunication Solutions for Cooperative Systems

Since the fundamental component of cooperative systems is wireless telecommunication systems much of the research focuses on this area. Since 2009 there has been dedicated telco technology DSRC 5.9 GHz for ITS solutions communication in Europe, but studies have also been investigating the potential to use other technologies, such as WiFi (IEEE 802.11 e/n/p), mobile data services (EDGE, HSDPA, LTE, LTE-A), WiMAX (IEEE 802.16 e/m), Bluetooth (IEEE 802.15.1) and MBWA (IEEE 802.20) to analyse properties of

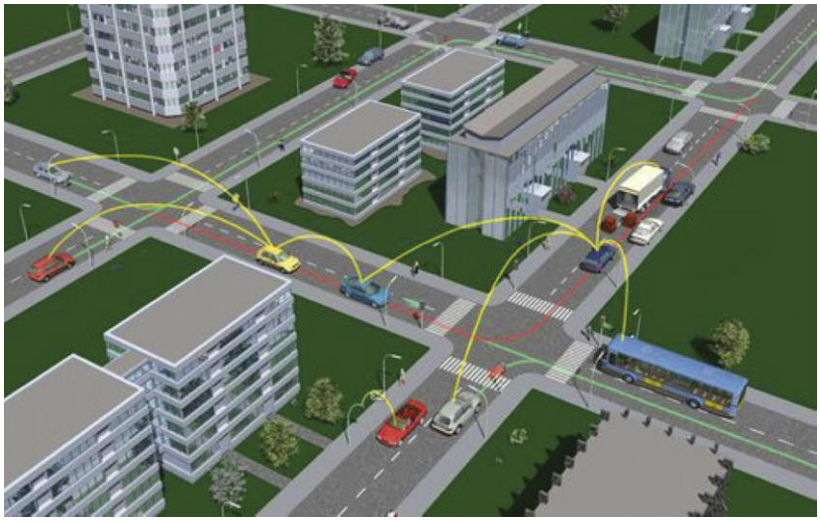

Figure 1: Cooperative systems basic scheme.
Source: www.car-to-car.org

these technologies and try to link them to the performance requirements of the cooperative applications.

Intelligent Truck Parking Application

One example of a cooperative system is the Intelligent Truck Parking Application (ITP), an R&D project currently under development at the Czech Technical University in Prague. This project is funded by the Technology Agency of the Czech Republic (project number TA02031411).

The goal of the system is to predict occupation of truck parking places on highways and immediately inform drivers via C2X telco technologies to help them plan their rest stops in accordance with legislation (e.g. a regular stop after eight hours of driving). This system mitigates risk by reducing the chance of a driver being unable to find an unoccupied parking space and thus having to drive to the next rest area. Based on actual information about parking occupation the driver can plan the best place to park and rest.

The Intelligent Truck Parking System uses on-line traffic data from the Czech tolling system which is the main data source for the statistical model developed to predict truck parking lot availability. To achieve greater predictive accuracy historical data from the tolling system is incorporated with actual data from traffic sensors (e.g. loops, video-sensors, parking entrance detectors).

The tolling system in the Czech Republic covers more than 1,365 km of highways and motorways and provides traffic data about vehicles above 3.5 t almost in real-time (with a maximum delay of approximately five minutes) so the data provides a complete description of truck traffic flow on highways including abnormalities in the flow of traffic and the source of the abnormalities.

Achievement of a reliable and highly accurate Intelligent Truck Parking system is highly dependent on a number of crucial factors:

- The availability of top quality and reliable data sources describing traffic flow as the basis for a statistical model to predict parking lot availability.
- The availability of reliable information regarding available truck parking lots in an appropriate area.
- The transmission of information via different communication channels e.g. VMS (Variable Message Sign), Smart Phones, RDS-TMC (Radio Data System – Traffic Message Channel) or C2C (Car-to-Car) systems.
- Implemented ITS systems need to be supported and become an integrated part of the ITS architecture.

The statistical prediction model development for the Intelligent Truck Parking application is almost finished. The model has been constructed using historical data from the tolling system and data from transport surveys of selected highway parking places. The next step is to validate outputs of the model and calibrate it to improve its accuracy.

References:
[1] T. Zelinka, Z. Lokaj, M. Svitek: "Specific telecommunications solutions for ITS applications, Circuits, Systems, Signal and Telecommunications Conference", in proc. Applied mathematics in electrical and computer engineering, Cambridge, Harvard, USA, MA, 2012, p. 315-322, ISBN 978-1-61804-064-0.
[2] Z. Kocur et al: "Analysis of Influence of Disturbance Level on Data Transmission in Wireless Networks", in TSP 2010 – 33rd International conference on Telecommunications and Signal Processing [CD-ROM], Budapest: Asszisztencia Szervező Kft., 2010, p. 292-296. ISBN 978-963-88981-0-4.

Please contact:
Zdeněk Lokaj
Czech Technical University in Prague,
Faculty of Transportation Sciences,
Prague, Czech Republic
Tel: +420 224359611
E-mail: lokaj@fd.cvut.cz

Location Assurance and Privacy in GNSS Navigation

by Xihui Chen, Carlo Harpes, Gabriele Lenzini, Sjouke Mauw and Jun Pang

The growing popularity of location-based services such as GNSS (Global Navigation Satellite System) navigation requires confidence in the reliability of the calculated locations. The exploration of a user's location also gives rise to severe privacy concerns. Within an ESA (European Space Agency) funded project, we have developed a service that not only verifies the correctness of users' locations but also enables users to control the accuracy of their revealed locations.

Driving across unknown routes used to be stressful owing to the almost inevitable problem of getting lost. Thanks to Global Navigation Satellite Systems (GNSS) navigation devices, such problems are now in the past. These devices have revolutionized how we move: they guide us not only in unknown places but also within our own neighbourhoods. Our growing dependence on them to find our way is risky though: since the airwaves that carry GNSS signals are broadcast in the open air with a relatively weak strength, they are vulnerable to spoofing and meaconing attacks.

Spoofing interferes with GNSS signals and misleads a driver into calculating a different location. Instead of interfering, a meaconing attack intercepts, alters or delays GNSS signals.

Launching these attacks is illegal, but devices required for their implementation are easy for anyone to access. Although this may not bother many of us, for those who run businesses that depend on correct routing (transport of valuable goods), spoofing and meaconing represent serious threats.

To address these threats, researchers from the Interdisciplinary Centre for Security, Reliability and Trust (SnT) of the University of Luxembourg and engineers from 'itrust consulting' Luxembourg teamed up to execute a project funded by the European Space Agency (from 12/2010 to 11/2012). They designed and implemented a service that assesses the trustworthiness of a GNSS device's claim of being at a certain location [1]. The service, Localization Assurance Service Provider (LASP), runs on a trusted third party between location-based service providers and their users. The architecture where LASP runs is shown in Figure 1: all communications, except for those between GNSS and User Devices, rely on the Internet or the data mobile network infrastructure and they are supposed to be secure.

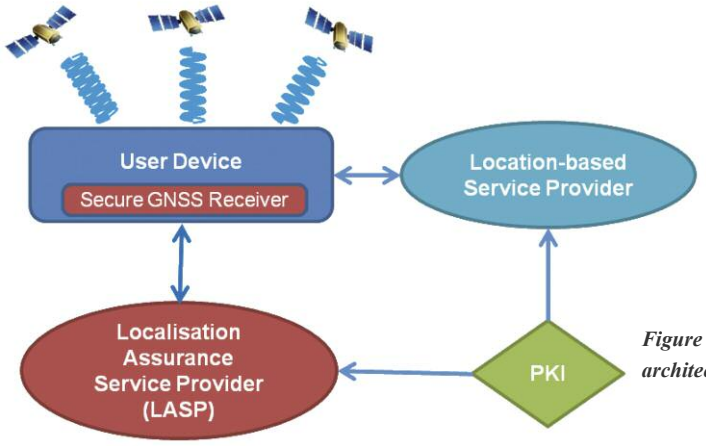

Figure 1: LASP service architecture

The LASP enacts a system of security checks, each check monitoring signal properties – such as signal strength, Doppler ratio, and clock bias [2] – that a GNSS navigation device has measured during localization. The outcomes of different security checks are intelligently processed based on a trust framework implemented with probabilistic conditional reasoning and subjective logic [3]. The results are then combined to obtain a value expressing to what extent a given localization can be trusted. This value, called localization assurance level, is embedded in a certificate issued by the LASP together with the putative user location. High assurance levels are given to localizations derived from untampered GNSS signals while low assurance levels to localizations that have been found to have inconsistencies likely due to spoofing.

In addition to the localization assurance, SnT researchers have studied solutions to protect location privacy. Thanks to a technique called selective blinding, users can control the accuracy of the location that is contained in a LASP certificate without invalidating the certificate. Location-based service providers can thus receive certified localizations that adhere to the need-to-know principle [4].

References:

[1] C. Harpes et al: "Implementation and validation of a localisation assurance service provider", in proc. IEEE NAVITEC'12, Noordwijk, The Netherlands, 2012

[2] D. Marnach et al: "Detecting Meaconing Attacks by analysing the Clock Bias of GNSS Receivers", Artificial Satellites, Journal of Planetary Geodesy 48(2), pp. 63-84, 2013

[3] X. Chen et al: "A trust framework for evaluating GNSS signal integrity", in proc. IEEE CSF'13, New Orleans, USA, 2013

[4] G. Lenzini, S. Mauw, J. Pang: "Selective location blinding using hash chains", in proc. SPW'11, Cambridge, UK, 2011

Please contact:

Carlo Harpes
itrust consulting, Luxembourg
E-mail: Carlo.Harpes@itrust.lu

Xihui Chen, Gabriele Lenzini, Sjouke Mauw, and Jun Pang
University of Luxembourg
E-mail: {Xihui.Chen, Gabriele.Lenzini, Sjouke.Mauw, Jun.Pang}@uni.lu

Paving the Way for Apps in Vehicles

by Avenir Kobetski and Jakob Axelsson

What will cars talk about when they start communicating with each other? Volvo and SICS Swedish ICT have launched a collaborative project to open the computer systems of cars for the market of apps.

In a world where cars are containing increasingly sophisticated software, it comes as no surprise when computer scientists cooperate with vehicle companies. During the last couple of years SICS has increased its efforts to work together with the vehicle industry in Sweden, within areas such as data mining, networked sensor systems, and software and systems engineering.

Building on its expertise in the field of the Internet of Things, SICS is currently working on a vision of the Internet of Cars (IoC), or simply federations of vehicles and other embedded systems that interact to provide new services to their users. The idea is to let vehicles exchange the information that they normally collect, for example the state of the road, traffic jams, the exact time of arrival, etc. Handled in a smart way, such information can contribute to making trips safer, more efficient and a lot more fun.

To reach the IoC vision, several advancements are needed. First, since it is not possible to foresee all future IoC applications, there should be a way to add intelligence to the vehicles after they have left the factory, in much the same way as apps are added to today's smartphones. However, since the safety of cars is more critical than that of typical cellphones, this requires an entirely new level of security and robustness.

In a project, initiated by SICS and Volvo, a framework for opening up computer systems of vehicles is being developed. This framework builds on Autosar, the leading automotive Electrical/Electronic (E/E) architecture standard, used in millions of cars throughout the world. As a result, Autosar is extended with a sandboxed environment, dedicated to apps. The interface between such an environment and the underlying functionality will be defined by the vehicle companies, allowing them to choose exactly which parts of the car's system can be influenced by the apps, and which should never be touched.

Currently, initial concepts for app installation and communication of data, both within the vehicle and with the outside world, have been developed [1] and simulated in a desktop environment. In the next couple of months, the simulation work will be transferred from a PC to embedded hardware. A couple of credit-card computers, Raspberry Pis, will be interconnected and used to simulate electronic control units (ECUs) of a real car. Next, a number of test apps will be developed and run in the ECU simulators to evaluate the vehicle app concepts.

Some app ideas that will be tested are related to the field of construction equipment, where Volvo sees great potential to use apps to coordinate construction work. An example is scheduling of the construction vehicles of different types, e.g. so that the hauler doesn't have to wait for the loader [2]. Also, coordination between construction equipment and private cars could lead to great cost reductions and better safety at road works sites.

Another example is to equip traffic lights at intersections with a decision app, and to let them act on speed information from the approaching vehicles to optimize the traffic flow through intersections. If the speed information is exchanged directly between vehicles, the car itself can detect a stop in the traffic ahead and adapt its speed to avoid a sudden halt, etc.

The second direction of SICS research towards the IoC vision is concerned with the business infrastructure that vehicle apps will create. New innovative companies will develop apps for cars, in a similar way to what occurred in the mobile industry. This will give rise to entirely new business ecosystems, holding both great opportunities and challenges.

For instance, ways of sharing information between different parties will need to be defined, so that an app developer can implement and test their software without full access to the overall product. The distribution of rights and responsibilities between the parties are also crucial. Who is liable in a situation where an incident occurs? Also, how should the streams of income be set up and divided among the parties?

To investigate these issues, SICS is currently conducting an empirical research project based on case studies and interviews, trying to identify the primary interfaces between stakeholders in a typical IoC-ecosystem [3]. This will be used to create a reference model, intended to provide guidance on how to organize future IoC ecosystems efficiently.

In the long run, the IoC concepts will affect the entire vehicle and transportation business. The opportunities are numerous, but the stakes are high. SICS will focus on carefully pushing both of the above research directions, with a first demonstration of app-enabled cars expected at the end of 2014.

Links:

http:// www.sics.se/projects/fresta-apps-in-vehicles
AUTOSAR consortium:
http://www.autosar.org/
Raspberry Pi community:
http://www.raspberrypi.org/

References:

[1] J. Axelsson, A. Kobetski: "On the Conceptual Design of a Dynamic Component Model for Reconfigurable AUTOSAR Systems", in APRES'13, 2013.
[2] D. Rylander, J. Axelsson: "Lean Method to Identify Improvements for Operation Control at Quarry Sites", in ISARC'13, 2013.
[3] E. Papatheocharous, J. Axelsson, J. Andersson: "Towards an Innovative Open Ecosystem Infrastructure for Federated Embedded Systems Development", in SESOS'13, 2013.

Please contact:

Avenir Kobetski or Jakob Axelsson
SICS Swedish ICT
Tel: +46 72 238 92 70, +46 72 734 29 52
E-mail: avenir.kobetski@sics.se, jakob.axelsson@sics.se

Mobile Opportunistic Traffic Offloading for Map-Based ADAS Applications: the MOTO Project Approach

by Raffaele Bruno, Andrea Passarella, Leandro D'Orazio and Filippo Visintainer

Map-based Advanced Driver Assistance Systems (ADAS) is one of the key applications for future "Intelligent Car" environments. Unfortunately, network capacity limitations may hinder their efficiency and large-scale adoption. The FP7 MOTO project investigates offloading on vehicle to vehicle (V2V) opportunistic networks as a suitable way to overcome such limitations, thus enabling ADAS and other data-hungry future vehicular applications.

Modern ADAS applications largely take advantage of maps with detailed information to support the users' driving behaviour (curve speed warning, passing/overtaking assistant, eco-routing/eco-driving). These maps are more accurate than those used for simple navigation and contain very high definition details, such as enhanced geometry, road curvature, height, slope, speed limits, lane information, etc. Since real-time updates are required, car manufacturers support the storage of these maps in central servers, where they are updated periodically, and can be downloaded in real time through mobile communication technologies.

A drawback to this model is that during moderate to high vehicular traffic conditions, the bandwidth available to individual users may be limited, even when modern cellular technologies (i.e. Long-Term Evolution, LTE) are used. Despite extremely high peak rates, the net throughput available to LTE users is not expected to exceed a few Mbps [1]. In moderate to heavy vehicular traffic conditions, when ADAS network traffic is supported by only a few LTE Base Stations (eNodeBs), the ADAS applications together with other data-hungry applications served by the same eNodeBs are very likely to saturate this capacity.

The FP7 MOTO project is a recently started EU project under the "Future Networks" Objective. MOTO investigates how to exploit mobile opportunistic networking solutions to offload traffic from congested cellular infrastructures, and considers ADAS applications in vehicular environments as one of the key use cases. An opportunistic network [2] is a mobile self-organizing network formed, in this case, by vehicles that communicate directly through wireless technologies (in the case of vehicular applications, the MOTO project focuses on the IEEE 802.11p standard in ad hoc mode). Unlike the conventional Mobile Ad hoc Networking (MANET) paradigm, in opportunistic networks the existence of a simultaneous end-to-end multi-hop path between senders and receivers is not taken for granted. Instead, messages are routed dynamically towards the final destination(s), by exploiting communication opportunities that arise thanks to the nodes' mobility (e.g., two vehicles coming close enough to establish a direct single hop wireless connection). Whenever such an event occurs, each node immediately evaluates the suitability of the other to carry locally stored messages closer and closer to the final destination and, if appropriate, hands over messages. This approach avoids problems typical of mobile ad hoc networks, such as disconnections, network partitions, and topology instability.

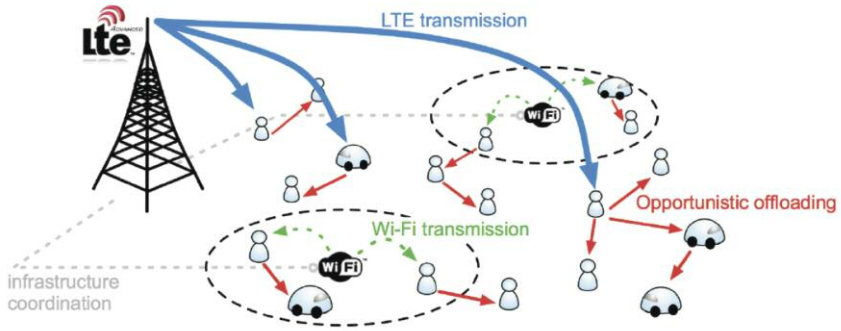

Figure 1: General offloading scenario of the MOTO project

The opportunistic networking paradigm has recently been gaining momentum as a solution for offloading traffic from congested cellular networks in the case of heavy data traffic concentrated in well-identified physical areas [3]. The main idea behind MOTO (see Figure 1) is that offloading can be most beneficial when a large number of users located in a given physical area need to download (almost) the same information at (almost) the same time from central servers. This is exactly the scenario of ADAS applications, as drivers moving in a congested area need the same maps containing the same information at almost the same time (although multicasting is typically not viable, due to the dynamicity of the vehicular traffic and the absence of synchronization among vehicles). Applying offloading to support ADAS applications is, thus, one of the use cases and scenarios considered

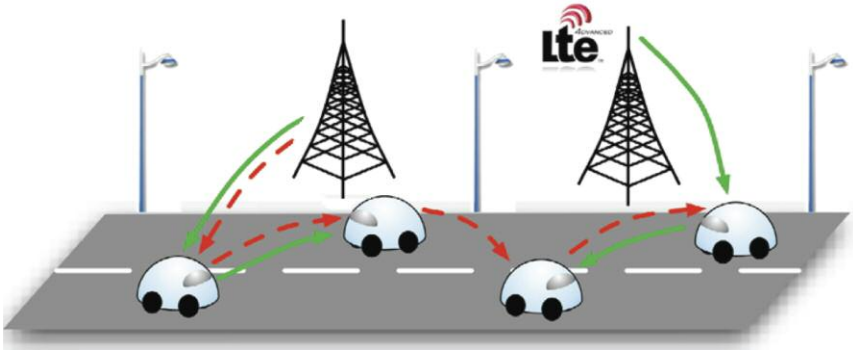

Figure 2: MOTO offloading for ADAS applications

by MOTO (see Figure 2). Relying exclusively on LTE, vehicles would need to establish a dedicated connection to the central servers, from which to download the multimedia files required by the ADAS application. The physical co-location of vehicles means that this network traffic will have to be sustained by a very limited number of eNodeBs, and consequently the congestion resulting from ADAS traffic (and possibly other "background" multimedia traffic generated by the users) over the LTE network may be excessive. To avoid this problem and allow ADAS applications to be effective, MOTO solutions select a subset of vehicles (seeds) which can download the maps from the central servers through the LTE network. The maps are then disseminated through V2V communications using 802.11p technology to the other interested users, exploiting mobile opportunistic data dissemination protocols. As a last-resort option, individual users may resort to the LTE network when the maps have not been received by a stated deadline. By exploiting MOTO services, drivers should receive the detailed maps required by the ADAS application, while drastically reducing the congestion on the LTE network (thus avoiding blocking and lock-out problems). At the current stage (first six months in the project), we have defined in detail the use cases for offloading through opportunistic networking in vehicular scenarios. The system-level architecture required to support offloading from LTE to opportunistic networks is about to be released. Moreover, we are running extensive simulations to quantify the limit of LTE-only solutions in supporting ADAS applications, and the benefit coming from a solution based on offloading.

While ADAS applications are one of the specific focuses of the MOTO project, the same offloading techniques may directly or indirectly help a range of other applications and services that future intelligent cars are likely to offer. For example, opportunistic networks formed by vehicles can natively support data dissemination between vehicles, for traffic alert applications, or infotainment applications, without overloading the LTE network. Offloading should also indirectly help safety-critical applications as the LTE network will not be congested due to the traffic of non-safety-critical applications, and could more easily sustain time-sensitive traffic. At the same time, non-safety-critical applications will not be blocked, but can still be supported through the opportunistic part of the network.

The MOTO project is coordinated by Thales (France) and also involves CNR (Italy), CRF (Italy), UPMC (France), AVEA (Turkey), FON (UK) and Intecs (Italy).

Link:
http://www.fp7-moto.eu

References:
[1] L. Korowajczuk (Editor): "LTE, WiMAX and WLAN Network Design, Optimization and Performance Analysis", Wiley, July 2011.
[2] L. Pelusi, A. Passarella, M. Conti: "Opportunistic Networking: data forwarding in disconnected mobile ad hoc networks", IEEE Communications Magazine, vol. 44, no. 11, November 2006, pp. 134-141
[3] J. Whitbeck et al: "Push-and-Track: Saving Infrastructure Bandwidth Through Opportunistic Forwarding", Elsevier Pervasive and Mobile Computing, vol. 8, issue 5 (August 2012), pages 682-697

Please contact:
Raffaele Bruno, Andrea Passarella, IIT-CNR, Italy
E-mail: raffaele.bruno@iit.cnr.it, andrea.passarella@iit.cnr.it

Leandro D'Orazio
Centro Ricerche Fiat, Italy
E-mail: leandro.dorazio@crf.it

Filippo Visintainer
Centro Ricerche Fiat, Italy
E-mail: filippo.visintainer@crf.it

Autonomous Traffic Warning System with Car-to-X Communication

by Josef Jiru

Of the 2.36 million traffic accidents recorded in Germany in 2011, traffic congestion and tailgating caused the most injuries. Because they are usually positioned at one to two kilometer intervals, conventional traffic control systems do not adequately warn drivers of sudden danger situations. When positioned between the systems, drivers must rely solely on their own perception of the situation. To address this issue, Fraunhofer ESK is working with Ruetz Technologies and TRANSVER to develop an autonomous traffic warning and information platform that can more quickly detect critical situations and provide precise local warnings through finer granularity.

The combination of various sensors enables much faster and more reliable detection of critical situations on high-speed roads through autonomous stations on the edge of the road (road side units or RSU). Fraunhofer researchers designed the system which utilizes a combination of stationary sensors, such as radar or measuring loops, and information relayed from vehicles.

The system analyzes the data and identifies potentially dangerous situations such as the sudden appearance of road construction sites, traffic back-ups, the presence of objects on the road or inclement weather. Drivers are informed of the precise location of the danger through warning lights integrated in the reflector posts and directly via car-to-x communication.

The data gathered by the system must be processed in real-time. The main task

involves merging the data generated by various sources in order to create a consistent overview of the traffic situation. Matching the information from multiple sensors covering the same area ensures a high level of data quality at the source. Furthermore, local as well as remote data is aggregated depending on the current network capacity and needs of the application scenario, and then forwarded. The system's modular design enables the transparent use of different sensor technology.

From sensor measurement to warning light

Traffic data, such as the position and speed of individual vehicles, is recorded within a dense grid, merged if needed, and then transmitted via wireless technology to a higher-level control unit (remote terminal unit or RTU). The RTU uses the accumulated information to ascertain the current traffic situation in real-time. If the system identifies a problem that could lead to a dangerous situation for traffic heading in a given direction, drivers are given a heads-up via warning lights that will blink with a standardized pattern. If appropriately equipped, vehicles can also directly receive standardized DEN messages and display traffic warnings. This raises the level of awareness among traffic participants in this particular grid. Direct car-to-x messages are transmitted directly via ETSI ITS-G5 technology and give the drivers more detailed information on the position and nature of the danger ahead, e.g. a broken down vehicle or the end of a traffic jam. Fraunhofer ESK is creating a heterogeneous network concept for communication between vehicles, infrastructure and RTUs in which the data are aggregated in an adaptive manner. The concepts and processes are validated in a simulation environment and then implemented with in-house developed software and hardware on a test route.

ezCar2X Framework

To easily generate these kinds of driver assistance systems, Fraunhofer ESK developed ezCar2X, a flexible software framework that creates the key components for rapidly developing prototype applications for use in vehicle-to-environment networking. This enables vehicle manufacturers, suppliers and road infrastructure operators to quickly implement new ideas and test them in a near-real environment.

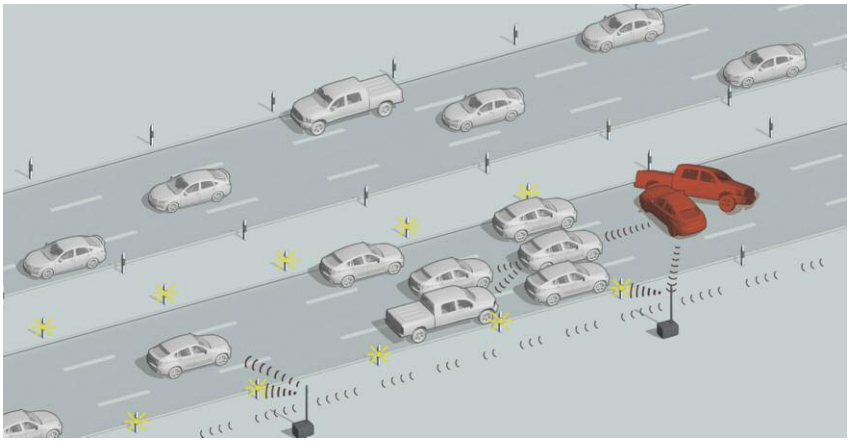

Figure 1: Accident on a three-lane highway. One of the accident vehicles transmits a warning to a road side unit (RSU) and to vehicles approaching from behind, which then forward the warning further. Beginning at the accident site, roadside warning lights are activated to visually warn vehicles approaching from behind.

The Fraunhofer ESK car-to-x framework ezCar2X features components for the following functions:

- ETSI ITS-compliant communication services, including ITS G5 and GeoNetworking
- CAN bus interface for fast integration into the vehicle
- Interfaces to external sensors, including GPS, speed or object tracking
- Digital map integration
- Facilities such as CAM and DENM management and precise determination of the vehicle's own position
- Interfaces for connecting various HMI devices such as tablets and touch screens.

The framework is a collection of C++ libraries that are available for Linux and Windows systems. Extensive API documentation and various examples speed up the learning curve. The software components were designed with a low degree of interdependency to enable them to be combined as needed. Unused functions do not create an additional load on resources at runtime. The consistent use of abstract concepts and known design patterns leads to a high degree of modularity. Individual functions can be swapped and replaced by other implementations if needed without having to modify the remaining system.

The project is funded by the Bavarian Ministry of Economic Affairs, Infrastructure, Transport and Technology.

Link:
http://www.esk.fraunhofer.de/en/projects/VIP.html

References:

[1] K. Roscher, et al: "ezCar2X: a modular software framework for rapid prototyping of C2X applications", 9th ITS European Congress, Dublin, Ireland, 4-7 June 2013
[2] T. Steiner, K. Roscher, J. Jiru: "Cooperative Glare Reduction Using V2V Radio Technology", 5th International Symposium on Wireless Vehicular Communications: WIVEC2013, 2-3 June 2013, Dresden, Germany

Please contact:

Josef Jiru
Fraunhofer Institute for Communication Systems ESK, Germany
E-mail: josef.jiru@esk.fraunhofer.de

Understanding V2X Communication Dynamics Through Complex Network Science

by Nicholas Loulloudes, George Pallis and Marios D. Dikaiakos

With the proliferation of wireless hardware that supports communication standards designed exclusively for the vehicular environment (IEEE 802.11p / WAVE), the wide deployment of Vehicular Networks (VANETs) is one step closer to realization. However, due to the inherent dynamics of mobility, the spatio-temporal topological features of such networks remain largely unknown, thereby putting at stake the quality of service provided by those intelligent applications for which VANETs were envisioned. The study of V2X communication dynamics from the aspect of Complex Network science can provide crucial insights to these features.

Our work stems from previous studies of large-scale, real-world systems (i.e. Internet topology, biological and social networks), which have led to important observations with significant influence [1,2]. During the design phase of applications such as safety message dissemination, congestion avoidance and infotainment that rely on V2X communications, certain key questions should be answered. For instance, when performing message routing, "What is the size and how well is the VANET connected in the arrow of time?"; when performing periodic 1-hop broadcasting, the question is "Does the topology support message spread with minimal rebroadcasts hence reducing latency, overhead and packet collisions?"; when performing multi-hop communication, "How far in the topology can packets travel?"; when the network is fragmented, "What are the properties of these clusters?".

Using publicly available, highly realistic vehicular mobility traces [3] in a large-scale urban environment (~34 km^2), we take per-second snapshots of the network over two hours, by considering the exact geographic position of ~53,000 nodes and the propensity for the establishment of an ad hoc wireless link given the communication range of V2X hardware. From accurate digital maps, we obtain the structure of the underlying road topology to establish a LOS/NLOS communication paradigm. Then we model the VANET topology as a graph, where vehicles correspond to vertices and links to edges. We analyze the structure and evolution of the communication graphs in variable traffic conditions and settings of V2X-enabled vehicle market penetration (P). Additionally, we extract the position of 7,000 signalized intersections in the area of study, which could potentially function as RSUs, and study

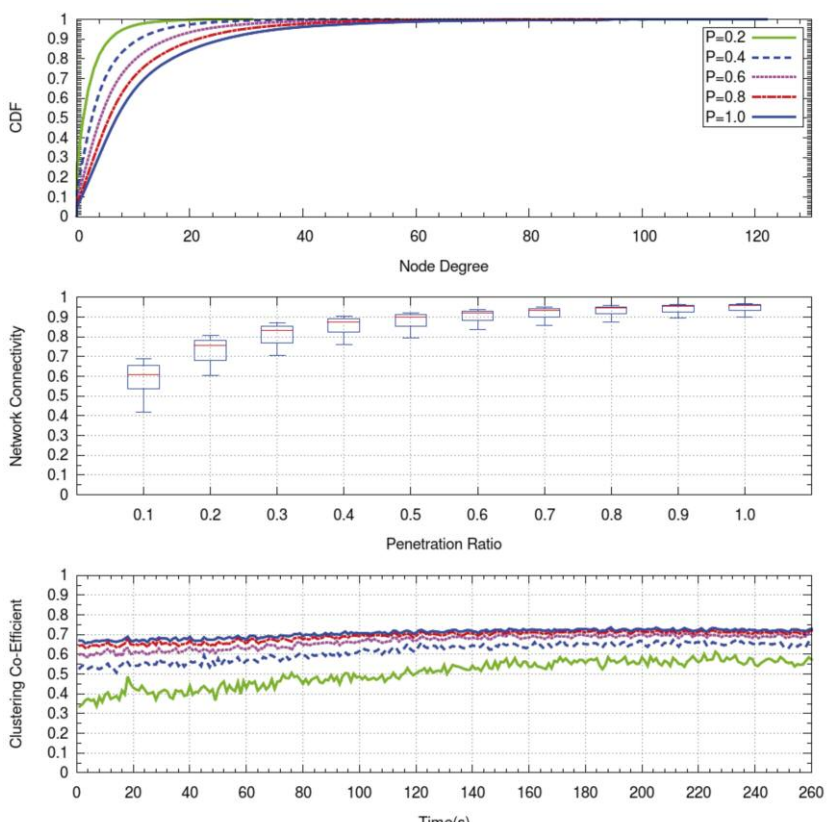

Figure 1:
Top: Cumulative Distribution Function of Node Degree
Centre: Network Connectivity vs. Penetration Ratio
Bottom: Clustering Co-Efficient vs. Time

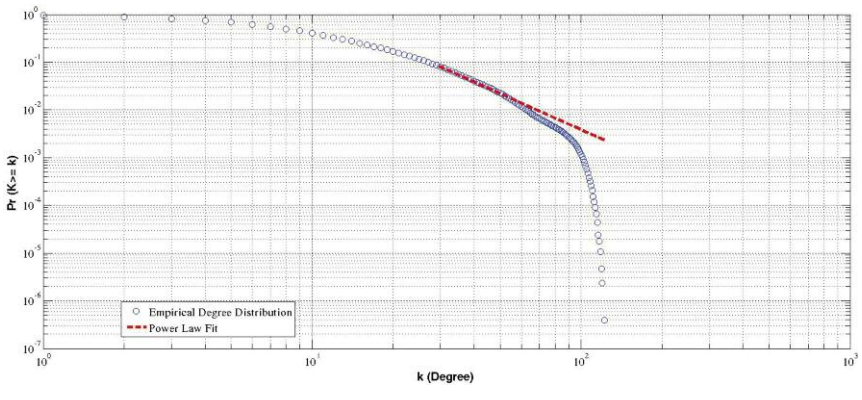

Figure 2: Empirical Node Degree Distribution and Power-Law Fit

their effects. Overall we analyse 436,000 graphs.

Our initial findings suggest that the VANET size (number of edges) grows linearly with increases in P. By examining the network connectivity (Figure 1a) we observe that even for values as low as P=10%, on average 62% of vehicles in the area of study were connected either with a direct or multi-hop link. This is counterintuitive since one would expect that in the case of a low P there would be a low probability that an encounter between two random vehicles could result in the establishment of a link. Furthermore, the geographical dispersion of the V2X-enabled vehicles, and consequently the distance between any established links, would prevent their fusion into multi-hop communication paths. When utilizing only 5% of the available RSUs, the minimum connectivity increases from 42% to 57%, showing that such stationary nodes can extend considerably the communication potential of vehicles. Studying the aggregate node degree distribution over the two hours (Figure 1b), we observe that VANETs are heterogeneous, where there can coexist several isolated vehicles and also others with over 100 neighbours. Nevertheless, we make an important observation in that the degree distribution does not follow a Power Law for any P, making VANETs not scale-free. Figure 2 depicts the empirical degree distribution (P=100%), which exhibits a cutoff at around 120 neighbours, in contrast to the power law

Figure 3: Geographic Coverage of Giant Cluster in the VANET

fit that stretches on and on without a cutoff. Although at first sight vehicles with 120 neighbours look like "hubs", they only appear at large network sizes and their links represent less than 1% of the network.

The network diameter exhibits very high variability, occasionally having shifts in excess of 20 hops, within only a few seconds. Moreover, the diameter follows the average degree of separation up to ~130 hops. The degree of separation increases almost proportionally to the geographic network size for all values of P, hence along with the fact that the degree distribution doesn't follow a power-law, we are led to the conclusion that VANETs do not exhibit small-world properties.

The VANET is partitioned in several clusters for all values of P. At low P it exhibits several small clusters, while high P favours the formation of at least one giant cluster (>0.1*network size) and a few smaller ones surrounding it. As P increases over 50% and a particular vehicle density threshold is reached, no new clusters are created. Newly arriving vehicles join one of the ~500 existing clusters. The average clustering co-efficient (Figure 1c) remains stable ~68% for P>=40%. This is a good indication of the connectivity of the nodes within the cluster, enabling message dissemination to a large audience with minimal broadcasts. Giant clusters such as that depicted in Figure 3 (in red), cover ~20% of the geographic area, however their shape and coverage changes rapidly in time. Effectively, given their strong internal connectivity, giant clusters enable the persistent dissemination and temporal maintenance of information over an extended geography.

References:
[1] D.J. Watts, S.H. Strotgatz: "Collective Dynamics of Small-World Networks", Nature 1998
[2] A. Barabasi, R. Albert, "Emergence of Scaling in Random Networks", Science 1999
[3] TAPAS-Cologne dataset: http://sourceforge.net/apps/mediawiki/sumo/index.php?title=Data/Scenarios/TAPASCologne

Please contact:
Nicholas Loulloudes
Dept. Computer Science,
University of Cyprus
Tel: +357 22892676
E-mail: loulloudes.n@cs.ucy.ac.cy

The VESPA Project: Driving Advances in Data Management for Vehicular Networks

by Thierry Delot and Sergio Ilarri

The importance of Vehicular Ad Hoc Networks (VANETs) is expected to increase significantly in the next few years: According to a recent survey by ABI Research, about 62% of new vehicles will be equipped with vehicle-to-vehicle (V2V) communications by 2027. This opens up a wide range of opportunities to develop different types of services for drivers, related to safety, entertainment, and information. For example, vehicles can communicate with to others information about events such as accidents, traffic jams, emergency vehicles requiring right of way, the availability of parking spaces, etc.

The use of short-range V2V communications to transmit data to neighbouring vehicles offers several advantages: they are free of charge (facilitating the cooperation needed to establish data sharing mobile ad hoc networks), they enable quick interactions among vehicles (avoiding the use of remote centralized servers), they are suitable for geographic routing (delivering messages to an area rather than to specific vehicles), they do not require the deployment of a support infrastructure, they favour privacy (by avoiding data centralization and limiting the spatial scope where messages are delivered), etc. However, the typical V2V communications of vehicular networks also present signifi-

Figure 1: Searching for a parking space

Figure 2: Analyzing a navigation route and nearby events

cant challenges for the data management community, requiring new data management and query processing techniques [1]. For instance, since two vehicles that want to exchange information will usually only be within communication range for a short time, a procedure is needed to limit the amount of data to be transmitted. To this end, evaluating the relevance of data is a key. This is also important to avoid distracting the driver by presenting him/her irrelevant information.

In our research, we focus on the development of the system VESPA (Vehicular Event Sharing with a mobile P2P Architecture) [2, 3], designed for vehicles to share information in inter-vehicle ad-hoc networks. We started this project in 2007 as a joint collaboration between the LAMIH research center at the University of Valenciennes (France) and the Distributed Information Systems (SID) group at the University of Zaragoza (Spain), and the project also now has collaborators in other institutions, such as Inria Lille Nord Europe and Télécom SudParis (France). Since then, we have developed several prototypes for smartphones and evaluated our approaches using simulations.

We focus on several data management issues. First, we consider the problem of evaluating the relevance of data items exchanged in the network (to decide which data should be disseminated and which data should be shown to the driver), using techniques based on geographic computations in the Euclidean space as well as techniques exploiting road network information available in digital road maps. Second, we propose suitable data dissemination protocols based on the relevance of the data items. Third, we propose data aggregation approaches that can help the vehicles to extract and share knowledge about the environment by summarizing data items (e.g., to identify areas where it is easier to park or areas where accidents are more likely). Fourth, we propose solutions for the problem of sharing information about scarce resources such as available parking spaces; this is important because communicating the information to many vehicles without control could easily lead to fruitless competition among drivers for the same resource. Fifth, we tackle the problem of processing queries that need to access remote data; in relation to this topic we have analyzed the challenge of routing the query results to the vehicle submitting the query and we are currently studying the potential of mobile agent technology for routing queries and results to the appropriate nodes. The exchange of multimedia information (e.g., to provide pictures about accidents or parking spaces to other vehicles) is also within our research interest. Finally, we are studying the potential benefits of applying other data management and semantic techniques. We envision a future where intelligent vehicles will be equipped with V2V communications and perform multiple data management tasks.

Links:

Website of VESPA: http://www.univ-valenciennes.fr/ROI/SID/tdelot/vespa/.

V2V Penetration in New Vehicles to Reach 62% by 2027, ABI Research: http://www.abiresearch.com/press/v2v-penetration-in-new-vehicles-to-reach-62-by-202 [Last access: May 8, 2013].

References:

[1] T. Delot, S. Ilarri: "Introduction to the Special Issue on Data Management in Vehicular Networks", Transportation Research Part C: Emerging Technologies, 23: 1-2, Elsevier, 2012

[2] N. Cenerario, T. Delot, S. Ilarri: "A Content-Based Dissemination Protocol for VANETs: Exploiting the Encounter Probability", IEEE TITS, 12(3): 771-782, 2011

[3] T. Delot, S. Ilarri, N. Cenerario, T. Hien: "Event Sharing in Vehicular Networks Using Geographic Vectors and Maps", Mobile Information Systems, 7(1): 21-44, IOS Press, 2011.

Please contact:

Thierry Delot
University of Valenciennes, France
E-mail: tdelot@univ-valenciennes.fr

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

Intelligent Vehicles: Complex Software-Based Systems

by Javier Ibanez-Guzman and Christian Laugier

Motor vehicles are becoming complex-networked mobile computers. Modern vehicles include numerous networked microprocessor-based Electronic Control Units (ECUs) ensuring multiple vehicle functionalities that include safety critical functions. As a comparison, a Boeing 787 Dreamliner requires about 6.5 million lines of software code to operate its avionics and on board support systems. By 2009, it was estimated that a luxury vehicle should run on close to 100 million lines of software code on 70 to 100 networked ECUs [1]. Currently, the rapid introduction of sensor-based driving assistance systems, digital maps for navigation plus vehicle connectivity means a rapid increase in the use of software for safety related functions leading to platforms under full computer control and hence autonomous driving [2].

The degree of motor vehicles' dependence on software is currently increasing exponentially. Components are sourced from 1st and 2nd tier suppliers, and are assembled by vehicle open end manufacturers (OEMs). Further, there are growing interdependencies between components. For example, an autonomous emergency braking system (AEB) is connected to a perception component, which operates using vehicle embedded data. The detected object is sent to a controller and the brake actuator is activated as well as the vehicle's longitudinal control. Different software dependent components interact and must operate concurrently; most of them originate from different suppliers, operate at different frequencies and handle different types of data. Consequently it can be difficult to integrate such functions. A vehicle OEM needs to ensure safe operation, which imposes important design constraints and a complex validation process. For safety related systems, performing full tests remains a challenge given the degree of risk involved in certain situations, for example, AEB applied to pedestrian safety.

Advances in intelligent vehicle technology should lead to vehicles under full computer control within the next few years. Most major vehicle OEMs have research programmes in this area. This potential was demonstrated by the technological push made by Google Inc. with their 'Google Cars' [2]. They have converted existing passenger vehicles into fully autonomous vehicles, demonstrating that commercial hardware for such a capability is available with autonomous capabilities attained through the implementation of advanced software [2]. Autonomous vehicles are 'systems of systems' where different sensor, actuation and decision making functions operate concurrently and react as a function of the environment the vehicle traverses. Figure 1 shows a typical architecture for the software-based components of an Intelligent Vehicle. It does not include all the mechanisms for controlling vehicle accessories or engine control, etc. Each of the systems shown is software dependent, most have to deal with different levels of uncertainty due to sensor limitations, vehicle response and incomplete information, as well as the actions of other entities sharing the vehicle work space. Intelligent Vehicles are only one component of much larger systems, namely Intelligent Systems where connectivity facilitates the interaction between these platforms and their ecosystem, e.g. road infrastructure sensors, traffic information, other transportation systems, etc. That is, there will be a continuum interaction with these elements through the 'Internet of Things' paradigm.

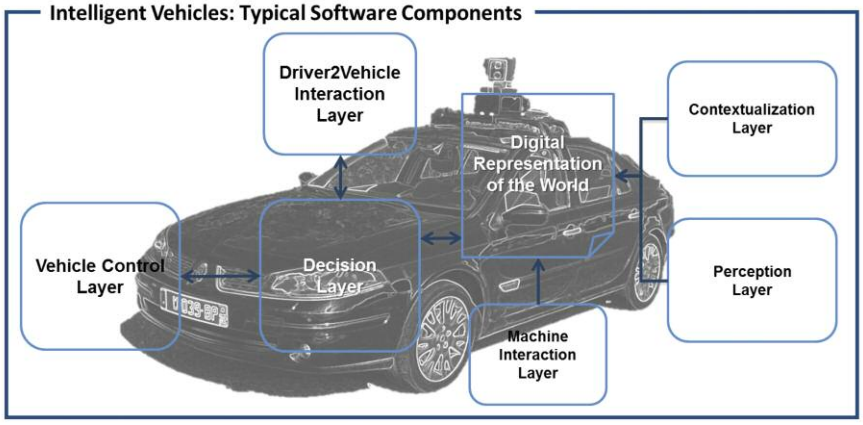

Figure 1 : Typical software components

As software becomes more complex, there is the need for systematic development and validation methods, and these need to be beyond the classical sensor-based and model-based approaches. There is the need to guarantee high integrity levels in a theoretical manner, that is to use analytical and formal methods that ensure that computer controlled vehicles are inherently safe.

Vehicles operate in uncertain environments: despite purpose-designed road infrastructure and traffic laws, unexpected situations occur resulting in accidents which are mainly caused by human driver error. Sensor-based algorithms are used to reduce driver error and algoithms for autonomous driving handle unexpected events when driving in standard traffic conditions. To enhance safety, the risk of an accident occurring has to be evaluated by the on board decision-making mechanisms. This can be estimated by identifying the difference between the driver's intention and expectation of the driver's

reaction with respect to the spatio-temporal relationship between the subject vehicle and other entities sharing the same road segment.

For example, when addressing risk assessment at road intersections to command driving assistance systems, a software-based solution would use data available in current standard vehicles and navigation systems as well as the sharing of information between vehicles using V2V wireless communications links [3]. The problem is formulated as a Bayesian inference problem, the solution demonstrated experimentally and through simulation means [3]. However, it remains difficult to demonstrate that the proposed solution is inherently safe. That is, unlike other domains, it is difficult to make any theoretical demonstration regarding the attainable level of safety.

If Intelligent Vehicles are to become computer controlled, guaranteeing their reliability and robustness is a major challenge. To lessen the effects of software complexity the automotive industry has formed the "Automotive Open System Architecture" (AUTOSAR), a consortium whose purpose is to standardize basic software functionality, leverage scalability to multiple platforms, ensure software transferability, etc. There are also associated norms that address partially different issues. However, there is a very strong need for software oriented validation methods and underlying theories. If vehicles are to run autonomously, their validation on punctual use cases and extensive testing might not suffice.

Vehicle connectivity enables the sharing of information between vehicles and with the infrastructure. It allows interaction through various communication channels with other services and with the Internet. However, it makes vehicles vulnerable to hacking, spoofing, etc. Thus, the development of encryption methods under the constraint of preserving user privacy is currently underway.

References:
[1] R.N. Charette: "This Car Runs on Code", in IEEE Spectrum, February 2009, http://spectrum.ieee.org/green-tech/advanced-cars/this-car-runs-on-code, accessed May 2013
[2] J. Ibanez-Guzman, C. Laugier, J-D Yoder, S. Thrun: "Autonomous Driving: Context and State-of-the-Art, in Handbook of Intelligent Vehicles", Ed. A. Eskandarian, pp 1271-1310, Springer, New York 2011
[3] S. Lefevre, J. Ibanez Guzman, C. Laugier: "Risk Assessment at Road Intersections: Comparing Intention and Expectation", in proc. IEEE Intelligent Vehicles Symposium IV2012, Alcalá de Henares, June 2012.

Please contact:
Javier Ibanez-Guzman
Tel: +33 1 768 57576
E-mail:
Javier.ibanez-guzman@renault.com

Modularity Analysis of Automotive Control Software

by Yanja Dajsuren, Mark G.J. van den Brand and Alexander Serebrenik

A design language and tool like MATLAB/Simulink is used for the graphical modelling and simulation of automotive control software. As the functionality based on electronics and software systems increases in motor vehicles, it is becoming increasingly important for system/software architects and control engineers in the automotive industry to ensure the quality of the highly complex MATLAB/Simulink control software. For automotive software, modularity is recognized as being a crucial quality attribute; therefore at Eindhoven University of Technology in the Netherlands we have been carrying out industrial case studies on defining and validating the modularity of Simulink models.

This research is part of the Hybrid Innovations for Trucks (HIT) project, which is carried out in a consortium of automotive manufacturer, suppliers, and research institutes, and aims to deliver a significant contribution in realizing up to 7% CO_2 emission reduction and up to 7% fuel saving for long-haul vehicles. Achieving this goal necessitates definition of proper quality techniques to enable the development of more efficient control software. In conventional vehicles, software is used in the multimedia, comfort, and safety systems. However, in (hybrid) electric vehicles, more complex control software for engine, after-treatment, and energy management systems for reducing fuel consumption and harmful emissions are being developed. Automotive control software is being developed using model-based design tools like MATLAB/Simulink and Stateflow. Simulink models can consist of several dozens of subsystems, hundreds of building blocks, and many hierarchical levels, which result in complex and difficult to maintain models. Although there are existing techniques and tools such as Mathworks Automotive Advisory Board (MAAB) guidelines and Model Advisor from Mathworks available to ensure guideline conformance and correctness, techniques for assessing quality of Simulink models are limited.

For automotive software, modularity is recognized as being paramount since changing or reusing non-modular software is very costly. MathWorks provides quality related tools like "Modeling Metric Tool" and "sldiagnostics" to quantitatively measure the content of a Simulink model (Stateflow model as well), to improve the productivity and quality of model development, e.g. model size, complexity, and defect densities. Other quality metrics e.g. for measuring instability, abstractness, and complexity of Simulink models have been intro-

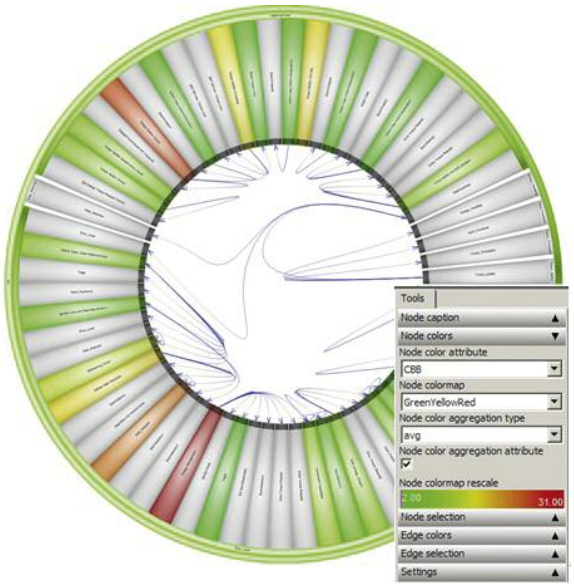

Figure 1: Visualization of system modularity with the help of SQuAVisiT tool [2]. This tool gives early feedback about system modularity, making it cheaper and easier to reuse and maintain than traditional techniques. Illustration: TU/e.

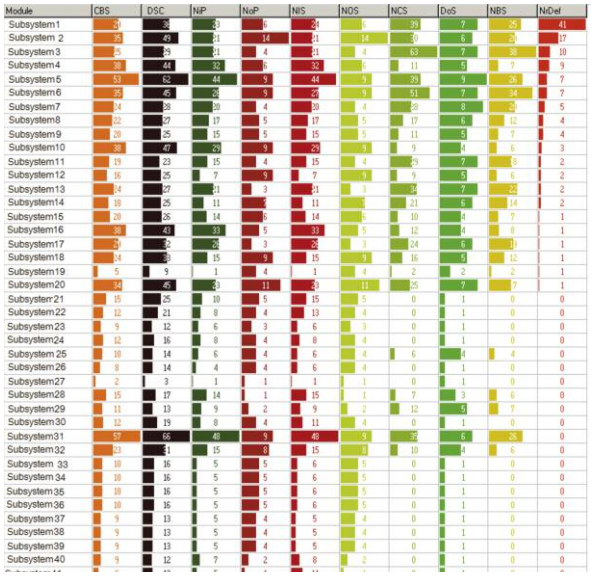

Figure 2: Visualization of modularity metric values and number of faults of the subsystems. This can help the domain experts locate easily the most problematic subsystems and their respective metric values. Illustration: TU/e.

duced, however the validation of the metrics is not provided. Therefore, we focused on modularity as one of the main quality attributes for Simulink models. According to the ISO/IEC 25010 international quality standard, modularity is defined as the degree to which a system or computer program is composed of discrete components such that a change to one component has minimal impact on other components. Following the Goal-Question-Metrics (GQM) approach, we defined a modularity metrics suit consisting of nine coupling and cohesion-related metrics for Simulink models and validated it with experts' evaluation [1].

We identified three independent metrics based on the statistical analysis and identified a relation between modularity metrics and presence of errors. We have observed that high coupling metric values and high number of hierarchical levels (subsystem depth) frequently correspond to presence of errors. We developed a Java tool to measure the defined metrics on Simulink models. The tool uses a Java parser for Simulink MDL files of the ConQAT open-source tool. We are extending our tool to measure not only modularity metrics but also other key quality attributes such as complexity, reusability, testability, and maturity of the Simulink models and extend the visualization of software quality with SQuAVisiT (Software Quality Assessment and Visualization Toolset) toolset [2].

Currently, our metrics tool interfaces with the SQuAVisiT toolset as it is a flexible tool for visual software analytics of multiple programming languages. Figure 1 illustrates an example visualization of the modularity aspect of the industrial application that we studied (the subsystem names here are blurred due to confidentiality reasons). Subsystems of a Simulink model are illustrated as nested rectangles in the outer rings of the radial view. The relations between (basic) subsystems, such as input and output signals, are shown as curved arrows in blue. The colours on subsystems are used to visualize values of modularity metrics. The green subsystems show the modular and red ones show subsystems which require attention to improve their modularity. Figure 2 illustrates another example of quality visualization. The first column displays the list of subsystems and the rest of the columns list all the metric values and last column lists the number of faults of the subsystems. Subsystems are sorted by descending number of fault value. This can help the domain experts locate easily the most problematic subsystems and their respective metric values. However, the effectiveness of using different visualizations of Simulink model quality remains a future work.

For more than a decade automotive Architecture Description Languages (ADLs) have been considered as one of the main solutions to represent system and functional architecture and to facilitate communication between different stakeholders. Based on our evaluation of a set of (automotive) ADLs like EAST-ADL, AADL, and SysML against automotive specific modelling requirements, definition of automotive architectural quality is still an open issue [3]. To this end, further work is needed to investigate quality metrics for automotive architectural models.

Link: http://www.win.tue.nl/~dajsuren

References:

[1] Y. Dajsuren, M.G.J. van den Brand, A. Serebrenik, and S. Roubtsov: "Simulink models are also software : Modularity assessment", proc. of the Ninth International ACM Sigsoft Conference on the Quality of Software Architectures (QoSA 2013), Vancouver, Canada, 2013.

[2] M. G. J. van den Brand, S. Roubtsov, and A. Serebrenik: "SQuAVisiT: A flexible tool for visual software analytics", in proc. of the European Conference on Software Maintenance and Reengineering, pp. 331-332, Washington, DC, USA, IEEE Computer Society, 2009.

[3] Y. Dajsuren, M. G. J. van den Brand, A. Serebrenik, and R. Huisman: "Automotive ADLs: a study on enforcing consistency through multiple architectural levels", in proc. of Quality of Software Architectures (QoSA), pp. 71–80. ACM, 2012.

Please contact:

Yanja Dajsuren
Eindhoven University of Technology, The Netherlands
Tel: +31 40 247 5052
E-mail: y.dajsuren@tue.nl

Electro Vehicle Car-Sharing System – A Simulation Model for Big Cities

by Dmitry Rozhdestvenskiy and Petr Bouchner

We address the problem of the introduction of car sharing systems for electric vehicles within big cities through simulation of such systems. Simulation is an important step in the early stages of the development of new technologies. In contrast to standard simulation techniques, this project deals with simulation through introduction of a virtual online system and building up a model of real consumers' behaviour based on their opinions from "pseudo real" experience.

Simulation is a strong tool which allows researchers to test a hypothesis with different initial conditions and at the same time, in the case of a project with a big outlay, it helps management staff to understand the feasibility of a proposed solution without risking loss of investment capital.

Electric Vehicle (EV) Car Sharing (CS) is such a case. Even a small project with 36 stations and 72 EV in operation requires an investment of more than €26 million (estimated by authors for Prague city), so the question of the popularity of the service and the subsequent load (amount of users) for each station becomes crucial.

Online simulation with virtual experience offers a promising solution. To predict and simulate the human decision-making process and to predict the potential number of customers, we need a strong statistical background. This project offers a different approach: instead of dealing with soft systems and trying to simulate human behaviour, we are letting individuals make their own decisions through a "virtual" service, available on the Internet. This service provides a strong background for analysis and serves as an input for further offline simulations. At the same time it should help inhabitants and visitors of big cities, such as Prague, to understand the benefits of this proposed service and become familiar with EV technology, possibly reducing the impact of the phenomenon of range anxiety.

The simulation software was created with a predefined price policy and predefined CS distribution over the region. This information was obtained from a study that had been performed previously at Czech Technical University [1]. It is possible to update the software to reflect different price models.

Figure 1: http://elektromobilita.fd.cvut.cz/ web interface.
a) Login page
b) User input acquisition
c) Route results visualization
d) Simulation of the movement
e) Details of the calculation
f) Virtual questionnaire

To obtain relevant data about consumer decision making in a broader transport context, the simulation offers locally available modes of transport within a selected region, including public transport and taxi services, and calculates time consumption and price for all of them. "Walk on foot" is included as a transport option.

The energy consumption model is introduced to calculate power demand on the infrastructure and developed from the known principles of mechanics and the vehicle parameters. It combines friction resistance, air resistance, acceleration resistance, power required to overcome road slope, electrical loads/losses from "all other" electrical loads in the vehicle from user-related systems, such as climate control, external lights, and entertainment sound systems. Within the calculation, the consumption model takes into account current weather forecast, elevation profile and velocity profile of the route (start/stop at the intersections). The velocity profile is calculated based on the response from Google direction service and depends not only on the maximum speed within a part of the route but also on the individual's driving style which is determined based on the information provided by consumer.

Simulation provides the user with the ability to set the desired location and destination address (if needed, waypoints can be added) through a GUI (Figure 1). Based on this, together with external information, it calculates price and time consumption for three modes of transportation for the defined route. Simulation software also calculates and displays the state of all vehicles in the system (battery charge, odometers etc.) and represents their current locations. If a vehicle is currently in use by someone else its movement is displayed on a map.

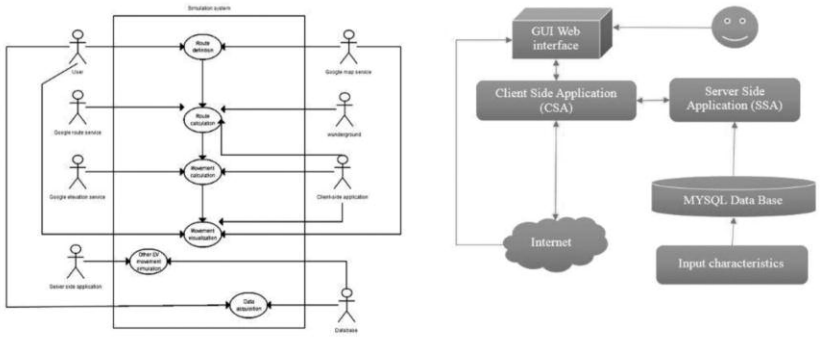

www.elektromobilita.fd.cvut.cz
(a)Use case Diagram (b) Rich picture

Figure 2: Use case diagram and rich picture of proposed simulation.

After the calculation is finished, the customer has to choose one mode of transportation. When this decision is made, information about the route is sent and saved in MySQL database. This can be used later for offline simulation.

If a customer chooses the EV car sharing solution, simulation of the movement based on real-time motion planning is displayed on the web interface (in the future it will be possible for the movement to change dynamically during simulation). This provides a real time simulation of EV movement along the route including "windshield" view, indications of the EV's power consumption, its current speed and route instructions (car navigation style). After the simulation finishes the customer is informed about details of his/her journey via set of performance graphs (Figure 1, bottom left).

If customers choose public transport, they are provided with detailed information about the route in textual form which is also represented on the map. If customers choose a taxi, they are simply redirected to a web page through which they can order a taxi. A use case diagram and rich picture of the proposed simulation is represented in Figure 2.

After the simulation software was launched, several tests were performed to understand its functionality and to determine the response of the EV consumption model to different initial conditions such as weather, elevation profile and driver behaviour. The slope of the road has the biggest influence, with an increased slope resulting in up to a 30% increase in energy consumption for a path of a given distance and a given number of crossings. At the same time, weather and driver behaviour can have an influence of up to 10%.

The software is stable and ready for mass usage, but there are still some tasks to be done in further studies, such as the introduction of proper battery charge models, dealing with precision of elevation profiles, introduction of an offline simulation of customer behaviour based on existing data, and an expansion of the population of test drivers, including potential city visitors.

Link:
http://elektromobilita.fd.cvut.cz/index (for Google Chrome)

References:
[1] M. Lapáčková: bachelor thesis (in Czech), Faculty of Transportation Sciences, Czech Tech. Univ. in Prague, 2012
[2] C Walsh et al: "Electric Vehicle Driving Style and Duty variation Performance Study"
[3] R.A. Barnitt et al: "Analysis of Off-Board Powered Thermal Preconditioning in Electric Drive Vehicles"

Please contact:
Dmitry Rozhdestvenskiy and Petr Bouchner, Czech Technical University in Prague, Faculty of Transportation Science, Department of Vehicles, Czech Republic, Prague,
E-mail: rozhddmi@fd.cvut.cz, bouchner@lss.fd.cvut.cz

The Car that Looks Like Me: Similarity Cues can Increase Trust in the Self-Driving Cars of the Future

by Frank Verberne, Jaap Ham and Cees Midden

In the future, cars will be better at driving themselves than humans are. These cars will be useless, however, if humans do not trust them. Researchers from Eindhoven University of Technology, the Netherlands, are studying how to increase trust in automation technology in cars. In a human-technology interaction, we investigate whether technology can use similar trust enhancing mechanisms to those used by humans to build trust in each other, such as emphasizing similarity. Will technology that thinks, acts, and looks similar to its user will be trusted more than its non-similar version?

In the near future, fully autonomous cars will be widely available and will be able to drive more safely, at greater speeds, and more sustainably than any human can. However, if most drivers refuse to relinquish control, perfectly capable cars will be gathering dust, due solely to the mismatch between technological possibilities and technological acceptance.

Research has shown that people need to have sufficient trust in technology before they are willing to use it [1]. At Eindhoven University of Technology we study how to increase trust in these

technologies. While some scholars focus on experience with new technologies to increase confidence in them, we focus on using mechanisms that humans also use, such as emphasizing similarity to another person. Therefore, we create similarities between technologies and their users to increase trust even before the actual use of the technologies. More specifically, three types of similarity cues are used: cognitive, behavioural, and appearance similarity.

In a first set of studies, the effect of cognitive similarity on trust was investigated. Participants were first asked to rank four driving goals (energy efficiency, speed, comfort, and safety) from one to four, one being the most important goal for them, four being the least important. There was no optimal ranking suited for everyone, every participant chose their own ranking. Then they were presented with a description of an Adaptive Cruise Control system that either shared or did not share their driving goal ranking. In the latter case, the system had the reverse ranking of that of the participant. Participants had to indicate their trust and acceptance of the system. Results showed that participants trusted the system more when it shared their ranking versus when it did not [2]. Thus, sharing goals (cognitive similarity) leads to greater trust in an Adaptive Cruise Control.

In a second set of studies, the effect of behavioural similarity on trust was investigated. In human-human interactions, humans (unconsciously) mimic each other's body posture. This mimicry enhances liking and strengthens bonds between people, even between strangers [3]. In this set of studies, we studied whether mimicry could be used by a virtual agent to increase trust in the agent. Such an agent could appear in a display in future self-driving cars, to provide a digital face to the automation technology driving the car. In our lab, participants played a risky game with a virtual agent that either mimicked or did not mimic them. Participants' head movements were measured using an orientation tracker, and in the mimicry condition, participants' own head movements were copied by the virtual agent with a delay of several seconds. In the non-mimicry condition, the virtual agent used the head movements of the previous participant. In the risky game, participants were presented with ten different routes and had to decide to either plan the route themselves (safe option), or let the agent plan the route for them (risky option). Results showed that participants gave more routes to a mimicking agent versus a non-mimicking one. Thus, mimicry (behavioural similarity) leads to greater trust in virtual agents.

Figure 1: PhD student Frank Verberne testing a driving simulator for future experiments (photo: Bart van Overbeeke).

In the current set of studies, the effect of appearance similarity on trust is being investigated. Previous research has shown that people had greater trust in an individual, based on a photo, when that photo was morphed with their own photo [4]. That is, individuals whose faces look similar to a person's own face are trusted more. In a first study, we investigate whether a virtual agent with a face similar to that of the participant is trusted more by that participant than a virtual agent that has a dissimilar face. First, we took pictures of all our participants and we used those pictures in FaceGen (a face modeller program) to create a 3D virtual head of all participants. Next, one standard head was morphed with all the participants' virtual heads with a 50-50 blend. Then, participants played risky games with a virtual agent that had a custom head. For one half of the participants the custom head was their own 50-50 blend (containing 50% of their own virtual head, self similar), for the other half the custom head was someone else's 50-50 blend (containing 50% of the virtual head of another participant, self non-similar). We expect that participants trust self-similar agents more than self-non-similar agents. The study is currently in progress. In future studies, we plan to extend the effects of these similarity cues to risky scenarios in a driving simulator.

Working together with DAF and Delft University of Technology, we will be able to test our ideas in more realistic settings in the future. The results of this line of research so far suggest that in the future, self-driving cars could use similarity to their users to gain their trust and persuade them to hand over the steering wheel.

Link:
http://hti.ieis.tue.nl/node/3344

References:
[1] J.D. Lee, K.A. See: "Trust in automation: Designing for appropriate reliance", Human Factors, 46, 2004
[2] F.M.F. Verberne, J. Ham, C.J.H. Midden: "Trust in smart systems: Sharing goals and giving information to increase trustworthiness and acceptability of smart systems in cars", Human Factors, 54, 2012
[3] T.L. Chartrand, J.A. Bargh: "The Chameleon effect: The perception-behavior link and social interaction", JPSP, 76, 1999
[4] L.M. DeBruine "Facial resemblance enhances trust" Proc R. Soc. Lond. B., 2002.

Please contact:
Frank Verberne
TU/e, The Netherlands
Tel: +31402475250
E-mail: f.m.f.verberne@tue.nl

Digital Road Authority for Coordination between In-Car Navigation Systems and Traffic Control Centres

by Rob van der Mei, Maaike Bots and Frank Ottenhof

The use of in-car technology and real-time traffic information has been growing spectacularly over the past few years, and this growth is not likely to come to an end in the near future. However, despite their enormous benefits for our modern society, today's navigation systems take into account important safety and environmental aspects to a limited extent at best. Moreover, the enormous diversity of the traffic-management applications and in-car systems available today often leads to contradictory routing advices and in many cases even to network instability. This raises the need for proper coordination between in-car navigation systems and traffic-control centres. CWI is developing data fusion models to estimate highway traffic. In the future these can be used as input for traffic forecasts.

A powerful means to coordinate the wide variety of traffic data is to implement so-called digital road authorities (DRAs). A DRA is a virtual traffic manager that monitors the performance of the road network, and undertakes action when needed. DRAs are often implemented in large-scale, congested areas with many actuators and sensors. To use DRAs, the road network is typically sub-divided in regions, each of which is managed by a dedicated regional DRA. These regions, in turn, are subdivided in smaller areas. In this way, the road network can be seen as a virtual tree of national networks, regional networks, area networks, and so on. Similarly, these networks are supervised by digital road supervisors, digital road region supervisors and digital road area supervisors.

DRAs may be coupled to all measurement and control systems within their network, not only to roadside equipment (like sensors, ramp meters, cameras, radars, detection loops) and VRI's, but also on in-car systems and applications. The DRA constantly monitors whether the requirements on reachability, safety, pollution and traffic throughput are met, and undertakes action when needed. To this end, the DRA for example may ask its neighbouring DRAs to decrease the influx of traffic. In case of overload in the whole network, it may occur that the neighbouring DRAs cannot help out. In that case, the DRA may contact a next higher layer DRA. And so on.

A successful example of DRA technology is the TrafficLink software [1] offered by Trinité Automation [2].

Societal importance and emergency situations

The proper use of DRA technology has far-going societal consequences. DRAs not only provide effective means to relieve road congestion and pollution, but may even save lives in emergency situations where every second counts. In the emergency medical services, like ambulance services, coupling dedicated sensor-based emergency apps to DRAs may (1) be an effective means to determine the fastest route to the emergency scene, by automatically switching to coordinated traffic light schedules in order to minimize travel time over a given route, (2) provide a means to detect road accidents more promptly, saving valuable time, and (3) enable car drivers on the same route to receive warning messages via their telephone or in-car system.

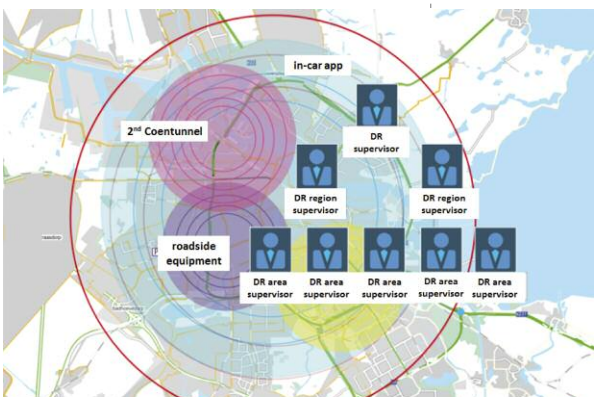

Figure 1: Illustration of the hierarchy of DRAs and heterogeneity of traffic-monitoring systems

Prediction of traffic jams

The use of DRAs may also provide a way to predict, and hence take proper action to avoid, traffic jams. For example, by coupling weather forecasting services to DRAs, the occurrence of road congestion in specific areas may be predicted, and dynamic rerouting and coordination of traffic light schedules may be enforced to provide relief.

Research

In order to fully exploit the enormous possibilities offered by DRA to enhance mobility and save lives, a variety of research challenges are to be faced, including the following:

1. how to accurately estimate the traffic on highways and local roads for different types of data from different sources, like including detectors, cameras, Bluetooth- and GPS-devices,
2. how to accurately estimate travel times from A to B on the basis of the heterogeneity of data from different sources, and
3. how to proactively manage traffic on the basis of traffic forecasts to relieve/avoid congestion.

To cope with the heterogeneity of data sources, we are currently developing and validating quantitative models using techniques from data fusion in order estimate traffic and congestion on highways and local roads. These models serve as an input to travel-time models which, in turn, are used to develop dynamic congestion control mechanisms based on traffic forecasts.

References:

[1] http://www.trafficlink.nl
[2] http://www.trinite.nl.

Please contact:

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

Frank Ottenhof
Trinité Automation, The Netherlands
E-mail: f.ottenhof@trinite.nl

Research activities and innovative developments in European research institutes

Imperia: The 21st Century Eco-Friendly Sports Car

by François Henry

The Imperia is back on track in the 21st century, thanks to its intelligent hybrid propulsion management developed by Green Propulsion, a spin-off company from Université de Liège, Belgium. Clever electronics gives the new Imperia GP roadster outstanding sports car performance with a carbon footprint of 50 g/km.

Founded in 2001, Green Propulsion, a company that specializes in developing clean vehicles, has quickly become a leader in its field. The revival of the Imperia brand represented a unique opportunity to create a lightweight, ultra-high-performance hybrid car from scratch.

After a half-century in the oblivion, the Belgian make was revived by several motor sports aficionados who helped Green Propulsion shine up the glorious image of what remains of the last Belgian make of cars.

The PowerHybrid concept

The Imperia GP roadster is motorized by the PowerHybrid system, controlled by the HMS (Hybrid Management System) electronic brain which continuously balances energy transfer between the thermal engine, electric motor, battery pack, and wheels.

The power train is based on three energy components: a four cylinder 1.6 litre direct injection turbo engine generating 147 kW (200 bhp), a 110 kW (150 bhp) axial flux electric motor, and a lithium-polymer battery pack that stores 11 kWh. The energy sources are a classical fuel tank and a 3 kW battery charger.

Starting on electric power only, its range is some 60 km, but the Imperia's range is unlimited the moment it switches to the thermal engine once the batteries are drained. In hybrid mode, the batteries can also be recharged preventively when the vehicle is in motion. While driving, the driver can at any time select the type of propulsion required (electric or hybrid) by flicking a switch. Using another control, the driver decides whether to run the batteries down completely (so as to recharge them on the mains at the destination) or to keep a reserve for use upon entering urban traffic.

PowerHybrid

The innovative hybrid management strategy minimizes total CO_2 emissions from well to wheel. To compare pollution from thermal and electrical power sources, one must consider emissions from each stage of the energy supply system. Decisions regarding operating modes and battery charging automatically result from this minimization strategy.

For the thermal engine, this means considering not only direct emissions from the exhaust pipe but also fuel extraction, transport and refining emissions. A total -direct plus

Figure 1: The Imperia GP roadster at the Francorchamps race track

indirect- CO_2 emissions map has been pre-calculated from the engine's fuel consumption.

For the electric part, each electron taken out of the battery will need an equivalent charge at the end of the trip. Battery drain depends not only on the electric motor demand but also on many other conditions, such as battery state of charge and temperature. The charging AC power needed at the mains is proportional to the battery drain multiplied by an efficiency factor. An average CO_2 emission factor from the electricity production site is determined (in kg CO_2/kWh) per country. For the fuel emission computation, secondary emissions from the power plant attributed to fuel extraction, transport and refining are added.

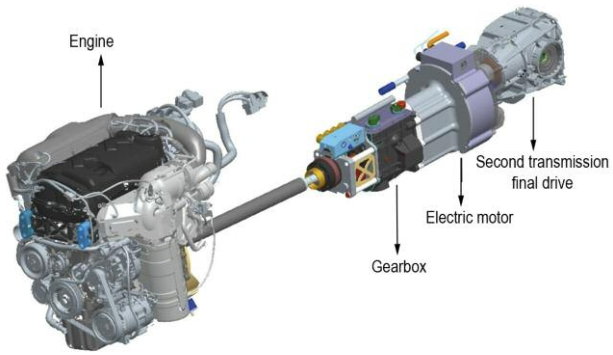

Figure 2: The Imperia hybrid powertrain

In the vehicle, a CO_2 emissions map is calculated for both motors. The emissions function in line with the vehicle power request is thus known at any vehicle speed and for each motor. Based on this function, a proprietary optimization algorithm determines the ideal power split between the fossil fuel engine and the electric motor. The driver may override this automatic hybrid management system if s/he knows a recharging station is close enough to allow a higher battery discharge rate.

Results

The Imperia GP roadster is one of the first in its market segment to offer a hybrid plug-in mode: a car that can be recharged via a 220V power socket. A full charge is obtained in 4.5 hours, for an unbeatable cost of € 0.034 (daytime rate) or € 0.016 (night rate) per km.

The PowerHybrid design minimizes well-to-wheel pollution from primary energy. Consequently, the Imperia GP roadster emits very little CO_2 overall, and even zero direct emissions in electric mode. Following European regulation R101 for plug-in hybrids, the Imperia GP's mixed consumption -petrol and electricity- barely reaches 1.9 l/100km plus 11.5 kWh/100km and the CO_2 emissions are below 50 g/km. In a more rigorous well-to-wheel philosophy, the environmental aims are also completely fulfilled with carbon emissions of 110 gm/km well-to-wheel or around a third as much as the Porsche Carrera S with similar performance.

The first Imperia consumers will take delivery at the end of the car's first production year in December 2013. Visit www.imperia-auto.be for details.

Results according to European regulation R101:

- Gas Mileage: 1.9 l/100km
- Electricity consumption: 11.5 kWh/100km
- Carbon Emission: 50 g/km

References

Y. Toussaint, F. Henry: "PowerHybrid: An innovative powertrain concept uniting sportiness and environmental awareness", in proc. of European Electric Vehicle Congress 2012, Brussels, 2012.

J. Nzisabira, et al.: "Eco-efficiency methodology of Hybrid Electric Vehicle based on multidisciplinary multi-objective optimization", in proc. of ULG-LTAS, 8th World Congress on Structural and Multidisciplinary Optimization, 2009.

Please contact:

François Henry, Green Propulsion R&D/Test engineer
Tel: +32 42476011
E-mail: f.henry@greenpropulsion.be

An Innovative Active Suspension System for Autonomous Vehicles: A Safe and Comfortable Ride and Good Handling

by Katerina Hyniova

Between 2001 and 2012, the Josef Bozek Research Center of Combustion Engines and Automobiles at the Czech Technical University in Prague, Czech Republic, has been developing innovations in suspension technology. The research team has designed a unique suspension system that uses a linear electric motor as a controlled actuator and provides the desired forces between wheel and car body. Many experiments on energy management in the system have also been undertaken. In order to verify various control strategies and to test ways of optimizing energy consumption, a unique one-quarter-car test stand has been designed and constructed.

The suspension system is an important aspect of car design because it influences both the comfort and safety of passengers. Two major performance requirements of any automotive suspension system are to provide a comfortable ride and good handling when random disturbances from road unevenness and variable cargo act upon the running vehicle. Passenger comfort can be interpreted as attenuation of sprung mass acceleration or as peak minimization of sprung mass vertical displacement, while good handling can be characterized as attenuation of unsprung mass acceleration. Effort devoted to passive suspension design is ineffective because improvements to ride comfort are achieved at the expense of handling and vice versa. Instead, the best net result can be achieved by active suspension, i.e. an additional force (see Figure 1) can act on the system and simultaneously improve both of these conflicting requirements. Another important goal of the control design is to maintain robustness of the closed loop system.

We focused largely on one-quarter- and one-half- vehicle models as control objects [3], and a one quarter-car model was developed for controller design and simulation. The model (Figure1) consists of unsprung mass m_w (wheel) and sprung mass m_b(vehicle body taken as one quarter of the total body mass), conventional passive suspension system (spring k_1 and damper b_1) and active suspension force F_a placed in parallel to the passive suspension system. "k_t" represents tyre stiffness, z_r denotes road displacement (disturbance), z_b displacement of the body mass and z_w wheel displacement. Contrary to commonly used hydraulic or pneumatic actuators [1] the research team used a unique solution – a linear electric motor controlled to obtain a variable mechanical force F_a for the car suspension. The main advantage of such a solution is the ability to generate desired forces between the wheel and the car body, giving the car sprung mass good insulation from the road surface disturbances. We

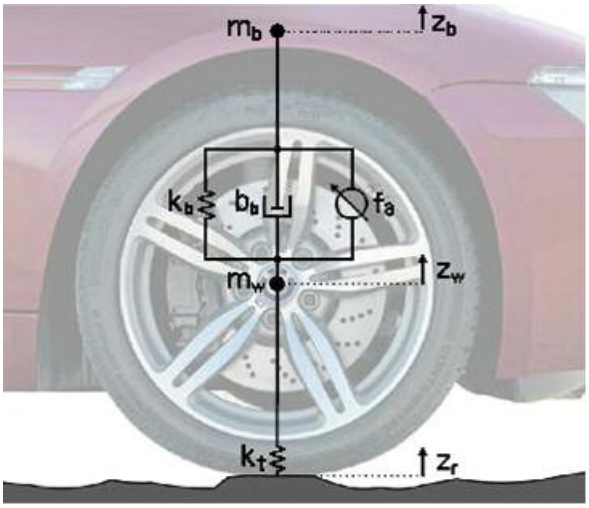

Figure 1: One-quarter-car model

learned that, under certain circumstances, the use of linear motors as actuators enables transformation of mechanical vibration energy to electrical energy, which can be accumulated in supercapacitors and used later when required. In this way, it is possible to reduce or even eliminate demand on the external power source. In contrast to traditional drives that use rotational electric motors and lead screws or toothed belts, the direct drive linear motor exhibits the property of contactless transfer of electrical power according to the laws of magnetic induction. The electromagnetic force is applied directly without the intervention of mechanical transmission. Low friction and no backlash result in high accuracy, high acceleration and velocity, high force, high reliability and a long lifetime, which enable not only effective usage of

Figure 2: Experimental test stand

modern control systems, but also facilitate efficient control of suspension vibration.

Design of an appropriate controller is a complex issue. The research team tested commonly used high authority control concepts, such as Linear Quadratic Gaussian control, feed-forward control concept, fuzzy control as well as H-infinity robust control. Most of the proposed concepts don't consider robustness of the controllers, although this aspect is fairly important and well-founded. For example, changes in ambient conditions can lead to changes in the dynamics of the vibration system, which can result in substantial deterioration of the feedback controller's performance.

The same one-quarter-car configuration has been used for real experiments and verification. The mechanical configuration of the test stand is shown in Figure 2: Under the tyre there is placed another linear electric motor that uses an input experimental signal to generate road displacement (road deviations) under the running wheel.

Various active suspension controllers have been developed via Matlab, implemented into dSpace and connected to the test stand system. As discussed in [2] and [3], the most satisfactory responses of the suspension system have been gained using H-infinity robust control.

This research has been supported by MŠMT project No.LN00B073 "Josef Božek's Research Center of Combustion Engines and Automobiles and by the project "INGO LG110039".

Link: http://www3.fs.cvut.cz/web/index.php?id=934

References:

[1] R.S. Sharp, et al.: "Road Vehicles Suspension Design – Review", VSD 16(3), 1987
[2] K. Hyniova et al.: "On Experiments Taken by The Active Shock Absorber Test Stand", AMECE'13, pp.105 – 109, 2012
[3] K. Hyniova, et al: "H-Infinity Controlled Actuators in Automotive Active Suspension Systems", ESDA, 2008

Please contact:

Katerina Hyniova
Dept. of Digital Design, Faculty of Information Technology
Czech Technical University in Prague
E-mail: hyniova@fit.cvut.cz

Wire-Speed Regular-Expression Scanning at 20 Gbit/s and Beyond

by Jan van Lunteren and Christoph Hagleitner

Network intrusion detection systems and emerging analytics applications for business intelligence rely on fast pattern-matching functions to scan data in real time. This type of scan operation has become increasingly challenging because of the growing number of regular expressions that have to be supported (e.g. to detect new types of intrusions) in combination with the continuous increase in network link speeds, which currently are on the order of multiple tens of gigabits per second.

At IBM Research – Zurich, together with our colleagues in Haifa and the US, we have designed a new programmable pattern-matching accelerator called RegX that is capable of simultaneously scanning thousands of regular expressions against multiple data streams at wire-speed for state-of-the-art network links. Millions of active streams can be scanned in an interleaved fashion by storing and retrieving the complete scanner state when switching between streams. The accelerator supports dynamic and incremental modification of the internal data structures, enabling updates to the pattern set without interruption of the scan operation.

The RegX accelerator combines the simple processing model of deterministic finite automata (DFA) with non-deterministic automata (NFA)-like mechanisms to realize high scan rates and obtain excellent storage efficiency. The latter is particularly important for dealing with the well-known "state-explosion" problem that occurs when certain combinations of regular-expression patterns with specific overlap conditions are mapped on the same DFA, resulting in extremely large data structures. One of these mechanisms is a lane concept that allows a flexible allocation of available scanner resources (programmable state machines) to a time-interleaved processing of multiple streams, and enables combinations of complex patterns that can cause state explosion to be separated by distributing them over multiple engines (a lane in this context comprises all the resources that are assigned to the scanning of a single data stream). A second mechanism is based on the extension of state machines with dedicated processing elements that allow the splitting of (individual) complex patterns into smaller pieces that can be scanned independently.

The micro architecture of the accelerator provides a range of flexible dispatch and execution options, allowing instructions to be executed, for example, upon the occurrence of a specific character or class of characters (e.g., a digit) in the input stream, upon the detection of one or multiple (sub)patterns, or upon fulfillment of a certain condition (e.g., for testing the distance between two detected (sub)patterns). Instructions are typically based on simple hardware primitives that are executed in a single cycle, and are exploited in a RISC fashion by an intelligent compiler stack to obtain high

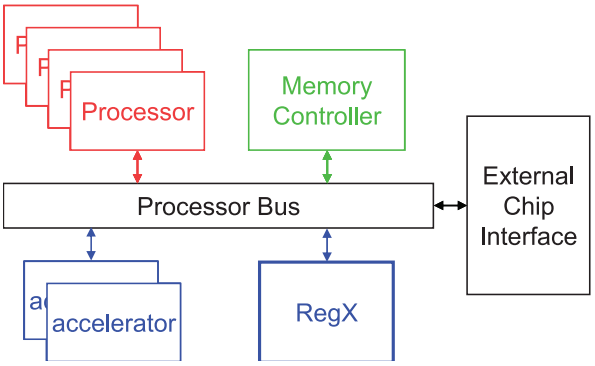

Figure 1: Processor-bus-attached RegX accelerator

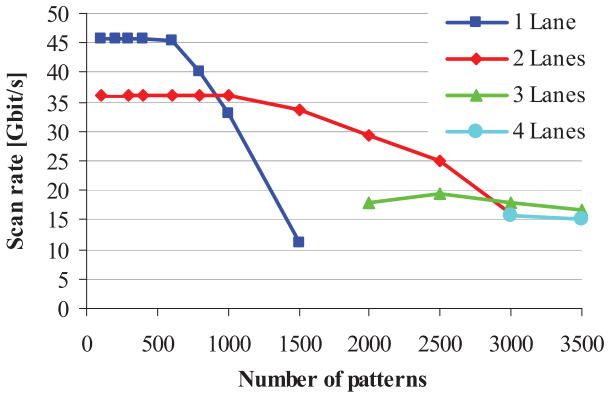

Figure 2: Measured scan performance for various lane configurations and pattern counts

scan performance. RegX employs a special cache concept that is adapted to the characteristics of pattern-matching workloads to obtain hit rates of 99% and higher. This concept involves an L1 cache that is partially managed in hardware in a conventional way and partially managed by a software application that exploits hardware-based profiling to optimize the placement of the cache lines.

RegX has been implemented as part of the IBM Power Edge of Network™ (PowerEn) system-on-a-chip (SOC) in 45-nm SOI technology, consuming 15mm^2 out of the total chip area of 410mm^2 and running at a clock frequency of 2.3 GHz. It has been integrated as a processor-bus-attached accelerator providing high-speed pattern-matching functions to applications executed on general-purpose cores as shown in Figure 1. A unique feature of RegX that sets it apart from other pattern-scanning engines is its unparalleled memory efficiency that allows it to handle much larger and more complex regular-expression sets at a constant scan rate that is independent of the input characteristics when operating out of dedicated memory (128 KB per scan lane, 512 KB in total). This renders RegX less susceptible to denial-of-service attacks than other scanning accelerators. The accelerator design achieves a peak scan rate of 9.2 Gbit/s for single streams, and a theoretical aggregate scan rate of 73.6 Gbit/s when processing 8 streams in parallel. The latter cannot be entirely reached on the PowerEn™ implementation due to limitations on the data transfer bandwidth to the RegX accelerator, which reduces the peak scan rate to about 50 Gbit/s.

The RegX accelerator supports pattern collections that exceed the size of the dedicated memory, by storing part of the data structures in main memory, which, however, involves a significant access penalty on the order of 400 cycles. The hardware/software managed cache concept guarantees in this case a graceful degradation of the scan performance. This is illustrated in Figure 2, which has been adapted from [1] and shows measured scan rates for configurations involving between one and four scan lanes and for pattern sets of varying sizes. The initial flat portions of the graphs correspond to the case that the (compiled) patterns fit into the dedicated RegX memory, resulting in a constant scan rate. The "graceful" declines in scan rates occur when the number of patterns exceed the size of the dedicated memory and a part of the data structure has to be stored in the main memory. By increasing the number of scan lanes that process a given input stream, the total amount of dedicated memory that is available for a scan operation can be increased, so that larger numbers of patterns can be scanned at a constant flat rate. The latter, however, comes at the cost of a reduced aggregate scan rate because fewer different data streams can be scanned in parallel when multiple scan lanes operate on a single stream. For typical network intrusion detection workloads, scan rates were measured on the PowerEn™ hardware in the range from 15 to 40 Gb/s.

Reference:

[1] J. van Lunteren et al: "Designing a Programmable Wire-Speed Regular-Expression Matching Accelerator," Micro-45, Vancouver, CA, December 2012, pp. 461-472

Please contact:

Jan van Lunteren, IBM Research GmbH, Switzerland
E-mail: jvl@zurich.ibm.com

DRMSim: A Routing-Model Simulator for Large-Scale Networks

by Aurélien Lancin and Dimitri Papadimitriou

The expansion of Internet topology, which comprises thousands of Autonomous Systems (AS), has resulted in a number of important research challenges. The Border Gateway Protocol (BGP), which is used to make core routing decisions on the Internet, starts to show its limitations in terms of the number of routing table entries it can store locally, update in a timely fashion and dynamically exchange. Because it is impractical to deploy newly designed routing protocols on the Internet a large-scale, simulation is an unavoidable step to validate their properties. However, the increasing routing information processing (CPU) and storage (memory) introduces similar challenges for the simulation of state-full routing. For this purpose, we introduce DRMSim a Dynamic Routing Model simulator of routing models on large-scale networks.

Today's Internet routing system relies on the Border Gateway Protocol (BGP). This path-vector routing protocol that matches both technical and economic aspects of Internetwork routing requires storage and timely updates of individual routing states at each node, i.e., it belongs to the class of adaptive stateful routing. The Internet backbone currently consists of more than 45,000 abstract nodes called Autonomous Systems (AS) and about 450,000 nodes called routers (part of the default-free zone). The number of AS increases at a rate ranging from 10 to 15 % per year, but there is no accompanying increase in the average shortest path length. However, as the Internet AS topology grows, the number of relationships between AS of a similar degree (peering or tangential relationships) increases faster than the number of relationships between low-degree to high-degree nodes (customer-provider or radial relationships). In addition, to address prefix de-aggregation practices (for traffic-engineering purposes), this excess of tangential links results in an increase in the number of BGP routing paths even though the network size itself doesn't increase. As a consequence, BGP shows its objective performance limits in terms of:

- Memory space required to locally store the routing table entries while minimizing the stretch in the routing paths it produces;
- Communication cost and processing of routing table entries (resulting from the dynamics of the routing information exchanges, which in turn delay the routing system convergence time).

For this purpose, the three-year EULER exploratory research project (Grant No.258307), part of the Seventh European Framework Programme, which began in October 2010, aims to explore novel dynamic routing schemes adapted to the Internet's short-term dynamics (driven by transient topological changes and traffic engineering decisions) and long-term evolution. These dynamic routing schemes are designed to meet the fundamental trade-off between memory space (routing table storage), routing path stretch, and adaptation cost (communication and computational complexity).

In order to validate the functionality and the performance properties of these routing schemes, we have designed and developed an efficient routing model simulator called DRMSim. This simulator comprises the following features and properties:

- Discrete-event routing model simulator developed in Java (running on any Java 1.5 platform or higher);
- Modular design: new topology generators, routing models and performance metrics as well as dynamicity models can be easily added to the simulator;
- Compared with other network simulators, it maintains for every node the local knowledge of the whole network; this technical challenge is addressed by means of efficient data structures.

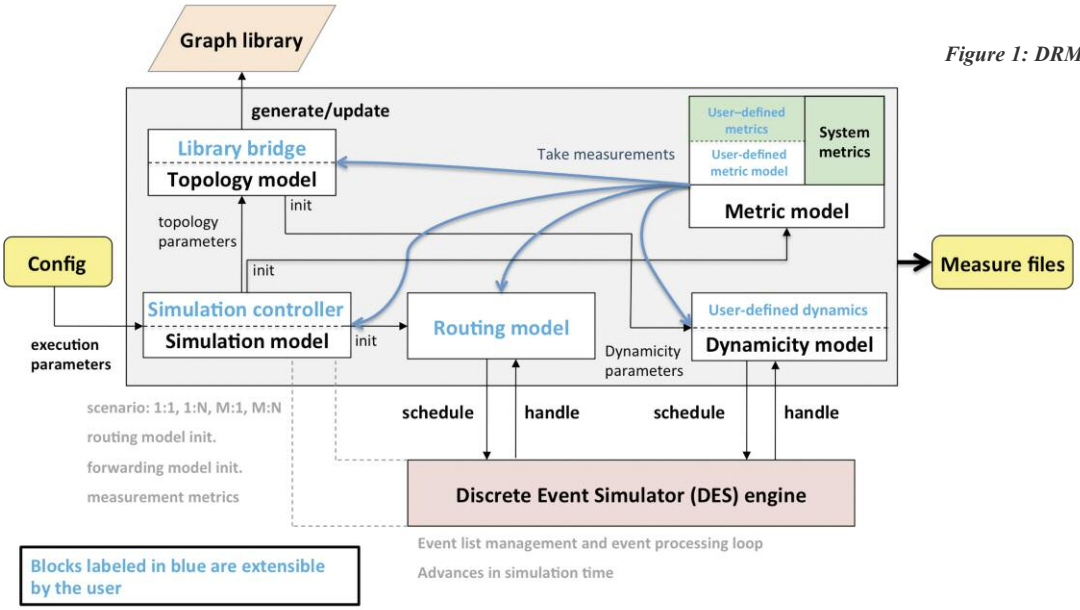

Figure 1: DRMSim Architecture

DRMSim also includes:
- Topology model generators reproducing the properties of the Internet topology (such as Brite, Inet, and GLP);
- The capability to import external topologies (e.g., CAIDA maps);
- A set of standard routing models (such as distance-vector and source routing), compact routing, as well as existing Internet routing schemes such as BGP, allowing researchers to compare the performance of their routing scheme on topologies including up to 16,000 nodes;
- Store and load routing tables to execute simulation from conditions known to the experimenter;
- A metric computation module which automatically computes routing performance metrics such as the routing path stretch, the size of the routing tables, the communication cost (number of routing messages exchanged), etc.

DRMSim is under active development: geometric routing and greedy routing models will soon complement the set of routing models it supports.

Link: http://drmsim.gforge.inria.fr

Please contact:
Aurélien Lancin, Inria and Univ. Nice Sophia Antipolis CNRS, I3S, UMR 7271, France
E-mail: aurelien.lancin@inria.fr

Dimitri Papadimitriou, Alcatel-Lucent Bell, Belgium
E-mail: dimitri.papadimitriou@alcatel-lucent.com

LuxDrops – User-Friendly Management of Confidential Data in the Cloud

by Pascal Bauler, Daniel Heinesch, Antoine Schlechter and Fernand Feltz

Cloud adoption is slowed down by a lack of confidence in the proposed cloud platforms. We first summarize the research activities on confidential and user-friendly data management for Cloud environments. Then we present Luxdrop, a solution that offers secure and transparent cloud solutions by combining security aspects with advanced tracing and visualization possibilities.

Despite the momentum of Cloud Computing and the predicted growth rate of this market segment, the domain faces controversial discussions and security issues. By moving IT solutions to the Cloud, geographical and organizational borders vanish. As a consequence, cloud customers have to precisely analyse their IT solutions and especially their data management in order to decide on the required levels of confidentiality. This analysis is even more complicated if the Cloud provider and customer are under different legislations. The complex legal context and the underlying security concerns are the major causes of the slowdown in cloud adoption. Research activities at the Centre de Recherche Public – Gabriel Lippmann address these concerns by designing a user-friendly management platform for confidential data in the Cloud.

A recent study by the Fraunhofer Institute [1] identifies three mandatory core features to manage data objects in the Cloud: Data sharing between multiple users, Data replication with online and offline access to the data, and automated versioning of the data. Recent customer communities and expert studies [2], [3] identify data related security concerns as a major concern in cloud adaption. In this context, the data objects can be classified into three different classes depending on their privacy and criticality:
- Private data, which are only accessible by the owner. The data are typically private information which should never be accessible by a third party,
- Open data on the other hand are accessible by everyone and no privacy constraints are applicable,
- Confidential data can be shared with well identified, trusted persons but are not designated for large scale distribution.

This classification is applicable to physical documents and objects in the same way as to digital data managed in a cloud-based data store. In a classical use case, private documents are mainly stored in a personal safe with access limited to the owner of the safe. Public documents are largely available and accessible to everyone typically within libraries or book stores. Confidential data are also stored in a safe but access to this safe is granted to selected individuals. However any access to this safe is properly managed, monitored and traced.

Independently of the above mentioned classification every data owner has to decide on the level of trust to place in the safe provider. In practice, the user has to decide whether or not to hand-over a duplicate key to the safe provider.

In the context of the LuxDrops research project, these different classes of data are mapped to cloud storage. The management of public data in cloud based environments is trivial and requires no further discussion. On the other hand the management of private and confidential data raises several questions which have to be addressed. As summarized in [1] the key concern is mainly user authentication. This concern can be partially addressed through strong passwords in the context of private data. However, a password based approach is not really secure when working with confidential data shared among trusted persons. Shared passwords can be easily leaked and security breaches cannot be easily identified and traced. As a consequence, specific issues have to be addressed when managing shared confidential data in the Cloud:
- Email based user authentication might not be sufficient to manage confidential data
- Legal requirements applicable to the cloud provider and customer have to be considered
- Access management, monitoring and tracing are only partially available and no appropriate visualisation tools are available to correctly manage the storage space and identify issues.

The LuxDrops platform implements the core features of a storage platform as mentioned by the Fraunhofer Institute. In contrast to existing solutions, the LuxDrops platform offers advanced sharing, replication and versioning features specially adapted to manage confidential data:

- Reliable user authentication is guaranteed through LuxTrust certificates, which are heavily used by the Luxembourg banking industry. Other types of certificates can be easily integrated on request.
- The storage platform can be deployed on-premise or in the context of local Cloud offerings. By locating cloud provider and customer in a same country, legal aspects are easier to handle.
- Innovative visualization tools improve the traceability and transparency of the overall solution. Key visualization features are: an intuitive multi-dimensional search engine and a meta-data browser to visualize access rights. A first design of the meta-data browser shows the access rights at file/folder or at user/groups level and also provides statistical overviews.
- Depending on the level of trust put in the cloud provider the customer can activate client side data encryption.

The LuxDrops solution aims to increase confidence in Cloud based storage solutions. A first version of the LuxDrops prototype will be made available during the second quarter of 2013 and will be evaluated by selected pilot customers.

Screenshot of the LuxDrops application

References:

[1] M. Borgmann, et al: "On the Security of Cloud Storage Services", Fraunhofer SIT – Technical Reports, 2012

[2] E. Sweeney: "Industry Recommendations To Vice President Neelie Kroes On The Orientation Of A European Cloud Computing Strategy", 2011

[3] D. Bigo, G. Boulet, et al: "Fighting cyber crime and protecting privacy in the cloud", European Parliament, Policy Department C: Citizens' Rights and Constitutional Affairs, 2012.

Please contact:

Pascal Bauler, CRP – Gabriel Lippmann, Luxembourg

E-mail: bauler@lippmann.lu

InGeoCloudS: A Cloud-Based Platform for Sharing Geodata Across Europe

by Claudio Lucchese, Raffaele Perego, Salvatore Trani, Makis Atzemoglou, Benoit Baurens and Dimitris Kotzinos

InGeoCloudS is a project funded by the European Commission under the CIP-Pilot actions program. The goal of InGeoCloudS is to design and build a cost-effective platform for data publication and services, able to promote collaboration and innovative development in the European environmental and geological/geophysical science domain.

The Inspired Geodata Cloud Services (InGeoCloudS) project, coordinated by AKKA Informatique et Systèmes (France) was launched in February 2012 with the aim of establishing the feasibility of using a Cloud-based approach for the publication and use of geodata across Europe. The initiative seeks to leverage the economies of scale achievable for a multi-consumer consortium and its ubiquitous availability of access for the geographically distributed end-users of the European institutions in the environmental field. The purpose is to demonstrate that a Cloud infrastructure can be used by public organizations to provide more efficient, scalable and flexible services for creating, sharing and disseminating spatial environmental data. InGeoCloudS is exploiting this concept based on the work of eight partner institutions from five different countries (including ERCIM members CNR, Italy and FORTH, Greece); some partners are IT enterprises and some are public data providers, covering hydrogeology and natural hazards applications. The project roadmap entails two main steps: Pilot1, which is currently available to project partners, and Pilot2 that will open up the services to a broader audience in summer 2013. The whole set of services will be available for free for the duration of the project.

Environmental science, like other fields, both benefits and suffers from an avalanche of data. From an IT resource management perspective, hardware and network resources have become a critical bottleneck and major cost item. One of the main characteristics of the Cloud is that its resources are virtually unlimited in terms of storage and computing power – it scales transparently and in a semi-automated manner, while offering up-to-date underlying technology – and it offers a pay-as-you-go/pay-per-use delivery model with potential reduction of traditional IT infrastructure costs.

From a scientific point of view, rising data quantity and quality has not been accompanied by an equivalent increase in visibility, accessibility and sharing. The 2007 Infrastructure for Spatial Information in the EC (INSPIRE) Directive established rules for geographic and environmental data to ensure that the geodata are consistently available, interpretable and usable across European regional and

state boundaries. The Directive requires the use of established standards and online availability of geodata. InGeoCloudS services intend to support and respect all obligations stemming from the INSPIRE Directive and to facilitate data providers in fulfilling their obligations.

The InGeoCloudS platform is structured in three main software layers. The most basic one provides applications with transparent access to an elastic distributed file system, a GIS-enabled database server, and to some compute facilities for managing on-demand jobs and elastic services. The second layer provides higher-level services including elastic pool of map servers, database synchronization mechanisms, and also an advanced data publication service which allows data providers to publish OGC services and their corresponding datasets by simply uploading their data in some standard

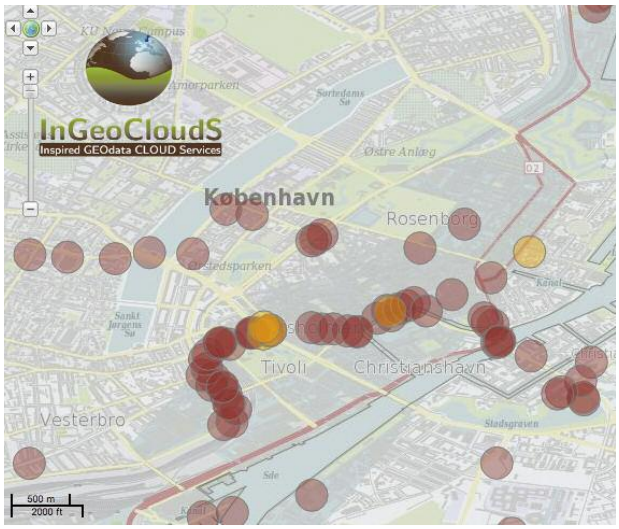

Figure 1: InGeoCloudS Portal: visualization of Copenhagen groundwater data

format and exploiting the geo-publication services of the InGeoCloudS infrastructure. The last layer includes the Web portal and an elastic web servers pool hosting the front-end of geo-applications.

One of the most innovative InGeoCloudS services is the Linked Open Data based data and metadata management service. The InGeoCloudS consortium designed an interoperable Geo-Scientific Observation Model (GSOM) that allows the mapping of different data sources into a harmonized data/metadata representation space. Thus the potential users of the platform would have the ability to publish their data as Linked (Open) Data and link them to existing data either inside or outside of the platform. By exploiting this model and RDF-based technologies, InGeoClouds opens the doors to cross-country queries as well as cross-discipline analysis, e.g. facilitating correlation of data between various scientific domains.

Three main use cases are currently hosted by the platform: (i) Geohazard management in Slovenia, which provides an early landslide warning system on the basis of past landslides, rainfall forecasts and geology for 14 Slovenian municipalities. (ii) Measurement of earthquake post-effects in Greece, with automated notification services for interested parties implemented by the Earthquake Planning and Protection Organization (EPPO). (iii) Integration of generalized services for groundwater management, showing areas where high concentrations of pesticides can be met or displaying samples chemical analysis. These use cases are owned by geological institutes in Denmark, France and Greece (GEuS, BRGM and EKBAA/IGME) and illustrate how pan-European exploitation of data can be achieved.

We hope to make InGeoCloudS a sustainable platform capable of attracting data providers from diverse horizons. Every owner of geoscientific data (not just geologists) is expected to find the platform and its services attractive, facilitating the publishing of data to larger and broader audiences and developing new services economically. Consortium members – and possibly other partners – will offer both scientific and necessary IT skills in this endeavour, far beyond the project's timespan.

Link:
http://www.ingeoclouds.eu/

Please contact:
Claudio Lucchese
ISTI–CNR, Italy
Tel: +39 050 3152468
E-mail: claudio.lucchese@isti.cnr.it

Electromagnetic Fields for Neuron Communications

by Rié Komuro and Ilangko Balasingham

Are electromagnetic fields harmful to the brain? We are continuously exposed to electromagnetic (EM) fields generated by electronic devices such as cell phones. The effects on human bodies by EM waves, especially ones with the frequency range between 3 kHz and 300 GHz, called radio frequency (RF), are of great interest. Reported effects of RF fields on living systems are widely variable; harmful, negligible, or nil. Although no definitive evidence has been found, it is generally considered that exposure to low energy RF waves could be a risk to human health. However some beneficial effects were also reported.

Our research group belongs to the Department of Electronics and Telecommunications at NTNU in Trondheim, Norway and has been working on the effects of EM fields on neurons for about three years. The work was inspired by some results obtained by a group at the Florida Alzheimer's Disease Research Center in 2010 [1]. The research group anticipated that long-term cell phone use could damage the brain, especially the region controlling memory. The same physical environment as cell phone use for two one-hour periods a day was established, and genetically modified mice mimicking Alzheimer's disease (AD) as well as healthy mice were put there. After several months of EM wave exposure, the experimental results showed some positive effects for

both groups of mice, contrary to what was expected before the experiment; the healthy mice increased cognitive performance, and the AD mice reduced the type of protein which is accumulated in the brain developing AD. This result interests us in the effects of EM fields to the brain; why and how the pathological changes occurred.

Each of the basic building blocks in the central nervous system consists of neurons. The membrane potential at each neuron rapidly rises and falls. This spike-shape event is called an action potential, and repetitive spikes play an important role in communication among neurons. They can be described as nodes of a network to process and transmit information through electrical and chemical signals. The role of message propagation is fundamental when we deal with neurodegenerative diseases, such as AD. Diseases of this type cause damage to the communication network of the brain and the links between neurons.

Our goal is to understand the fundamental mechanism of the network of the neurons on the cellular level, especially the component related to memory. We attempt to establish a scientific explanation for how and why neurons and networks can function as RF wave receivers and demodulators for given RF transmitter characteristics. An induced EM field has thermal and non-thermal components. The non-thermal effects are due to the extremely low frequencies, that is, modulation frequencies of the information signal. The schematic picture showing the process of propagation of EM waves to a neuron is presented in Figure 1. The generated EM waves near the head penetrate the skull and reach neurons in the brain. It is considered that the induced fields could stimulate each neuron and affect the action potential.

Our work is carried out using "NEURON", a simulator that models individual neurons and networks of neurons. Simulations of RF fields inside the skull have been already done. However, these data cannot be directly used as inputs for simulations on the neuron level because the scale of neurons is much smaller than that of a skull; without adjustment, the simulated RF would make exactly the same impact upon all neurons in the brain. In practice, each neuron can be affected differently depending on the distance from the source.

For the first step, we have been exploring the effects of RF fields on a single neuron with different types of waves as an input current. When the fundamental mechanism is well-studied, we will move on to simulation with multiple neurons and then eventually a network of neurons at the region playing the role of consolidation of memory. In order to confirm the validity of simulation results, experiments are also required. We are currently concentrating on the theoretical part, but it is necessary to collaborate with neuroscientists in the future. The detailed outline of the project is introduced in [2].

Our research results could potentially help development of some medical treatments in the future. The World Alzheimer's Report 2009, issued by Alzheimer's Disease International, estimated that 35.6 million people worldwide would be suffering from the disease in 2010, and the number was estimated to nearly double every 20 years, to 115.4 million in 2050. Development of safe and effective treatments

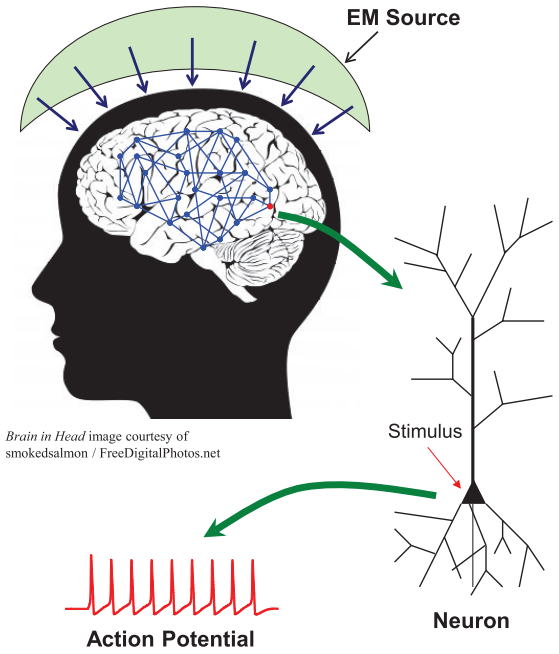

Figure 1: propagation of RF waves to a neuron. The generated EM waves near the head propagate to the brain through the skull. Neurons in the network could be stimulated, and the action potential occurring on the membrane potential of neurons could be affected.

is, therefore, an urgent matter. Surgical implant of a pacemaker-like device into the brain of a patient with mild AD has already started in the US. Elucidation of the mechanism of the neuronal network could enable scientists to invent new types of treatment; for example, it might be possible to develop a nanomachine which stimulates the region of the brain that is malfunctioning owing to a neurodegenerative disease from the outside of or inside the body and reconstruct the network [3].

Link: http://www.neuron.yale.edu/neuron/

References:

[1] G. W. Arendash et al.: "Electromagnetic field treatment protects against and reverses cognitive impairment in Alzheimer's disease mice", J Alzheimers Dis, 19(1), 2010.
[2] F. Mesiti et al.: "On the modeling and analysis of the RF exposure on biological systems: a potential treatment strategy for neurodegenerative diseases", Nano Commun Networks, 3(2), 2012.
[3] F. Mesiti and I. Balasingham: "Nanomachine-to-neuron communication interfaces for neuronal stimulation at nanoscale", IEEE J Sel Areas Commun, 2013 (in press).

Please contact:

Rié Komuro, NTNU, Norway
Tel: +47 7359 3670
E-mail: rie.komuro@ntnu.no

A Patented Position Determination Solution

by Per Kreuger

More dynamic and less resource-intensive: Researchers at SICS have developed a new solution to determine mobile phone positioning.

Communication is becoming increasingly mobile, but tracking the location of individuals within networks is quite a difficult task. As researchers at SICS began analyzing the problems surrounding mobile phone positioning, they realized the solution was right in front of their eyes—in the network itself. The network has contact with mobile devices in a variety of situations, and with the new method from SICS, the locations can be tracked within the network. The proposed method uses many short traces to map the collective behaviour of network users without storing longer individual trajectories.

The biggest win is that this method makes manual configuration of network parameters unnecessary: with statistics and optimization, it can be done automatically. The demand for this kind of technology is only going to grow, especially with the introduction of more heterogenous networks.

An important and difficult configuration problem

Configuring the mobility management entities (MMEs) that calculate the location of a mobile phone is a resource-intensive task. Cellular networks currently rely on fixed collections of cells – tracking areas – for user equipment localization. Tracking areas are manually configured and maintained; this is therefore an important configuration and resource management problem. The imperfections of the configuration results are particularly apparent for people moving collectively over tracking area borders, for example when a conversation breaks off and the internet connection is lost on a subway.

It was with this background that, a few years ago, Ericsson turned to SICS to work together to develop a better method and a tool to decide how the MMEs should be configured and where the boundaries between these tracking areas should be.

Knowing how people move is essential to determine the best placement for the boundaries. When we considered this, we came to an important realization: the network sees a lot of movements and generates the information necessary to set the parameters itself. A more radical approach would therefore be to dynamically and autonomously reconfigure the tracking areas online and make them local to each cell.

The research project has been successful and has now resulted in a new patent. A completely new mechanism has been designed based on a distributed algorithm to disseminate mobility information between nodes in the network, a local probabilistic model estimated from this data and an optimization mechanism to implement efficient incremental paging in each area.

Taking the solution forward

Ericsson has now taken over the rights to the patent [1] and Per Kreuger hopes to see the technology up and running soon. Configuring the network is time-consuming and expensive for the operators and need for network solutions like this one is growing in pace with the fact that we're more connected and we use cellular networks for more task and in new contexts. Our solution shows that configuration efficiency can be improved considerably [2,3], and it's currently an excellent example of the benefits of Probabilistic Network Management.

Mobility is only one of the areas in which SICS is currently developing Probabilistic Network Management (PNM) approaches. The Probabilistic Network Management paradigm specifies non-deterministic methods in which decisions are based on probabilistic objectives and richer statistical monitoring information, rather than on strict performance guarantees and measurements. Compared to current network management technology, PNM approaches provide new effective means of resource-efficient and flexible network management solutions.

Link:
http://www.sics.se/projects/probabilistic-network-management

References:
[1] Å. Arvidsson, D. Gillblad, P. Kreuger: "Tracking user terminals in a mobile communication network", Patent PCT/EP2011/060090}, June 2011, Ericsson AB http://patentscope.wipo.int/search/en/detail.jsf?docId=WO2012171574
[2] P. Kreuger, D. Gillblad, Å. Arvidsson: "Zero configuration adaptive paging (zCap)", IEEE 76th Veh. Technol. Conf., Québec, 3-6 Sept. 2012, 978-1-4673-1881-5, http://www.ieeevtc.org/vtc2012fall/
[3] P. Kreuger, D. Gillblad, Å. Arvidsson: "zCap: A zero configuration adaptive paging and mobility management mechanism", Journal of Network Management, forthcoming.

Please contact:
Per Kreuger, SICS Swedish ICT
Tel: +46 (0)8 633 15 22, E-mail: piak@sics.se

SoFWIReD – Southampton Fraunhofer Web and Internet Research & Development

by Sepideh Chakaveh

The SoFWIReD project brings together two ERCIM institutes, University of Southampton and Fraunhofer Gesellschaft to determine the future of Web and Internet Science.

The information age as we know it, has its roots in a handfull of enabling technologies most of all the World Wide Web for the provision of global connectivities. The emergence of a Web of Data in terms of the publication and analysis of UK Open Government Data has provided new insights about the impact of the Web in Society [1].

The second most important technology in this development has been the emergence of streaming processes based on new and innovative compression methods such as MP3 where audio & video contents could be accessible to everyone on the web.

The University of Southampton where Sir Tim Berners-Lee the inventor of the WWW holds a chair in Computer Science, together with Fraunhofer Gesellschaft, Europe's largest research organization for applied research and the inventors of MP3, have jointly launched a common project: SoFWIReD.

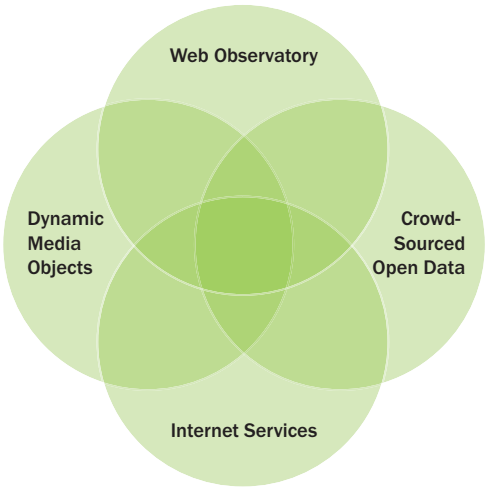

Research topics add'ressed by SoFWURed

The SoFWIReD-project aims to find answers to research questions such as:

- What are the socio-economic reasons as to why individuals participate in a collective endeavour?
- What legal frameworks govern (or should govern) the resources that are created?
- What is the psychology of identification with an online collective community?
- What role is there for policy-makers to engage in and facilitate collaborative endeavour?
- How can collective intelligence emerge, given the different languages used by different genders, races, classes, and communities?

The SoFWIReD-project is developing a comprehensive, interoperable platform for data and knowledge driven processing of open data and will investigate aspects of Collective Intelligence. The insights generated in the project will form the basis for supporting companies in the collective intelligence transition by consulting, organisational development and software solutions.

The project addresses the following research topics:

Web Observatory

The goal is to establish an observation mechanism capable of detecting or predicting information cascades on the Web. The idea is to identify and collect data on the Web, to analyse how it impacts business activity, and to develop mechanisms and tools that will enable interpretation and analysis [2].

Crowd-Sourced Open Data

The goal is to leverage the value of open data for e-business by exploring the value of open data in supply chains and logistics, and to deploy mechanisms for public engagement in collecting and leveraging open data related to e-government. This refers to sharing concepts on a digital collaboration environment to encourage and facilitate new and innovative "brainstorming" processes. New concepts may emerge due to the contributions of individuals, professional or otherwise who are able to share their ideas, as it is not limited to professionals, but rather the general public who wishes to become involved through this collaboration and social dissemination fashion which is only possible using the web.

Dynamic Media Objects

Media objects that are consumed in the web such as audio & video content are mainly streamed. It is recently established that these objects are inherently dynamic and are in stark contrast to static objects such as text files. The focus of this part of the SoFWIReD project is to understand the nature and behaviour of dynamic media objects and their life cycle. Also some generic technical & legal interactions concepts and the relevant impacts on the Web such as privacy and trust issues are to be addressed in here [3].

Internet Services

This aims at developing advanced, secure and trustworthy services that can help companies to improve the flow of information as well as to adjust their structure and organisation to remain competitive in a mobile and dynamic Web environment.

Each of these four research themes represents a major direction for innovation with significant impact on Web technology at large. They result from considering the practical demands of both the research community and customers in industry, alike.

Link: http://sofwired.org/

References:

[1] T. Berners-Lee, W. Hall, J. Hendler, N. Shadbolt, D. Weitzner: "Computer Science: Enhanced: Creating a Science of the Web", Science 313 (5788): 769–771

[2] W. Hall, N. Shadbolt: "Web Science and the Web of Linked Data: Opportunities in the Information Age", in Self-Innovation Forum, 11 November 2010, Shenzhen, China.

[3] S. Chakaveh, W. Vonolfen: "Edited Web Collective Intelligence Project Proposal", Fraunhofer Gelleschaft – ICON Programme Documentations

Please contact:

Sepideh Chakaveh
Chair of Media Technology & Edutainment Working Group
Web & Internet Science Research Group
Electronics & Computing Department
University of Southampton
E-mail: sc2@ecs.soton.ac.uk

Preserving Linked Data

by Carlo Meghini, Anna Molino, Francesca Borri and Giulio Galesi

The PRELIDA project aims at building bridges across the Digital Preservation and Linked Data communities, raising awareness of already existing outcomes of Digital Preservation in the Linked Data communities, while at the same time posing new research questions for the preservation domain.

In January 2013 the European Commission launched the project PRELIDA – Preserving Linked Data, a two year Coordination Action of the VII Framework Programme. The main goal of PRELIDA is to build bridges between the areas of Linked Data and Digital Preservation, with two principal objectives: making the Linked Data community aware of the existing results of the Digital Preservation community, and identifying the issues and problems raised by the need to preserve Linked Data which pose new research challenges. To achieve these goals, the project will target stakeholders of the Linked Data community (e.g. data providers, service and technology providers, as well as end user communities), who have not been traditionally targeted by the Digital Preservation community.

The activities of PRELIDA are motivated by the recent emergence of a whole new industry implementing services on top of large data streams, and the impact of this economic sector – known as the "data economy" – may soon exceed the current importance of the software industry, since the sheer amount of data offered and consumed on the Internet will steadily increase by orders of magnitude, generating the potential for many new types of products and services. For instance, governments and organizations will only make their data available in open form on the Linked Data cloud if there are assurances that it will be properly maintained, with particular emphasis on quality and permanent access. It thus becomes crucial to be able to guarantee the integrity, accessibility and usability of Linked Data over the long-term, and these are precisely the objectives of Digital Preservation.

Unfortunately, so far there has not been much interaction between the Linked Data and the Digital Preservation groups. However, interest in both the adoption of linked data by the digital preservation community, and the recognition of preservation as an important challenge for linked data is now growing rapidly. This is one of the reasons why the PRELIDA consortium includes a society of organizations in the area of Digital Preservation (APA) plus two key members of the Linked Data area (University of Innsbruck and University of Huddersfield), and is coordinated by the ISTI institute of the Italian National Research Council which has expertise in both areas.

The main outcomes of PRELIDA will be a State of the Art report on Linked Data and their preservation needs, and a Road Map for addressing the new challenges that preserving linked data poses. The Road Map should drive scientific and technological developments in this field, as well as future research programmes that the Commission may decide to fund.

The ambitious targets of PRELIDA will be achieved through a number of different means. The first crucial instrument is the establishment of a continuous working group, bringing together key researchers and stakeholders from both communities. The principal task of the working group is to identify key sectors within the two areas, working out the particular challenges that Linked Data pose to the long-term preservation problem. To accomplish this task, use cases representing examples of long-term access to Linked Data will be developed by key stakeholders, and a Technology/Research observatory will be set up in order to identify the most significant actors working on Linked Data and Digital Preservation challenges.

The Working Group members will be invited to three workshops. During the Opening Workshop participants will concentrate on the current state of Digital Preservation solutions, presenting and discussing the preservation needs of the Linked Data community. The focus of the Midterm Workshop will be decided on the basis of the interactions and findings of the working group, while the main aim of the final Consolidation and Dissemination Workshop is the presentation of a preliminary roadmap, as well as the discussion of the key findings for research communities, relevant industries, potential stakeholders, and policy makers.

An online infrastructure will be provided with the purpose of creating a network that will support continuous interaction between the consortium and the working group members. Moreover, liaisons will be established with other research projects and organizations working in the relevant areas.

Finally, postgraduate students and young researchers with knowledge from both fields will be invited to two summer schools, where they will acquire thorough knowledge of the state of the art of both communities. The first school will be held together with the European Semantic Web Conference (ESWC) Summer School, where speakers from the Digital Preservation area will be invited to present preservation solutions and discuss challenges. In the second year, a school dedicated to the topic of preserving Linked Data will be held in conjunction with the Consolidation and Dissemination workshop.

In conclusion, PRELIDA will facilitate the establishment of a scientific, technological and user group community that can be expected to outlast the duration of the project.

Link: http://www.prelida.eu/

Please contact:
Carlo Meghini, ISTI-CNR, Italy
E-mail: carlo.meghini@isti.cnr.it

MediaEval 2012 Evaluation Campaign

by Gareth J. F. Jones and Martha Larson

MediaEval is an international multimedia benchmarking initiative offering innovative new tasks to the multimedia community. MediaEval 2012 featured tasks incorporating social media search, analysis of affect and location placing of images.

Workshop participants. Photo: Xincaho Li

MediaEval is an international multimedia benchmarking initiative that offers innovative new content analysis, indexing and search tasks to the multimedia community. MediaEval focuses on social and human aspects of multimedia and strives to emphasize the "multi" in multimedia, including the use of speech, audio, tags, users, and context, as well as visual content. MediaEval seeks to encourage novel and creative approaches to tackling these new and emerging multimedia tasks. Participation in MediaEval tasks is open to any research group who signs up. MediaEval 2012 was the third evaluation campaign in its current form, which follows on from VideoCLEF track at CLEF 2008 and CLEF 2009.

MediaEval 2012 offered six main tasks coordinated in cooperation with various research groups in Europe and elsewhere. The following tasks were offered in the 2011 season:

- Placing Task: This task required participants to assign geographical coordinates (latitude and longitude) to each of a provided set of test videos in two sub-tasks: placing anywhere in the world, precise location in a known city. Participants could make use of metadata and audio and visual features as well as external resources.
- Spoken Web Search Task: This task involved searching for audio content within audio content using an audio content query. It addresses the challenge of search for multiple, resource-limited languages with application in low-literacy communities in the developing world.
- Affect Task: This task required participants to deploy multimodal features to automatically detect portions of movies containing violent material. Violence is defined as "physical violence or accident resulting in human injury or pain". Any features automatically extracted from the video, including the subtitles, could be used.
- Social Event Detection Task: This task requires participants to discover events and detect media items that are related to either a specific social event or an event-class of interest. Social events of interest were planned by people, attended by people and the social media captured by people.
- Tagging Task: The task required participants to automatically assign tags to Internet videos using features derived from speech, audio, visual content or associated textual or social information. This year the task focused on labels that reflect the genre of the video.
- Visual Privacy Task: Participants were required to explore methods to obscure human faces so as to make them unrecognisable in digital imagery with application in situations where persons may be captured in a video frame, but may wish to protect their privacy. will need to propose methods whereby human faces occurring in digital imagery can be obscured so as to render them unrecognisable.

MediaEval 2012 also introduced the idea of Brave New Tasks, as activities with smaller participant groups as incubators of potential main tasks for future years. The MediaEval 2012 Brave New Tasks were: User Account Matching, Search and Hyperlinking and MusiClef: Multimodal Music Tagging.

The MeviaEval 2012 campaign again culminated in a workshop that was held at Fossabanda Santa Croce in Pisa, Italy from 4 to 5 October. The workshop brought together the task participants to report on their findings, discuss their approaches and learn from each other. MediaEval participation increased again in 2012 with a total of 54 papers appearing in the Working Notes, and 60 participants attending the workshop. In addition to organizer and participant presentations, the workshop features invited presentations by Jana Eggink, BBC Research and Development, London and Nicola Ferro, University of Padova, co-ordinator of the Promise Network of Excellence. The workshop concluded with a meeting of task organizers and other interested researchers that consisted of presentations and discussions of task proposals for MediaEval 2013. An exciting development of the workshop was the increasing collaborations between task participants arising from informal breakout discussions, which are now leading to further experiments with MediaEval datasets and submissions of joint papers to international conferences. The Working Notes proceedings from the MediaEval 2012 workshop have again been published by CEUR workshop proceedings.

MediaEval 2012 received support from a number of EU and national projects and other organizations including: AXES, Glocal, WeKnowIt, Chorus+, Quaero, IISSCoS, Technicolor, IBM India and CMU.

MediaEval 2013

The MediaEval 2013 campaign is currently in progress and participants will be presenting results of their work at the MediaEval 2013 Workshop in Barcelona from 18 to 19 October, just before the ACM Multimedia 2013 conference.

Further details of MediaEval are available from the MediaEval website.

Links:
http://www.multimediaeval.org
Online proceedings: http://ceur-ws.org/Vol-927/

Please contact:
Martha Larson, TU Delft, The Netherlands
Tel: +31 15 278 5812, E-mail: M.A.Larson@tudelft.nl

Event report

TransForm School on Research Directions in Distributed Computing

by Panagiota Fatourou

The TransForm School on Research Directions in Distributed Computing (SRDC 2013) was held in Heraklion, Crete Island, Greece, on June 10-14, 2013 at the premises of the Foundation for Research and Technology-Hellas (FORTH). It was organized by Prof. Panagiota Fatourou in the context of the EC-funded Marie Curie Initial Training Network, called TransForm.

The school included a series of talks by twenty renowned researchers covering hot topics on current research in the general area of distributed computing with emphasis on multi-core computing, synchronization protocols, and transactional memory. Additionally, it provided discussion sections on future research directions in the field. The school also hosted short presentations by graduate students and early-stage researchers where they had the chance to discuss their current research work.

The school was very successful with about 65 participants in total. Out of them, 23 were well-known researchers (both in academia and the industry) from all over the world, and the rest were early stage researchers and graduate students from 16 different Universities and Institutions located in Europe, Israel, Canada or the US. The school served as a major dissemination and training activity of TransForm and significantly contributed towards the dissemination of advanced scientific knowledge on the field and the promotion of international contacts among key scientists from academia and industry.

More information:
http://www.ics.forth.gr/carv/transform/srdc/

Call for Participation

AITA 2013 – 12th International Workshop on Advanced Infrared Technology and Applications

Turin, Italy, 10-13 September 2013

The main purpose of the AITA workshop series is to provide an international forum to present and discuss current trends and future directions in Infrared Technology, mainly for civilian applications. The workshops also aim at fostering the creation of a permanent network of scientists and practitioners for easy and immediate access to people, technologies and ideas. The events have been successful so far not only because of the high scientific quality of the communications but also for the friendly atmosphere that always characterizes AITA workshops.

In AITA 2013 edition, special emphasis will be given to the following topics:
- Advanced technology and materials
- Smart and fiber-optic sensors
- Thermo-fluid dynamics
- Biomedical applications
- Environmental monitoring
- Aerospace and industrial applications
- Astronomy and Earth observation
- Non-destructive tests and evaluation
- Systems for cultural heritage
- Image processing and data analysis
- Near-, mid-, and far infrared systems

Considering the relevance of technological transfer on AITA topics, this year a technical seminar entitled "Energy certification and diagnostics in buildings: regulatory framework and evaluation tools" is also organized: the seminar is aimed at those who wish to become familiar with IR thermography applications but will also highlight the most recent advances in the field.

More information:
http://ronchi.isti.cnr.it/AITA2013
E-mail: aita@isti.cnr.it

Call for Participation

W3C Workshop on Publishing using the Open Web Platform

Paris, Centre Georges Pompidou, 16-17 September 2013

The goal of this W3C Workshop is to bring together major players, including publishers, standardization organizations, technology developers, booksellers, accessibility organizations and others to identify areas where work is needed to make the Open Web Platform suitable for commercial publishing, especially in print, all the way from authoring through to delivering the printed product and beyond.

This is one of a series of W3C Workshops in the areas of publishing, and the first to focus on the complete publishing workflow and on issues particular to producing printed products.

Workshop topics inlcude:
- "XML First" XML/XHTML/HTML from authors
- Revision control and change tracking for the Web
- Content management, version tracking and workflow
- Accessibility: applying WCAG 2.0 and ATAG 2.0 to print
- Formatting to print using CSS
- Internationalization: publishing is for everyone. What changes to CSS or HTML might be needed for multilingual texts, including parallel texts and study materials with glosses?
- Print on demand: color management, ink control, specifying media, binding, trimming, finishing...
- Multiple output formats: are CSS media queries enough? What about alternate content, image replacement, subsetting?

To ensure productive discussions, the workshop is limited to 80 attendees. Participation is free and open to W3C members and non-members.

Deadline for submitting position papers is 15 July 2013.

More information:
http://www.w3.org/2012/12/global-publisher/

Call for Participation

Second Plenary Meeting of the Research Data Alliance

Washington DC, USA,
16-18 September 2013

The Research Data Alliance (RDA) vision for global research data infrastructures involves overcoming barriers in realizing the importance of data sharing for next century science. RDA provides the framework to accelerate international data-driven innovation and discovery by facilitating research data sharing and exchange, use and re-use, standards harmonization, and discoverability. This will be achieved through the development and adoption of infrastructure, policy, practice, standards, and other deliverables and this work is primarily undertaken by the Working and Interest Groups.

Scientists, research data infrastructure providers, research data practitioners, policy-makers and research data stakeholders from all over the world will gather at RDA's second bi-annual plenary meeting. The meeting will provide existing and new and aspiring RDA community members an opportunity to conduct business and make progress on their plans and deliverables. The plenary will be a forum to demonstrate the value of RDA and to receive feedback from the broader research and policy communities. RDA members and plenary participants will get an update on both the governing and technical activities since March 2013 (RDA Launch).

Plenary sessions will include Keynotes from internationally renowned visionaries, including the US Government, reports from currently active Working and Interest Groups, a forum for affiliate organizations, and general RDA business meetings. All this coupled with parallel breakout sessions for the Working and Interest Groups, other ad-hoc groups and a dedicated session for RDA Newcomers on how to get involved, propose new working or interest groups and interact with the RDA members.

More information:
http://rd-alliance.org/future-events

Call for Participation

Collaboration meets Interactive Surfaces: Walls, Tables, Tablets and Phones

A One-day Workshop co-located with ITS 2013, St Andrews, Scotland, UK, 6-9 October, 2013

The vast screen real estate, which is provided in large-scale interaction environments presents novel ways to visualize and interact with data-rich models. In parallel to this technological revolution, interactive surfaces have also become widespread in different sizes and devices. This workshop proposes to bring together researchers who are interested in improving collaborative experiences through the combination of multiple interaction surfaces with diverse sizes and formats, ranging from large-scale walls, to tables, tablets and phones.

Topics:
Topics to be covered include, but are not limited to:

- Design and evaluation of collaborative environments with interactive surfaces, either remotely or co-located
- Collaborative applications on interactive surfaces for concrete domains
- Communication, cooperation and coordination as well as social protocols
- Interactive surfaces to enhance spatial perception of content and/or support navigation during collaboration activities
- Issues when moving from desktop-based collaboration to large-scale walls, tabletops and touch-based mobile devices
- Physical navigation and collaborative sense making
- Integration of different devices and surfaces for collaboration
- Collaboration paradigms and user interface designs that address enhancement of collaborative activities using interactive surfaces and tabletops.

More information:
https://sites.google.com/site/collaborationsurfaces/

Call for Papers

5th Joint Virtual Reality Conference (JVRC 2013)

Campus Paris Saclay,
11-13 December 2013

The 5th JVRC conference gathers the 19th Eurographics Symposium on Virtual Environments (EGVE – http://www.eg.org/) and the 10th Conference and Exhibition of the European Association of Virtual Reality and Augmented Reality (EuroVR – http://www.eurovr-association.org/).

As previous issues, JVRC 2013 will bring together people from industry, commerce, research including technology developers, suppliers and all those interested in virtual reality, augmented reality, mixed reality and 3D user interfaces to exchange knowledge and share experiences of new results and applications, live demonstrations of current and emerging technologies and form collaborations for future work.

Topics
Topics of interest include any topics related to "human factors issues" (user studies and evaluation trials, presence and cognition, 3D user interfaces, 3D interaction metaphors, self-representation and embodiment, virtual humans, etc.), "technologies" (VR system architecture, collaborative and distributed VR, augmented Reality (AR) and mobile devices, mixed Reality (MR), advances in display technologies, etc.), applications (industrial applications, aerospace and transport, construction and architecture, manufacturing and engineering, medical and rehabilitation, etc.).

Deadlines for Submissions

- 17 July: Full and Short scientific papers
- 26 July: Special abstract for collaborative JVRC-ICAT demos
- 6 September: Short papers for industrial tracks, posters and all demos, exhibition proposals

More information:
http://jvrc2013.sciencesconf.org/

Florin Stoican and Sorin Olaru

Set-theoretic Fault-tolerant Control in Multisensor Systems

Fault-tolerant control theory is a well-studied topic but the use of the sets in detection, isolation and/or reconfiguration is rather tangential. The authors of this book propose a systematic analysis of the set-theoretic elements and devise approaches which exploit advanced elements within the field. The main idea is to translate fault detection and isolation conditions into those conditions involving sets. Furthermore, these are to be computed efficiently using positive invariance and reachability notions. Constraints imposed by exact fault control are used to define feasible references (which impose persistent excitation and, thus, non-convex feasible sets). Particular attention is given to the reciprocal influences between fault detection and isolation on the one hand, and control reconfiguration on the other. Co-author Florin Stoican was an ERCIM postdoctoral fellow at NTNU in 2011/2012.

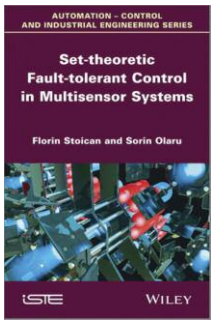

ISBN: 978-1-84821-565-8; hardcover; 176 pages;
May 2013, Wiley-ISTE

Benoît Michel

Digital Stereoscopy – Scene to Screen 3D Production Workflow

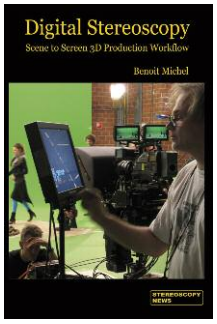

Digital Stereoscopy provides a large number of technical details on Digital Stereoscopy, making it a reference work for professionals in the field, with explanations of the most commonly used software, hardware, and standards. To make the book as useful as possible, an abundance of sidebars explaining the most important concepts in each chapter and a glossary of the most frequently used 3D-specific terms have also been included. Over 250 black-and-white diagrams and pictures complement the text, some of them accompanied by QR-code links to full-color images on the web.

Intended audience

This book is for anyone who wants to get involved in stereoscopic imagery or those who already work in 3D but want to hone their skills. Broadcasting students are not forgotten: The first chapters give the basic principles needed to discover 3D imagery, while the subsequent chapters contain the practical information that they will need to put them into practice.

http://www.stereoscopynews.com/digitalstereoscopybook
ISBN: 978-1-480-15709-5; paperback, 354 pages, March 2013

Stefania Gnesi and Tiziana Margaria

Formal Methods for Industrial Critical Systems: A Survey of Applications

Today, formal methods are widely recognized as an essential step in the design process of industrial safety-critical systems. In its more general definition, the term formal methods encompasses all notations having a precise mathematical semantics, together with their associated analysis methods, that allow description and reasoning about the behavior of a system in a formal manner.

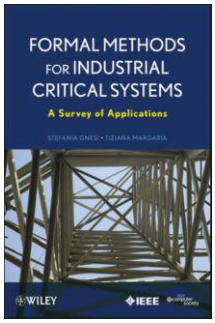

Growing out of more than a decade of award-winning collaborative work within ERCIM, Formal Methods for Industrial Critical Systems: A Survey of Applications presents a number of mainstream formal methods currently used for designing industrial critical systems, with a focus on model checking. The purpose of the book is threefold: to reduce the effort required to learn formal methods, which has been a major drawback for their industrial dissemination; to help designers to adopt the formal methods which are most appropriate for their systems; and to offer a panel of state-of-the-art techniques and tools for analyzing critical systems.

ISBN: 978-0-470-87618-3; Paperback; 292 pages
March 2013, Wiley-IEEE Computer Society Press

Job opportunity

Research positions at the Warsaw Center of Mathematics and Computer Science

The Warsaw Center of Mathematics and Computer Science (WCMCS) is a research consortium consisting of the Faculty of Mathematics, Informatics and Mechanics of the University of Warsaw, and the Mathematical Institute of the Polish Academy of Sciences. The Center was appointed the Polish Leading National Research Center in mathematical sciences by the Ministry of Science and Higher Education in 2012. The WCMCS offers a number of research positions with a few hours teaching duties for the outstanding scientists whose field of research is compatible with those of the research groups in the units of the Center. PhD students from other universities can apply for PhD internships in April and November, young researchers can apply for postdoctoral positions in March and October, while senior scientists can apply for the position of a leader of the research project at any time. The rules of the program and the information regarding the application procedure can be found on the WCMCS website.

More information: http://wcmcs.edu.pl

European Laboratory on Big Data Analytics and Social Mining

The "SoBigData" Lab is a recently created research initiative in Pisa, Italy, with the mission to perform advanced research and analyses on the emerging challenges posed by big data, namely the digital breadcrumbs of human activities continually sensed by the ICT systems that people use. This is a joint initiative among different Italian institutions with the aim to aggregate multi-disciplinary competencies for addressing together the scientific challenges needed to set the ground for social mining science and technology. Fields of research include "mobility data mining", "social network analysis", and "privacy, security and trust". The Lab is currently supported by Institute of Information Science and Technologies of CNR, Institute of Informatics and Telematics of CNR, Department of Computer Science, University of Pisa, Region of Directorate General Presidency EU Liaison Office, CNR Department for Engineering, ICT, Energy and Transportation. The initiative will be launched by a bootstrap workshop held in Pisa on 18 July 2013.

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

National ICT Awards 2013 for Python and SIG

Programming language Python and company Software Improvement Group (SIG) were awarded a Dutch National ICT Award on 13 June. Both winners have their origins at CWI. Python was awarded the COMMIT/ Award for the most valuable product resulting from ICT research that is over ten years old. Python was developed in the early 1990s by Guido van Rossum, who worked at CWI at the time. The jury said that this award is a recognition of the persistent pioneering work that resulted in worldwide use of the language today. Examples of users are Google, game producers, Walt Disney and NASA. Van Rossum's former group leader Dick Bulterman (CWI) received the award on his behalf. SIG received the ICT Milieu Award for their application energy profile register. Founded in 2000 as a CWI spin-off company, SIG received the prize for their application that determines the energy profile of software-hardware combination. The jury's report said: "The application energy profile register has an enormous potential (…). It makes the true cost of software use visible by coupling energy meters to hardware. It initiates a thought process. This makes it the fair winner of the ICT Milieu Award 2013."

Dick Bulterman (CWI, right) receives the COMMIT/ Award on behalf of Guido van Rossum from jury chair Gerard van Oortmerssen (left). Source: Nederland ICT.

It is with great sadness that we remember **Professor Michael Wilson**, Scientific Computing Department, STFC, who died suddenly and unexpectedly at the end of April. Michael joined RAL (Rutherford Appleton Laboratory) in 1986 after having been a researcher in psychology at Cambridge. He led a team of researchers who worked on user interface design and visualisation of scientific data. He also worked on projects concerning knowledge-based assistance to querying and explanation, and heterogeneous distributed information systems. His work on time-based media for the web was a major contribution to the SMIL language which later became the basis for the Multimedia Messaging System (MMS) used universally on mobile phones. He was also involved in work on web ontology languages, OWL and SKOS. Later, Michael became increasingly involved in promoting science and technology, managing the UK and Ireland office of the World Wide Web Consortium and leading the communications activities for e-Science. In his most recent role Michael was on secondment to the Department of Business Innovation and Skills, leading the work of the e-Infrastructure Leadership Council Secretariat. Within ERCIM Michael was well-known as an active member of the Executive Committee and project participant, contributing greatly to financial management, governance and the change to the new structure of ERCIM.

Michael will be sadly missed by colleagues. Our thoughts are with his family and friends at this time.

Cryptographer Ronald Cramer appointed Fellow of IACR

Ronald Cramer from CWI in Amsterdam and Leiden University was appointed Fellow of IACR. This was announced by the International Association for Cryptologic Research, IACR. The selection committee praised the mathematician for his contributions to the development of modern cryptography. He received the title "for fundamental contributions to cryptography, for sustained educational leadership in cryptography, and for service to the IACR". Cramer is the first researcher active in the Netherlands to receive this prestigious award. The ceremony takes place during the 33rd CRYPTO conference in August 2013 in Santa Barbara, Ca., USA. Cramer is, amongst others, known for the Cramer-Shoup encryption – which is adopted by an international ISO standard – and for the hash proof systems building on this research.

More information: http://homepages.cwi.nl/~cramer/
http://www.iacr.org/fellows/2013/cramer.html

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/

Swiss Informatics Research Association
c/o Professor Abraham Bernstein, Department of Informatics,
University of Zurich, Binzmühlestrasse 14, CH-8050 Zürich
http://www.sira.s-i.ch/

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/

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

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

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/