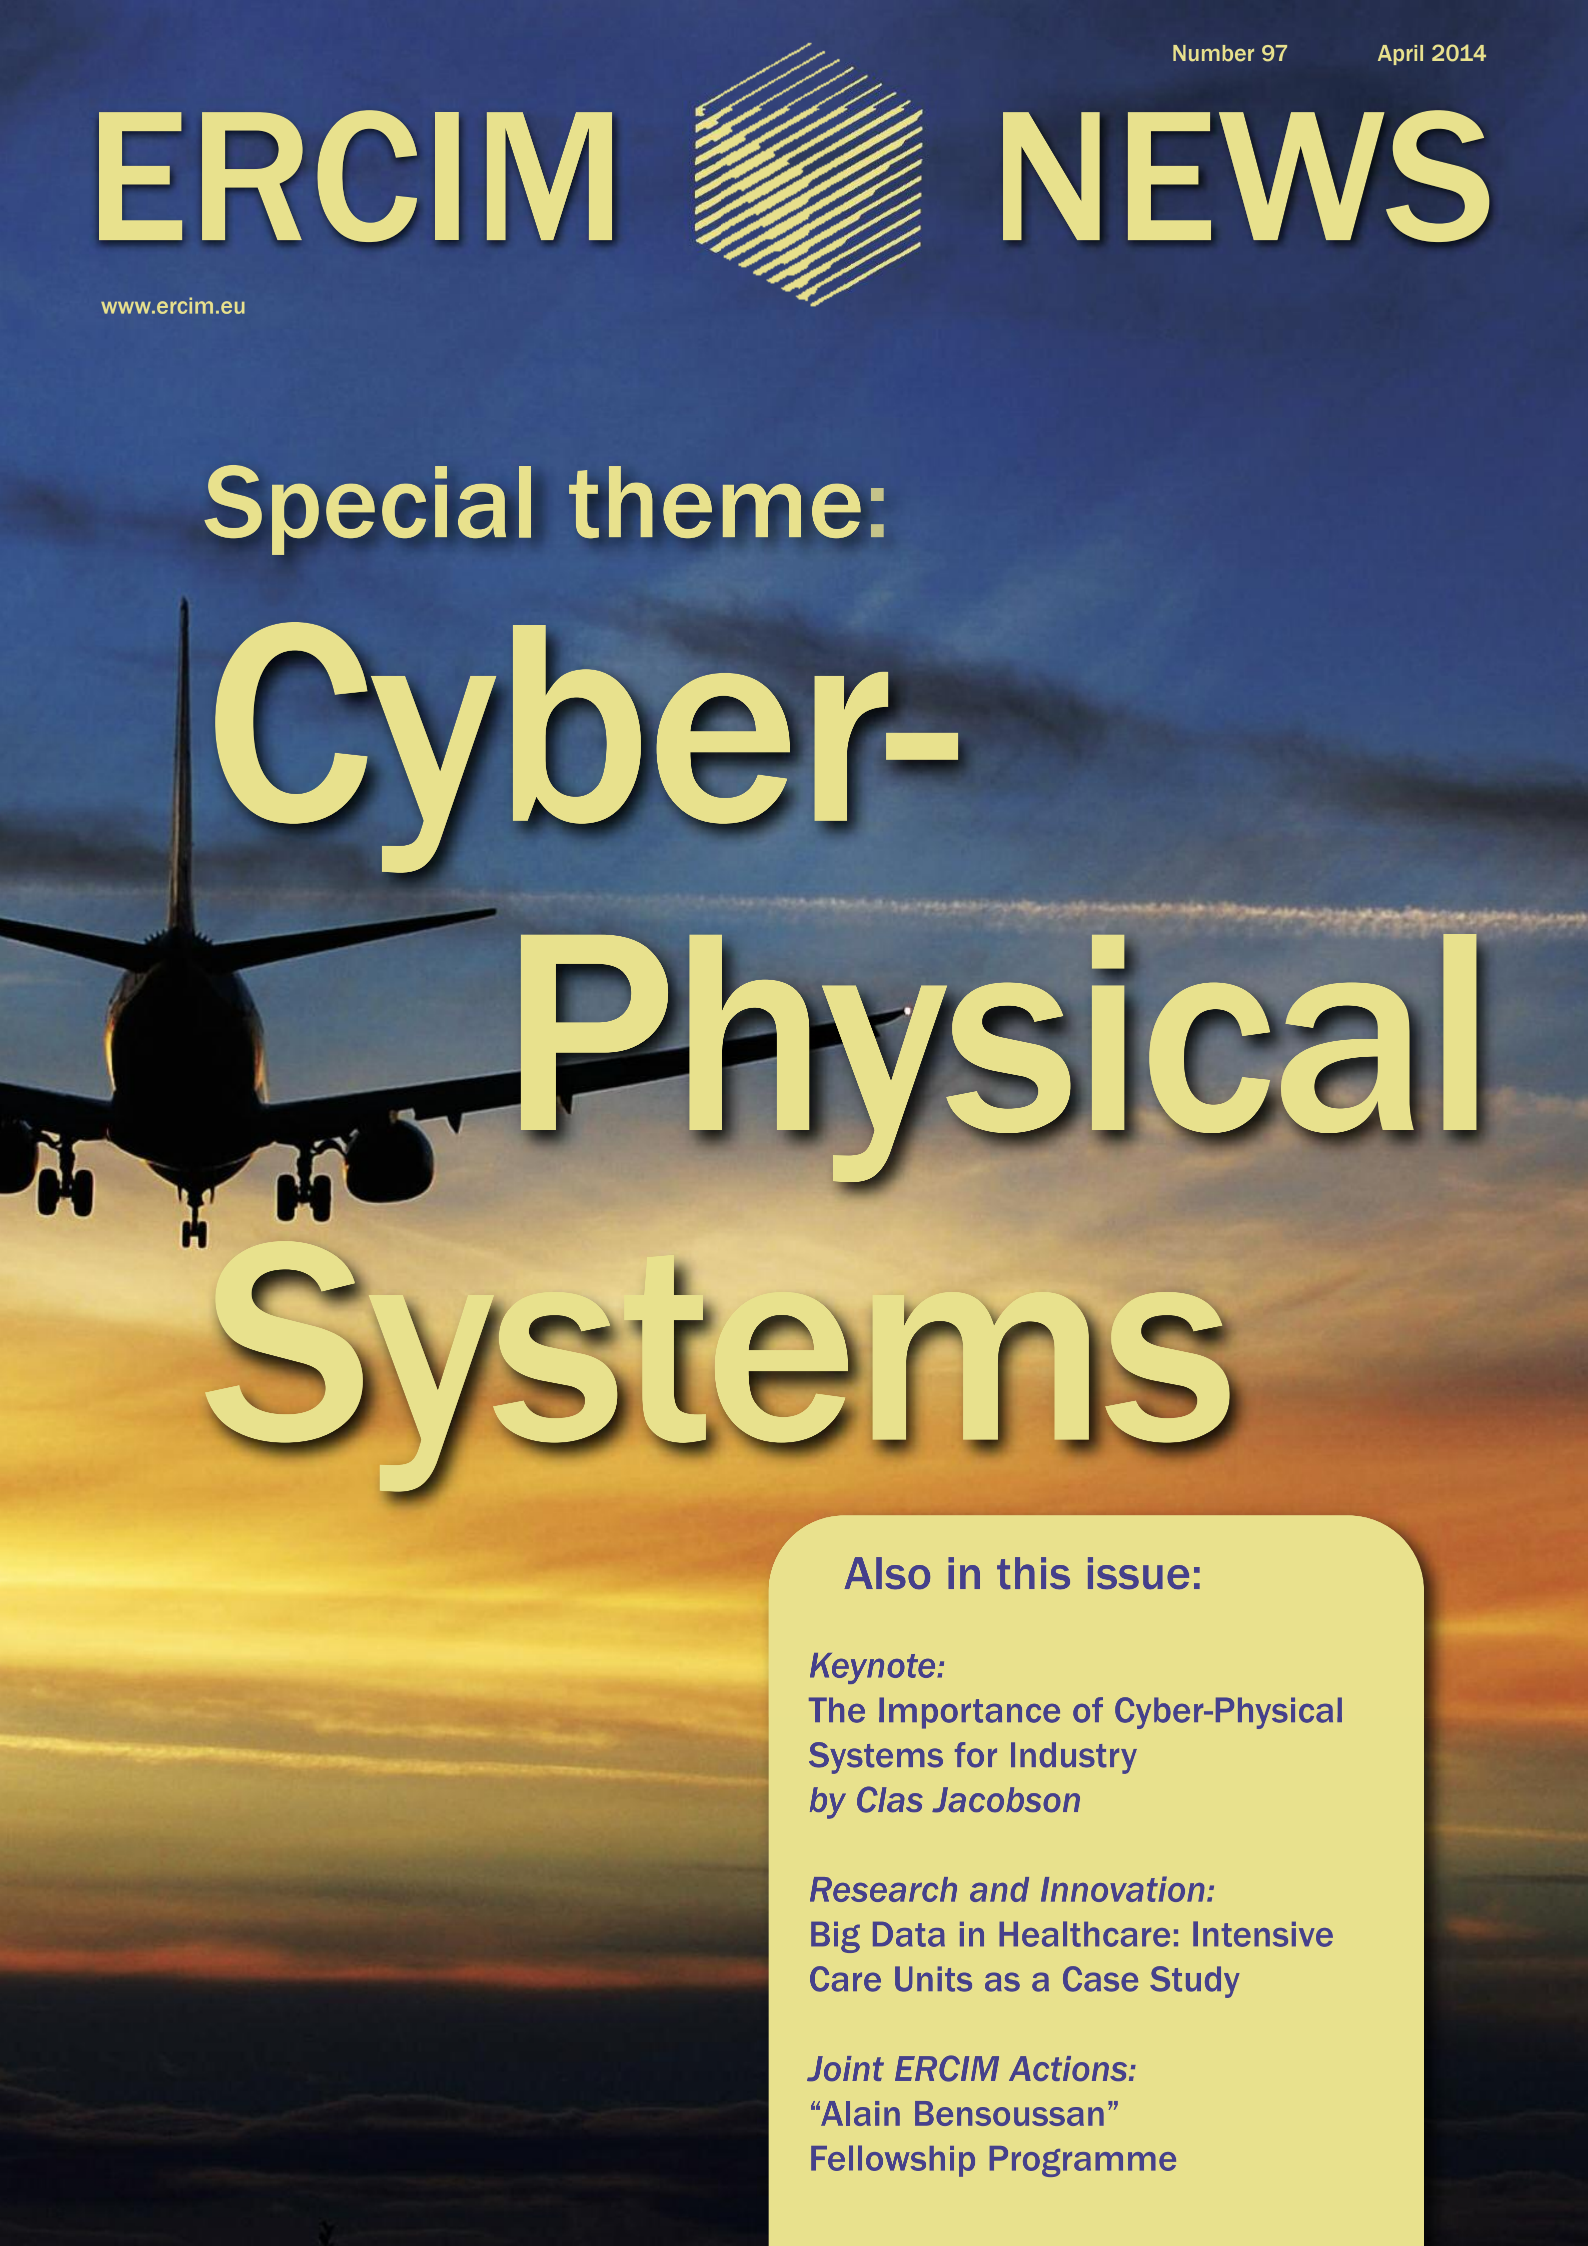

Number 97 April 2014

ERCIM NEWS

www.ercim.eu

Special theme:

Cyber-Physical Systems

Also in this issue:

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of about 6,500 printed copies and is also available online.

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

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

ERCIM News online edition
The online edition is published at http://ercim-news.ercim.eu/

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

Next issue
July 2014, Special theme: Smart Cities

Contents

KEYNOTE

JOINT ERCIM ACTIONS

SPECIAL THEME

The special theme section "Cyber-Physical Systems" has been coordinated by Maria Domenica Di Benedetto, University of L'Aquila, Center of Excellence DEWS, Françoise Lamnabhi-Lagarrigue, CNRS, Laboratoire des Signaux et Systèmes, and Erwin Schoitsch, AIT Austrian Institute of Technology/AARIT.

RESEARCH AND INNOVATION

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

EVENTS/IN BRIEF

*Clas Jacobson,
Chief Scientist for the
United Technologies
Systems & Controls
Engineering, USA*

The Importance of Cyber-Physical Systems for Industry

by Clas Jacobson

The presence and importance of Cyber Physical Systems is intended as the orchestration of networked computational resources with multi-physics (e.g., mechanical, chemical, electrical) systems in industry cannot be overemphasized. The engineering problems faced daily is managing dynamics, time, and concurrency in heterogeneous (interconnected) systems where the amount and complexity of intelligence (the cyber part) is growing rapidly and where software implementations are a major portion of system design, validation and ultimately verification. Research and education in this field is of strategic importance for business for years to come. Industry and Government in United States have posed CPS at the center of the engineering research agenda since 2007 when the President's Council of Advisors on Science and Technology (PCAST) highlighted CPS as the "number one" Priority for Federal Investments in Networking and Information Technology). Since 2010, the European research and industrial community has focused on CPS as paradigms for the future of systems. Acatech (German National Academy of Science and Engineering) developed an Integrated Research Agenda for Cyber-Physical Systems published in 2011.

To witness the importance placed on CPS by US industry, UTC together with Applied Material, Global Foundries, IBM, Intel, Micron, Raytheon and Texas Instruments is participating in an important research initiative in the US under the DARPA-SRC-SIA Semiconductor Technology Advanced Research Network (STARnet): The TerraSwarm center (http://www.terraswarm.org/) which is one of the six centers sponsored by the program consists of nine universities, 92 students, 22 faculty researchers, and 19 industry associate personnel. TerraSwarm, headquartered at University of California at Berkeley, is addressing the huge potential—and associated risks—of pervasive integration of smart, networked sensors and actuators into our connected world. The center aims to enable the simple, reliable, and secure deployment of a multiplicity of advanced distributed sense and control applications on shared, massively distributed, heterogeneous, and mostly uncoordinated platforms through an open and universal systems architecture.

IBM and UTC are the founding members of another initiative: the industrial cyber-physical systems center (iCyPhy, http://www.icyphy.org/), a research consortium built as a partnership between the University of California at Berkeley (UC Berkeley), the California Institute of Technology (Caltech), and the member companies. iCyPhy has been formed to identify and develop new engineering techniques that will make it easier to successfully build products and services that combine complex software, hardware and mechanical components. The iCyPhy team aims to identify repeatable, standardized and measurable processes to ensure high-performing systems, and to uncover new product development methods that can help companies reduce costs, increase reliability and innovate more quickly.

The interest demonstrated by industry in general is strong, but for UTC it is even stronger. United Technologies Corporation (UTC) is a global conglomerate with interest in Elevators (Otis) and Climate, Controls and Security, and in aerospace businesses (Pratt&Whitney gas turbine engines, Sikorsky rotorcraft and Aerospace Systems including air management, electric power distribution and landing systems). In addition to the research initiatives above, the UTC Institute for Advanced Systems Engineering at UConn with $10 million in planned investments, is a significant investment in educating the next generation of engineering leaders. The Institute focuses on methods for the design of CPS of interest to business such as aircraft, buildings and other highly engineered solutions that include numerous intelligent components and sub-systems. As we look to the future, this expertise will be crucial for engineers who want to design and create the most advanced products and solutions.

CPS are essential for the future of the system industry worldwide and collaboration at all levels, from practicing engineers to product architects, from tool makers to technology providers, from service to research, is necessary.

Call for Nominations

Cor Baayen Award 2014

The Cor Baayen Award is awarded each year to a promising young researcher in computer science and applied mathematics.

The award consists of 5000 Euro together with an award certificate. The selected fellow will be invited to the ERCIM meetings in autumn. A short article on the winner, together with the list of all candidates nominated, will be published in ERCIM News.

Nominees must have carried out their work in one of the 'ERCIM countries': Austria, Belgium, Cyprus, Czech Republic, Finland, France, Germany, Greece, Hungary, Italy, Luxembourg, Norway, Poland, Portugal, Spain, Sweden, Switzerland, The Netherlands and the United Kingdom. Nominees must have been awarded their PhD (or equivalent) after 30 April 2011.

The award was created in 1995 to honour the first ERCIM President.

Detailed information and online nomination form:
http://www.ercim.eu/activity/cor-baayen-award

ERCIM "Alain Bensoussan" Fellowship Programme

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

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

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

Conditions

Applicants must:

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

Application deadlines

30 April and 30 September

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

ERCIM Working Group Workshops in 2014

The purpose of an ERCIM Working Group is to build and maintain a network of ERCIM researchers in a particular scientific field. The working groups are open to any researcher in the specific scientific field. Their main activities are the organization of workshops and the preparation of common project proposals. The meetings and workshops planned for 2014 inlcude:

- ERCIM/EWICS/ARTEMIS Special Session of the ERCIM Dependable Embedded Systems Working Group at Euromicro SEAA on "Teaching, Education and Training for Dependable Embedded and Cyber-physical Systems", Verona, Italy, 27-29 August 2014
  http://esd.scienze.univr.it/dsd-seaa-2014/?page_id=200
- 2014 Security and Trust Management workshop co-loacated with the European Symposium on Research in Computer Security (ESORICS), Wroclaw, Poland, 7-11 September 2014
  http://esorics2014.pwr.wroc.pl/
- ERCIM/EWICS/ARTEMIS Workshop of the ERCIM Dependable Embedded Systems Working Group on Dependable Embedded and Cyber-Physical Systems and Systems-of-Systems at SAFECOMP 2014, Florence, Italy, 8-12 September 2014
  http://www.safecomp2014.unifi.it/decsos
- FMICS 14 - 19th International Workshop on Formal Methods for Industrial Critical Systems, September 11-12, 2014, Florence, Italy
  http://fmics2014.unifi.it/
- MUSCLE International Workshop on Computational Intelligence for Multimedia Understanding, Paris, 1-2 November 2014
  http://iwcim.isep.fr/index.html
- 7th International Conference of the ERCIM WG on Computational and Methodological Statistics, jointly held with CFE 2014, Pisa, Italy 6-8 December 2014
  http://www.cmstatistics.org/ERCIM2014/

More information:
http://www.ercim.eu/activity/workgroup
http://www.ercim.eu/activity/f-events

Introduction to the Special Theme

Cyber-Physical Systems

by Françoise Lamnabhi-Lagarrigue, Maria Domenica Di Benedetto and Erwin Schoitsch

Embedded systems - some visible, others integrated into every day equipment and devices – are becoming increasingly pervasive, and increasingly responsible for ensuring our comfort, health, services, safety and security. In combination and close interaction with the unpredictable real-world environment and humans, they become “Cyber-Physical Systems” (CPS), which act independently, co-operatively or as “systems-of-systems” composed of interconnected autonomous systems originally independently developed to fulfill dedicated tasks. Some of these systems may be older legacy systems.

The elements of the physical system are connected by the exchange of material or energy, while the elements of the control and management system are connected by communication networks which sometimes impose restrictions on the exchange of information, with software as a critical core element. Dependability - particularly resilience, safety and security - is a more complex issue than ever before for embedded distributed systems in the conventional sense. Prototype systems are the electrical grid, a power plant, an airplane or a ship, a manufacturing process with many cooperating elements, e.g., robots, machines, warehouses, conveyer belts, a large processing plant with many process units, a building with advanced distributed HVAC control, etc.

From the German Agenda CPS, Intermediate Results, Acatech, 2010: “Cyber-physical systems typically comprise embedded systems (as parts of devices, buildings, vehicles, routes, production plants, logistics and management processes etc.) that

- use sensors and actuators to gather physical data directly and to directly affect physical processes
- are connected to digital networks (wireless, wired, local, global)
- use globally available data and services
- possess a range of multi-modal human-machine interfaces (dedicated interfaces in devices, or unspecific interfaces accessed through browsers, etc.).”

The extreme importance of CPS for industrialized countries was highlighted in the August 2007 Report of the President's Council of Advisors on Science and Technology (PCAST) presenting a formal assessment of the Federal Networking and Information Technology R&D (NITRD). This report led the National Science Foundation to create the Cyber-Physical Systems Program in 2009, a major cross-cutting initiative at NSF. CPS-related conferences have been integrated into a major annual event, CPS WEEK, which rotates between the United States and Europe.

During the Workshop on Control of Cyber-Physical Systems held at the University of Notre Dame London Centre, 20-21 October 2012, organized by Panos Antsaklis, Vijay Gupta and Karl Henrik Johansson, the following definition was adopted: “Cyber-Physical Systems (CPS) are physical, chemical, biological and engineered systems whose operations are monitored, coordinated, controlled and integrated by a computing and communication core.” Desired characteristics or descriptive words of well-designed and engineered Cyber-Physical Systems include: coordinated, distributed, connected, heterogeneous, robust and responsive, providing new capability, adaptability, resiliency, safety, security, and usability, with feedback loops including often humans and the environment and related systems. This intimate coupling between the cyber and physical will be manifested across a broad range of length scales, from the nano-world to large-scale wide-area systems of systems. Correspondingly, CPSs will often have dynamics at a wide range of time-scales, such as at the discrete clock scale for some computational aspects to multi-day or even year-long time scales for system-wide properties and evolution.

Applications with enormous societal impact and economic benefit will be created. Cyber-Physical Systems will transform how we interact with the physical world just as the Internet transformed how we interact with one another. After several FP7 related projects or supporting actions, CPSs are now a targeted research area in Horizon 2020 (http://ec.europa.eu/programmes/horizon2020/) and public-private partnerships such as ECSEL (Electronic Components and Systems for European Leadership (http://ec.europa.eu/digital-agenda/en/time-ecsel), which integrates the former ARTEMIS (http://www.artemis-ju.eu), ENIAC (http://www.eniac.eu) and EPoSS (http://www.smart-systems-integration.org), which are dedicating great effort to this area of research. The importance of “Systems of Cyber-Physical Systems” is highlighted by EU-Roadmap activities such as the CPSoS project (see article on page 21), developing a strategic policy document “European Research and Innovation Agenda on Cyber-Physical Systems of Systems” (http://www.cpsos.eu/).

The design of such systems requires understanding the joint dynamics of computers, software, networks, physical, chemical and biological processes and humans in the loop. It is this study of joint dynamics that sets this discipline apart. Increasingly, CPSs are autonomous or semi-autonomous and cannot be designed as closed systems that operate in isolation; rather, the interaction and potential interference among smart components, among CPSs, and among CPSs and humans, requires coordinated, controlled, and cooperative behaviour. Very difficult challenges are posed for control of CPS owing to a variety of factors, such as very broad time and length scales, the presence of network communication and delays, coordination of many components (with an associated increased risk of component failure as the number of components grows to be very large), model reduction, tractability, etc.

Cyber-physical systems (of systems) cannot be designed and managed using theories and tools from only one domain, see the FP7 NoE HYCON2 deliverable D3.4.1 “Roadmap: From distributed and coordinated control to the management of cyber-physical systems of systems”. The behaviour of the large coupled physical part of the system must be modelled, simulated and analysed using methods from continuous systems theory, e.g. large-scale simulation, stability analysis, and the design of stabilizing control laws. On the other hand, methods and tools from computer science for the modelling of distributed discrete systems, for verification and testing, assume-guarantee methods, contract-based assertions etc. are indispensable to capture both the behaviour on the low level (discrete control logic, communication, effects of distributed computing) and global effects, in the latter case based on abstract models of complete subsystems. Logistic models as well as models and tools for performance analysis of discrete systems will be useful for system-wide performance analysis. Finally, theories from physics, e.g. structure formation in large systems, and from economics and social science (market mechanisms, evolution of beliefs and activity in large groups) may also prove to be useful.

The emphasis of this special issue is on an integrated perspective of real-time computing, communication, dynamics, and control of CPS, particularly of non (extra-) functional properties. A first set of contributions are dedicated to the representation and foundations of CPS including:
- Analysis, estimation, synthesis, design, and verification of CPS
- Platforms for smart infrastructure connecting the three layers : the Cloud layer, the Middle layer and the Physical layer
- Dependability (safety, security, reliability etc.) and resilience of CPS, including the Systems of Systems aspects
- Specification formalisms, including languages and software tools
- Real-time computing and resource-aware control for CPS
- Networks and protocols for CPS

Several contributions are dedicated to specific applications:
- Smart Electric Grid
- Water Management
- Autonomous Vehicles and Smart Transportation
- Next-generation Port and Air Traffic Management
- Security
- Smart Medical Technologies

As stated in the previous cited NoE HYCON2 Deliverable, “the size of cyber-physical systems of systems and their ‘multimodality’ or hybrid nature consisting of physical elements as well as quasi-continuous and discrete controls, communication channels, and local and system-wide optimization algorithms and management systems, implies that hierarchical and multi-domain approaches to their simulation, analysis and design are needed. These methods are currently not available. In individual domains, e.g. dynamic modelling and simulation, verification of discrete systems, design of controllers for guaranteed system stability on different system levels, and optimization of flows across the system, further progress can be expected that will have a high impact on the engineering of systems of systems. However, the simultaneous use and the integration of heterogeneous models and tools to capture system-wide properties reliably and with firm guarantees are currently completely open issues. The critical properties of cyber-physical systems of systems go beyond what can be analysed and designed systematically today: dynamic reconfiguration of complex systems, large-scale dynamics, waves of events or alarms, interaction of autonomous, selfish systems, and coupling of physical and computational elements via communication channels.”

Please contact:
Maria Domenica Di Benedetto
University of L’Aquila, Center of Excellence DEWS, EECI, Italy
E-mail: mariadomenica.dibenedetto@univaq.it

Francoise Lamnabhi-Lagarrigue
CNRS-Laboratoire des Signaux et Systèmes, EECI, France
E-mail: lamnabhi@lss.supelec.fr

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

Cyber-Physical Systems: Theoretical and Practical Challenges

by Ezio Bartocci, Oliver Hoeftberger and Radu Grosu

Cyber-Physical Systems (CPS) are the next generation of embedded ICT systems that are becoming pervasive in every aspect of our daily life. In this article we discuss some of the theoretical and practical challenges that we are currently facing in this area.

Most modern computing devices are ubiquitous embedded systems employed to monitor and control physical processes: cars, airplanes, automotive highway systems, air traffic management, etc. In the past, research on embedded systems tended to focus on the design optimization problems of these computational devices. In recent years, the focus has shifted towards the complex synergy between the computational elements and the physical environment with which they interact. The term Cyber-Physical Systems (CPS) was coined to refer to such interactions. In CPS, embedded computation and communication devices, together with sensors and actuators of the physical substratum, are federated in heterogeneous, open, systems-of-systems. Examples include smart cities, smart grids, medical devices, production lines, automotive controllers, and robotics.

Here we discuss some of the research problems that we are addressing on both the theory and practical application scenarios of CPS, such as automotive and medical systems.

Theoretical foundations and challenges

The behaviour of CPS is characterized by the nonlinear interaction between discrete (computing device) and continuous phenomena (the physical substratum). For this reason, research on hybrid systems plays a key role in modelling and analysing CPS. CPS are usually spatially distributed and they exhibit emergent behaviours (i.e. traffic jams, cyber-attacks), which result from interactions among system components, and which cease to exist when specific components are removed from the systems. Owing to their ubiquity and impact on every aspect of our life, one of the greatest challenges of this century is to efficiently predict the emergent behaviours of these systems. The complexity of their models, however, often hinders any attempt to exhaustively verify their safe behaviour.

An alternative method is to equip CPS with monitors and to predict emergent behaviours at runtime. This approach makes CPS self-aware, opening up new approaches to designing systems that can dynamically reconfigure themselves in order to adapt [1] to different circumstances. However, monitoring introduces a runtime overhead that may alter the timing-related behaviour of the system under scrutiny. In applications with real-time constraints, overhead control strategies may be necessary to reduce the overhead to acceptable levels by, for example, turning on and off the monitoring. Gaps in monitoring, however, introduce uncertainty in the monitoring results. Hence, our current research [1] also focuses on efficient techniques to quantify this uncertainty and compute an estimate of the current state of the system.

Automotive scenario

The extensive integration of sensor networks and computational power into automotive systems over recent years has enabled the development of various systems to assist the driver during monotonous driving conditions, and to protect the passenger from hazardous situations. This trend will inevitably lead to autonomous vehicles. In order to plan actions and reliably negotiate traffic, these vehicles need sensors capable of fault tolerant observation of their environment. Additionally, vehicle to vehicle (V2V) and vehicle to infrastructure (V2I) communication technology – also known as V2X communication – will be integrated into future automotive systems. V2X communication allows exchange of information between vehicles and roadside units about position and speed of vehicles, driving conditions on a particular road, accidents, or traffic jams. Information exchange thereby allows

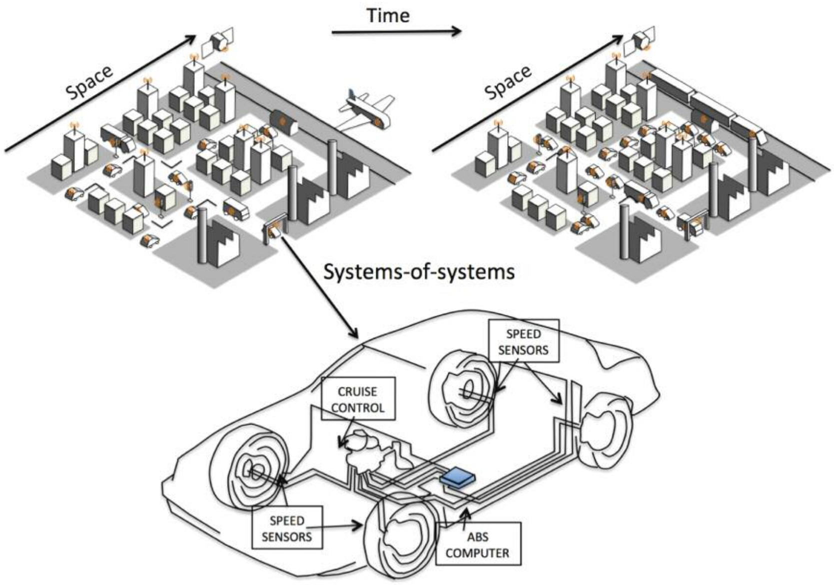

Figure 1: Multi-level CPS in the automotive domain.

traffic load to be distributed among several roads during rush hour, as well as preventing accidents and multiple collisions and sending automated emergency calls. Figure 1 schematizes different levels of an automotive system-of-systems consisting of sensor networks within a car, and the interaction of vehicles on a higher system level.

Fast error detection, fault tolerant system designs and new planning strategies are required to cope with the increasing failure rates of microchips owing to continuous shrinking of devices, as well as reliance on unreliable sources of information (e.g., information sent by other vehicles). Some of these problems can be solved by knowledge-based techniques, such as autonomous reconfiguration and substitution of faulty subsystems and components by using system ontologies [2].

Medical cyber-physical systems

Medical CPS refers to modern medical technologies in which sophisticated embedded systems equipped with network communication capabilities, are responsible for monitoring and controlling the physical dynamics of patients' bodies. Examples include proton therapy machines, electro-anatomical mapping and intervention, bio-compatible and implantable devices, and robotic prosthetics. Malfunctioning of these devices can have adverse consequences for the health of the patient. The verification, validation and certification of their reliability and safety are extremely important and still very challenging tasks, owing to the complexity of the involved interactions. The modelling and efficient simulation of the patient body will play a key role in the design and validation of Medical CPS and in the development of personalized treatment strategies. To this end, our research has largely focused on modelling and analysis techniques for cardiac dynamics to predict the onset of arterial and ventricular fibrillation. In [3] we show that with a normal desktop with GPU technology, it is possible to achieve simulation speeds in near real-time for complex spatial patterns indicative of cardiac arrhythmic disorders. Real-time simulation of organs without the need for supercomputers may soon facilitate the adoption of model-based clinical diagnostics and treatment planning.

Link:
http://ti.tuwien.ac.at/cps

References:

[1] E. Bartocci et al.: "Adaptive Runtime Verification", in proc. of RV 2012, LNCS, vol. 7687, Springer, pp. 168-182, 2012

[2] O. Höftberger, R. Obermaisser: "Ontology-based Runtime Reconfiguration of Distributed Embedded Real-Time Systems", in proc. of ISORC 2013

[3] E. Bartocci et al.: "Toward real-time simulation of cardiac dynamics", in proc. of CMSB '11, New York, NY, USA, pp. 103-112, 2011.

Please contact:

Radu Grosu
Vienna University of Technology
E-mail: radu.grosu@tuwien.ac.at
http://ti.tuwien.ac.at/cps/people/grosu

Ezio Bartocci
Vienna University of Technology
E-mail: ezio.bartocci@tuwien.ac.at
http://www.eziobartocci.com

Oliver Hoeftberger
Vienna University of Technology
E-mail:
oliver.hoeftberger@tuwien.ac.at

Modelling, Analysis and Co-Design of Wireless Control Networks

by Maria Domenica Di Benedetto and Alessandro D'Innocenzo

In the context of the EU FP7 Network of Excellence HYCON2, DEWS, a Centre of Excellence established at the University of L'Aquila in 2002, is developing a unifying mathematical framework that takes into account the joint dynamics of physical systems, control algorithms and wireless communication protocols. In addition, the framework provides an enabling technology for Cyber-Physical System (CPS) design and implementation.

The growing relevance of Cyber-Physical Systems (CPS) is due in part to the development of technologies that are essential for their effective deployment, namely: embedded systems and wireless communication networks. In particular, Wireless Sensor and actuator Networks (WSN) have greatly facilitated and enhanced the automated, remote and intelligent monitoring and control of a large variety of physical systems. These networks consist of many tiny, inexpensive and low-power devices, each incorporating sensing, actuation, processing, and wireless communication capabilities. These devices are now used in a number of application domains from industrial and building automation, to environmental, wildlife, and health monitoring, and disaster relief management. In these applications, computation and communication devices represent the enabling technology for WSN, while some of the most advanced WSN applications currently available are used for control tasks. Wireless Control Networks (WCN) are WSN applied to control.

The opportunities offered by WCN, and the increasing computing power of control nodes, are accompanied by some tough challenges. To advance the state-of-the-art in modelling, analysis and design of WCNs, classical control problems have to take into account the joint effect of Control, Computation and wireless Communication, evidencing the need to solve new theoretical and implementation problems. Despite some initial fundamental scientific achievements, the convergence of the "3Cs" has not yet had a significant impact on society and industry. In some cases, methodologies have been derived for general models, but their computational complexity results are intractable, i.e. they require the solution of NP-hard or even undecidable problems. In other cases, methodologies have been developed with an affordable computational complexity, but for simplified models that neglect the most relevant issues arising in real implementations - e.g., critical non-idealities, communication protocol dynamics and constraints, limited resources - rendering them practically useless. In general, scientific research on WCNs models the network non-idealities as aggregated network performance variables or disturbances, neglecting the dynamics introduced by the communication protocols.

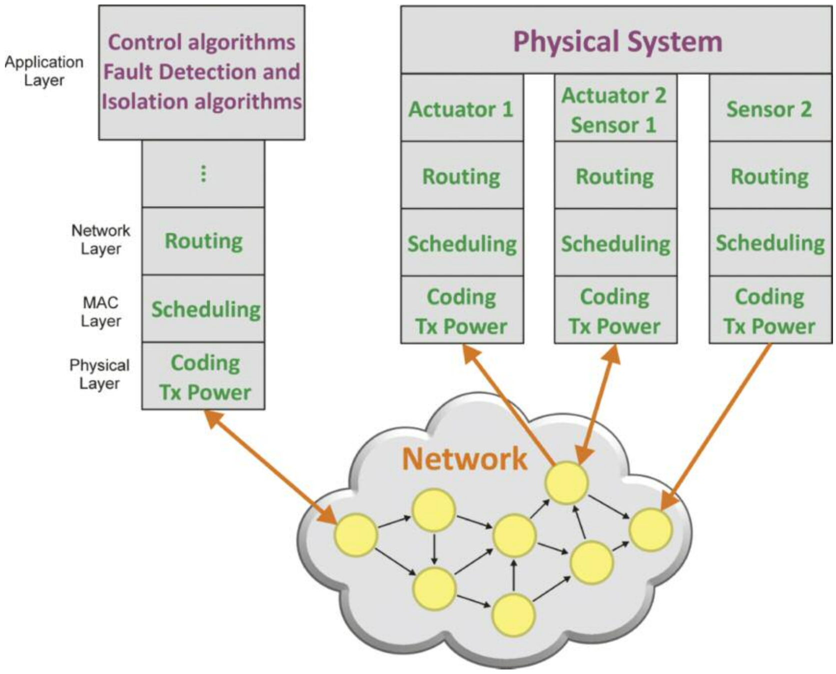

Figure 1: Control loop over wireless networking protocols.

In the context of the EU FP7 NoE HYCON2, multidisciplinary research at the DEWS Centre of Excellence, University of L'Aquila, Italy, is developing a unifying mathematical framework that takes into account the joint dynamics of physical systems and communication protocols (see Figure 1). By formally quantifying the interactions within the heterogeneous environment typical of WCN, we expect that powerful formal methods and tools, developed over decades by different (and often disconnected) scientific communities, can now be combined to solve problems arising in a WCN.

In collaboration with Prof. George J. Pappas, University of Pennsylvania, USA, and Prof. Karl H. Johansson, Royal Institute of Technology, Sweden, we proposed to model as a Switching Linear System the joint dynamics of a physical system, an embedded controller and of the MAC (scheduling) and Network (routing) ISO/OSI layers of a time-triggered communication protocol over a shared multi-hop communication network [1]. In particular, our framework makes it possible to model recently developed wireless industrial control protocols such as WirelessHART and ISA-100.

On the basis of this mathematical framework, we considered WCNs subject to failures and/or malicious attacks in the communication nodes [2]: we addressed and solved the problem of designing a set of controllers and communication protocol parameters, making it possible to detect and isolate failures on-the-fly and consequently apply an appropriate controller to stabilize the closed-loop system. We leveraged both algebraic and graph theoretic methods to derive necessary and sufficient conditions both on the physical system (i.e., its dynamics) and on the communication protocol (i.e., on topology, scheduling, routing and network coding) that guarantee stabilizability and allow failures to be detected and isolated.

In collaboration with Prof. Raphaël Jungers, Université Catholique de Louvain, Belgium, we recently began to study the necessary and sufficient conditions for stabilizability for a WCN switching linear model. Our goal has been to leverage the particular algebraic structure induced by the networking communication protocol in order to improve theoretical understanding of the dynamics at stake in these systems. This has enabled us to design tailored controllers whose performances are better than for classical switching systems. We have also tried to distinguish situations where exact algebraic methods are still applicable from situations where the time-varying delays make the system hard to control with exact methods, obliging the researcher to resort to conservative Lyapunov-like methods.

Future activities involve the extension of our unifying mathematical framework to accurately model all the ISO/OSI layers of the communication protocols for WCNs, as well as the development of novel control methodologies and algorithms for this framework.

Links:
http://dews.univaq.it/
http://www.hycon2.eu/
http://www.hartcomm.org/protocol/wihart/wireless_technology.html
www.isa.org/isa100

References:
[1] R. Alur, A. D'Innocenzo et al.: "Compositional Modeling and Analysis of Multi-Hop Control Networks", IEEE TAC, Special Issue on WSN, 56(10):2345-2357, 2011.
[2] A. D'Innocenzo, M.D. Di Benedetto, E. Serra: "Fault Tolerant Control of Multi-Hop Control Networks" IEEE TAC, 58(6):1377-1389, 2013.
[3] R.M. Jungers, A. D'Innocenzo, M.D. Di Benedetto: "Feedback stabilization of dynamical systems with switched delays", 51st IEEE CDC, pp. 1325-1330, 2012.

Please contact:
Alessandro D'Innocenzo
University of Aquila, Italy
E-mail:
alessandro.dinnocenzo@univaq.it

Observers for Nonlinear Systems with Delayed Measurements

by Tarek Ahmed-Ali, Iasson Karafyllis and FrançoiseLamnabhi-Lagarrigue

Designing observers for nonlinear systems with communication constraints is a crucial challenge for Cyber-Physical Systems (CPS). Small gain techniques are used to derive new kinds of robust sampled-data observers for wide classes of nonlinear systems with delayed measurements.

The design of observers for systems with sampled and delayed measurements has been attracting increasing attention. The interest is greatly motivated by many industrial applications. From a practical viewpoint, an observer is a software sensor using a model and online measurements, and providing a real-time estimation of physical quantities which cannot be measured online. Those estimates can then be used for system control (dynamical output feedback) or monitoring (fault diagnosis). The most used observer algorithm is the famous Kalman filter, and its convergence is well established for linear systems. Its variant known as Extended Kalman Filter [1] can be used for nonlinear systems, but with the problems of gain tuning and stability guarantee. After the Luenberger [2] result for linear time-invariant systems, various other works have addressed this problem for nonlinear systems.

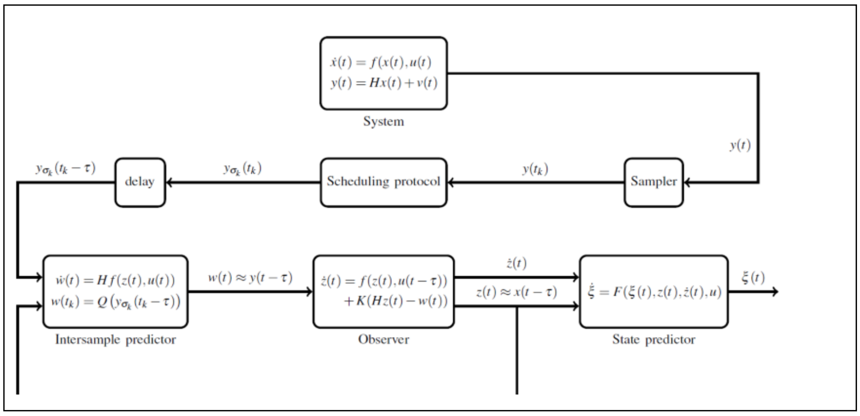

Figure 1: Structure of the observer z(t) given the system x(t).

System

$$\dot{x}_1 = -x_1^3 + x_2$$
$$\dot{x}_2 = -x_2^3 + sat(u)$$
$$y = x_1 + v$$

Observer

$$\dot{z}_1(t) = -z_1^3(t) + z_2(t) + k_1(z(t), x_1(t) + v(t), u(t))$$
$$\dot{z}_2(t) = -z_2^3(t) + sat(u(t)) + k_2(z(t), x_1(t) + v(t), u(t))$$
$$\hat{x}(t) = z(t)$$

where $v \in L^{\infty}_{loc}(\Re_+;\Re)$ denotes the measurement error and

$$k(z,y,u) := \begin{bmatrix} L_1 \\ L_2 \end{bmatrix}(z_1 - y), \text{ for all } (z,y,u) \in \Re^2 \times \Re \times \Re \text{ with } z_1^2 + z_2^2 \leq \sqrt{10}$$

$$k(z,y,u) := \begin{bmatrix} L_1 \\ L_2 \end{bmatrix}(z_1 - y) - \frac{\varphi(z,y,u)}{z_1^2 + z_2^2}\begin{bmatrix} z_1 \\ z_2 \end{bmatrix}, \text{ for all } (z,y,u) \in \Re^2 \times \Re \times \Re \text{ with } z_1^2 + z_2^2 > \sqrt{10}$$

with

$$L_1 = -p\frac{10402 + 9p}{2(32 - p^2)} \quad ; \quad L_2 = -\frac{1}{32 - p^2}\left(\frac{p^2}{64}(10400 + 9p) + 1\right), \; p \leq \frac{\sqrt{2}}{5376}$$

$$\varphi(z,y,u) := \max\left(0, -z_1^4 + (z_1 + sat(u))z_2 - z_2^4 + \frac{1}{8}\left(z_1^2 + z_2^2\right) + (z_1 - y)g\left(z_1^2 + z_2^2\right)(L_1 z_1 + L_2 z_2)\right)$$

Figure 2 (left): The System.
Figure 3 (right): Designed observer z(t) for estimating online x(t).

Two main approaches are described in the literature: one based on optimization, and one on analysis. In the latter approach, the most common idea is to characterize classes of systems which can be turned into another one, by change of variables (or immersion), allowing for a simpler observer design. The approach thus takes advantage of explicit structures, but remains dependent on the appropriate forms. Our work falls within this second approach. Very frequently we meet a system with: nonlinear characteristics, partial measurements, sampled measurements, measurements with errors, measurements with delays or uncertain sampling schedule. In fact, we rarely meet a system without one of these "annoying features"! The main question is then: How can we design a global exponential observer for such a system?

Our team recently proposed new robust observers with respect to measurement errors and perturbations of the sampling schedule which use robust global exponential state predictors [3]. The global exponential sampled-data observers design is accomplished by using a small gain approach. Sufficient conditions are provided, which involve both the delay and the sampling period. The structure of the proposed observer, illustrated in Figure 1 can be described as follows:

- A hybrid sampled-data observer is first used in order to utilize the sampled and delayed measurements and provide an estimate of the delayed state vector,
- The estimate of the delayed state vector is used by the robust global exponential predictor. The predictor provides an estimate of the current value of the state vector.

For example, for the system x(t) in Figure 2, the designed observer z(t) is given in Figure 3, and Figure 4 shows

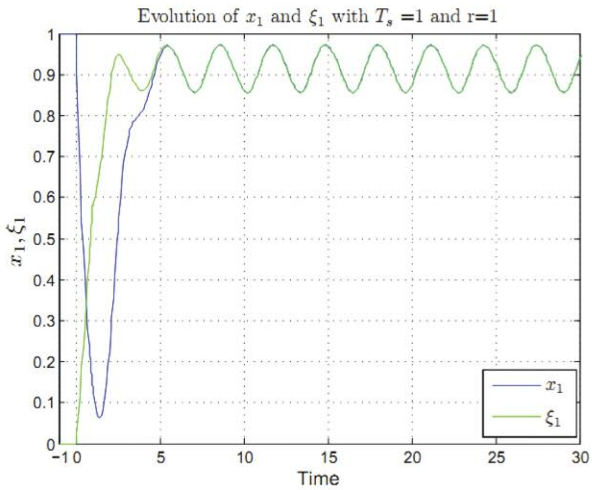

Figure 4: Time evolution of the first component of the observation error (x(t) – z(t)).

the exponential convergence of the observation error (x(t) – z(t)) to zero.

Future work will consider the application of recently proposed control schemes for which the input is delayed, the measurements are sampled and delayed and only an output is measured (the state vector is not available). These control schemes will consist of an observer for the delayed state vector, an inter-sample predictor for the output signal, an approximate predictor for the future value of the state vector and the nominal feedback law applied and computed for the predicted value of the future state vector.

References:

[1] R. E. Kalman: “A New Approach to Linear Filtering and Prediction Problems”, Transactions of the ASME - Journal of Basic Engineering, vol. 82, 1960, p. 35-45

[2] D. G. Luenberger: “An Introduction to Observers”, IEEE Transaction on Automatic Control, vol. 16, 1971, p. 596-602

[3] T. Ahmed-Ali, I. Karafyllis, F. Lamnabhi-Lagarrigue: “Global Exponential Sampled-Data Observers for Nonlinear Systems with Delayed Measurements”, Systems and Control Letters, 2013, 62(7) 539-549.

Please contact:

Tarek Ahmed-Ali Lagarrigue
ENISICAEN-GREYC UMR 6072, France
E-mail: tarek.ahmed-ali@ensicaen.fr

Iasson Karafyllis, National Technical University of Athens, Greece
E-mail: iasonkar@central.ntua.gr

Françoise Lamnabhi-Lagarrigue
CNRS-EECI SUPELEC, France
E-mail: lamnabhi@lss.supelec.fr

Platform-Aware Design of Embedded Controllers

by Dip Goswami, Twan Basten and Samarjit Chakraborty

Correctness, implementation efficiency and good quality of control (QoC) are essential for embedded controllers in cyber-physical systems. Awareness of the implementation platform plays a key role in achieving these goals. Recent results from Eindhoven University of Technology (TU/e) and Technische Universität München (TUM) report substantial improvements in the design of embedded controllers.

Embedded Control

Applications in domains such as automotive, avionics, professional printing, electron microscopy, semiconductor lithography, and medical imaging closely interact with physical processes. This interaction between hardware/software (cyber) components and physical processes in this class of cyber-physical systems (CPS) is regulated using feedback controllers. In an embedded implementation (see Figure 1), a feedback controller is realized by one or more tasks (e.g., tasks for sensor reading, control input computation and actuation) that are mapped onto one or more interconnected processors. Embedded platforms and control applications are typically constrained in terms of computation and communication resources, size, power, required latency and throughput, etc. Traditionally, the controller design and its implementation are done in isolation either with idealized assumptions such as instantaneous computation and communication, equidistant sampling, guaranteed message delivery, etc. or with worst-case assumptions on the controlled physical processes, the plant. In today’s constrained embedded realizations for CPS, this may lead to inefficient design solutions, suboptimal QoC, and even violations of functional correctness when the idealized assumptions are not met. To address these problems, recent research at TU/e and TUM advocates the joint design of controllers and their implementation platforms. One key idea is to allocate platform resources to embedded controllers based on the state of the plant.

State-based Resource Allocation

Timing of the feedback signals in a control loop plays a crucial role in QoC. Through time-triggered communication (e.g., TDMA) with perfectly synchronized processors and communication bus (see Figure 1), it is possible to allocate dedicated communication slots. This leads to a predictable communication delay which can potentially be translated into superior QoC. Therefore, time-triggered communication is in high demand in safety-critical application domains. Pure time-triggered communication is, however, typically bandwidth consuming and hence expensive in terms of resource usage, while the availability of time-triggered bandwidth is limited. On the other hand, event-triggered communication (e.g., CAN) induces a priority-driven sharing of bandwidth and, therefore, is bandwidth efficient. It results in variable and occasionally large communication delays. Delay variability in feedback loops

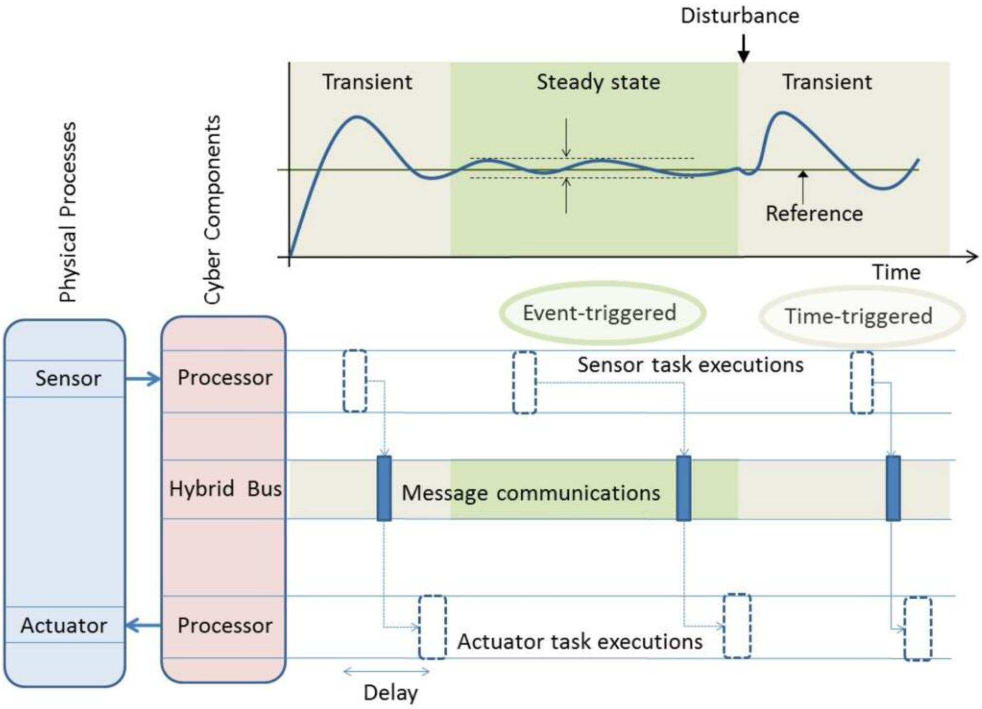

Figure 1: Control over hybrid communication bus

degrades QoC and may even lead to violations of functional correctness.

A hybrid communication protocol offers both time-triggered and event-triggered communication. The wide variety of functional and timing requirements in modern distributed systems (e.g., automotive E/E architectures) with mixed-criticality applications makes hybrid protocols (e.g., FlexRay in automotive systems) an attractive communication medium. By appropriately exploiting the advantages of both time-triggered and event-triggered paradigms, it is possible to achieve higher QoC within given constraints on time-triggered bandwidth.

The idea of state-based resource allocation is to map the feedback signals in a control loop dynamically to time-triggered and event-triggered implementations depending on the state of the plant being controlled [1]. Multiple control loops may share common time-triggered bandwidth where each loop has its own event-triggered bandwidth; only one loop can use the time-triggered bandwidth at any given point in time. The question then is which loop should have access to the time-triggered bandwidth. Upon the occurrence of a disturbance (i.e., an unpredicted change in one or more physical parameters), the system undergoes a state change, referred to as a transient. During this transient state, a control loop needs better quality of feedback (i.e., shorter delay, less jitter) for faster disturbance rejection. Hence, a loop which is in a transient state requests access to the time-triggered bandwidth. On the other hand, when the system is in steady state, it uses event-triggered bandwidth. A number of experiments conducted at TUM indicate significant improvement in QoC for a given time-triggered bandwidth [1]. The gains depend on the setup, in particular the used sharing policy and the criticality levels of the control loops.

Research Challenges

The need for correct and efficient embedded controllers with good QoC for cyber-physical systems necessitates further developments. Control theory needs to be extended to take into account the platform constraints of the embedded electronics [2]. Platform design needs the development of QoC-aware resource-allocation strategies [3]. But, most importantly, we need a general theoretical and conceptual framework for multi-objective and multi-domain design with techniques from both control theory and real-time embedded systems.

Links:
http://www.rcs.ei.tum.de/en/research/
http://www.es.ele.tue.nl/

References:
[1] Masrur et al.: "Timing analysis of cyber-physical applications for hybrid communication protocols" DATE 2012, IEEE
[2] Heemels et al.: "An introduction to event-triggered and self-triggered control" CDC 2012, IEEE.
[3] Aminifar et al.: "Control-quality driven design of cyber-physical systems with robustness guarantees" DATE 2013, IEEE.

Please contact:
Twan Basten
TU/e, the Netherlands
Tel: +31 40 247 5782
E-mail: a.a.basten@tue.nl

Samarjit Chakraborty
TUM, Germany
Tel: +49 89 289 23551
E-mail: samarjit@tum.de

Designing Cyber-Physical Systems for Smart Infrastructures: The Rainbow Platform

by Andrea Giordano, Giandomenico Spezzano and Andrea Vinci

Although Cyber-Physical System (CPS) technologies are essential for the creation of smart infrastructures, enabling the optimization and management of resources and facilities, they represent a design challenge. The Rainbow platform has been developed to facilitate the development of new CPS architectures.

The complexity of Cyber-Physical Systems (CPS) [1], resulting from their intrinsically distributed nature, the heterogeneity of physical elements (i.e. sensors and actuators), the lack of reliability in communications, the large-scale [2] and variability of the environments in which they are employed, makes data analysis and operation planning a very complex task. The Rainbow platform has been designed to address some of these issues. Rainbow hides heterogeneity by providing a virtual object concept, and addresses the distributed nature of CPS by introducing a distributed multi-agent mechanism. Rainbow aims to keep computation close to the sources of information (i.e., the physical devices), thus reducing the need for remote communication. Furthermore, it addresses the dynamic adaptivity requirements of CPS by fostering the use of swarm intelligence algorithms [3]. Rainbow also exploits the Cloud to carry out heavy computational tasks when needed.

An approach currently used when implementing a CPS architecture involves two layers: a physical and a remote (cloud) layer. The physical layer sends data to a remote server, which processes them and computes an operational plan. The remote server then sends the sequence of operations that must be executed to each device on the physical layer. The reasoning is performed in the remote layer. However, this solution cannot be applied when there are constraints on the response times, i.e., when a system needs to react rapidly to critical events that may compromise its integrity and functionality. Communication lags and remote processing can cause intolerable delays.

Rainbow relies on a new integrated vision that allows the design of a large-scale networked CPS based on the decentralization of control functions and the assistance of a remote layer to optimize their behaviour. Decentralization is obtained using a distributed multi-agent system in which the execution of a CPS application is carried out through cooperation among agents. This enables the exploitation of swarm intelligence techniques, where a complex global behaviour results from the interactions of simpler entities, acting solely on the basis of local knowledge and without the need for a global coordinator. These techniques introduce some appealing properties, such as adaptivity, fault tolerance and self-configuration.

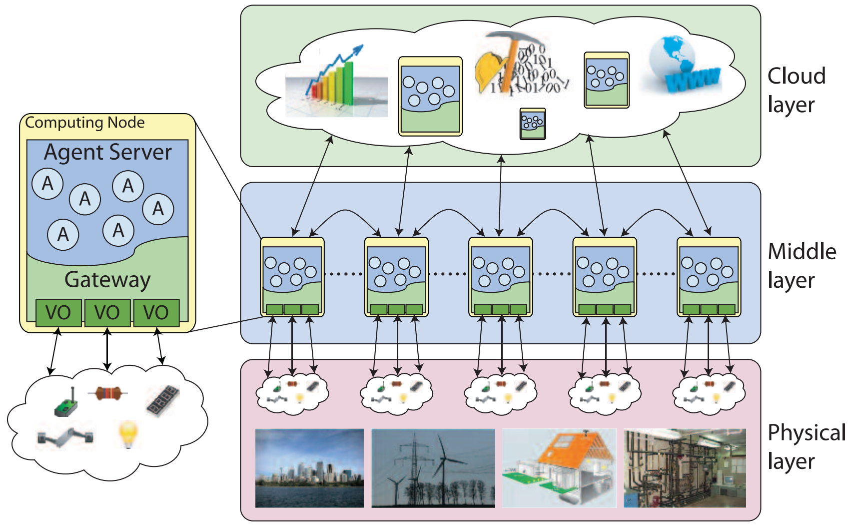

Figure 1: The Rainbow Platform

Rainbow is a three layer platform composed of a Cloud layer, a Middle layer and a Physical layer.

Sensors and actuators immersed in the physical environment are directly connected to computing nodes, which are single-board computers like Raspberry PI or BeagleBoard, where they are represented as virtual objects (VOs). The VO concept aims to hide the heterogeneity of the physical devices, in terms of capabilities, functionalities and communication protocols. VOs offer a transparent and ubiquitous access to the physical devices via a well-established API, and allow agents to connect directly without needing to consider proprietary drivers or needing to address fine-grained technological issues. The networking computing nodes represent the middle layer of the Rainbow platform.

Each computational node that hosts VOs also contains an agent server. VOs and Agent Servers are co-located in the same computing nodes in order to guarantee that agents directly exploit the physical part through VO abstraction. Instead of transferring data to a central processing unit, we transfer processes (injecting fine-grain agents on the nodes) toward the data sources. As a consequence, less data needs to be transferred over a long distance (i.e. toward remote hosts) and local access and computation is fostered in order to achieve good performance and scalability.

The upper layer of the Rainbow platform concerns the Cloud part. This layer addresses all those activities that cannot be executed in the middle layer, such as tasks that require high computational resources or when a historical data storage is mandatory. On the contrary, all the tasks where real time access to the physical part is required are executed in the middle layer. The data analysis executed by the remote server is used by the middle layer to optimize its operations and behaviour.

Rainbow has been developed at the Institute of High Performance Computing and Networking (ICAR-CNR) within the RES-NOVAE project, "Buildings, Roads, Networks, New Virtuous Targets for the Environment and Energy", funded by Italian Government. RES-NOVAE aims to implement new solutions for Smart Cities, and the Rainbow platform will be employed in many domains including power grids, water management, traffic control and environmental monitoring.

Link:
http://smartcityexhibition.it/it/partners/expo-partner/progetto-res-novae

References:
[1] A. Lee: "Cyber Physical Systems: Design Challenges", in proc. of the 11th IEEE Symposium on Object Oriented Real-Time Distributed Computing, IEEE Computer Society Washington, DC, USA, 2008.

[2] I. Stojmenovic: "Large scale cyber-physical systems: Distributed actuation, in-network processing and Machine-to-Machine communications", in proc. of the 2nd Mediterranean Conference on Embedded Computing (MECO), pp. 21-24, 2013.

[3] E. Bonabeau, M. Dorigo, G. Theraulaz: "Swarm Intelligence: From Natural to Artificial Systems" , New York, NY: Oxford University Press, Santa Fe Institute Studies in the Sciences of Complexity, Paper: ISBN 0-19-513159-2,1999.

Please contact:
Andrea Giordano, Giandomenico Spezzano, Andrea Vinci
ICAR-CNR, Italy,
E-mail: giordano@icar.cnr.it, spezzano@icar.cnr.it, vinci@icar.cnr.it

Cloud-Based Industrial Cyber-Physical Systems

by Armando Walter Colombo and Stamatis Karnouskos

The domain of industrial systems is increasingly changing as it adopts emerging Internet based concepts, technologies, tools and methodologies. The rapid advances in computational power, coupled with the benefits of the Cloud and its services, has the potential to give rise to a new generation of service-based industrial systems whose functionalities reside in-Cloud (Cyber) and on-devices and systems (Physical).

As we move towards an infrastructure that is increasingly dependent on monitoring of the real world, timely evaluation of data acquired and timely applicability of management (control), several new challenges arise. Future factories [1] are expected to be complex System of Systems (SoS) that will empower a new generation of applications and services that, as yet, are impossible to realise owing to technical and financial limitations. The European IMC-AESOP project is a visionary undertaking by key industrial players such as Schneider Electric, SAP, Honeywell and Microsoft, investigating the applicability of cloud-based Cyber-Physical Systems (CPS) [2] and Service Oriented Architectures (SOA) [1] in industrial systems, such as monitoring of paper production machines by FluidHouse (Finland) and lubrication monitoring and control systems in mining industry by LKAB (Sweden).

The Factory of the Future (FoF) [1] will rely on a large ecosystem of systems in which large scale collaboration will take place. Additionally, it is expected that CPS will harness the benefits of emerging Cloud Computing, including resource-flexibility and scalability. This not only has the potential to enhance the functionality of CPS, but will also enable a much wider consumption of CPS's data and services. The result will be a highly dynamic flat information-driven infrastructure that will empower the rapid development of better and more efficient next generation industrial applications whilst simultaneously satisfying the agility required by modern enterprises.

The first step in the infrastructure evolution was to empower the individual devices with Web services, and enable them to participate in a service-based infrastructure [1]. This is achieved by enabling them to: (i) expose their functionalities as services, and (ii) empower them to discover and call other (web) services to complement their own functionalities. The next step is to take advantage of modern capabilities in software and hardware, such as the Cloud and the benefits it offers. CPS have two key parts integrated in balance: the physical part for interacting with the physical environment (e.g. composed of sensors and actuator constellations), and the cyber part, which is the software part managing and enhancing the hardware capabilities of the CPS as well as its interaction with the cyber-world. The prevalence of the Cloud and its benefits enables us to expand the cyber-part of the CPS and distribute it on-device and in-Cloud. This gives rise to a new generation of Cloud-Based CPS that may revolutionize the Industrial Systems domain.

Cloud-based CPS and SOA may lead to a new information-driven infrastructure. Today, plant automation systems are composed and structured by several domains viewed and interacting in a hierarchical fashion mainly following the specifications of standard enterprise architectures. However, with the empowerment offered by modern SOA, the functionalities of each system or even device can be offered as one or more services of varying complexity, which may be hosted in the Cloud and composed by other (potentially cross-layer) services [3]. Hence, although the

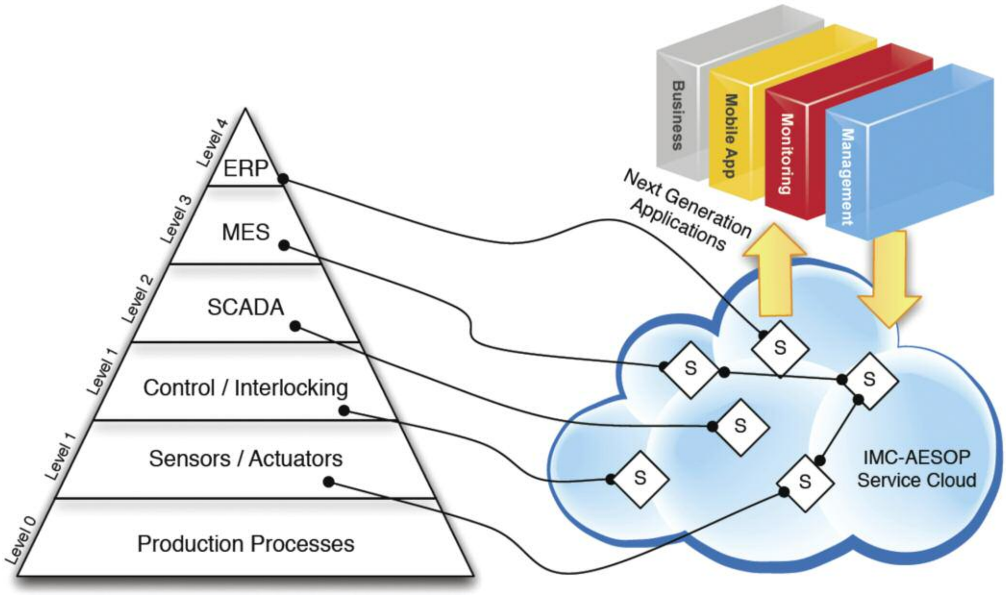

Figure 1: Industrial Automation Evolution: Complementing the traditional ISA-95 automation world view (pyramid on the left side) with a flat information-based infrastructure for dynamically composable services and applications (right side).

traditional hierarchical view coexists, there is now a flat information-based architecture that depends on a big variety of services exposed by the CPS [2], and their composition. Next generation industrial applications can now be rapidly composed by selecting and combining the new information and capabilities offered (as services in the Cloud) to realise their goals. The envisioned transition to the future Cloud-based industrial systems is depicted in Figure 1.

The Cloud enabled CPS have profound implications for the design, development, and operation of CPS. Although the device-specific part, i.e. the cyber (on-device) and physical part are still expected to work closely together and provide the basic functionalities for the CPS, the in-Cloud cyber part may evolve independently. Due to its nature, the in-Cloud part will require connectivity of the CPS with the Cloud where added-value sophisticated capabilities may reside. In contrast, the on-device cyber-part may consider opportunistic connections to the Cloud, but in general should operate autonomously and in-sync with the physical part. The nature of the functionalities as well as the degree of their dependence on external resources, computational power, operational scenarios, network connectivity etc. will be the key factor for hosting them on-device or in-Cloud. The IMC-AESOP project has proposed an architecture and services [3] that are based on this vision.

Links:
IMC-AESOP: http://www.imc-aesop.eu
Upcoming Book on IMC-AESOP vision and results, ISBN 978-3-319-05623-4: http://www.springer.com/engineering/production+engineering/book/978-3-319-05623-4
Cyber-Physical Systems: Uplifting Europe's innovation capacity, European Commission CPS event: http://www.amiando.com/cps-conference.html

References:
[1] A.W. Colombo, S. Karnouskos: "Towards the factory of the future: A service-oriented cross-layer infrastructure", in: ICT Shaping the World: A Scientific View, vol 65-81, European Telecommunications Standards Institute (ETSI), John Wiley and Sons, 2009

[2] Acatech: "Cyber-Physical Systems: Driving force for innovation in mobility, health, energy and production", tech. rep., acatech – National Academy of Science and Engineering, 2011, http://kwz.me/fJ

[3] S. Karnouskos et al.: "A SOA-based architecture for empowering future collaborative cloud-based industrial automation", in proc. of IECON 2012, Montréal, Canada.

Please contact:
Armando Walter Colombo
Schneider Electric Automation, Hochschule Emden/Leer, Germany
E-mail: Armando.Colombo@schneider-electric.com, awcolombo@technik-emden.de

DEECo: Software Engineering for Smart CPS

by Tomáš Bureš, Ilias Gerostathopoulos and Rima Al Ali

The functionality of smart complex networked Cyber-Physical Systems (CPS) is increasingly reliant on software. Software dominates to such an extent that smart CPS can be classified as software-intensive systems [1] – systems in which software is by far the most intricate and extensive constituent. The complexity of the software is underpinned by the fact that smart CPS are inherently distributed and need to combine collaborative behaviour with autonomicity, self-awareness and self-adaptation. The DEECo framework addresses the holistic development of such systems.

The complexity of software in smart Cyber-Physical Systems (CPS) means that software cannot be developed by ad-hoc means; instead, there is a need for systematic software engineering methods to reduce the development complexity and increase reliability and robustness by using appropriate software models and abstractions. A distinct challenge of smart CPS is that their software architecture undergoes continuous modifications: Components appear and disappear as CPS devices enter/exit the system, components form and dissolve cooperation groups as they start/finish a particular joint activity, and communication links are established/released depending on the actual availability of network connectivity. This is illustrated by the firefighter scenario (taken from one of our case studies), where firefighters (captured as CPS components) coordinate within and across mission sites and take advantage of stationary and mobile nodes (also captured as CPS components) existing in their vicinity to communicate with and sense their environment (see Figure 1).

To fulfill the requirements of smart CPS, researchers at Charles University in Prague, have developed the DEECo framework [2] for building smart CPS. This framework consists of a component model, a high-level design method, analyses, and a runtime framework – all tailored to the needs and specifics of smart CPS. It offers a synergy between component-based systems, ensemble-based systems, agent systems and control systems.

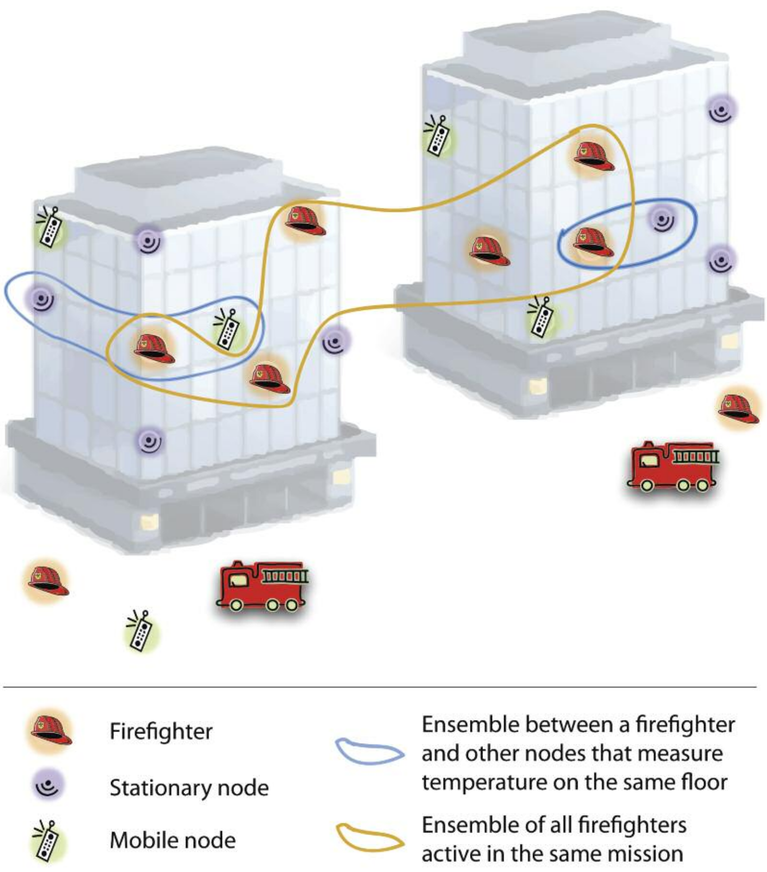

Figure 1: Use of components to model a smart CPS consisting of firefighters and other mobile and stationary nodes in the vicinity.

DEECo Framework

Component model - A component in DEECo [2] is the basic software unit of development and deployment. It constitutes state (referred to as knowledge) and behaviour, materialized into processes. Each process executes cyclically similarly to a real-time task. To achieve component interaction, components are dynamically assembled into collaboration groups called ensembles (e.g. a firefighter and all temperature sensors on the same floor, all firefighters active at a mission site). Within an ensemble, the components interact in terms of implicit knowledge exchange, which is handled by the execution environment.

High-level design - Invariant Refinement Method (IRM) [3] is a design method for DEECo-based systems. It captures high-level goals and requirements in terms of invariants, which describe the desired state of the system-to-be at every time instant (thus strengthening component autonomy and reliability). Invariants are maintained by component coordination. As a design activity, top-level invariants are iteratively decomposed into more fine-grained sub-invariants, essentially yielding a detailed contractual design of system implementation – either in terms of local component behaviour, corresponding to a component process, or in terms of component interaction, corresponding to an ensemble.

Analysis - Thanks to its well-defined semantics, DEECo allows timing properties to be quantified via static analysis and simulations. These include end-to-end response time and estimating level of inaccuracy of perceived knowledge depending on the communication latency and physical model of sensed values (given as differential equations).

Runtime Environment - The framework is backed by jDEECo, an implementation of DEECo in Java. This runtime environment includes scheduling of component processes, dynamic grouping of components into ensembles, and distributed knowledge exchange. Technically, jDEECo employs gossip-style network communication to uniformly address IP-based, as well as peer-to-peer broadcast-style WPAN networks.

Case Studies and Evaluation

To date, DEECo has been a successful part of the EU FP7 IP project ASCENS and employed in a number of case studies, including intelligent vehicle navigation, emergency coordination, and ad-hoc cloud deployment. Case studies have confirmed that DEECo represents a significant development simplification while preserving robustness and dependability properties of the designed system.

Future Work

In the next research and development steps, we intend to focus on enhancing the framework of efficient communication means in situations of limited connectivity, handling uncertainty of knowledge, and verification in the presence of dynamicity.

Links:
http://www.d3s.mff.cuni.cz
https://github.com/d3scomp/JDEECo

References:
[1] K. Beetz and W. Böhm: "Challenges in Engineering for Software-Intensive Embedded Systems", in Model-Based Engineering of Embedded Systems, Springer, 2012, 3–14
[2] T. Bures et al. "DEECo – an Ensemble-Based Component System", in proc. of CBSE'13, ACM, 2013, 81–90
[3] J. Keznikl et al. "Design of Ensemble-Based Component Systems by Invariant Refinement," in proc. of CBSE'13, ACM, 2013, 91–100.

Please contact:
Tomáš Bureš
Charles University in Prague, Czech Republic
E-mail: bures@d3s.mff.cuni.cz

The Software-Defined Network Revolution

by Marco Canini and Raphaël Jungers

Thanks to the introduction of Software-Defined Networking (SDN) [1], it is becoming possible, for the first time since the early days of computer networks, for operators to design and implement their own software in order to operate and customize the network to specific needs. But this comes with challenges involving many different disciplines.

Computer networks, such as the Internet, are playing an increasingly important role in our society. Despite this, they are too often designed and managed using ad-hoc techniques that are not as mature as other engineering and scientific disciplines. This is largely a consequence of the dependence on proprietary technologies and lack of open interfaces. Today's networks consist of various types of switches, routers and packet-processing middleboxes. Switches and routers are architecturally composed of two components: a data plane and a control plane. The data plane handles line-rate forwarding of packets that arrive at the device. The control plane handles the logic needed to configure the forwarding rules in the data plane. Traditional switches and routers have proprietary firmware and control logic implementations. This not only inhibits vendor interoperability, but also hampers flexibility, as operators cannot introduce new control plane functionality or protocol into a switch.

With the recent emergence of Software-Defined Networking (SDN), the networking industry is on the verge of a profound transformation. SDN is an emerging paradigm that aims to alleviate these issues by decoupling the data and control planes. SDN outsources the control of packet-forwarding switches to a set of software controllers running on a server cluster. This enables operators to run third-party software or create their own to build networks that meet their specific, end-to-end requirements, such as fine-grained policy enforcement, Quality of Service, and, of course, efficiency [2]. In SDN, control applications operate on a global, logically-centralized view, achieving a higher level of abstraction for managing networks. This view enables simplified programming models to define a high-level policy (i.e., the intended operational behaviour of the network) that can be compiled down to a collection of forwarding rules and installed in the data plane. This is in stark contrast to today's manual, error-prone, ad-hoc approaches.

As such, SDN promises to radically change the way networks are managed and operated. Indeed, for the first time since the early days of computer networks, it is becoming possible to enable network operators to design and implement their own network-controlling software and develop new functionality and services for end-users. A pioneering initiative in this direction is the OpenFlow protocol [3], which allows a software controller to dynamically manage the state and the behaviour of network elements.

However, while offering unprecedented potential, SDN is still in its infancy and much research effort is needed to better define its foundations: What are the abstractions that promote a simple yet expressive and efficient network-programming model? How will these abstractions yield efficient implementations that support rapid reactions to unplanned network events? Is it practically feasible to devise efficient abstractions that may still lend themselves to automated verification of integrated network systems? For instance, is it possible to prove that a network system satisfies key properties that would

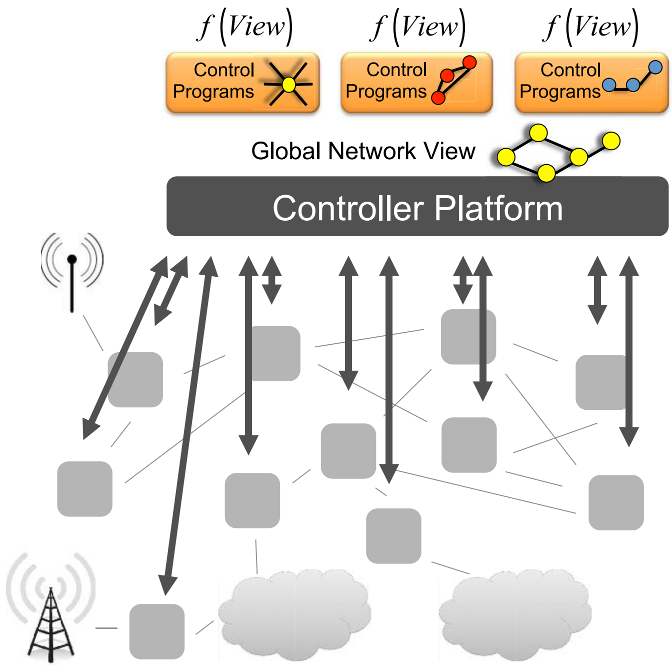

Figure 1: Example Software-defined network architecture: the control plane is separated from the data plane and control applications operate on top of a global (possibly virtualized) network view.

ensure smooth network operation? How to design a logically centralized (i.e., physically distributed) control platform that adequately and cost-effectively provides the required levels of availability, responsiveness and scalability? What control functionality can we keep distributed, what should we centralize? How can we achieve optimal dynamical resource allocations and fairness among users? Also, what methodologies can we devise to troubleshoot these systems when things in practice do not work according to theory (e.g., due to hardware malfunctioning)? Finally, what strategies should we use to deploy SDN into existing networks? We believe that these questions require us to apply principles and techniques from a large body of different disciplines, including Operations Research, Control Theory, Distributed Systems, Programming Languages, Software Engineering, Security, and Embedded Systems.

As such, SDN systems share many challenging technical problems with other Cyber-Physical Systems (CPS): (i) they have to cope with human and unpredictable demand, (ii) they are akin to a Systems-of-Systems structure comprising a two-tier distributed system, in which each node has heterogeneous processing capabilities and a large part of their computational resources are embedded (e.g., router firmware, NPUs, FPGAs), and (iii) constraints from the physical world are strong and various (e.g., link bandwidths and latencies, or failures due to random errors or malicious behaviours). On top of that, SDN systems constitute an opportunity for other CPS that might need to rely on a programmable and modular network of computing and forwarding resources. Perhaps even more than for other CPS, SDN seems to have only one limitation to our ability to control them: our own imagination.

At UCLouvain, we have built a team centred on the new SDN paradigm in order to leverage its possibilities and develop the new generation of computer networks. Our goal is to build a rigorous and systematic methodology for designing and operating SDN systems, and we are convinced that the only feasible approach is to combine advances in Electrical Engineering, Computer Science and Applied Mathematics, ranging from computer networking to software and requirements engineering, optimization and control theory. Our research group includes faculty, students, and post-docs from various backgrounds in these disciplines. More details on this project are available at our website: http://sites.uclouvain.be/arc-sdn/. PhD and Post-doc positions are available within our team.

Link:
http://sites.uclouvain.be/arc-sdn/

References:

[1] N. Feamster, J. Rexford, E. Zegura: "The road to SDN: An intellectual history of programmable networks," ACM Queue 11, 12 (December 2013)

[2] "Ethane: taking control of the enterprise," in Proceedings of ACM SIGCOMM, 2007

[3] N. McKeown et al.: "OpenFlow: enabling innovation in campus networks," ACM SIGCOMM Comput. Commun. Rev. 38, 2 (March 2008), 69–74.

Please contact:

Marco Canini and Raphaël Jungers
Université catholique de Louvain, Belgium.
E-mail: marco.canini@uclouvain.be, raphael.jungers@uclouvain.be

Securing Interconnected Cyber-Physical Systems through Strategic Information Sharing

by Florian Skopik and Thomas Bleier

Cyber-attacks are becoming increasingly sophisticated, targeted and coordinated. Consequently, new paradigms are required for detecting attacks in critical cyber-physical systems, such as the smart grid. Many attack detection tasks are currently performed within individual organizations, and there is little cross-organizational security information sharing. Information sharing is a crucial step to acquiring a thorough understanding of large-scale cyber-attack situations, and is necessary to warn others against threats.

The smooth operation of critical infrastructures, such as telecommunications or electricity supply is essential for our society. In recent years, however, operators of critical infrastructures have increasingly struggled with cyber security problems [1]. Through the use of ICT standard products and increasing network interdependencies [2], the surfaces and channels of attack have multiplied. New approaches are required to tackle this serious security situation. One promising approach is the exchange of network monitoring data and status information of critical services across organizational boundaries with strategic partners and national authorities. The main goal is to create an extensive situational awareness picture about potential threats and ongoing incidents, which is a prerequisite for effective preparation and assistance in large-scale incidents.

The Challenges of Sensitive Information Sharing

In practice, security information sharing is usually accomplished via ad-hoc and informal relationships. Often, national Computer Emergency Response Teams (CERTs) assume the role of a contact point for coordinating and aggregating security incidence reports. However, the information that is provided is usually not targeted to particular vertical industry sectors, such as power grids. We suggest that sector-oriented views, along with rich information and experience reports are required to make such platforms more effective. Furthermore, there is a crucial trade-off to be considered: existing platforms require information to be verified centrally (in order to avoid hoaxes); therefore, the speed of information distribution suffers. Timeliness of information is very important when protecting against aggressive attackers and zero-day exploits. Consequently, there is a need for new standards that employ suitable direct sharing models, which allow the targeted exchange of specific information about discovered vulnerabilities of ICT systems utilized in critical infrastructure control systems, as well as current threats (such as new SCADA (supervisory control and data acquisition)-targeted malware) and recent incidents. The application of these standards further implies the existence of a federated trust and reputation model to address reservations of users, and to attract a critical mass of users. This is also in line with the objectives of the recently introduced European "Network and Information Security" (NIS) directive [3]. NIS explicitly recommends the implementation of national cyber security centres, which are not only informed about the security status of the national critical infrastructure providers, but also play a coordinating role in the prevention of, or protection from attacks.

The Project "Cyber Incident Information Sharing" (CIIS)

The research project "Cyber Incident Information Sharing" (CIIS) aims to develop methods and technologies for the exchange of information on cyber incidents to better defend against cyber-attacks and to streamline the analysis of current threats. CIIS assumes that preventive mechanisms and resilient architectures for critical infrastructures are in place. If, despite these measures, an attack is successful, a number of counter measures are being studied (see Figure 1). In particular, a novel anomaly detection mechanism is being developed, which utilizes large-scale event processing and correlation to detect anomalous system behaviour. This anomaly detection approach is specifically designed to enable incident information sharing on top in a privacy-preserving way. In order to draw conclusions from exchanged data about

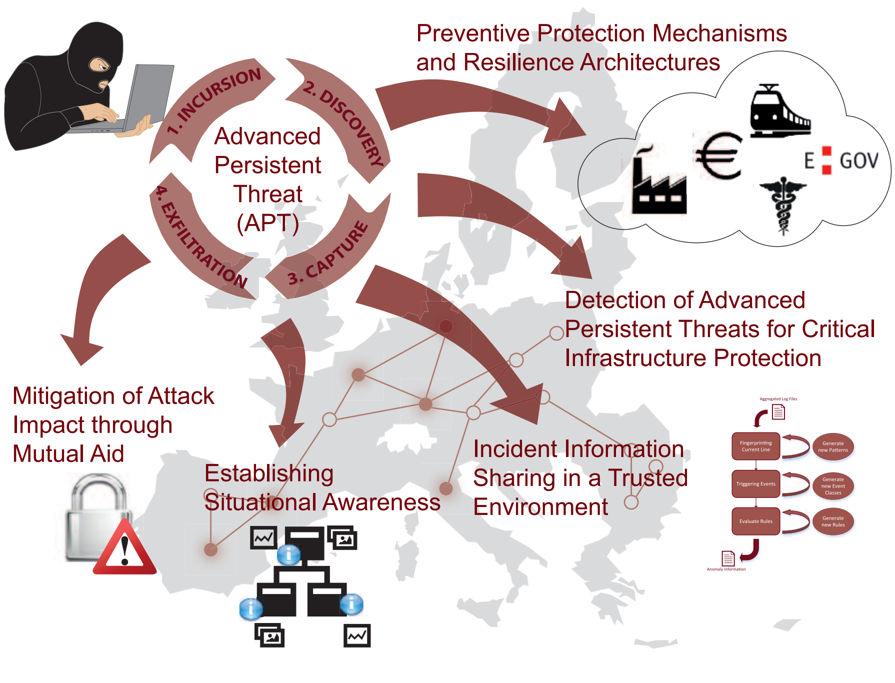

Figure 1: Overview of CIIS Research Efforts

current threats and to enable an assessment of current risks and concrete actions to be deployed, further tools for interactive analysis and visualization and establishing situational awareness for the technical operations personnel are being developed. Exchanging best practices to deal with cyber-attacks and providing aid to affected organizations mitigates the impact of successful cyber-attacks and ensures fast recovery.

CIIS Project Consortium

In order to attain these ambitious goals and finally ensure the wide applicability of developed tools, the project consortium consists of a vital mix of experts with a thorough knowledge of the area of security (Austrian Institute of Technology), visualization and situational awareness (VRVis Forschungs GmbH), and critical infrastructure operation (T-Systems Austria, Energie AG). Additionally, the involvement of experts in the field of social sciences (Institute for the Sociology of Law and Criminology, Netelligenz e.U.) is essential to illuminate issues of trust and privacy in the exchange of sensitive information and to investigate incentives to participate in such initiatives. In addition to the development of scientific methods, the proper demonstration of the applicability in a real-world environment is of paramount importance in order to test and evaluate the planned system. Pilot cases are supported on the one hand through operators of critical infrastructures, on the other hand by national authorities (Ministry of the Interior, Ministry of Defence and Sports).

Moreover, CIIS consortium members are actively involved in international initiatives, such as the Multinational Capability Development Campaign (MCDC) which enables beneficial collaborations across Austria's borders. CIIS is a two-year national project running from 2013 to 2015 and is financially supported by the Austrian security-research program KIRAS and by the Austrian Ministry for Transport, Innovation and Technology (BMVIT).

Link: http://www.ait.ac.at/ciis

References:

[1] R. Langner: "Stuxnet: Dissecting a cyberwarfare weapon", Security & Privacy, IEEE, 9(3), 49-51, 2011

[2] S. M. Rinaldi: "Modeling and simulating critical infrastructures and their interdependencies", IEEE HICSS, 2004

[3] Commission Proposal for a Directive concerning measures to ensure a high common level of network and information security across the Union; http://ec.europa.eu/digital-agenda/en/news/commission-proposal-directive-concerning-measures-ensure-high-common-level-network-and, 07/02/2013.

Please contact:

Florian Skopik and Thomas Bleier
Austrian Institute of Technology, AIT/AARIT, Austria
E-mail: florian.skopik@ait.ac.at, thomas.bleier@ait.ac.at

A European Roadmap on Cyber-Physical Systems of Systems

by Michel Reniers and Sebastian Engell

We are developing a strategic policy document "European Research and Innovation Agenda on Cyber-Physical Systems of Systems"

Cyber-Physical Systems of Systems (CPSoS) is a 30-month Support Action, funded by the EU (under FP7), that provides an exchange platform for Systems of Systems (SoS) related projects and communities. It focuses on the challenges posed by the engineering and the operation of technical systems in which computing and communication systems interact with large complex physical systems ([1], [2], [3]). Its approach is simultaneously integrative - aiming to bring together knowledge from different SoS related communities - and applications driven. It will integrate the different approaches to the design, analysis and control of systems of systems that are pursued by different communities in theory and application, and relate the methods and tools proposed for dealing with SoS to key application domains which are important for Europe's competiveness as well as for the well-being of its citizens. The project will conduct an in-depth examination of application-specific issues, capture cross-industry and cross-application findings and propose new avenues for SoS analysis, design and control. This will contribute to the development of a science of systems of systems and a European R&I agenda on SoS, involving different scientific communities, application domain experts, end-users and vendors.

The final outcomes of the project will be:

- Analysis of industrial and societal needs and of the state-of-the art of tools, theories, and methods for cyber-physical SoS, and identification of synergies, open issues and promising transdisciplinary research directions,
- Definition of a research agenda for cyber-physical systems of systems and their engineering needs,
- Contribution to coordination of SoS related projects, building up a network of key researchers and application domain experts in the area, and raising public awareness of the impact of research on systems of systems engineering, analysis and control.

The contributions will be summarized in a strategic policy document "European Research and Innovation Agenda on Cyber-physical Systems of Systems", supported by a set of in-depth technical papers. The project will prepare the EU stakeholders for extracting a competitive advantage from the recent

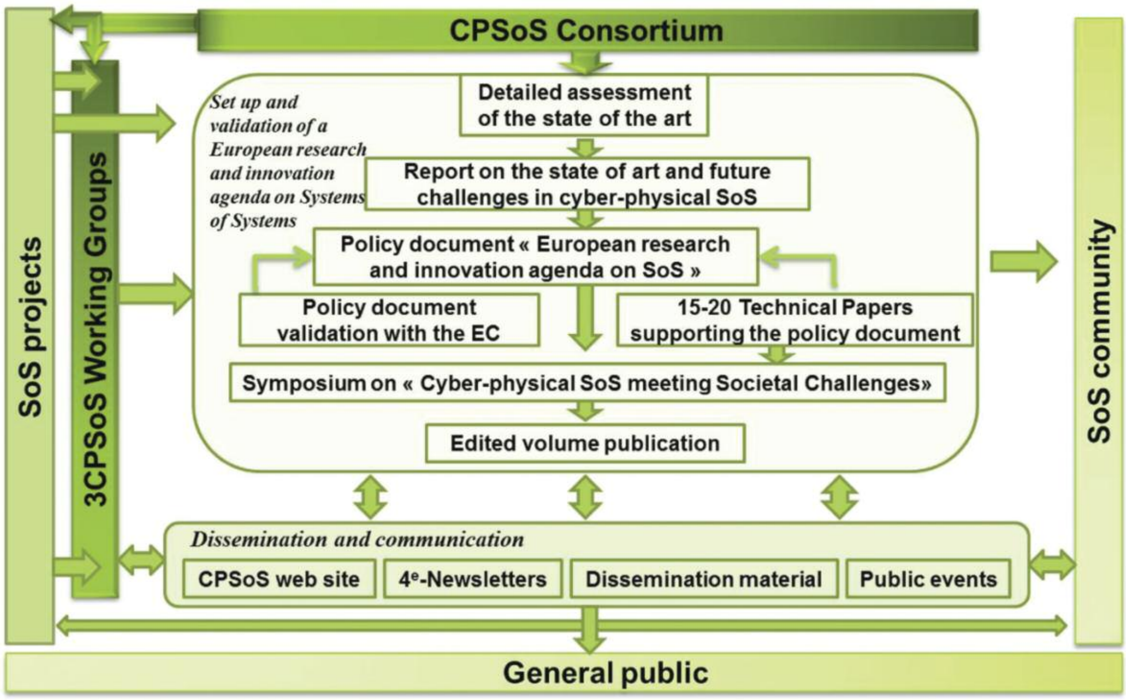

Figure 1: Workflow of the CPSoS project.

and the future developments in the area of SoS. The project's approach is summarized in Figure 1, which shows several of the steps taken in the project to achieve the above outcomes.

Three working groups have been created on:
1. Systems of Systems in Transportation and Logistics (Chair: Prof. Haydn Thompson, Haydn Consulting),
2. Physically Connected Systems of Systems (Chair: Prof. Sebastian Engell, TU Dortmund),
3. Tools for Systems of Systems Engineering and Management (Chair: Prof. Wan Fokkink, TU Eindhoven).

On 31 January 2014, a first meeting of the members of the working groups was held in Düsseldorf, Germany. The discussions addressed the following main questions:
- What are typical use cases of cyber-physical systems of systems?
- What are the main difficulties encountered in the engineering, realization and operation of CPSoS?
- What are the specific demands and challenges for advanced methods and tools for CPSoS engineering and operation?
- What are the most important open research questions for CPSoS over the next five years?

Numerous examples for cyber-physical systems of systems were presented and discussed, and the three Working Groups provided prioritized lists of future topics for research and development in the area of CPSoS. These will be discussed with domain experts, within the consortium, and elaborated in more detail with the help of the members of the working groups in order to provide a first draft of a strategy document by June 2014.

A first draft of the roadmap paper will be presented to the European Commission in July 2014. The roadmap paper will then be further discussed in workshops with external participation in the autumn of 2014 to obtain feedback and further input. A Workshop on Tools and Methods for Cyber-Physical Systems of Systems is planned to take place on September 12, 2014, co-located with iFM 2014 in Bertinoro, Italy.

More information, including the composition of the consortium, the members of the working groups, papers of importance to the domain and discussion papers and final outcomes can be found on the project website.

Researchers and industrial stakeholders with an interest in cyber-physical systems of systems are invited to contribute in the following ways:
- send suggestions via the short questionnaire (http://www.cpsos.eu/?page_id=462)
- send any other suggestions to the project coordinator Prof. Sebastian Engell (via Web site, "Contact" page, or via email s.engell@bci.tu-dortmund.de)
- subscribe to the project events' announcements and newsletters.

The CPSoS project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement No 611115.

Link: http://www.cpsos.eu

References:
[1] M.W. Maier: "Architecting Principles for System of Systems", Systems Engineering, Vol. 1, No. 4, 1998 pp. 267-284

[2] A. Lee: "Cyber Physical Systems: Design Challenges", in proc. of the 11th IEEE Symposium on Object Oriented Real-Time Distributed Computing, IEEE Computer Society Washington DC, USA, 2008, pp. 363-369

[3] J. Fitzgerald et al.: "Model-based Engineering for Systems of Systems: the COMPASS Manifesto", COMPASS Technical Report, 2013.

Please contact:
Sebastian Engell
CPSoS project coordinator
Technische Universität Dortmund, Germany
E-mail: s.engell@bci.tu-dortmund.de

Linking Design Disciplines in Cyber-Physical System Development: The DESTECS/Crescendo Technology

by John Fitzgerald, Peter Gorm Larsen and Marcel Verhoef

A new generation of innovative products is emerging that link computing elements – hardware and software – with physical processes. They surround us in our daily life and we are becoming dependent upon their embedded intelligence, as well as their interconnections. Successful design of these "Cyber-Physical Systems" (CPSs) requires close collaboration between all engineering disciplines involved, and rapid innovation is often required in order to meet short windows of opportunity under economically volatile circumstances. The DESTECS consortium has developed co-modelling technology to support cooperative working, and has delivered methods validated in industry case studies. These results are embodied in a new tool set that reduces the effort required to perform design iterations, and improves their impact, right from the outset of design.

Cyber-Physical Systems (CPSs) are challenging to develop because the cyber and physical parts have different kinds of semantic foundation, thus a communication hurdle exists between the engineering disciplines [1]. CPS methods and tools must bridge this gap in order to create a common basis for design and analysis. Various approaches have been suggested to overcome this challenge [2]. In the DESTECS project, we demonstrated how model-based techniques could support multidisciplinary design by bringing together Continuous Time (CT) models of physical systems with Discrete Event (DE) models of the embedded controller behaviour. Instead of requiring control and software engineers to adopt a new common design notation, we focused on harnessing models of physics and software within a common framework with a sound formal semantics, enabling the resulting co-model to be simulated as a whole. The project resulted in the development of the Crescendo technology [3], which enables stakeholders from different engineering backgrounds to continue using the methods that are familiar to them, while jointly engaging in holistic analysis and tradeoff across the disciplinary divide. This enables cross-discipline dialogue in the early stages of design, often exposing hidden (domain specific) assumptions that would otherwise be found during integration and test. This not only reduces cost and improves time-to-market and product quality, but it also promotes the impact analysis of design alternatives, providing a better basis for decision making during product development. In Crescendo, the coupling between discrete and continuous models is formalized in a "contract" that describes the state variables and events that are exchanged between the connected models, as well as the shared design parameters and model variables. All these artifacts are described and managed in a single place and automatically shared among models, which improves consistency. Model variables are set from a script, enabling engineers to undertake sensitivity analysis on the model by performing a parameter sweep or to perform fault resilience analysis by enabling or disabling parts of the model that represent error behaviours. Providing the co-simulation contract is maintained, the participating models can be modified at will, so that engineers can explore the design space,

Figure 1: The ChessWay Prototype

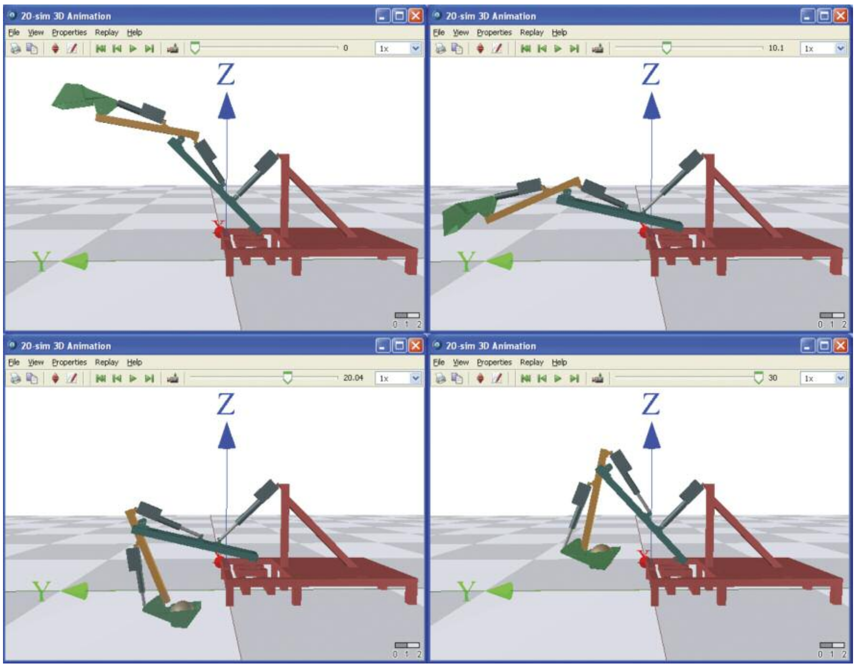

Figure 2: Co-simulation of the Verhaert Excavator model

for example examining different controller strategies on the same physical system model, or to subject a single controller to a set of physical system models with different fidelity.

Our collaborative modelling and simulation approach has been implemented in the Crescendo tool, which currently uses 20-sim, a commercially available tool, for the CT models of the physical system, coupled with DE controller models expressed using the VDM notation. 20-sim supports several textual and graphical methods to describe CT models which are automatically converted into sets of differential equations that are numerically solved. A library of domain-specific sub-models describes mechanical, pneumatic, hydraulic and electrical components. A 3-D mechanics editor allows a virtual mock-up to be created to support visualization. On the DE side, VDM is an object-oriented formal modelling technique that supports abstract data types, constants, functions, operations and additional integrity constraints such as invariants and pre- and post-conditions. The language enables the specification of distributed and embedded systems, with asynchronous operations, time and an explicit notion of the computation and communication architecture onto which the software is deployed. The cyber-side engineer can thus specify which software runs where, and assess the impact of that chosen deployment on overall performance and timing.

Several companies have applied and evaluated the technology. During the project, CHESS (NL) developed a personal transporter akin to the famous Segway, with a distributed safety monitor (Figure 1). Verhaert (B) developed a novel controller for a dredging excavator, enabling even novice operators to dig perfectly straight trenches (Figure 2). Neopost (NL) modelled a new system design for automatically folding stacks of documents and placing them in envelopes at high speed (Figure 3). Other smaller-scale applications based on industry-specified challenges included a tilting conveyor belt system, an aircraft flare dispensing unit and movement control for an interplanetary rover. The industrial case studies have shown that co-simulation is a valuable tool to support the multi-disciplinary design dialogue from the initial (conceptual) stages of product development. We used 20-sim and VDM-RT in DESTECS, but the Crescendo tool has an open architecture and is built upon a well-defined semantics that is essentially technology-agnostic.

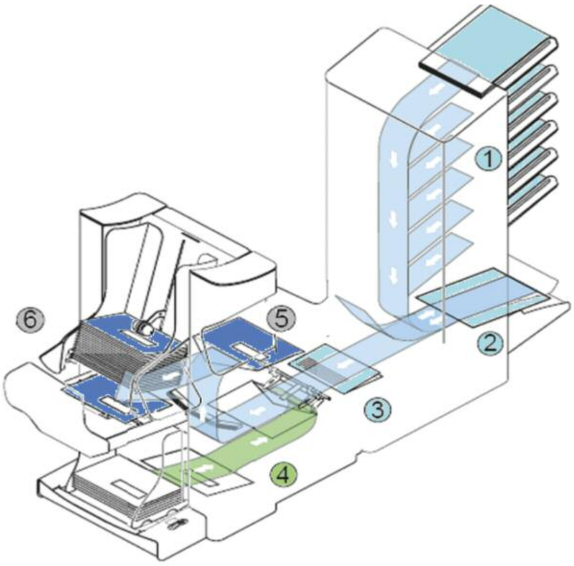

Figure 3: The Neopost document insertion system

Of course, simulation is not a "silver bullet", but our industry case studies showed that it is possible – and indeed very cost-effective – to create abstract and competent co-models of CPSs in early development phases. These models remain lightweight and compact, allowing developers to play "what-if" scenarios at relatively low cost. The insight gained by these simple experiments raises the confidence in the models quickly, because design decisions are continuously validated. Implicit choices and hidden assumptions are immediately exposed, replacing "gut feeling" by credible, objective and quantitative information that now can be assessed by all the designers involved, regardless of their "home" discipline. This typically gives direction, depth and momentum to the design effort because the potential of certain designs and the likelihood and impact of potential risks can be determined very rapidly, without large upfront investments. Because co-models can be subjected to advanced verification or form the basis of implementations, they have a utility far beyond initial multidisciplinary design.

Links:
http://www.destecs.org
http://www.crescendotool.org

References:
[1] T.A. Henzinger, J. Sifakis: "The Embedded Systems Design Challenge", FM 2006: Formal Methods, pp. 1-15, 2006

[2] E.A. Lee: "Cyber Physical Systems: Design Challenges" in proc. of ISORC 2008, pp. 363-369, 2008

[3] J. Fitzgerald, P. Gorm Larsen, M. Verhoef (eds): "Collaborative Design for Embedded Systems - Co-modelling and Co-simulation", Springer, April 2014, ISBN 978-3-642-54117-9

Please contact:
John Fitzgerald
Newcastle University, UK
E-mail:
john.fitzgerald@newcastle.ac.uk

Peter Gorm Larsen,
Aarhus University, Denmark
E-mail: pgl@eng.au.dk

Marcel Verhoef
Chess, the Netherlands
E-mail: marcel.verhoef@chess.nl

Cassting: Synthesizing Complex Systems Using Non-Zero-Sum Games

by Nicolas Markey

The design of complex systems such as those used in transportation, mobile communication and home automation, raises fundamental challenges: these systems are made of heterogeneous components that interact continuously with each other (collaboratively or as adversaries) and with their environments. Available methods (such as model-based verification, quantitative model-checking, or controller synthesis) only address specific aspects of complex systems. The FP7 project "Cassting" is developing a novel approach for analysing and designing complex systems in their totality, using non-zero-sum games to model their behaviours.

Formal methods for verifying computerised systems offer a powerful approach for validating the adequacy of the systems against their requirements. Based on deep theoretical results in logics and automata, model checking is one such approach. This approach involves modelling a system as a mathematical object (a finite-state automaton in the basic case), and applying algorithms for comparing the behaviour of that model against a logical property. Model checking has now reached maturity [1] through the development of efficient symbolic techniques, state-of-the-art tool support, and numerous successful applications to various areas.

Building on this framework, automatic synthesis techniques have recently been proposed. These techniques allow a model of a system (that is correct in construction) to be built automatically based on certain specifications. This model can then be refined to produce a validated implementation, thereby avoiding a trial-and-error approach of model checking. It is usual - and convenient - to rephrase the problem of automatic synthesis in the framework of game theory. In this metaphor, the system to be synthesized is then represented as a strategy, which should win against any behaviour of the system's environment, represented by the other player. The winning condition is used to express the correctness properties that we want the system to satisfy. This game-based approach to the synthesis of systems interacting with their environment has been greatly developed over the last twenty years, and is approaching maturity [2].

It should be noted, however, that up until now, formal methods have mostly focused on two-player zero-sum games, which only encompass the purely antagonist setting and may only model centralized control. Multiple-player games with non-zero-sum objectives have been considered, e.g., in algorithmic game theory, but they were strategic games, as opposed to the infinite-duration games on graphs that apply particularly well to synthesis.

Figure 1: Home automation is one of our target applications.

Within the Cassting project, we are designing a game-based framework for modelling and reasoning about cyber-physical systems. This is much more than just a minor extension of the existing framework: while the basic concept of using game models is preserved, radically new concepts and techniques have to be developed in order to build a powerful formalism, equipped with efficient algorithms and tools for modelling and synthesizing such complex systems. In particular, while winning strategies are the central concept in adversarial games, multiple-player games come with a full range of equilibria for globally characterizing the admissible behaviours of the players. Understanding how this can be adapted to infinite duration games representing cyber-physical systems is already an important part of our work [3]. Setting up a complete framework, with powerful formalisms and efficient algorithms for synthesizing the relevant strategies and the corresponding systems, is thus the main focus of our research. In order to fit to the systems we model, quantitative constraints (in particular on time or energy consumption) are a must-have feature of our modelling formalism: components of cyber-physical systems are often small independent devices with limited resources. Perturbations and imperfect communications are also part of our framework.

To validate our approach, our industrial partners have selected several case studies for us. These cases come from areas of smart buildings and smart grids, trying to optimize energy consumption and harvesting at different levels of scale. In the first case study, we aim at designing distributed, robust controllers for a floor heating system, in order to improve the efficiency of the currently implemented centralized controller. The second case focuses on a home-automation system, with a dynamic set of devices, including HVAC (climate-control) appliances, lighting, etc., with the aim of making the whole system easily reconfigurable. Finally, a third case study is concerned with optimizing energy harvesting and consumption of a whole set of houses equipped with solar panels and heat pumps, taking care of weather forecasts and varying energy prices.

The Cassting project includes five academic partners: CNRS/Laboratory Specification and Verification, France (Coordinator); University of Mons, Belgium; Université Libre de Bruxelles, Belgium; University of Aalborg, Denmark; RWTH Aachen, Germany. It also has two industrial partners: Seluxit, Denmark (an SME specialized in home automation) and Energi Nord, Denmark (an energy provider). The project started in April 2013, and will run for three years, with a total budget of approximately 2.7 million euros. Cassting is part of the Coordination Action FoCAS (Fundamentals of Collective Adaptive Systems).

Links:
Cassting project: http://www.cassting-project.eu/
FoCAS coordination action: http://www.focas.eu/

References:
[1] M. Edmund et al.: "Turing Lecture: Model Checking: Algorithmic Verification and Debugging", Communications of the ACM 52(11):74-84 ACM, 2009

[2] C-H. Cheng et al.: "MGSyn: Automatic Synthesis for Industrial Automation", in proc. of CAV'12, LNCS 7358, Springer, 2012

[3] J. Y. Halpern: "Beyond Nash Equilibrium: Solution Concepts for the 21st century" In K. R. Apt, E. Grädel (eds.), Lectures in Game Theory for Computer Scientists, chapter 8, Cambridge, 2011.

Please contact:
Nicolas Markey
CNRS & ENS Cachan, France
E-mail: nicolas.markey@lsv.ens-cachan.fr

Dependable System of Systems Engineering: the COMPASS Project

by John Fitzgerald, Steve Riddle, Paolo Casoto and Klaus Kristensen

In many areas of life, from buying holidays to organising the national defence, we are coming to depend on systems that are composed of several independently owned and managed, pre-existing systems. How can we engineer such Systems of Systems (SoSs) so that they merit the trust we place in them? A consortium of European and Brazilian researchers and practitioners are developing semantically sound languages and tools for Model-based SoS Engineering, integrated with well-established systems engineering practice. Industry case studies in areas including smart homes and emergency response have driven the requirements for our methods and tools, and have provided a basis for their evaluation.

Systems of Systems (SoSs) are synergistic collaborations between separately owned and managed systems that together deliver an emerging behaviour on which people come to rely. For example, the collective response of the emergency services in responding to a natural disaster comes as a result of the co-working of authorities that are normally – sometimes fiercely – independent. As we come to depend on the collective behaviour of these SoSs, we need to engineer them so that reliance on them is justified, even in the face of their capacity to change goals and functions.

Comprehensive Modelling for Advanced Systems of Systems (COMPASS) is an EU FP7 ICT project, now in its third year. Our vision is providing systems engineers with well-founded engineering notations, methods and tools that allow them to build models of SoSs and analyse their collective properties well ahead of commitments to costly design decisions.

Three specific challenges must be addressed to achieve the benefits of SoS engineering [1]:

1. Verification of emergent behaviour. We do not assume that emergent behaviour is unanticipated: it may be the intended outcome of building the SoS. Verification requires the composition of Constituent System (CS) properties.
2. Collaboration through contracts. Full information about each CS is unlikely to be visible to the SoS engineer. The CS must agree a contract specifying the range of behaviours that each CS can rely on (and guarantee). We need to be able to verify CS behaviour conforms to the contract.
3. Semantic heterogeneity. CSs may have a combination of components that are specified with discrete or continuous state or time, or at different levels of abstraction. This heterogeneity must be addressed in order to be able to verify emergent properties.

Throughout the project, there has been close collaboration between the technology developers and case study owners. This co-operative development has improved the COMPASS tool chain maturity and the industrial deployability of COMPASS technology.

An emergency response case study is developed by project partner INSIEL, who observe: "The COMPASS project is going to provide INSIEL several benefits from both methodological and technological points of view, the most significant being the ability to model, by means of executable CML models, properties emerging from the architecture of our System of Systems. Such requirement is, in fact, really critical for the specific domain where INSIEL's products work; each product needs to satisfy a set of constraints concerning quality of service and response time. In particular, by means of model checking, fault tolerance modelling and CML execution INSIEL will be able to evaluate if constraints will be respected at design time. CML modelling also allows us to analyze how SoSs will behave when a different context appears, or when the topology of the SoS may change."

A smart homes audio/video case study is developed by project partner Bang & Olufsen (B&O), who observe: "Significant results have been achieved in the field of SoS requirements engineering, SoS architectural level modelling and SoS model simulations. The requirements engineering methods and techniques have improved requirement consistency and stakeholder impact analysis for B&O's development organisation. A Streaming Architectural Framework (SAF) has been produced, which has enabled communication regarding integration challenges of the different CS's streaming architectures in the B&O SoS.

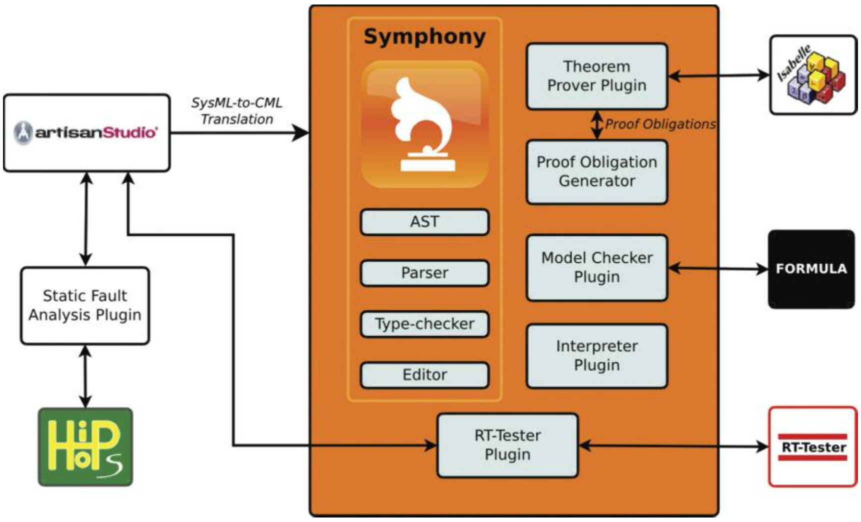

Figure 1: The COMPASS Toolset.

As part of early verification improvements, a formal CML (COMPASS Modelling Language) model for a new B&O distributed Leadership algorithm was developed. The algorithm enables desired emergent properties in the SoS. The correctness of the algorithm is vital for the end-users' experience in an SoS multi-product setup. B&O was able to find errors by analyzing the model, and correct the C++ implementation, before the algorithm was deployed in the products."

The identified specific challenges will continue to be the focus of further work in Cyber-Physical Systems of Systems.

Project partners are Newcastle University, UK; Aarhus University, Denmark; University of York, UK; Bremen University, Germany; Universidade Federal de Pernambuco, Brazil; Bang & Olufsen, Denmark; Insiel, Italy; Atego, UK

Link:
http://www.compass-research.eu

Reference:

[1] John Fitzgerald et al,"Model-based Engineering for Systems of Systems: the COMPASS Manifesto". COMPASS Technical Report, 2013. http://www.compass-research.eu/Project/Whitepapers/COMPASS_Manifesto.pdf

Please contact:

Steve Riddle
Newcastle University, UK
E-mail: steve.riddle@newcastle.ac.uk

Towards a Unified Theory for the Control of CPS: A Symbolic Approach

by Alessandro Borri, Maria Domenica Di Benedetto and Giordano Pola

We present a general framework for the efficient synthesis of control algorithms enforcing complex logic specifications on Cyber-Physical Systems (CPS).

Cyber-Physical Systems (CPS) are attracting particular attention in the research and industrial communities given their impact on the future of Internet of Things, Systems of Systems, wearable electronics, brain-machine interaction and swarm systems, and other fields that require the integration of different, complementary competences. We focus on the role of control and in particular of Networked Control Systems (NCS) [1] and embedded control software synthesis [2] in building the foundations of CPS.

Particular emphasis is given in NCS to the non-idealities affecting communication among plants and controllers. The relevant ones are quantization errors, variable sampling/transmission intervals, time-varying delay in delivering messages through the network, limited bandwidth, packet dropouts, and scheduling protocols. Given the generality of the NCS model, results for the analysis and control design are difficult to obtain. Researchers, thus, typically focus only on subsets of these non-idealities. Results available in the literature mostly concern stability and stabilizability issues [1]. However, emerging requirements for CPS address different and perhaps more complex control objectives, for example: obstacle avoidance, synchronization specifications, enforcement of limit cycles and oscillatory behaviour.

The computer science community has developed methodologies for the control design of discrete systems with complex logic specifications. These techniques are general enough to address issues arising from CPS applications. However, this methodology cannot be directly applied to NCS, which include continuous dynamics. The focus of this article is an approach to deal with this class of systems by applying “correct-by-design embedded control software synthesis”.

Central to this approach is the construction of symbolic models, which are abstract descriptions of continuous systems where a symbol corresponds to an "aggregate" of continuous states. Once a symbolic model is constructed, which is equivalent to or approximates the original continuous dynamics, then the techniques developed for discrete models can be applied. Several classes of dynamical and control systems admit symbolic models, see for example [2] and the references therein.

In particular, we considered a general model of NCS comprising all the above mentioned non-idealities in the communication infrastructure. The proposed model is flexible enough to include other relevant features of CPS, such as specific communication protocols, data compression and encryption in the message delivery, and scheduling rules in the computing units. Under mild assumptions on the boundedness of the time-varying delays involved in the NCS, and of the incremental forward completeness on the plant, a symbolic model is constructed and proven to approximate the given NCS in the sense of alternating approximate simulation [2], for any desired accuracy. While being sound, this result exhibits the following drawback: if a controller, designed on the basis of the symbolic model for enforcing a given specification, fails to exist, there is no guarantee that a controller, enforcing the same specification, does not exist for the original NCS model. When alternating approximate simulation is replaced by alternating approximate bisimulation [2], the above drawback is overcome: we showed that if the plant of the NCS enjoys the stronger property of global asymptotic stability, a symbolic model can be constructed which approximates the NCS in the sense of alternating approximate bisimulation for any desired accuracy.

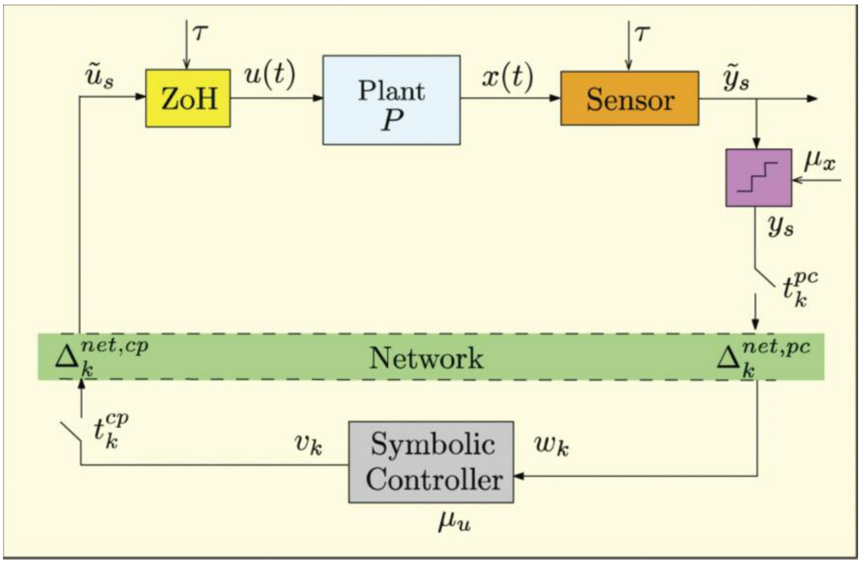

Figure 1: Example of control over networks in a cyber-physical system.

Once symbolic models are constructed, the symbolic control design of NCS can be addressed. We consider the following control problem: given a specification expressed in a fragment of Linear Temporal Logic, find a symbolic controller so that the closed-loop NCS system matches the specification with any desired accuracy. We solved this control problem by extending computer science results to metric transition systems. The resulting symbolic controller has been proven to be the approximate parallel composition of the symbolic model and of the specification. This approach is computationally intensive. We have developed efficient algorithms based on “on-the-fly” techniques to mitigate computational complexity.
The methodology proposed here is reported in detail in [3], which also includes an example of application of the techniques to the control design of a pair of nonlinear control systems sharing a common communication network. Since the requested symbolic controllers would exhibit a large space complexity (more than 10^{20} integers), we applied the aforementioned “on-the-fly” algorithm and designed symbolic controllers with a total space complexity of just 25,239 integers, still solving the original control problem. The closed-loop NCS obtained has been validated through the OMNeT++ network simulation framework, thus showing the effectiveness of the approach.

To the best of our knowledge, this is the first contribution where a general model of NCS is considered and where complex control problems are addressed. However these results are still theoretical and more work is needed to be relevant for industrial CPS.

Link: http://dews.univaq.it/

References:

[1] A. Bemporad et al. (Eds.): “Networked Control Systems”, LNCIS 406.
[2] P. Tabuada: “Verification and Control of Hybrid Systems: A Symbolic Approach”, Springer 2009.
[3] A. Borri et al.: “A symbolic approach to the design of nonlinear networked control systems”, HSCC’12.

Please contact:

Alessandro Borri,
IASI-CNR Biomathematics Laboratory, Rome, Italy
E-mail: alessandro.borri@iasi.cnr.it

Maria Domenica Di Benedetto and Giordano Pola
University of L’Aquila, Italy.
E-mail: {mariadomenica.dibenedetto, giordano.pola} @univaq.it

Information Assurance System in the Arrowhead Project

by Sandor Plosz, Markus Tauber and Pal Varga

A common communication language is the key element for interoperation of systems. Defining such a common language can be a formidable task, particularly in legacy and industrial systems. This problem is addressed within the Arrowhead project, which focuses on the automation of industrial systems with the goal of achieving energy efficiency and flexible use of energy.

The Arrowhead project consortium comprises 78 partners, including the Austrian Institute of Technology and AITIA International, Inc., as well as a number of manufacturers, energy providers, energy grid maintenance companies and communication system providers.

The investigated domains and use cases indicate that the integration of (legacy) sensors and M2M communication components into an automation system facilitates improved efficiency of operations.

The Arrowhead framework

An important task within the ARROWHEAD project is to design a generic framework to achieve interoperability of various sub systems, which often derive from legacy applications. The concept of this framework is based on service oriented architectures (SOA). For each system, an interface description must be provided (black box design) which can be made accessible to other systems through the core system functions. Each system may have multiple instances for different application areas, which are represented as System of Systems. System functions are derived from System of Systems Design Description (SoSDD) template. The structure of the framework is depicted in Figure 1.

A system may provide four core functionalities that are defined within the System of Systems Design (SoSD): Information Infrastructures, Information Assurance, System Management and Application related functionalities. These functionalities comprise services, such as service discovery, orchestration, authorization, and status, which are required to be defined for every system. Each partner involved in an individual use case publishes a list of functionalities that their (automation) system can provide and requirements for their system to accomplish particular goals that have been defined at the start of the project. Requirements may relate to dependability, safety, reliability, integrability, interoperability, architectural and communication. Based on these, the framework will be refined, and finally the procedure on how to integrate legacy systems into the arrowhead framework will be described and disseminated in project deliverables.

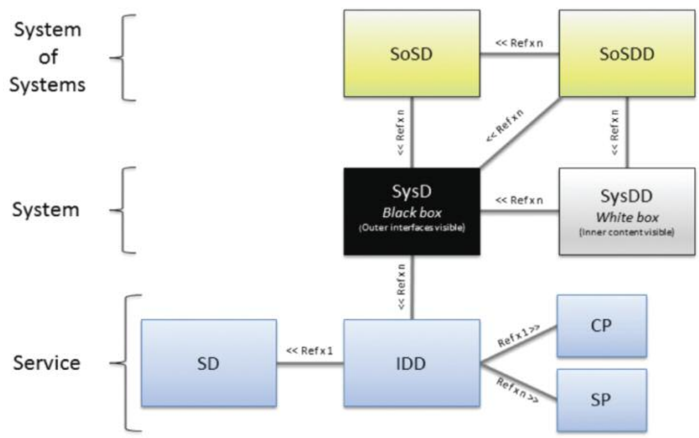

Figure 1: Structure of the Arrowhead framework.

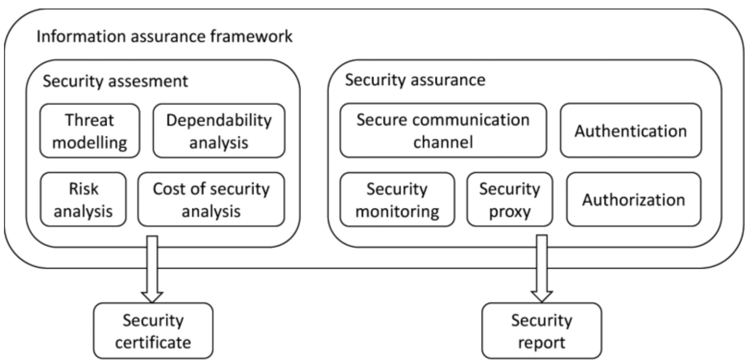

Figure 2: Arrowhead Information assurance framework.

Information assurance

In the information assurance core system a security assessment methodology is designed which makes the security assessment of systems and applications straightforward for the developers and system designers prior to integration into the Arrowhead framework. This method is based on threat modelling [1] and risk assessment [2] techniques which can be applied from the early stages of development. Additionally, applications can make use of dependability analysis methods, where additional factors are considered for evaluation, such as socio-technical aspects, safety and reliability and maintenance support. Despite their potentially large impact on security, socio-technical aspects - such as user expectations and goals, and how users use their systems - are often neglected [3].

Security requirements are outlined during application design. Important security assurance services will be developed within the information assurance core system. These include: key distribution for authentication, creating secure tunnels between systems for safe communication, certificate distribution security monitoring and security proxy for maintaining autonomy. Within the Arrowhead project we will initially collect the security requirements of different systems. Services will be developed accordingly to meet these requirements. Ways of integrating additional functionalities into the system will also be analysed and described.

An important use case for achieving energy efficiency within the inter-framework communication infrastructure is provided by cost-of-security analysis. Robust security solutions are accompanied by large processing demands in terms of speed and energy efficiency. In the security assessment methodology we address this issue by examining individual application requirements and common security assurance methods, which will enable us to suggest best methods for applications in terms of energy efficiency.

Main research question

Assuring security in the M2M communication domain requires much forward planning to make sure no interruption will be inflicted on a deployed service during operations, especially if novel IT systems are interoperating with traditional/legacy M2M components. Security assessment needs to be revisited to get the best of both the IT and M2M worlds. This planning requires the security assurance framework to be flexible, allowing security services to be updated or replaced, and new services to be installed on-the-fly. To achieve this, we need to detach security assurance from the applications themselves. This is a new approach which differs from traditional IT trends. We will address these issues as part of our work for the Arrowhead framework. This approach will enable application designers to concentrate on the task at hand without having to worry about interoperability, efficiency and security, which will be assured by well-tested systems through the Arrowhead framework.

Link:
http://www.arrowhead.eu

References:

[1] M N. Johnstone:, “Threat Modelling with Stride and UML”, in proc. of the 8th Australian Information Security Management Conference, 2010

[2] P. Saripalli, B. Walters: “QUIRC: A Quantitative Impact and Risk Assessment Framework for Cloud Security”, 2010 IEEE 3rd International Conference on Cloud Computing

[3] P. Masci et al.: “Towards Automated Dependability Analysis of Dynamically Connected Systems”, International Symposium on Autonomous Decentralized Systems (ISADS) 2011

Please contact:

Sandor Plosz
AIT Austrian Institute of Technology
Tel: +3614633402
E-mail: Sandor.Plosz.fl@ait.ac.at

Non-Asymptotic Estimation for Online Systems: Finite-Time Algorithms and Applications

by Jean-Pierre Richard and Wilfrid Perruquetti

Finite-time estimation is a way to perform real-time control and achieve specified time performance. The Inria project-team Non-A is focusing on designing such “non-asymptotic” algorithms; combining mathematics (algebra, homogeneity) with real-time applications (robotics, energy, environment, etc.).

For engineers, a wide variety of information is not directly obtained through measurement. Some parameters (e.g. constants of an electrical actuator, time delays in communication) or internal variables (torques applied to a robot arm, posture and localization of a mobile, detection of impacts or angles in biped walk, etc.) are unknown or are not measured. Similarly, more often than not, signals from sensors are distorted and tainted by measurement noises. To control such processes, and to extract information conveyed by the signals, one often has to estimate parameters or variables.

Estimation can concern parameters (identification), states (observation), derivatives (differentiation), inputs (left inversion) or noisy data (filtering). Consequently, a vast number of results have already been produced in this area, concerning either control or signal processing. However, unlike traditional methods, the majority of which pertain to asymptotic statistics, the particularity of Non-A is to develop algorithms which converge after a finite-time. This is summarized in the project’s name: “Non-Asymptotic estimation for online systems”.

The Non-A project-team is a joint action of Inria with Ecole Centrale de Lille, University of Lille 1 and CNRS (LAGIS UMR 8219). It also involves members from ENSEA Cergy, University of Reims and University of Lorraine. It is located at the Inria research centre Lille North-Europe and involves 23 people (Researchers, PhD students, Post-Doc, Engineers) from 13 countries. It was created in January 2011 in the continuity of a previous project ALIEN (led by Michel Fliess). After having successfully undergone its evaluation in 2013, Non-A is being supported by Inria until the end of 2017.

Why do we develop finite-time algorithms? In most fields of application, the time response constraint is crucial. Using our algorithms, computations are performed in real-time (i.e. as the application is running). Finite time convergence can benefit both the engineer who has to fulfill design specifications prior to certification, and the control researcher looking for a mathematical proof of their “separation principle”.

Application fields are plentiful and past results have concerned robotics, vehicle control, aeronautics, hydroelectric power plants, shape-memory or magnetic actuators, power electronics, secured communications, financial engineering and image/video processing. Currently, the team is focusing largely on the networked control systems, including WSAN (Wireless Sensors and Actuators Networks), but other applications range from circadian rhythms to high precision machining. Our "model-free control" also attracts various industrial contracts and has given birth to AL.I.E.N. SAS, a company created in Nancy.

Two complementary alternatives for finite-time estimation are considered. The first develops an algebraic framework for identification initiated in 2003 (Fliess, Sira-Ramirez, ESAIM COCV, 9, 151-168). Recent results concern the fast estimation of time-delay systems [1], of frequencies in noisy periodic signals [2], of modes and states in switched systems [1], of impulsive systems, as well as the differentiation of multivariate signals. The other develops the concept of homogeneous nonlinear systems and observers: roughly speaking, homogeneity [3] allows for extending local results, on a sphere, to global ones, in the whole space. Very recently we generalized this concept to time delay systems. Homogeneity applies to control, input-to-state stability and finite-time observers, and also to fixed-time stabilization [3], where the convergence time is fixed independently of the system's initial conditions (Figure 1).

Rather than being oriented to a specific domain of application, Non-A is a method-driven project (namely, algebra and homogeneity for fast estimation and control). However, one must not forget that applicability remains a guideline in all our research. In the long term we will be concentrating our engineering effort on the Internet of Things, and more particularly the self-deployment of WSAN. Various packages are being developed within the ROS environment, to assist with robot cooperation: e.g., cooperative SLAM, path-planning and path-tracking and localization with a reduced number of landmarks. In a more ad-hoc way, some work also has an environmental focus: biology (rhythms of bivalve molluscs for pollution detection), meteorology (sunshine prediction for solar energy management), fluid mechanics (limit flow control for energy saving).

Non-A is a partner of European grants FP7 HYCON2: "Highly-complex and networked control systems" and Interreg IVA 2 Seas SYSIASS: "Autonomous and Intelligent Healthcare System".

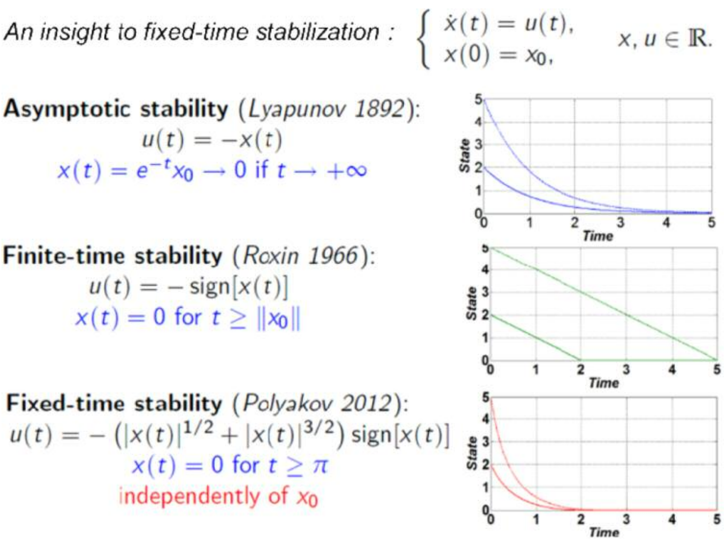

Figure 1: Increasing constraints on the convergence performance

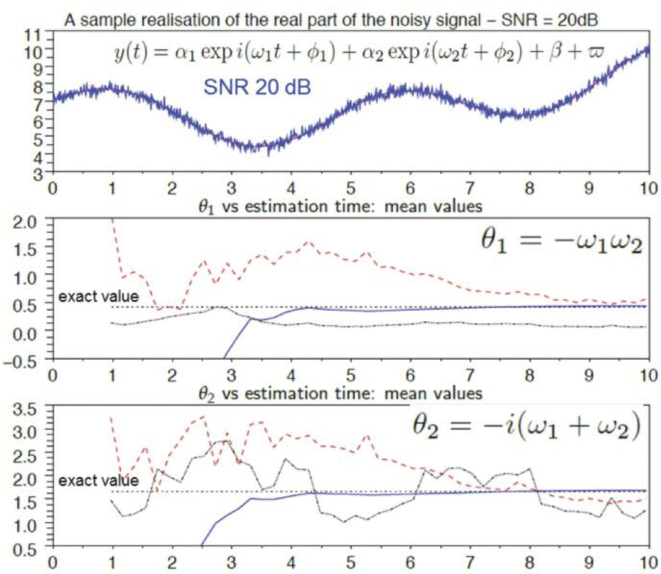

Figure 2: Frequency estimation for noisy signals, taken from [2] – Blue = our algebraic method, Black & Red = Prony-like methods.

References:

[1] L. Belkoura, J.P. Richard, M. Fliess: "Parameters estimation of systems with delayed and structured entries", Automatica, 45(5), 1117-1125, 2009; http://hal.inria.fr/inria-00343801
[2] R. Ushirobira et al: "Algebraic parameter estimation of a biased sinusoidal waveform signal from noisy data" in proc. of SYSID'2012, 16th IFAC Symp. on System Identification, Brussels, 2012, http://hal.inria.fr/hal-00685067
[3] E. Bernuau et al.: "On homogeneity and its application in sliding mode control", J. of the Franklin Institute, 2014, http://hal.inria.fr/hal-00942326

All the project's publications are available at http://www.inria.fr/en/teams/non-a/(section)/publications

Links:

http://www.inria.fr/en/teams/non-a
http://alien-sas.com/

Please contact:

Jean-Pierre Richard and Wilfrid Perruquetti
École Centrale de Lille, France
E-mail: jean-pierre.richard@ec-lille.fr, wilfrid.perruquetti@inria.fr

Soft Real-Time Scheduling Approaches in Embedded Control Systems

by Daniele Fontanelli, Luca Greco and Luigi Palopoli

Two trends can be recognized in the recent development of embedded control systems: the adoption of sophisticated sensing systems (which require large and highly variable processing times) and the proliferation of control applications that are deployed on the system. The combination of the two trends has caused an obsolescence of hard real-time design techniques, which allocate computing resources based on the worst case demand of the application. There exists an effective replacement, however, that allows us to reconcile performance guarantees and efficiency in resource management.

An ever growing number of control applications utilize sophisticated sensing devices that can extract multiple information from each measurement. For instance, visual sensors are in charge of device localization, people detection and the extraction of features of interest for control purposes (e.g., the lines delimiting a lane in an automated driving system).

The price paid for the adoption of such complex devices is a large and highly variable processing time. Indeed, in a "clean" environment, the amount of computation required to extract the relevant features can be very low, while it sky-rockets in the presence of a large number of artifacts or under dim illumination.

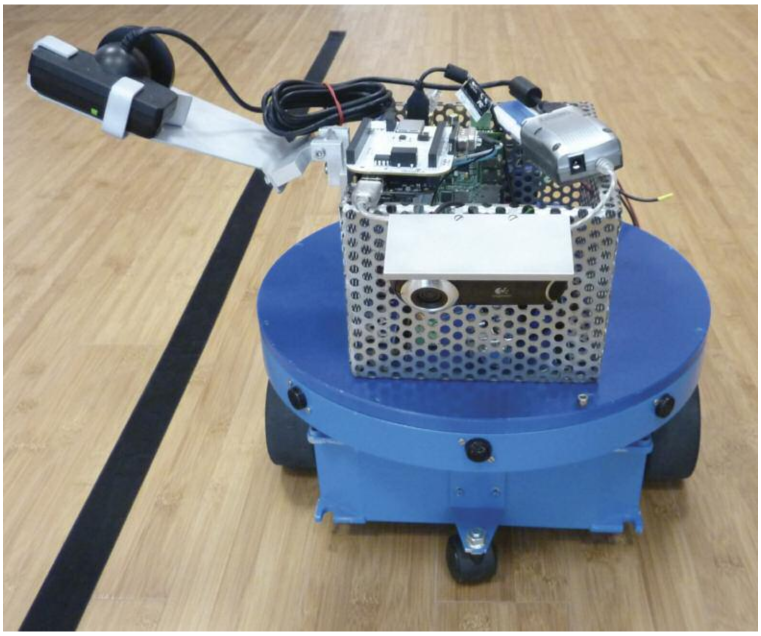

Figure 1: The actual robot adopted for the case study.

The need for sharing

Equally important is a different requirement of modern control systems: to maximize hardware sharing between the different control functions. In a modern car, for instance, the complexity of the communication infrastructure is currently measurable in kilometres of copper cables and hundreds of kilograms of weight. The same applies to Electronic Control Units (ECU) used to execute computations, whose number easily exceeds 100.

Positioning and interconnecting of these devices has a tremendous impact on the complexity of system engineering. Hardware complexity has reached the critical point where it becomes imperative for new or more advanced control functions to exploit the existing computation and communication infrastructure.

Limits of classic approaches

With strong fluctuations in the computing workload on one side, and the need for intense hardware sharing on the other, the designer is confronted with problems that can hardly be addressed by traditional design approaches. The classic mandate is to grant to each control function hardware resources commensurate to its worst case demand (e.g., a dark and cluttered scene for visual sensors). This way, it will be executed with regular timing and with constant delays, thus allowing control designers to formulate and enforce guarantees on control performance. The rigid application of this old paradigm to the new context naturally leads to a massive waste of resources (which lie unused for most of the time) and to a drastic reduction in hardware sharing opportunities. On the contrary, if the control function receives resources proportionate only to its average requirements, there can potentially be large intervals of time when the system is not properly controlled: performance and even stability could be lost!

A stochastic approach

The University of Trento and the University of Paris Sud, within the framework of the FP7 HYCON2 project (running from 2010 to 2014), has started a research activity which aims at reconciling performance guarantees and hardware sharing. The key idea underlying this research is that the limits of classic methods can be overcome by a suitable stochastic description of the computation requirements. Clearly the new approach calls for new performance metrics (e.g. Almost-Sure Stability and Second Moment Stability), which are themselves stochastic in nature. Broadly speaking, system trajectories are allowed to have stochastic fluctuations, but eventually they will behave properly (except for a small and statistically irrelevant set of them). The application of these notions to our context needed a formal mathematical model to describe the stochastic evolution of the

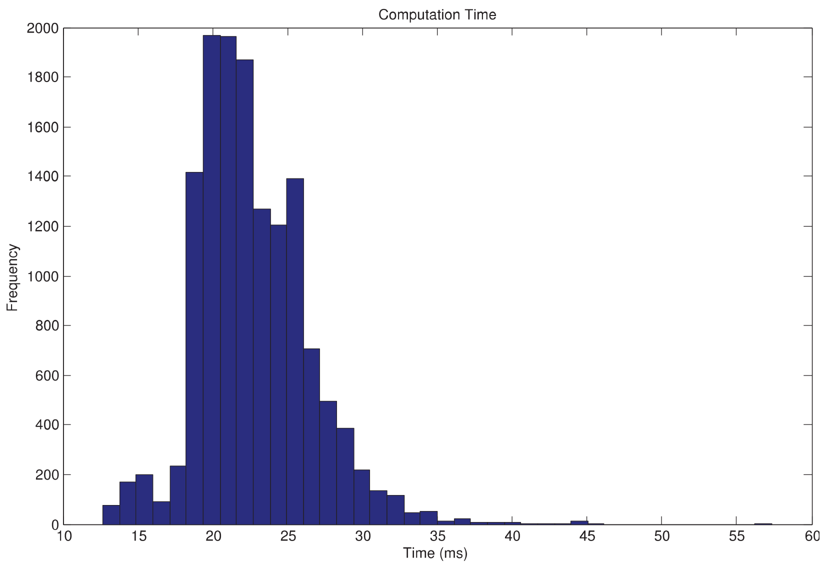

Figure 2: Probability mass function of the computation times observed during the experiments

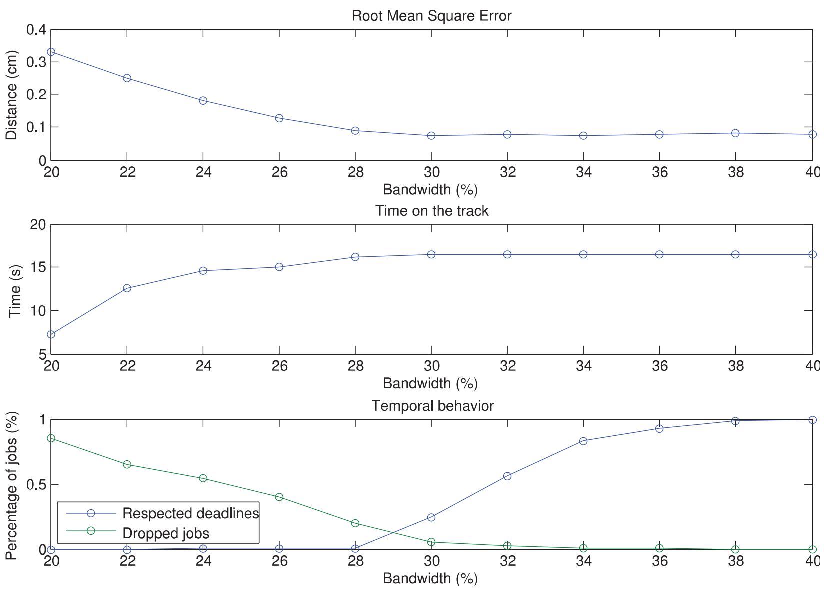

Figure 3: Experimental results achieved with different bandwidth values: root mean square of the deviation from the desired position (top); time the robot succeeds to remain on the track (middle); percentage of task executions that complete within a deadline or that are cancelled (bottom).

delays incurred in control computation. In [1,2] we met this requirement by combining a soft real-time scheduler [3] with an appropriate programming discipline. The former guarantees that the different computing tasks will have timely access to the computation resources, while the latter forces interactions between computers and environment to take place on well defined instants. Leveraging these ideas, it was possible to establish a connection between the stochastic control properties and the fraction of computing resources to allocate.

A concrete case study

We considered a mobile robot (Figure 1) that was required to follow a black line drawn on the ground. The variability in the scene background determines a similar variability in the computation time (Figure 2). In the worst case, the computation time is beyond 40ms; if the system is operated at 25 frames per second this corresponds to a worst-case utilization beyond 100%. With the classic approach, the control function would require the allocation of the entire CPU. Our methodology has revealed that 21% of CPU utilization is in fact sufficient to stabilize the system.

In Figure 3 we show the experimental results achieved with different bandwidth values. The top plot reports the integral of the root mean square of the deviation from the desired position, while the middle plot reports the time the robot succeeds to remain on the track (the duration of the experiment is 40s). The performance dramatically improves when we allocate more than the minimal 21%, but there is no apparent advantage in exceeding 30%. In the bottom plot, we report the percentage of task executions that complete within a deadline equal to the period and the percentage of executions cancelled because of time-out. In the classic setting these figures would be 100% and 0% respectively. In contrast, the plot reveals that the system can operate with provable performance far away from this regime with substantial resource savings.

Conclusion

The experiments here presented show that the stochastic scheduling approach is a viable solution for real-time control of systems, in which the desired control performance can be achieved even with limited computing resources.

References:

[1] D. Fontanelli, L. Greco, L. Palopoli: "Soft RealTime Scheduling for Embedded Control Systems", Automatica, 49:2330-2338, July 2013.

[2] D. Fontanelli, L. Palopoli, L. Greco: "Deterministic and Stochastic QoS Provision for Real-Time Control Systems", in proc. or RTAS 20111, Chicago, IL, USA, p. 103-112

[3] R. Rajkumar et al.: "Resource kernels: A resource-centric approach to real-time and multimedia systems", in proc. of the SPIE/ACM Conference on Multimedia Computing and Networking, 1998.

Please contact:

Daniele Fontanelli, Luigi Palopoli, Università di Trento, Italy
E-mail: daniele.fontanelli@unitn.it, luigi.palopoli@unitn.it

Luca Greco, L2S Supélec, Paris, France
E-mail: luca.greco@lss.supelec.fr

Requirements Engineering Patterns for Cyber-Physical Systems

by Christophe Ponsard, Jean-Christophe Deprez and Robert Darimont

Engineering Cyber-Physical Systems (CPS) is challenging in many respects, including at the Requirements Engineering (RE) stage. In addition to generic RE techniques that already encompass the system dimension of CPS, such systems can benefit from domain knowledge related to specific requirements, such as accuracy, dependability and security. This knowledge can be structured as a reusable requirements pattern library.

Cyber-physical systems are complex systems with intertwined computation, communication and physical elements, often including human interactions in the loop. The design of such systems goes well beyond traditional embedded system controllers and exhibits characteristics of multi-agent systems. This is quite challenging for the engineering requirements process as it must be able to deal with multiple agents either under (co-)design or located in the environment (including physical processes, devices and human interactions). It also requires software design to be mixed with physical system design, and these two design processes rely on different sets of modelling languages (e.g. logic for computational elements and differential calculus for physical elements). Moreover, CPS are increasingly manipulating critical infrastructure such as the power grid, transportation systems, medical devices, security systems, which means they must satisfy a whole set of challenging non-functional requirements about safety, security, privacy, usability, energy-efficiency and adaptability.

In the scope of the SYLIS CORNET project (involving the CETIC and Fraunhofer IESE Research Centres), it is important to help companies to deal with the design of such complex systems, especially from the requirements engineering phase (RE). From a RE point of view, the most adapted methods able to deal with such complexity are goal-oriented methods, such as KAOS, which is able to capture and reason about multiple agents both within the environment and the system under design of hardware/software/human natures. KAOS also enables capturing of global system properties that are achieved as a result of the global collaboration of all CPS elements. Furthermore, KAOS enables the refinement of high level system goals into specific requirements or expectations placed on each kind of agent [1].

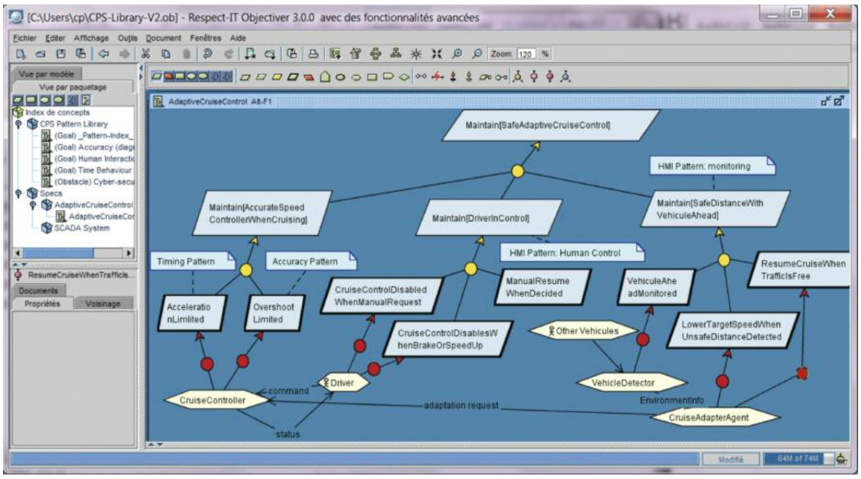

Figure 1: Structure of the CPS pattern library and instantiation to an adaptive cruise control system.

Based on those premises, in SYLIS, we revisited a set of key requirements found in a collection of CPS systems (such as [2]) with the aim of reusing them based on the successful technique of pattern (e.g. OO design patterns, security patterns, cloud patterns). The purpose of these patterns is to capture specific knowledge related to the design of CPS systems, such as:

- *Accuracy:* This is a key property of CPS systems and it can only be achieved through the proper collaboration of many components. Typically, the goal of feedback loops is to maintain a specific controlled variable at some target value level. On the dynamic side there are a number of typical behaviours to control systems that are generally managed by the control designer. These control characteristics should also be made visible to the requirements analyst to unveil possible impact on system or sub-system goals: for example: how fast should a target value be reached, what is the permitted overshot (occurring on step transition), what amplitude of oscillation is acceptable, etc.

- *Human interaction:* different patterns are possible: humans can be in a partial control loop (which should be adapted to the individual's reaction time); a human agent responsible for a specific requirement can be monitored by the system and their actions overridden by the system if the human fails at the duty. On the interface side, different modalities can be imagined with different kinds of ergonomics and cognitive loads.

- *Cyber-security:* CPS are composed of tightly interconnected elements which may be connected to the Internet either permanently or periodically. Consequently, exposing these elements to cyber threats with potentially high impact (e.g. SCADA systems controlling power grids, transportation system networks, connected cars, etc.). Reasoning based on anti-goals patterns on security threats is especially relevant to capture knowledge about how malicious agents can potentially attack a CPS, e.g. physical CPS element attack, denial of service, wrong command to controller. A number of security patterns can be enriched for the cyber-security context such as a secure controller or a secure communication channel.

The Objectiver tool is currently used to model our CPS patterns. The tool supports the full KAOS methodology and includes support for pattern management and instantiation. Figure 1 illustrates the structure of our pattern library model (package view on the left) as well as a simple instantiation to an adaptive cruise control system including multiple agents such as a car under control, a driver, and other cars in the environment (goal diagram on the right). We plan to deploy it in a more specific tooling, already successfully applied to the smartcard domain [3], and further enrich it. We also plan to release our pattern catalogue in a tool independent collaborative knowledge oriented website.

Links:
http://www.objectiver.com
https://www.cetic.be/SYLIS-1934
http://cloudpatterns.org

References:

[1] A. van Lamsweerde: "Requirements Engineering - From System Goals to UML Models to Software Specifications", Wiley, 2009.
[2] M. Heimdahl, et al.: "Modeling and Requirements on the Physical Side of Cyber-Physical Systems", Second International Workshop on the Twin Peaks of Requirements and Architecture, San Francisco, 2013
[3] N. Devos et al.: "Efficient Reuse of Domain-Specific Test Knowledge : An Industrial Case in the Smart Card Domain", in proc. of 34th ICSE, Zurich, June, 2012

Please contact:

Christophe Ponsard, CETIC, Belgium
E-mail: cp@cetic.be

Management of the Interconnection of Intermittent Photovoltaic Systems through a DC Link and Storage

by Alessio Iovine, Sabah Benamane Siad, Abdelkrim Benchaib and Gilney Damm

New connection constraints for the power network require more flexible and reliable systems, with robust solutions to cope with uncertainties and intermittence from renewable energy sources (renewables), such as photovoltaic arrays. The interconnection of such renewables with storage systems through DC links can fulfill these requirements. A "Plug and Play" approach based on the "System of Systems" philosophy using distributed and adaptive control methodologies is being developed.

Renewable energy can play a key role in producing local, clean and inexhaustible energy to supply the world's increasing demand for electricity. Photovoltaic conversion of solar energy is a promising way to meet the growing demand for energy, and is the best fit in several situations. However, its intermittent nature remains a real disability that can create voltage instability for large scale grids. In order to offer an answer to the new constraints of connection to the network (Grid-Codes) it is possible to consider a storage system; the whole system will be able to inject the electric power generated by photovoltaic panels to the grid in a controlled and efficient way. As a consequence, a strategy for managing energy in relation to the load and the battery constraints will be needed.

The approach is based on a "Plug and Play" philosophy: the global control will be carried out at local level by each actuator, according to the distributed control paradigm. The control objective is to develop a distributed method for stabilizing each part of the whole system, while performing power management in real time to satisfy the production objectives and assuring the stability of the interconnection to the main grid. Furthermore, this system may bring support (voltage and frequency) to the main grid by additional control schemes.

The whole control objective is then split in several tasks; the first is to increase the efficiency of the photovoltaic array, which should work in its maximum power point for delivering the maximum amount of energy. Maximum Power Point Tracking (MPPT) algorithms are used to compute the desired voltage that assures this maximum power production. A drawback is the uncertainty of doing this without irradiance and temperature sensors. An adaptive control scheme that is capable of achieving maximum power point tracking for changing environmental conditions can be used instead [1]. This system estimates parameters that depend on irradiance and temperature, increasing the robustness of the system and eliminating the need for a sensor.

A second step is to stabilize the DC/DC boost converter that connected the solar array to the DC network, calculating the optimal value of the duty cycle [2]. Lyapunov theory is used in this task, taking into account the parameters' estimator dynamics.

The focus then moves to the storage system and its connection to the DC network. Lead-acid batteries are commonly chosen as a storage system, and a DC/DC bidirectional converter is necessary to enable the two modes of functioning (charge and discharge) [2]. The regulation of the battery current and voltage consists of a cascade control; again, Lyapunov theory is used for providing stability. An algorithm that helps to manage power by defining priorities in relation to the load and the constraints of battery operation is also required.

With this structure, the DC grid is able to provide a continuous supply of good quality energy. Finally, the DC network must be interconnected with the AC grid, and must be controlled in such a way that it offers support (ancillary services) to the grid.

The final management system can be configurable and adaptable as needed. The battery is assimilated as a reservoir which acts as a buffer between the flow requested by the network and the flow supplied by the production sources, and its voltage is controlled by the DC/DC current converter. The three converters present in this system must, in a decentralized way, keep the stability of the DC network interconnecting all parts.

The whole system provides protection against faults and suppresses interference, and has a positive impact on the behaviour of the complete electrical system.

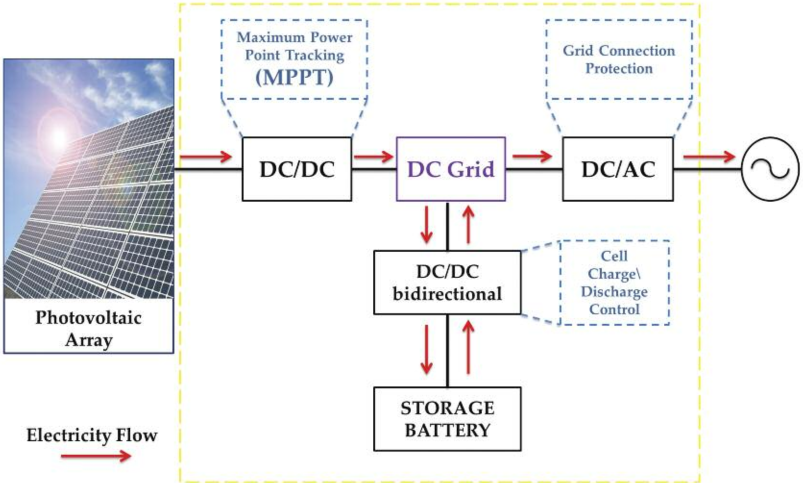

Figure 1: Photovoltaic system with storage.

Link: http://www.hycon2.eu/

References:

[1] F. Jaramillo-Lopez, G. Damm, G. Kenne, F. Lamnabhi-Lagarrigue: "Adaptive control scheme for maximum power point tracking of a photovoltaic system connected to the grid", 2013 European Control Conference (ECC), Zürich, Switzerland.

[2] S. Benamane Siad: "Gestion des flux énergétique dans un système photovoltaïque avec stockage connecte au réseau", Master thesis report, Alstom-Cnam, July 2012.

Please contact:

Alessio Iovine
University of L'Aquila, Italy
E-mail: alessio.iovine@graduate.univaq.it

Sabah Benamane Siad
Bouygues Energies et Services, France
E-mail:
s.siadbenamane@bouygues-es.com

Abdelkrim Benchaib
ALSTOM GRID, R&D Department, France
E-mail:
abdelkrim.benchaib@alstom.com

Gilney Damm
Université d'Evry Val d'Essonne, France
E-mail: gilney.damm@ibisc.fr

Multi-Terminal High Voltage Direct Current Networks for Renewable Energy Sources

by Miguel Jiménez Carrizosa, Yijing Chen, Gilney Damm, Abdelkrim Benchaib and Françoise Lamnabhi-Lagarrigue

The control of Multi-Terminal High Voltage Direct Current (MT-HVDC) networks is still an open problem. Communication within this wide-span network is a very delicate matter: the time delay and possible loss of information must be taken into account. We are aiming to develop and control a hierarchical stable structure that manages such a network.

Power systems' networks are among the largest and most complex physical man-made systems. They connect hundreds of millions of producers and consumers, cover continents and exhibit very complicated behaviours. Many poorly-understood phenomena result from the interactions between such a large number of devices and the large spatial dimension, and as a consequence their management and stabilization may be very challenging. At the same time, the dependence of modern societies on electricity is growing, with our reliance on pervasive computers, embedded systems and cellular devices. This dependence will most likely continue to grow if a lot of electric vehicles become a reality. For these reasons the impact of disturbances and blackouts has become astronomic.

Major changes are putting the power grid under great stress; in particular the large scale introduction of renewable energy sources is a major cause. The goal of reducing carbon dioxide emissions and creating a greener world will lead to an increased use of renewable energy sources (RES) in the form of wind farms or solar plants. Two distinctive features of these energy sources present an important challenge: most RES are dispersed over a very wide geographical area, and most RES are intermittent. Since further development of the classical Alternative Current (AC) Grid will not be possible owing to economic (in the offshore case) and environmental issues, the only foreseeable solution becomes the use of Direct Current (DC) links, and most probably a new DC off-shore Grid interconnecting

AC grids or other DC grids [1], constructing a new SuperGrid. Integrating these new technologies into the existing AC grid and introducing new tools for handling specific control problems related to the distributed control of large offshore wind farms interconnected by a DC network, as well as the connection of this DC network to the AC grid is currently under investigation.

Recently, multi-terminal HVDC (MT-HVDC) networks which consist of more than two converter stations have drawn more and more attentions. The MT-HVDC network offers a larger transmission capacity than the AC network and provides a more flexible, efficient transmission method. The main applications of MT-HVDC networks include power exchange among multi-points, connection between asynchronous networks, and integration of scattered power plants like offshore renewable energy sources. The control of a Multi-Terminal High Voltage Direct Current (MT-HVDC) network, as shown in Figure 1, is still an open problem. We are developing a hierarchical structure to manage such a network. The wide-span makes communication a very delicate matter. The time delay and possible loss of information must be taken into account when managing the network. It is therefore necessary to rely on lower level local only controllers [2], which will stabilize the system against changes and disturbances in the network.

Local controller

The lower level focuses on designing a local controller on each node in an MT-HVDC system, without using remote information. Its first objective is to make the converter accurately track its reference values which are provided by a higher level controller. In addition, in case of disturbances such as power imbalance, the system must be always kept stable and quickly converge to a new steady state condition. Finally, in a second step, the local controller should be driven by the higher level controller, which maintains sustainable operating conditions.

Most current research is based on Voltage Source Converters (VSC), and in particular with MMC (Modular Multilevel Converters) technology. The behaviour of a VSC terminal can usually be characterized by its DC and AC subsystems, which constitute its state vector (in particular the AC currents and DC voltage). Numerous control strategies for VSC-HVDC systems have been developed recently. Most, however (for example [3]), are based on intuitive assumptions without rigorous mathematical demonstrations.

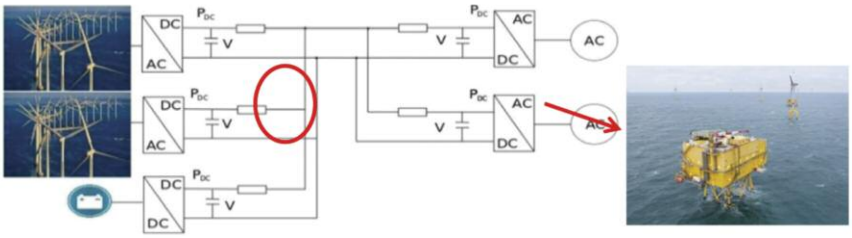

Figure 1: An MT-HVDC system (left) and an offshore converter station.

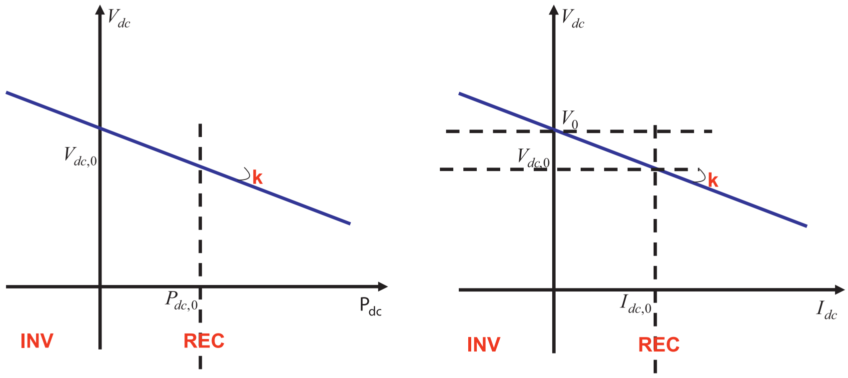

Figure 2(a) (left) and 2(b): Droop technique for primary control strategy.

Our group is investigating several control approaches to improve system performance via different control theories and rigorous model-based analysis.

There are three on-going approaches:
1. Nonlinear control based on feedback linearization: although some linear system properties can be applied to regulate the behaviours of the system, the main drawback is that the system can only operate in unidirectional mode.
2. Nonlinear control based on static and dynamic feedback linearization: this novel controller has been developed based on a static and on a dynamic feedback linearization structure in order to operate in bidirectional mode. For example, when the terminal absorbs the power from the DC grid it uses a static part; otherwise it uses the dynamic part. However, this method involves a switch system according to the power direction, which could present undesirable behaviours such as chattering and the appearance of large peaking values.
3. Globally stable passive control [1]: the main advantage of this controller is that global stability is guaranteed. In addition, the VSC terminal can operate more smoothly than the other approaches.

Primary (Droop) controller

Droop control reacts to a power imbalance, and can be explained as follows. When there is a power variation in one or several nodes, the other nodes will adjust the amount of generated/consumed power via droop control, by adjusting their voltages, see Figures 2(a) and 2(b). This control is implemented as a set of local proportional controllers with gains called the droop gain. These local gains are represented by the slope shown in Figures 2(a) and 2(b). This value, which depends on several variables, such as the size of each node and its available (primary) energy reserve, indicates how much power each generator provides when a variation occurs, in function of its capabilities.

Secondary controller

The main objectives of the secondary control are: to follow the tertiary level references; and to make a periodic power flow in order to guarantee the proper behavior of the system. The secondary control is also responsible for predicting and avoiding congestion in the lines.

When a disturbance occurs in the grid, it is the job of the secondary control to restore the equilibrium, taking into consideration the current energy reserves. To achieve this, communication between nodes is necessary. In this respect, the secondary control differs from the primary control, in which no communication occurs. The secondary control is a critical element in the proper functioning of the system because it takes into consideration disturbances as well as load forecast. It has information on the current and future short-term state of the system, and it creates the link between the physical system, in the time-frame of milliseconds, and the economic forecast that is in the order of hours. Its role is then to calculate voltage values to optimize the power flow to minimize transmission losses, for example, subject to all the constraints. These voltage values are then transmitted as references for the primary level.

To develop the secondary controller we have used Model Predictive Control (MPC), because it can provide an optimal solution taking into account these constraints whilst also considering forecasts, such as weather, load, or even prices.

In summary, the lower level controllers stabilize the system in the face of changes within the system as well as large disturbances, while higher level intercommunicating controllers steer the overall system to reliable and optimal operating conditions.

Link:
http://www.eeci-institute.eu/WINPOWER

References:

[1] M. Jiménez Carrizosa, et al.: "Bilinear and nonlinear control algorithms for a DC/DC converter for multi-terminal HVDC networks", IFAC World Congress 2014

[2] Y. Chen, et al.: "A globally stable passive controller for a VSC terminal in a multi-terminal HVDC grid", in proc. of ACC'14, Portland, USA

[3] J. L. Thomas , S. Poullain, A. Benchaib: "Analysis of a robust DC-bus voltage control system for a VSC transmission scheme", in proc. of the 7th Int. Conf. AC-DC Power Transmiss., 2001.

Please contact:

Gilney Damm
Laboratory IBISC, France
E-mail: gilney.damm@ibisc.fr

Abdelkrim Benchaib
ALSTOM GRID, FRANCE
E-mail:
abdelkrim.Benchaib@alstom.com

Françoise Lamnabhi-Lagarrigue
LSS-SUPELEC, France
E-mail: françoise.lamnabhi-lagarrigue@lss.supelec.fr

Smart Water Management: Key Issues and Promises from a Distributed Control and Observation Perspective

by Christopher Edwards, Prathyush Menon and Dragan Savic

The capability of modern utility distribution networks to exchange large amounts of data offers huge advantages in terms of efficiency of operation and optimization. On the other hand, the inherent communication network component exposes the overall system to the possibility of malicious cyber-attacks. The resilient operation of such distribution networks presents challenges to control theorists and practitioners alike.

New technologies are revolutionizing traditional water distribution systems. Currently, "smart methods" based on novel ICT technologies are being employed to deliver integrated supply-demand side management for improved efficiency. Although perhaps less well publicized than similar developments in the area of Smart Grids in power distribution networks, the aim of incorporating these new and smart methodologies into water distribution networks is to contribute to the delivery of a sustainable, low-carbon society, thus helping achieve progress towards the European 2020 targets on Climate and Energy. Figure 1 provides an example and an overview of a typical smart water network.

The primary source of information is real-time and near-real-time meter/sensor data, reporting pressures, flows and water levels at selected points throughout the water supply/distribution network This is based principally on instantaneous flow, pressure and/or level monitoring at major physical structures (i.e., intakes, treatment works, service reservoirs, pump stations, etc) delivered through a "supervisory control and data acquisition" system (SCADA). This information is augmented by flow and pressure measurements from meters at the entry to District Metered Areas (DMAs) along with selected pressure monitoring points elsewhere within the DMA, with data typically transmitted every 30-60 minutes, principally over a cellular communication network. Thus the intelligent water distribution network above comprises computing/monitoring facilities, valves/pumps (actuators), sensors (measurements), a physical commodity flow (water), reservoirs (source), consumers (sinks) and information flow (associated with actuator commands and sensors). See for example [1]. Such an infrastructure is often referred to as a Cyber-Physical-System (CPS) [2].

Modelling challenges

Models of a cyber-physical infrastructure can be developed based on the underlying physics of the problem: typically there are sources and sinks, and mass/flow/energy balance constraints. At a local level, components and subsystems can be modelled as individual lumped systems to be represented as generic, continuous/discrete time ordinary differential equations. The modelling of the interconnections between these local or node level elements, must capture both the exchange of information as well as the underlying topology of the physical and commodity flow realization, so that graph theoretic methods [3] can be exploited. Whilst such large-scale systems offer great potential and flexibility in terms of efficient operation, it is extremely important to have an effective monitoring mechanism in place so that failures/threats/risks can be swiftly identified, and suitable decisions can be made to contain and mitigate their effect. For example, timely detection of pipe burst events provides opportunities for water companies not only to save water but also money and energy, and

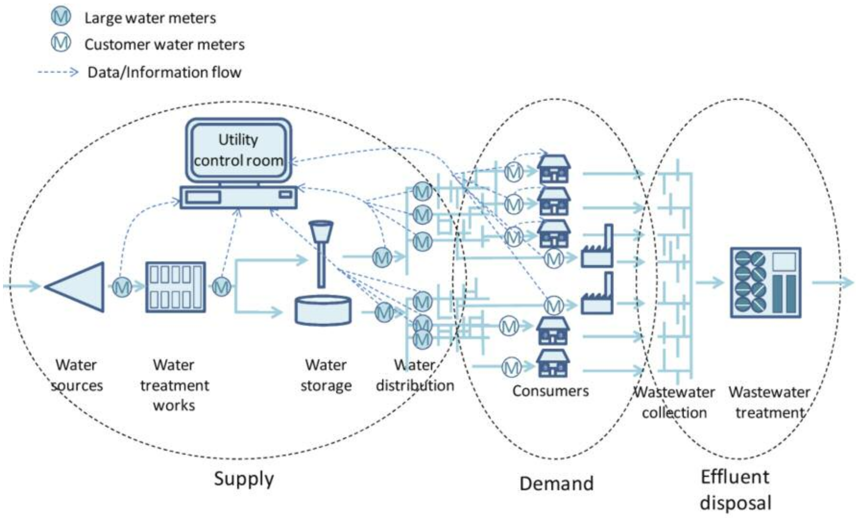

Figure1: Description of the architecture.

reduce their carbon footprint, and improve their operational efficiency and customer service.

Control Challenges

For efficient functioning of such infrastructure, in addition to optimizing normal fault-free operation, associated security aspects must also be taken into account. Hence the physical layer's control must be resilient and robust towards cyber-attacks and other malicious behaviour. A highly complex distribution network comprising many interacting, heterogeneous distributed systems, each under autonomous control and reacting to the environment and to cyber and/or physical signals of other systems, requires coordination/control – often in a decentralized or distributed manner - to accomplish a global objective/performance.

Monitoring challenges

One approach to developing monitoring schemes to detect failures/threats is to attempt to extend existing observer-based fault detection and reconstruction ideas to the CPS. In contrast to conventional observer problems, including a copy of the system being monitored internally requires a plethora of actuation and measurement signals from throughout the entire network. This is at best cumbersome and may not even be practical. Different architectures for the observer can be considered depending on which signals are realistically available. These could be classified as centralized (where global information is required), decentralized (only local information is required) and distributed (which involves an exchange of information amongst the observers themselves). Arguably, a scalable and distributed monitoring solution would be best suited to this problem. This is beneficial in terms of the sensor resources required, and hence the overall energy consumption.

Once a monitoring scheme has identified failures or cyber-attacks within the network, it is imperative to take pertinent timely decisions to isolate which of the sub-systems/elements is the root cause. Once isolated within the network, further decisions need to be made so that the operation of the remaining "healthy" part of the infrastructure experiences minimal impact and the effects of the failures/threat contained. However, since cyber as well as physical interconnections are present, an attack/fault on one sub-system may have an impact on the performance of other sub-systems and thereby globally within the network. Owing to possible redundancy within the network it may be possible to recover acceptable overall performance through reconfiguration and thereby achieving self-healing and providing resilience.

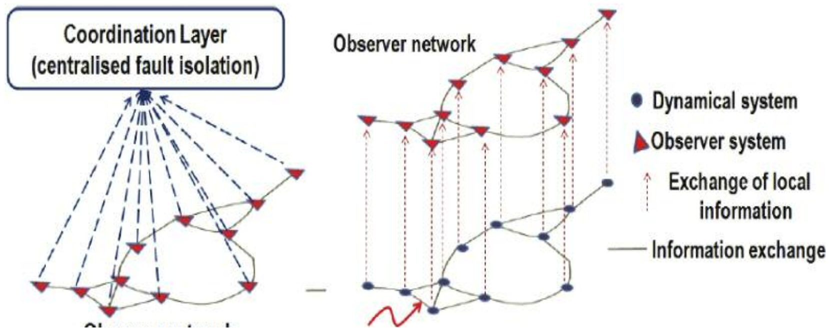

Figure 2: A simple example of a monitoring architecture.

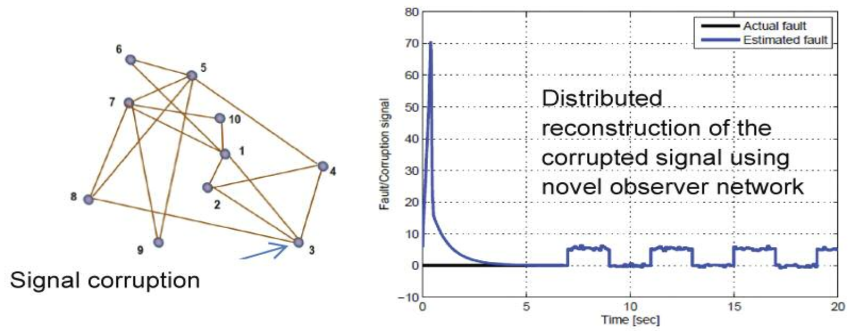

Figure 3: Signal corruption in the network and a reconstructed corruption signal.

References:

[1] D. A. Savić, J. K. Banyard (Eds.): "Water Distribution Systems", Thomas Telford Limited, 2011
[2] P. Derler, E. A. Lee, A. Sangiovanni-Vincentelli: "Modeling Cyber-Physical Systems", in proc. of the IEEE (special issue on CPS), 100(1):13-28, 2012
[3] M. Mesbahi, M. Egerstedt: "Graph Theoretic Methods in Multi-agent Networks", Princeton Series in Applied Mathematics, 2010.

Please contact:

Dragan Savic, Christopher Edwards, Prathyush Menon
University of Exeter, United Kingdom
E-mail: D.Savic@exeter.ac.uk, C.Edwards@exeter.ac.uk, P.M.Prathyush@exeter.ac.uk

SEAM4US: Intelligent Energy Management for Public Underground Spaces through Cyber-Physical Systems

by Jonathan Simon, Marc Jentsch and Markus Eisenhauer

Public transport operators suffer from the high energy consumption of their underground stations. Lighting, ventilation and vertical transport are crucial systems for the safety and comfort of passengers, but they represent the largest part of the non-traction energy required in metros. Intelligent control of these subsystems, however, can significantly reduce their energy consumption without impacting passenger comfort or safety or requiring expensive refurbishment of existing equipment.

Within the FP7 project SEAM4US, a system for intelligent energy management of public underground spaces has been developed, integrating cyber-physical systems to monitor the physical state of the station, passenger flow and energy consumption of all subsystems, as well as to control lights, fans and escalators. The system prototype has been deployed in the Passeig de Gràcia metro station in Barcelona. During its development, we have focused on building upon existing infrastructure that was not designed for energy efficiency while adding just as much additional technology as is necessary.

The SEAM4US system employs passenger density models and thermal models, integrating sensors and control algorithms in a model predictive control architecture that grants the optimal operation of the station plants under different occupancy and thermal conditions [1]. In this architecture, environmental parameters of the models are fed with sensor data, allowing the system to predict the environmental state depending on the chosen control policy. The results of this prediction determine the actual control output, while the effects of this control are also monitored to contribute to the calibration and learning process of the models.

A large wireless sensor network has been developed and deployed in the station, providing the thermal model with the current thermal state of the underground space. It has been designed to minimize energy consumption and maintenance work, specifically optimizing the operation of wireless sensor nodes in order to reduce battery replacement intervals, and providing self-diagnostics and self-configuration capabilities to ensure correct operation without human intervention. The positions of wireless sensor nodes and the frequency of sensor readings have been defined to allow the system to adapt to changes in energy consumption, user and environmental models. Existing CCTV cameras are enhanced with image processing to detect current and predict future levels of passenger density. Privacy of passengers is protected by performing all processing locally, filtering the image data to avoid recognisability of individuals, and transmitting only density levels in terms of integer numbers. In addition, an extensive network of smart meters has been installed to monitor energy consumption during the operation of the system and feeding it back to the model.

Control components have been developed to dynamically control the station lighting system, passenger transport, and ventilation. Based on a number of complex station models, combining input from the monitoring components with predictions about the state of the underground space, fan speed is adjusted proactively to optimize both air quality and energy consumption. Installation of a dedicated PLC and control logic allow for stepless control of the fan frequency while leaving the operation of this safety-critical system unimpeded by ensuring that the existing SCADA system, accessible by operators in the station or the Operations Control Center (OCC), always has priority and automatically takes over control if the SEAM4US system fails. By changing the escalator speed according to the predicted passenger density level, setting it to a slightly lower level for most of the time, escalator energy consumption can be reduced. Full speed is only required during peak times to ensure maximum capacity for passenger transport. Additional presence sensors ensure the safety of passengers, allowing the escalator to stop when no passengers are using it and to start when a person is approaching. As a complimentary approach, but integrated within the overall framework, a strategy for

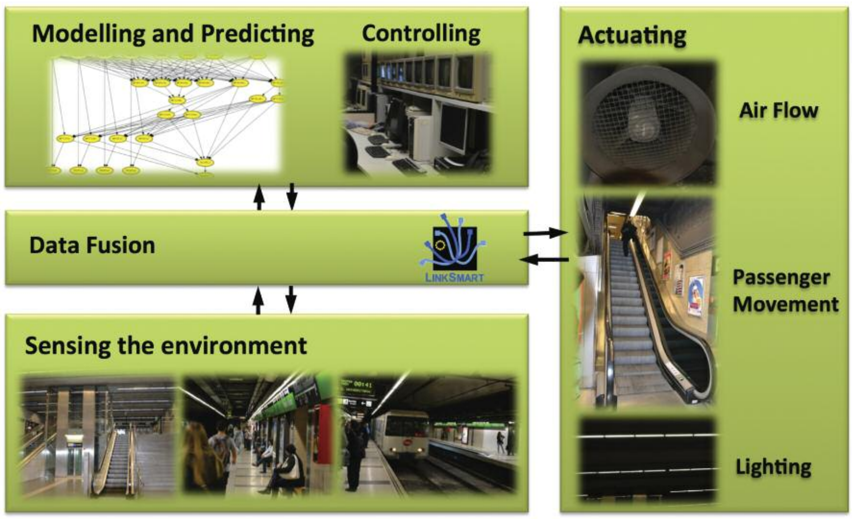

Figure 1: The SEAM4US approach

involving passengers in the energy-saving effort has also been developed by Fraunhofer FIT. By the use of a smartphone application, passengers are rewarded for taking the stairs instead of the escalators. This allows the SEAM4US system to run the escalators in energy-saving mode for longer periods of time. Dimmable LED lights have been deployed in the station, the brightness of which is controlled reactively according to current occupancy. While passenger safety requires more light when fewer people are present, luminosity can be reduced during peak times, however this level is obviously determined by legal and passenger comfort constraints. The system supervision component maximizes the reliability of the system and monitors the level of confidence of the system's predictions.

System development was coordinated by Fraunhofer FIT and is based on the LinkSmart middleware for networking heterogeneous devices and subsystems [2]. By means of its publish-subscribe mechanism, it enables efficient interfacing of architecturally independent components. In order to provide better support for energy efficiency applications of public transport spaces, it was enhanced with specific services for spatial modelling and handling of large amounts of event data. New components have been developed to integrate the existing station hardware as well as the newly installed sensors and actuators.

The result of this project is a complete prototypical solution for intelligent energy management of public underground spaces that integrates both existing and new infrastructure. The consortium consists of nine partners:
- Cofely Italia (Italy), energy service company – coordinator
- Marche Polytechnic University (Italy) – scientific coordinator
- Polytechnic University of Catalonia (Spain)
- Fraunhofer FIT (Germany), research institute – system development coordinator
- VTT (Finland), research centre
- University of Kassel (Germany)
- Almende (Netherlands), SME
- CNet Svenska (Sweden), SME
- TMB (Spain), metro operator.

Links:
http://www.fit.fraunhofer.de/en/fb/ucc/projects/seam4us.html
http://www.seam4us.eu
http://linksmart.eu

References:

[1] R. Ansuini, et al.: "Models for The Real-time Control Of Subway Stations", in proc. of BS2013
[2] M. Eisenhauer, P. Rosengren, and P. Antolin: "A Development Platform for Integrating Wireless Devices and Sensors into Ambient Intelligence Systems", in SECON 2009 workshops

Please contact:

Jonathan Simon, Marc Jentsch, Markus Eisenhauer
Fraunhofer FIT, Germany
E-mail:
jonathan.simon@fit.fraunhofer.de,
marc.jentsch@fit.fraunhofer.de,
markus.eisenhauer@fit.fraunhofer.de

A Comprehensive Port Operations Management System

by Christophe Joubert, Miguel Montesinos and Jorge Sanz

Galileo, the global navigation satellite system funded by the EU, will soon provide highly accurate and precise position measurements on Europe's roads. But the primary mode of international trade, the maritime industry - responsible for nearly 90 percent of world trade - still relies on outdated technology with limited precision capacity whilst being relatively expensive and inefficient, namely: Electronic Chart Display Information Systems (ECDIS) or paper charts in conjunction with GPS receiver and laser-based Berthing Aid Systems (BAS). Improvements in port traffic management and operational efficiency, as well as reductions in operating expenses, CO_2 emissions and the environmental impact of shipping can all be achieved by providing the necessary centimetre positioning/speed accuracy based on WiMAX and PDGNSS technologies. We have developed a set of built in devices and solutions on top of existing maritime embedded devices to form a safer integrated cyber-physical system for maritime transport.

The aim of the European project DOCKINGASSIST (FP7 - 286828, November 2011-October 2013) was to develop a novel wireless network based on IEEE 802.16-2009 [1], WiMAX, in order to create a centralized, cost-effective, real-time, accurate vessel location and monitoring system for harbours using a Precise Differential Global Navigation Satellite System (PDGNSS) [2] with increased accuracy to less than 10 cm, based on Real Time Kinematic (RTK) positioning (see Figure 1).

Within this project, we have been working on the continuous monitoring of position, speed and direction of vessels with high accuracy throughout all phases of the approach, allowing ships to be precisely located in the harbour zone, taking into account the dimensions of the ship. The system allows for enhancements to the trajectory through early adjustments from the rate of turn when entering in the port if required, as well as assisting in the approach towards the berth. The harbour is equipped with a differential GNSS base station and wireless technology offering a range of several kilometres, in order to send the DGPS corrections data to the vessels.

We specified the architecture and contributed to the development of the DockingAssist base station software system. The system allows a port to optimize the operational maritime activities related to the vessel flow within the service area of the port, integrating all stakeholders and all the relevant information systems. In particular, it integrates information available at the port, such as weather data and specific features with vessel positions provided by AIS systems and accurate vessel positions provided by the DockingAssist device. This complex system also includes a real time analysis engine (Decision Support System Module) that can be configured to detect anomalous events such as unauthorized vessel positions, early detection of vessel collision and challenging weather conditions.

Working in parallel with the Docking Assist project, the Prodevelop geospatial research team extended the DSS module to a real time analysis engine based on spatial information (Geo-Decision Support System Module) to detect geospatial events. This technology has been included in a port management GIS module that automates relevant operational events, such as anchoring, berthing/unberthing, pilot and tug operations, bunkering, entry and exit of areas such as the port service area, waypoints or inner harbour, port exit with pending requested anchoring, etc. Detected events are published in real-time to a web Geographical Information System (GIS) application [3], available to all connected clients, as shown in Figure 2.

The event display includes an event list, warning icons, map warnings, location of events on the map, integration with operations register, and query of events.

In order to process massive data streams coming from port and location sensors, in real time, the architecture makes use of the eXtreme Transaction Processing Platform (XTPP), message distribution middleware implementing Data Distribution Service (DDS) international OMG standard, Non-SQL repository, trajectory simplification, push notifications, eXtensible Messaging Presence Protocol (XMPP) IETF standard RFC 6120 and communication over WebSockets / Bidirectional-streams Over Synchronous HTTP (BOSH).

The results to date are very encouraging: "Posidonia Operations" has been built upon the experience acquired in Docking Assist to provide a solution for managing the sea-side operations of vessels. The system has built in and comprehensive integration capability with security management systems, building management systems, SCADA, AIS/VTS systems, gate operating systems, terminal operating systems, port community and data interchange packages.

In our research on the Docking Assist project we collaborated with several SMEs and research centres, namely ATEKNEA and ASCAMM (Spain), VTT (Finland), Port of Cork (Ireland), Marimatech (Denmark), Net Technologies (Greece) and Runcom Technologies (Israel).

This research is also part of the strategic goals and recommendations for the EU Maritime Transport Policy until 2018 (IP/09/84), namely to ensure the long term performance of the European maritime transport system as a whole. Our work is partially supported by the Spanish MEC INNCORPORA-PTQ 2011 program.

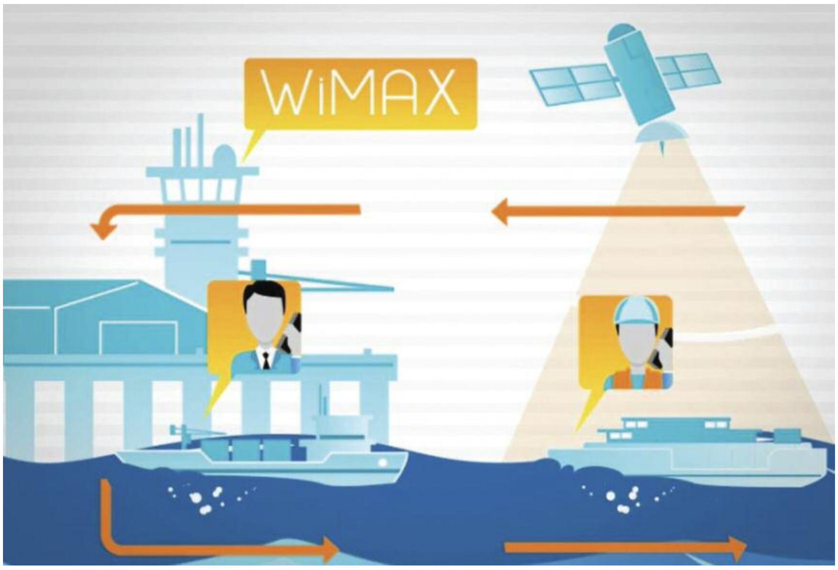

Figure 1: WiMAX-based PDGNSS location system improving port traffic management.

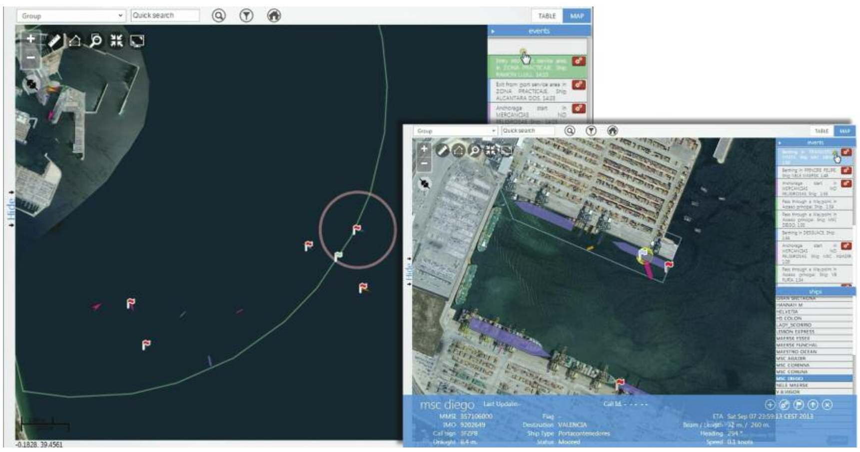

Figure 2: Event display to operators in the Posidonia Operations port system.

Links:
http://www.dockingassist.eu
http://www.prodevelop.es

References:

[1] IEEE 802.16 Working Group: "IEEE Standard for Local and Metropolitan Area Networks, Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems", IEEE Standard 802.16-2009, May 2009

[2] L. Jiang et al: "DockingAssist: A novel vessel navigation system design based on WiMAX and DGNSS", IEEE WPNC 2013, 1-6

[3] C. Joubert et al.: "Real-Time Visualization of MV/LV Energy Alarms on GIS Web Applications", ERCIM News 2013(92).

Please contact:

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

Safety Criticality Analysis of Multi-Agent Air Traffic Management Systems: A Compositional Hybrid Systems' Approach

by Elena De Santis, Maria Domenica Di Benedetto and Giordano Pola

We present a formal framework for analysing the impact of non-nominal operating modes on the safety of next generation ATM procedures under study in the SESAR 2020 Concept of Operation. Efficient complexity reduction algorithms are also derived for applying the proposed methodology to realistic large-scale ATM systems.

Air Traffic Management (ATM) is an important application domain exhibiting many features of CPS such as heterogeneity, complexity, and human operators in-the-loop.

Several disciplines are being used to assist ATM experts in the design of robust procedures; among these, of particular relevance is Resilience Engineering [3]. Resilience indicates that operations and organizations are capable of resisting a variety of demands and capable of recovering from variations, degradations, disruptions, and any condition affecting the stability of the operation or organization. In other words, Resilience Engineering addresses the design of joint cognitive systems, both in nominal and non-nominal conditions. However, because of the complexity of ATM joint cognitive systems, Resilience Engineering as applied to ATM is at an early stage of development. Formal mathematical models and analysis methods offer a key complementary approach that is needed to render Resilience Engineering effectively applicable to ATM systems. Making Resilience Engineering applicable to ATM was the main goal of the SESAR WP-E Research Project "Mathematical Approach towards Resilience Engineering in ATM (MAREA)".

Among the topics included in MAREA, of fundamental importance in the analysis of novel ATM procedures was the study of the impact of non-nominal operating modes on the safety of the overall system. To tackle this important problem we proposed in [2] a formal approach that is based on the use of a number of relevant hybrid systems' techniques, including compositional hybrid systems' modelling and hybrid observers' synthesis.

While aircraft dynamics is commonly described by differential equations, pilots' and air traffic controllers' behaviours are well modelled by finite state machines, whose discrete states and transitions mimic the procedure the agents are requested to follow. Hybrid systems' formalism, featuring both discrete and continuous dynamics, is characterized by an expressive power that we proved to be general enough to adequately describe ATM systems, both in nominal and non-nominal conditions. While nominal modes of operation are dictated by the procedure the agents are requested to follow, non-nominal modes may originate from several causes, including malfunction or disruption of technical devices or unpredictable behaviour of human operators in-the-loop. To study the impact of non-nominal operating modes on the safety of the ATM procedures we used the notion of critical observability [1]. This notion corresponds to the possibility of detecting whether the current discrete state of a hybrid system is in a critical set, representing unsafe, un-allowed or non–nominal situations. When a hybrid system exhibits this property, a hybrid observer can be constructed which automatically detects the criticality of the current discrete state.

Our first investigation of critical observability considered each agent of ATM systems in isolation. However, in multi-agent ATM scenarios, agents cannot be considered in isolation because their interaction is responsible for the occurrence of unsafe situations that cannot be captured in isolation. For this reason, we proposed a compositional hybrid systems framework that describes the behaviour of each agent as well as their interaction. This framework can be viewed as a generalization of the classical notion of serial composition to a multi-agent setting, and provides a systematic way to study critical observability of multi-agent systems. More precisely, models are first created of each agent with a hybrid system, then, by applying the compositional

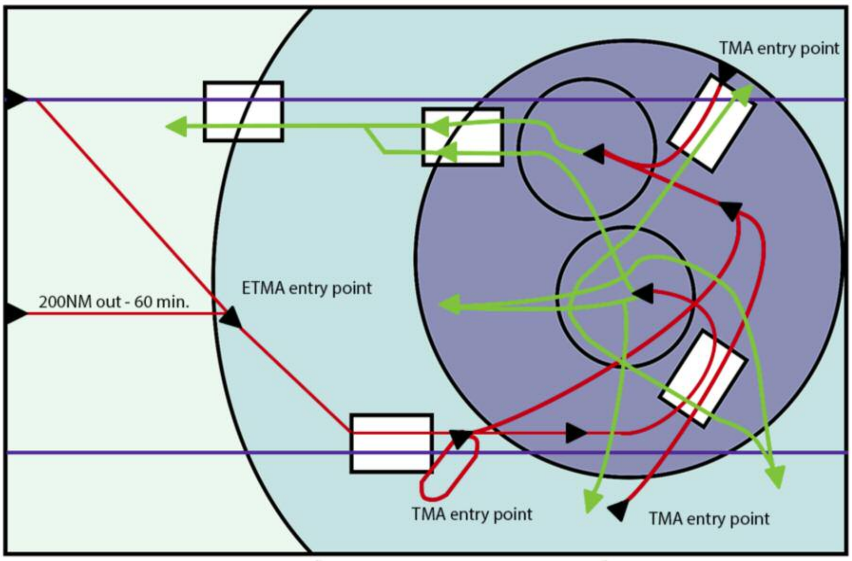

Figure 1: A possible scenario of the TMA T1 operation. Standard instrument departure routes are depicted in green, standard terminal arrival routes in red, and cruise routes at a lower flight level in blue.

rules, a unique hybrid system is obtained. A critical relation is then defined which describes the occurrence of safety-critical situations in the composed hybrid system. Studying safety in multi-agent ATM scenarios then translates to studying critical observability of the obtained (composed) hybrid system with respect to the critical relation.

Although formally sound, this approach is hardly applicable to realistic scenarios because of the large number of variables involved. To overcome these difficulties we proposed algorithms based on bisimulation theory, widely used in the area of formal methods to mitigate software verification.

We analysed the Terminal Maneuvering Area (TMA) T1 operation, a procedure selected within the MAREA consortium as a benchmark, exhibiting most relevant features arising in the novel SESAR 2020 Concept of Operation. We considered a scenario involving 25 agents, comprising more than 1.68 x 10^{18} discrete states. We showed that this procedure is not critically observable. This implies that there are safety-critical configurations which cannot be detected by pilots or air traffic controllers. In other words, in some situations, not only is the human operator's awareness of a safety critical situation incorrect, but furthermore, it cannot be improved before a safety-critical situation occurs. This analysis also proposed alternative solutions which ensure safety of the procedure.

We wish to thank Henk Blom and Mariken Everdij (NLR) for fruitful discussions on this paper.

Link:
http://dews.univaq.it/

References:
[1] E. De Santis et al: "Critical Observability of a Class of Hybrid Systems and Application to ATM systems", in "Stochastic Hybrid Systems: Theory and Safety Critical Applications", LNCIS 337
[2] E. De Santis et al: "Final modelling and analysis of SESAR 2020 ConOps", Report MAREA D4.4, http://complexworld.eu/wiki/File:D4.4-v1.0.pdf
[3] Resilience Engineering Perspectives, Preparation and Restoration, Vol. 2, Ashgate, England, 2009.

Please contact:
Elena De Santis, Maria Domenica Di Benedetto and Giordano Pola, Center of Excellence for Research DEWS, University of L'Aquila, Italy.
E-mail: {elena.desantis, mariadomenica.dibenedetto, giordano.pola} @univaq.it

Security in the Era of Cyber-Physical Systems of Systems

by Stamatis Karnouskos

Critical infrastructures are increasingly equipped with modern Cyber-Physical Systems (CPS) and Internet-based services that enhance their functionalities and operation. However, traditional security practices fall short when it comes to addressing the multitude of security considerations, not only at individual system but also at system-of-system level. The creation of recent sophisticated tools, such as Stuxnet, Duqu, Flame, and the Mask, is the prequel to a nightmare in a CPS-dominated future.

The existence and utilization of highly complex tools such as Stuxnet, Duqu, Flame, and the Mask in real-world attacks demonstrate that we have entered the era of sophisticated cyber warfare. The concern is that the Cyber-Physical Systems (CPS) [1] monitoring and controlling critical infrastructures such as the smart grid [2] may be susceptible to cyber-terrorism, and that even small criminally inclined groups would be able to create attacks with asymmetrical impact. Since the majority of the world's SCADA/DCS and PLC systems can be found in high-tech industrial facilities in Europe, US and Japan, it is imperative to invest in security as a process. Adequately addressing security in the cyber-physical system era, however, poses a significant challenge.

Attacks such as that of Stuxnet relied on a number of existing vulnerabilities, some of which dated back two years [3]. Updates should have been applied during that time, but owing to the "air-gap" isolation they were considered unnecessary. Additionally, many of the industrial infrastructures that employ CPS are long-lived with life times of 10+ years. This means updates are not always possible (for older systems), or are not implemented as often owing to the lengthier testing time and the fear of unwanted side effects. However, these poorly defended, poorly patched and poorly regulated systems will be the first ones that will be used as Trojan horses to attack the more modern systems with zero-day attacks. ``Don't touch a running system" may not apply in the CPS era.

Modern CPSs do not constitute a monolithic platform and are not developed by a single stakeholder. On the contrary, they consist of various hardware and software parts "glued" together to perform the required tasks. Hence the first problem that arises is how to trust the individual parts of the CPS and how to guarantee a deterministic behaviour. Addressing security only at hardware or software levels is not enough. The operational context also needs to be considered for safety and dependability reasons. Even if both are fully certified and addressed, there is still no guarantee that actions that compromise a CPS will not occur during its lifetime – hence adequate security measures also need to be taken in the operational context.

Trust is a fundamental issue to consider. As an example, CPS hardware components can be used as carriers of attacks and entry points to a system. Common attacks utilize “trusted” parts of a system, such as USB ports, the Ethernet card, the battery etc. to host and execute malicious code that bypasses the operating system’s guards. Digitally signed software should not be blindly trusted either. As an example, Stuxnet installed two kernel drivers that were digitally signed by valid certificates that were stolen from two different issuing companies. Real time online validation of certificates may limit the exposure window.

Software and hardware security are not the only issues to be considered; human users must be included in the process. Security clearance on people does not imply security on their accompanying assets. In the Stuxnet case [3], a trustworthy employee with an unknowingly rootkited laptop or an infected USB flash drive would be enough to spread the malware. This could be, for instance, a contractor carrying a personal device, who is assigned to do maintenance on a facility.

Lack of security-considerations at CPS development time may lead to insecure software with bugs that may result in unpredictable system behaviour or, even worse, a controllably malicious operational behaviour. Other pitfalls may also be possible, as demonstrated by the Stuxnet [3], which was able to take advantage of something that should never have existed in the first place, i.e., default hard-coded access accounts and passwords in industrial PLCs.

It is interesting that we place so much trust on CPS systems even when some of their complex operational stages may not be secure. Stuxnet impersonated the normal behaviour of the PLC, and any network management system or control room operator would have been unlikely to see a rogue PLC as its signals were faked [3]. As such, only the independently observed physical process and that reported by the Stuxnet-infected PLC data would mismatch. Hence, safeguards need to be in place, not only on individual CPS, but also on the processes in which they participate. This requires system-of-system wide behaviour monitoring and checks for anomalies. Heuristics for estimating behaviour deviation may provide hints, which should be assessed and analysed in conjunction with other metrics. This is challenging but probably achievable to some degree if the process is under the control of a limited number of stakeholders. However, in the envisioned widely collaborative CPS systems-of-systems this is a daunting task.

In CPS system-of-systems it will be difficult to do holistic code reviews, systematic testing and checks at design and runtime [2]. Hence, software “bugs” which may have a tangible impact on the physical world will happen more often, while their impact will be hard to assess. It is not clear how much effort will need to be invested in designing and integrating software for such complex system of systems versus testing it. Additionally, the range of qualifications that will be required by future engineers to perform these tasks will have to be broader and more in-depth which is challenging. Automatic tools that do the model checking as well as detect potential safety-critical issues on large scale multi-dimensional applications will be needed.

CPSs control real-world infrastructures and thus have a real-world impact. Dependability in CPS and their ecosystems will be the key factor for their application in critical systems; it will determine to what extent our core critical infrastructure will be vulnerable in the future. The CPS era is in need of solutions that will support it at device, system, infrastructure and application level [1]. This includes the whole lifecycle from cradle-to-grave of its components and services. This is a grand challenge and includes multi-disciplinary engineering, modelling, emergent behaviour, human interaction etc. Finally, it has to be kept in mind that security is a multi-angled process in which vulnerability and risk analysis may dictate what is an acceptable level. Solutions focusing asymmetrically on particular aspects, whilst neglecting others, may give a false sense of safety and security, which will be shattered by reality as Stuxnet, Duqu, Flame, and the Mask have recently demonstrated.

References:

[1] “Cyber-Physical Systems: Driving force for innovation in mobility, health, energy and production”. Tech. rep., Acatech – German National Academy of Science and Engineering, 2011

[2] S. Karnouskos: "Cyber-Physical Systems in the SmartGrid", in IEEE 9th International Conference on Industrial Informatics (INDIN), Lisbon, Portugal, 2011

[3] S. Karnouskos: "Stuxnet Worm Impact on Industrial Cyber-Physical System Security", in 37th Annual Conference of the IEEE Industrial Electronics Society (IECON), Melbourne, Australia, 2011.

Please contact:

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

Optogenetics to Unravel the Mechanisms of Parkinsonian Symptoms and Optimize Deep Brain Stimulation

by Antoine Chaillet, Diane Da Silva, Georgios Detorakis, Christophe Pouzat and Suhan Senova

By allowing a precise stimulation of targeted neurons through light impulses, optogenetics is revolutionizing our understanding of the brain. Its recent use in primate deep brain structures promises unprecedented data to model the mechanisms underlying Parkinson's disease symptoms and to improve their treatment by electrical stimulation.

Neuronal synchronization plays a central role in brain functioning. It is linked to memory, cognition and movement path generation. Abnormal synchronization in certain cerebral zones can also lead to pathological states such as Parkinson's disease, essential tremor, or epilepsy. In particular, Parkinsonian symptoms are known to be linked to persistent beta band (13-30Hz) oscillations in deep brain structures called basal ganglia.

Deep brain stimulation (DBS) is a symptomatic treatment of several synchronization-related neurological diseases [1]. It consists of electrical stimulation of deep brain structures through permanently implanted electrodes. In most existing DBS techniques, the stimulation is a square signal (1 to 5V amplitude at around 120 to 180Hz) whose precise parameters are derived primarily by trial-and-error on each patient. It operates in an open loop: no cerebral information or models of the dynamics involved are exploited.

Despite its strong success and impressive efficacy in most cases, current DBS treatment still suffers from considerable limitations, including: the spread of electrical stimulation within non-targeted brain tissues; and the low level of knowledge about the exact functioning of DBS. More importantly, a number of side effects have been reported, such as cognition issues, depression, speech disturbances, and balance impairment. Moreover, the permanent stimulation of deep brain structures leads to an overuse of embedded electrical resources, which in turn imposes further surgical operations for battery replacement. Finally, the trial-and-error selection of the present DBS parameters has been effective because of the almost immediate effects of DBS on Parkinsonian motor symptoms; new therapies utilizing DBS technology will not allow such a tuning (for instance, the beneficial effects of stimulation can take weeks to manifest in dystonia or obsessive compulsive disorders).

These limitations seem strongly linked to the disproportionate signal amplitude and the open-loop nature of the current DBS signals. Control engineering tools may contribute to a better understanding of basal ganglia circuitry and the adaptation of DBS signals to real-time cerebral activity. These are the keys for the development of more adaptive and parsimonious stimulation policies.

The recent technique of optogenetics provides unprecedented neurophysio-

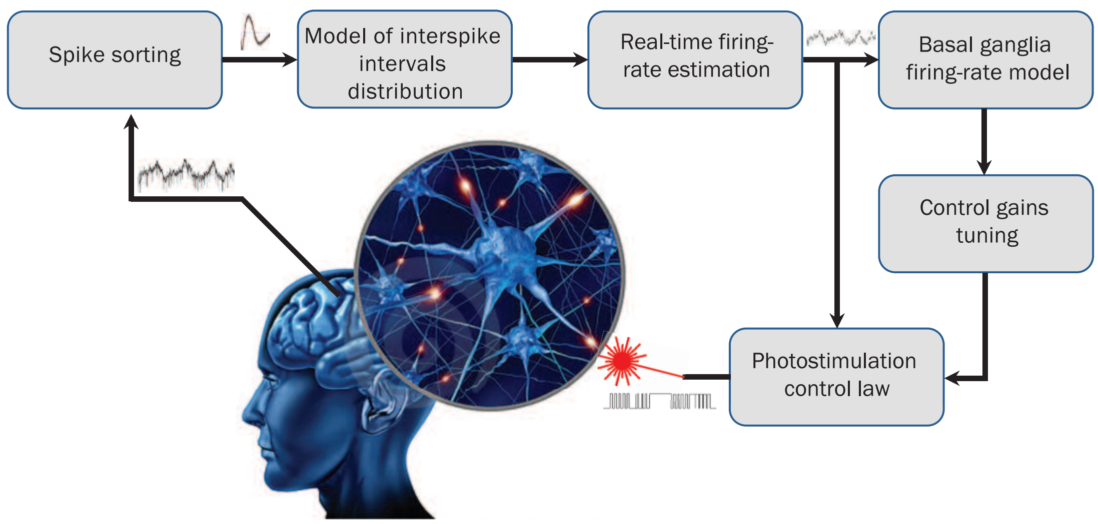

Figure 1: Schematic description of the SynchNeuro strategy. Electrical measurements of single spike recordings in specific brain structures, thanks to permanently implanted electrodes, allow to estimate the firing rate of the whole neuronal population after spike sorting and probabilistic identification. The estimated firing rate is then used off-line to identify the dynamical properties of the neuronal activity and to deduce the optimal tuning of the control gains. It is also used on-line to provide a real-time photostimulation feedback, using optogenetics, to reduce the magnitude of pathological oscillations.

logical data, which may improve our understanding of basal ganglia interaction. It relies on a gene transfer technology to make targeted neurons sensitive to photostimulation by visible wavelengths after the selective transfer of genes coding for the expression of photosensitive ionic channels on neuronal membrane [2](Boyden et al. 2005). Thus, simple pulses of intense light, for instance through implanted optical fibers, can induce the response of the photosensitized selected neurons only, while not affecting neighbouring neurons. This technique enables an accurate spatiotemporal neuromodulation at the level of the millisecond and of one targeted neural network. Combined with electrode measurements, optogenetics offers unprecedented possibilities to dissect basal ganglia mechanisms in health and disease, and thus to optimize DBS treatment.

The SynchNeuro project, started in 2013 and co-financed by the French ANR and the Digiteo research network, proposes a multidisciplinary research effort to investigate methodologies for the analysis and control of neuronal synchronization and their application to the treatment of neurological diseases by closed-loop DBS, based on optogenetics data. The aim is to understand the underlying neurological phenomena and to develop closed-loop real-time DBS controllers that intelligently and autonomously alter pathological oscillations.

More precisely, four objectives will be addressed. First, a firing-rate model of basal ganglia under photostimulation will be developed based on in vivo primate data. This modelling and identification will be performed on both healthy and Parkinsonian primates. In order to interpret the neurophysiological data acquired by optogenetics and in vivo extracellular recordings, the raw data needs to be processed and analysed. This involves the implementation of a spike sorting method [3] and the development of a stochastic model for parameter identification. Those parameters will later be used to develop a more precise model of the neural populations, including the specificity of photostimulation. Second, based on this firing-rate model, analytical conditions for the pathological oscillations onset will be developed using tools from nonlinear control theory. Third, innovative closed-loop DBS signals that adapt in real-time to the recorded brain activity will be designed in order to attenuate the pathological oscillations. These DBS signals are expected to be more efficient and more respectful than existing open-loop DBS strategies. Finally, both the population model and the developed DBS signals will be validated numerically and in vivo. The overall setup is reported in Figure 1.

By adapting control techniques and theoretical background to the specific constraints imposed by DBS (model imprecision, limited measurements and actuation, nonlinearities), decisive improvements in terms of adaptation, side effects and energy consumption can be expected. The optimization of the DBS technique is a perfect illustration of the gains that can be achieved through a closer interaction between control engineering and neuroscience.

Acknowledgement: R. Aron-Badin (CEA MIRCen) and P. Hantraye (CEA MIRCen), L. Greco (L2S) and S. Palfi (APHP H. Mondor hospital) are also partners of the SynchNeuro project and have contributed to this article. This work has received support from the French ANR project SynchNeuro, the French Digiteo project NEUROSYNCH, the French institute iCODE, and the European FP7 Network of Excellence HYCON2.

References:

[1] A.L. Benabid et al.: “Long-term suppression of tremor by chronic stimulation of the ventral intermediate thalamic nucleus”, The Lancet, 337(8738), pp.403–406, 1991

[2] E.S. Boyden et al.: “Millisecond-timescale, genetically targeted optical control of neural activity”, Nature Neuroscience, 8(9), pp.1263–1268, 2005.

[3] C. Pouzat, O. Mazor, G. Laurent: “Using noise signature to optimize spike-sorting and to assess neuronal classification quality”, journal of neuroscience methods, 122, pp.43–57, 2005.

Please contact:

Antoine Chaillet
L2S, University Paris Sud 11,
Supélec, France
E-mail: antoine.chaillet@supelec.fr

Big Data in Healthcare: Intensive Care Units as a Case Study

by Ioannis Constantinou, Andreas Papadopoulos, Marios D. Dikaiakos, Nicolas Stylianides and Theodoros Kyprianou

Traditional medical practice is moving from relatively ad-hoc and subjective decision making to modern evidence-based healthcare. Evidence is based on data collected from: electronic health record (EHR) systems; genomic information; capturing devices, sensors, and mobiles; information communicated verbally by the patient; and new medical knowledge. These resources produce a collection of large and complex datasets, which are difficult to process using common database management tools or traditional data processing applications.

Currently, the data generated in the process of medical care in intensive care units (ICUs) are rarely processed in real-time, nor are they collected and used for data analysis. The existence of user-friendly platforms for accessing and analyzing such massive volumes of data could leverage an era of medical knowledge discovery and medical care quality improvement. The current absence of such platforms is due partly to the difficulty of accessing, organizing, and validating voluminous health care data produced with high time-varying arrival rates by various types of sensors, and notably to the variability of proprietary software solutions in use. The need for fast computationally-intensive scientific analysis of the data generated in ICUs thwarts the use of traditional databases and data processing applications and demands the development of cloud based solutions.

Our platform, shown in Figure 1, is able to collect data from proprietary used software solutions, such as the Clinical Information System (CIS) or via the Intensive Care Window (ICW) framework [1]. Data are produced by specialized medical equipment, or - in the case of laboratory test results and medication records - manually entered to the hospital database system by healthcare professionals. Monitors are capable of monitoring about 95 vital signs measured by devices such as patient monitors, patient ventilators, smart infusion pumps and laboratory test databases. In particular:

- The ICU patient monitor device (Philips IntelliVue MP 70) measures about 30 vital signs such as Heart Rate, EGG and SpO2.
- The patient ventilator device (Puritan Bennett 840 ventilator) measures about 20 parameters and vital signs such us ventilator mode, minute volume and tidal volume.
- The smart infusion pump allows the persistent storage of two values: (a) the instant flow of drugs at the time they have been validated by a caregiver (typically a nurse) and (b) the total amount of drugs administered during the last hour.

Table 1 illustrates the measurement rates of each of the vital sign type, laboratory test and medication. Vital signs are measured more frequently than medication and laboratory

Description	Capture rate
Patient Monitor 31 Vital Signs (ECG, Heart Rate, SpO2, etc)	Every 30 seconds
Patient Ventilator 20 Vital Signs and parameters (ventilator mode, minute volume, tidal volume, etc)	Every 30 seconds
27 Laboratory Tests (glucose, urea, creatinine, potassium, etc)	Thirteen of the 27 tests are captured 5 times/day. The remaining 14 are captured 24 times/day.
19 Drugs (Thiopental, acetylcesteine, Furosemide, dexamethasone, etc)	Hourly

Table 1: Capture rates for different vital signs, lab test and drug data.

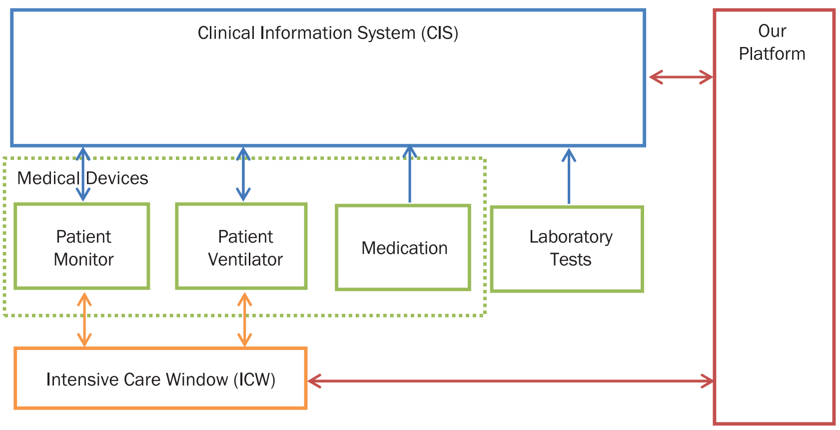

Figure 1: Platform interconnection diagram. Medical devices store data to CIS automatically. Healthcare professionals manually enter laboratory test and medication data to the system. Our platform is interconnected to CIS and ICW for automated data retrieval.

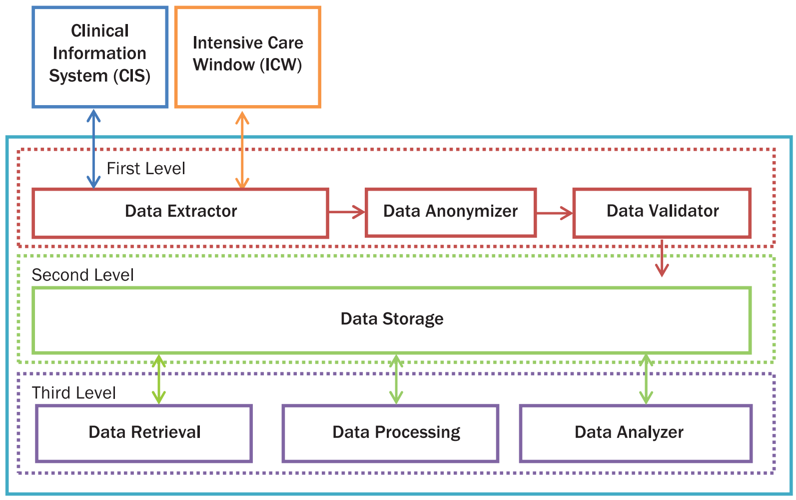

Figure 2: The future platform architecture. The platform consists of seven modules and three levels. The first level consists of Data Extractor, Data Anonymizer and Data Validator. The Data Extractor module is responsible for the interconnection between CIS and ICW. The Data Anonymizer is responsible for anonymizing critical private patient data. The Data Validator discards invalid values of vital signs, laboratory test and medication based on strictly predefined rules. At the second level is the storage module. The Data Storage module is a cloud NoSQL database. The third level is responsible for data retrieval, data processing and data analysis. The Data Retrieval module is an interface for retrieving data from the cloud. The Data Processing module is responsible for data processing and, finally, the Data Analyzer is responsible for the data analysis. The Data Analysis module will provide the powerful functionality of R statistical package.

tests. In a small 20-bed ICU, 2.8 million records with a capacity of about 11.6 MB are stored in the platform's data repository per day, which is equivalent to an annual data production of about 1.058 billion records with capacity of about 4.2GB.

Using data collected through our platform, we: (a) evaluated 12 different Glucose Variability (GV) indexes for mortality prediction; and (b) deployed a multivariate logistic regression model, enriched with patient characteristics and laboratory tests, such as age, gender, BMI, CRE, etc, in order to raise the prediction accuracy [2]. Specifically, data collected through our platform, for about 1000 patients for a period of 12 months, have been considered and analyzed using Matlab. We concluded that the majority of the GV indexes exhibit high accuracy for predicting mortality, with GVI achieving the best accuracy (72%). Furthermore, our model, based on multiple patient characteristics, achieved an 82% success rate in predicting mortality.

Our goal is to integrate our current platform to a cloud-based platform for storing and analyzing ICU medical data. Our future platform would comprise three levels (Figure 2). The first level is responsible for data extraction, anonymization and validation [1]. The second level is responsible for data storage and the third level includes the tools and the interfaces for retrieval, processing and analysis of the data. We plan to extend the third level by integrating an R (a powerful open source statistical analysis software) connector for Hadoop, optimizing and expanding the various components as well as testing it on ICU patients' vital signs for real time analysis and pattern matching.

References:

[1] N. Stylianides et al.:"Intensive Care Window: Real-Time Monitoring and Analysis in the Intensive Care Environment", IEEE Transactions on Information Technology in Biomedicine, Vol. 15, pp. 26-32, 2011.
[2] S. Kokkoris, et al.: "The ability of various glucose variability metrics for mortality prediction in critically ill patients: a comparative study", 27th Annual Congress CCIB – Barcelona, ESICM, 2013.

Please contact:

Marios D. Dikaiakos
University of Cyprus
E-mail: mdd@cs.ucy.ac.cy

SEMEOTICONS: Face Reading to Help People stay Healthy

by Sara Colantonio, Daniela Giorgi and Ovidio Salvetti

The European project SEMEOTICONS offers a fresh perspective on educational programs and lifestyle intervention for the prevention of cardiovascular diseases. SEMEOTICONS will develop a multisensory platform, in the form of a mirror, to be integrated into daily-life settings. The mirror will be able to detect and monitor over time facial signs correlated with cardio-metabolic risk, and give personalized advice to end-users on how to improve their habits.

The face is the pre-eminent channel of communication among humans: it is a mirror of status, emotions, mood. This is the base principle of Medical Semeiotics, which looks at the face as a revealer of the healthy status of an individual, through a combination of physical signs (e.g., skin colour, subcutaneous fat) and facial expressions. Medical Semeiotics dates back as far as Aristotle's times. Nevertheless, it is still used by medical practitioners today, who rely on their ability to read the face of patients to steer the medical examination and decide on diagnostic investigations.

SEMEOTICONS, funded by the EC (FP7), began in November 2013 with the objective of moving Medical Semeiotics to the digital realm. SEMEOTICONS proposes the translation of the semeiotic face code into computational descriptors and measures, which can be automatically extracted from videos, images, and 3D scans of the human face. In particular, SEMEOTICONS addresses signs related to cardio-metabolic risk as cardiovascular diseases are the leading cause of mortality worldwide.

SEMEOTICONS also aims to bring semeiotic analysis closer to everyday life, ie, from the surgery to the home, the gym, the pharmacy, etc. This will enable people to self-assess and self-monitor the status of their well-being over time. A multisensory platform in the form of an interactive smart mirror, called the Wize Mirror, will be developed. The Wize Mirror is designed to fit into daily-life settings, by maximizing non-invasiveness and natural interaction modalities, according to the Ambient Intelligence paradigm. It will seamlessly integrate contactless sensors (3D optical sensors, multispectral cameras, gas detection sensors, microphones) with a user-friendly touch-screen interface (Figure 1).

The heterogeneous data collected by the Wize Mirror – videos, images, 3D scans, and gas concentration of individuals standing in front of the mirror – will be processed by dedicated algorithms, which will extract a number of biometric, morphometric, colorimetric, and compositional descriptors [1]. The descriptors will be integrated to define a virtual individual model and an individual wellness index to be traced over time. The Wize Mirror will also offer suggestions and coaching messages, with a personalized user guidance, aimed at the achievement and the maintenance of a correct life-style.

The empowerment of individuals, in terms of their ability to self-monitor their status and improve their life-style, is expected to have a great impact on the reduction of health expenditure and disease burden. It is well-known that the cost of health systems grows exponentially with the aging of the population, together with the widespread use of complex, and sometimes inappropriate, diagnostic procedures. Prevention is the best strategy to limit the spread of cardio-metabolic diseases. SEMEOTICONS offers a fresh perspective on educational programs and lifestyle intervention.

SEMEOTICONS will address significant scientific and technological challenges. It promises touch-less data acquisition via a low-cost platform; temporal and spatial data synchronization; real-time processing of multimodal data; and intelligent solutions to map iconic facial signs to repeatable, reliable computational measures correlated with risk factors. The project consortium has been built to strike a good balance between innovation and technology-driven research. It comprises six research organizations (with competence in both ICT and medical areas), two industrial partners, and two SMEs, from seven European countries (France, Greece, Italy, Norway, Spain, Sweden, United Kingdom). The project coordinator is the Istituto di Scienza e Tecnologie dell'Informazione (ISTI) of the National Research Council of Italy (CNR), located in Pisa.

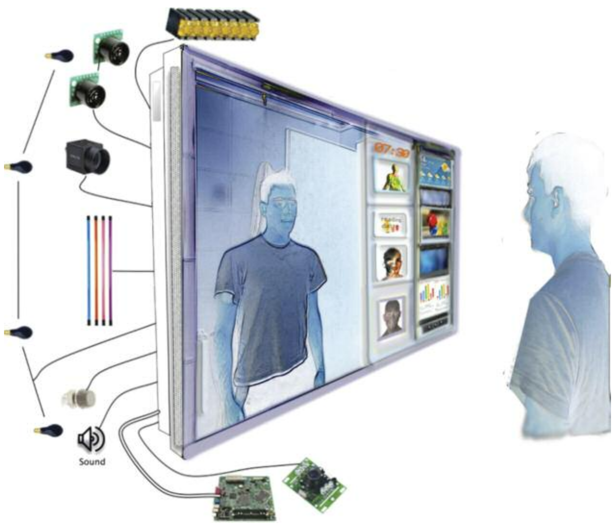

Figure 1: The Wize Mirror is a multisensory platform which collects videos, images, 3D scans of the human face and gas concentration signals, looking for signs correlated with cardio-metabolic risk.

Link: http://www.semeoticons.eu/

Reference:

[1] E. Vezzetti and F. Marcolin: "3D human face description: landmark measures and geometrical features", Image and Vision Computing 30 (2012), 698 - 712

Please contact:

Sara Colantonio - ISTI-CNR, Pisa, Italy
Tel.: +39 050 621 3141
E-mail: sara.colantonio@isti.cnr.it

DYMASOS - Dynamic Management of Physically Coupled Systems of Systems

by Radoslav Paulen and Sebastian Engell

We are developing new tools for the management of large physically-interconnected production and distribution systems.

The well-being of the citizens of Europe depends on the reliable and efficient functioning of large interconnected systems, such as electric power systems, air traffic control, railway systems, large industrial production plants, etc. Such large systems consist of many interacting components, such as generation plants, distribution systems, and large and small consumers. Those subsystems are usually managed locally and independently, according to different policies and priorities. The dynamic interaction between the locally managed components gives rise to complex behaviour and can lead to large-scale disruptions, such as power black-outs. Large interconnected systems with autonomously acting sub-units, such as these, are called "systems of systems".

DYMASOS addresses the issues associated with managing systems of systems, where the elements of the overall system are coupled by flows of physical quantities, such as electric power, steam or hot water, materials in a chemical plant, gas, potable water, etc. DYMASOS, initiated in October 2013, is a three-year research project funded by the European Commission under the 7th Framework Programme for Research & Technological Development in the "Information and Communication Technologies" theme.

DYMASOS focuses on the theoretical approaches for dynamic management of physically coupled systems of systems driven by and validated on real-world use cases and supported by eventual deployment of prototyped methodologies and toolsets. The DYMASOS consortium comprises academic partners, industrial partners, small and medium enterprises and a large provider of engineering and technological solutions.

The project will develop new methods for the distributed management of large physically connected systems with local management and global coordination. The explored approaches involve:
- Modelling and control of large-scale systems analogously to the evolution of the behaviour of populations in biological systems,
- Market-like mechanisms to coordinate independent systems with local optimization functions,
- Coalition games in which agents that control the subsystems dynamically group to pursue common goals.

The developed approaches will be validated in large-scale simulations of case studies provided by the industrial partners or already involved supporters from the industry. From the validation, general conclusions will be drawn about the suitability of the proposed distributed management and control mechanisms for certain classes of systems of systems. This will provide guidelines for the design of evolving large physically coupled systems of systems with respect to the interplay of local autonomy and global management. These guidelines, developed in close collaboration with industrial

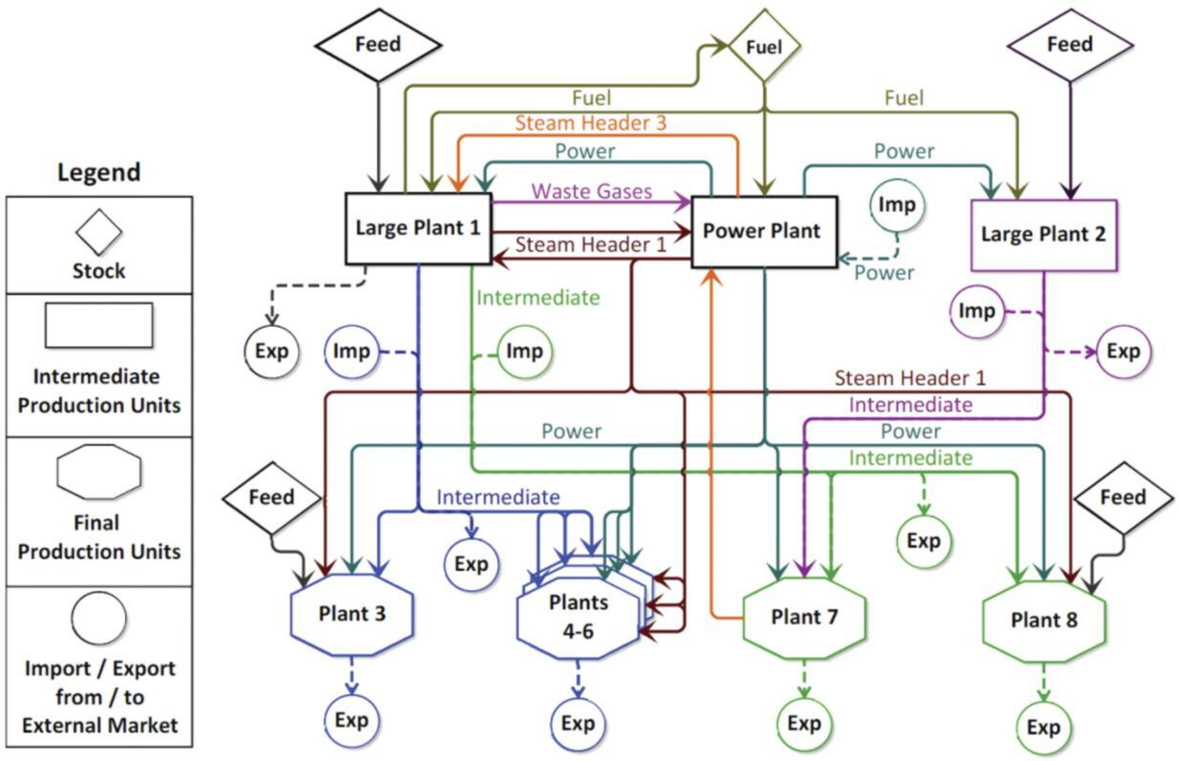

Figure 1: A schematic representation of a physically-coupled system of systems using a chemical production site as a case study.

partners, will influence the next generation of software and hardware development for optimal management of large interconnected systems.

The industry-driven case studies of real applications that are considered in the project include:
- The balancing and stabilizing of the electricity distribution grid of the city of Koprivnica that is a highly-automated complex system involving multiple electricity producers and consumers (households, charging stations of electrical vehicles, etc.).
- The dynamic management of charging of electric vehicles in the city of Malaga that involves the management of the network of electrical vehicles and power station with vehicle-to-grid charging capabilities.
- The energy management of a large chemical production site that seeks the right balance for an operation of the power plant (steam network) to satisfy the demands of the production subunits and to minimize the amount of imported energy resources.
- The cross-plant process management in an integrated chemical production complex. This case study is concerned with managing a network of parallel reactors that share a flare system and limited cooling power.

For the usability of the results, it is essential that the engineering, implementation and long-term maintenance of such advanced solutions are taken into account. The project will develop prototypical tools that facilitate development, testing, deployment and maintenance of these solutions. A key element of these engineering tools is an open simulation platform that enables dynamic management and control methods to be tested with accurate models of the overall system and the connection to external models that are available in the application domains.

The expected technical outcomes of the project are:
- Innovation in distributed management methods for complex interconnected systems,
- Progress in methods for the rigorous analysis and validation of systems of systems,
- Improvements in the management of electric grids and of large production complexes,
- Tools for the engineering of management systems for systems of systems,
- Identification of technology gaps in advanced management and coordination methods.

The developed coordination methods will lead to improved system stability and lower resource consumption in industrial production, and in electric power generation and distribution. This will result in a reduction of the CO_2 emissions, increased competitiveness of European industry and lower prices for consumers. DYMASOS is thus contributing to the goal of a greener and more competitive Europe.

Link: http://www.dymasos.eu

Please contact:
Sebastian Engell
Technische Universität Dortmund Germany
E-mail: Sebastian.Engell@bci.tu-dortmund.de

3DHOP: A Novel Technological Solution for the Development of Online Virtual Museums

by Marco Potenziani, Marco Callieri and Roberto Scopigno

The interactive visualization of 3D content on the Web has gained momentum in recent years, thanks to the release of ad-hoc technology as part of HTML5. However, so far, few products fully exploit this opportunity for the easy and quick embedding of high-resolution 3D models in a Web page. 3DHOP (3D Heritage Online Presenter) is an advanced solution for easy publishing of 3D content on the Web.

3DHOP allows anyone with basic webpage-creation skills to setup an HTML page displaying a scene composed by high-resolution 3D models. This free tool is essentially a viewer that can stream and render effectively 3D data over the net, letting the user interact with high-resolution 3D content (see Figure 1). It is oriented towards the Cultural Heritage (CH) field, where the need to handle highly detailed digital representations of cultural artefacts is often mandatory; however, it may also be useful in other contexts requiring high performance and interactive visualization of high-resolution 3D content.

3DHOP has been developed by the Visual Computing Lab, ISTI-CNR, in the framework of several EU projects addressing the management of large collections of 3D digital models: the EC IP project "3DCOFORM", the EC NoE "V-MusT" and the EC pilot project "3D-ICONS". It is the result of 10 years of R&D focused on the design and implementation of easy-to-use tools that enable CH professionals to present, visualize and enrich high-quality digital 3D models (the first important result of which was the Inspector system [1]).

The rationale underlying 3DHOP was the lack of tools to manage the interactive visualization on the Web of high-resolution 3D geometries in a simple way. While there are various products (both free and commercial) that can re-create complex virtual scenes made of simple 3D entities, they are not suitable for the management of complex 3D geometry,

Figure1: A collection of 3D models presented with 3DHOP (an example of layout)

```
presenter.setScene({
  meshes: {
    "Gargoyle": {
      url: "models/multires/gargo.nxs"
    }
  },
  modelInstances: {
    "Model1": {
      mesh: "Gargoyle"
    }
  },
  trackball: {
    type: TurnTableTrackball,
    trackOptions: {
      startPhi: 0.0,
      startTheta: 0.0,
      startDistance: 2.5,
      minMaxPhi: [-170, 170],
      minMaxTheta: [-10.0, 50.0],
      minMaxDist: [0.5, 3.0]
    }
  }
});
```

Figure2: Declarative programming in 3DHOP: the virtual scene is defined using different text fields, easily modified by the user with different values to build a similar small scene (fully textual approach).

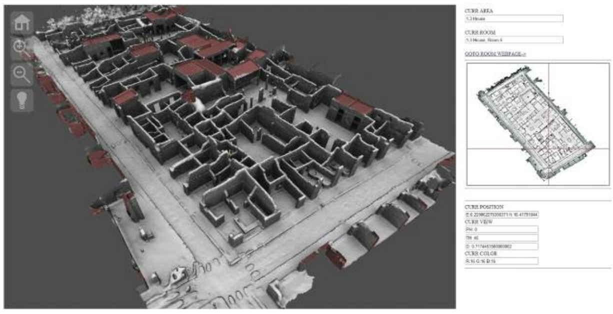

Figure3: Handling a large 3D mesh (about 20M triangles) in 3DHOP using the Nexus technology.

and are mainly aimed at expert Computer Graphics (CG) developers.

3DHOP can be used by anyone with simple webpage-creation skills, without the need for specific CG expertise; the goal was to be easy-to-use and easy-to-learn. For this reason, 3DHOP uses general web programming mechanisms (declarative programming). The various components of the visualization page (the 3D models, the Trackball, the Canvas, the Scene) are treated like HTML entities in the same way as the rest of the structural elements of the web page, and are declared and customized using simple parameters. A developer can just pick a starting example and modify it by simply changing some parameters of the entities (see Figure 2).

3DHOP has been created using standard Web technologies (like JavaScript and HTML5), without relying on external components; it can run on all the most widely used browsers (Chrome, Firefox, Opera and soon IE) and all the principal OS (Win, Mac-OS and Linux) without the need for plugins. It must be stressed that although 3DHOP has been designed for the Web, it only needs a Web browser to work; this means that it can also run locally, to implement museum kiosks, for example. Unlike similar tools, the 3DHOP viewer is totally client side, and does not require a specialized server.

3DHOP currently supports two types of 3D models: single-resolution (PLY and soon OBJ) and multi-resolution (NEXUS). In particular, the multiresolution Nexus technology (another achievement of ISTI-CNR), supports the streaming of high-resolution 3D meshes over the HTTP protocol, enabling the exploration of very large models (millions of triangles, like in Figure 3) on commodity computers with standard internet connections. 3DHOP is based on WebGL (the standard API for the CG on the Web developed as the JavaScript equivalent to OpenGL|ES 2.0), and on SpiderGL, a JavaScript utility support library for WebGL (again developed by the Visual Computing Lab [2]), which provides an API with the typical structures and algorithms for real-time rendering in the development of 3D web applications.

The evolution of 3DHOP will be focused on the integration of visualization components for other multimedia layers, like images, audio, video, etc. Work is already underway for the management of large 3D terrains and RTI (Reflectance Transformation Imaging) images [3].

With 3DHOP, we hope to offer a novel and versatile solution for publishing 3D content on the Web, facilitating deployment and access to high-resolution 3D resources by non-experts in 3D programming. This new technology, particularly suitable for Online Virtual Museums, should encourage the global dissemination and deployment of high-quality 3D cultural heritage content via the Internet.

Links:
3DHOP: http://vcg.isti.cnr.it/3dhop
Nexus: http://vcg.isti.cnr.it/nexus
SpiderGL: http://vcg.isti.cnr.it/spidergl

References:
[1] M. Callieri et al.: "Virtual Inspector: a flexible visualizer for dense 3D scanned models", IEEE Computer Graphics and Applications, 2008
[2] M. Di Benedetto et al.: "SpiderGL: a JavaScript 3D Graphics Library for next-generation WWW", Web3D Conference, 2010
[3] G. Palma et al.: "Dynamic Shading Enhancement for Reflectance Transformation Imaging", ACM Journal on Computing and Cultural Heritage, 2010.

Please contact:
Marco Potenziani, Marco Callieri, Roberto Scopigno, ISTI-CNR, Italy
E-mail: marco.potenziani@isti.cnr.it, marco.callieri@isti.cnr.it, roberto.scopigno@isti.cnr.it

The Entity Registry System: Publishing and Consuming Linked Data in Poorly Connected Environments

by Christophe Guéret and Philippe Cudré-Mauroux

Sixty-five percent of the world's population is deprived of ICT-enhanced data-sharing because of poor access to digital technology. The ExaScale Infolab and Data Archiving and Networked Services (DANS) have designed a new framework, the "Entity Registry System" (ERS), and produced a reference implementation making it possible to publish and consume Linked Open Data in poorly connected environments.

There is no longer any doubt about the social and economic value of accessible, open data. Leveraging the pervasive presence of the Web, Linked Data principles and Web-based portals are promising technologies for implementing data-sharing applications. But the trivial architectural assumptions around data connectivity and availability of large-scale resources keeps ICT-enhanced data sharing out of reach for about 65% of the world's population. The majority of the world's population is digitally disconnected from the rest of the world, and are thus deprived of (Linked) Open Data, and its associated benefits [1].

Work on the Entity Registry System (ERS) started in 2012 as a joint project between the ExaScale Infolab and Data Archiving and Networked Services (DANS) with a generous grant from Verisign Inc. This work is placed in the context of the World Wide Semantic Web initiative aimed at ensuring that the Semantic Web becomes a reality for everyone (see http://worldwidesemanticweb.org). Out of the three challenges - infrastructures, interfaces and relevancy - that the World Wide Semantic Web initiative encompasses, ERS tackles the first one.

In contrast to established economies and city dwellers, developing economies and rural populations typically do not have a high speed and stable access to the internet. They also lack access to Web-based data portals where they could publish their new data and consume third-party datasets. The need for sharing open data in such environments is real, however. Monitoring environmental hazards, online medical help, emergency response, and remote education are but a few examples of data sharing scenarios that are essential to the development of emerging countries [1,2]. Technology is increasingly available to such populations, though still in limited forms: solar-powered access to low-speed Internet, local mesh networks, community radio, usb-sticks to carry data, low-cost computers, basic mobile phones, etc. In this context, our goal is to design an approach to publish and consume Linked Open Data using such technologies. More specifically, we decided to focus on existing and relatively cheap equipment such as XO-1 laptops, RaspberryPis, and mesh networks.

In a nutshell, ERS creates global and shared knowledge spaces through series of statements. For instance, "Amsterdam is in the Netherlands" is a statement made about the entity "Amsterdam" relating it to the entity "the Netherlands". Every user of ERS may want to either de-reference an entity (for instance, by asking for all pieces of information about "Amsterdam") or contribute to the content of the shared space by adding new statements. This is made possible via "Contributors" nodes, one of the three types of node defined in our system. Contributors can interact freely with the knowledge base. They are responsible for publishing their own statements but cannot edit third-party statements. Every set of statements about a given entity contributed by a single author is wrapped into a unique named graph to avoid conflicts and enable provenance tracking. In a typical ERS deployment, the end-user machines (eg, the XO-1 laptops) work as Contributor nodes. Two Contributors in a

Figure 1: An example of ERS deployment that could be in a school. The four XO laptops are Contributors using the RaspberryPi as a Bridge to asynchronously exchange data with another school.

closed P2P network can freely create and share Linked Open Data. In order for them to share data with another closed group of Contributors, we introduce "Bridges". A Bridge is a relay between two closed networks using the internet or any other form of direct connection to share data. Two closed communities, for example two schools that are willing to share data, can each set up one Bridge and connect these two nodes to each other. The Bridges will then collect and exchange data coming from the Contributors. These bridges are not Contributors themselves, they are just used to ship data (named graphs) around and can be shut-down or replaced without any data-loss. Lastly, the third component we define in our architecture is the "Aggregator". This is a special node every Bridge may push content to. As its name suggests, an Aggregator is used to aggregate entity descriptions that are otherwise scattered among all the Contributors. When deployed, an aggregator can be used to access and expose the global content of the knowledge space or a subset thereof.

An ERS deployment may contain any number of Contributors, Bridges and Aggregators depending on the data-sharing scenario at hand. The system is, however, always "offline by default", relying only on Contributors as persistent data stores, and taking advantage of data flows that are biased towards locally relevant and potentially crowd-sourced data. Figure 1 shows a complete test deployment using four Contributors (XO-1 laptop) and a bridge (RaspberryPi model B) connected through a mesh network.

ERS has been developed in Python and Java. It uses CouchDB and CumulusRDF to store the RDF data serialized as JSON-LD. The code of our reference implementation is freely available under an open source license (see http://ers-devs.github.io). We are currently working on extending it to support further data sharing scenarios and to make it usable on more hardware.

Links:
http://ers-devs.github.io
http://dans.knaw.nl
http://exascale.info
http://worldwidesemanticweb.org

References:
[1] F. Estefan, C. Spruill, S. Lee: "Bringing Open Contracting Data Offline", Open Contracting Blog, 2013, [Online], http://www.open-contracting.org/bringing_open_contracting_data_offline

[2] T. Unwin: "ICT4D: Information and Communication Technology for Development", Cambridge University Press, 2009, ISBN 9780521712361.

Please contact:
Christophe Guéret
Data Archiving and Networked Services (DANS), The Netherlands
E-mail: christophe.gueret@dans.knaw.nl

Spin-Up: A European Project Aimed at Propelling University Spin-off Growth

by Manuel Au-Yong-Oliveira, João José Pinto Ferreira, Qing Ye and Marina van Geenhuizen

The "Spin-Up" project investigated what sort of entrepreneurship training and coaching program will contribute to the development of key entrepreneurial skills.

Entrepreneurial endeavours with roots in academia, also known as University spin-off firms (USOs), tend to have slow growth rates compared to corporate spin-offs and other high-technology start-ups. The project, called "Spin-Up" - reflecting its aim to improve the growth (or "spin upwards") of USOs - addressed the question: "What sort of entrepreneurship training and coaching program will contribute to the development of key entrepreneurial skills, both technical and behavioural, essential to enable and leverage university spin-off growth?" "Spin-Up" involved a European Consortium of organizations, comprising: Advancis Business Services and INESC TEC (both in Portugal), Leaders2Be and the Technical University of Delft (the Netherlands), and the Lappeenranta University of Technology (Finland). The University of Porto and the Technology Transfer Centre at the University of Lodz also participated. The main findings of the project were presented [1] at the European Conference on Entrepreneurship and Innovation (ECIE) 2013, in Brussels.

The research, which took place between 2011 and 2013, was funded by the European Commission – Lifelong Learning / Erasmus Program (Enterprise and University cooperation). Ninety-nine USO firms from four countries (Finland, Poland, Portugal and the Netherlands) made up the sample of firms from which in-depth data was gathered, in face-to-face interviews or via questionnaires answered electronically. This data was then processed and the results published in Van Geenhuizen and Ye's Spin-Up Research Report [2], this report thus representing the main thrust of the research effort, which revealed that USOs' growth is hampered by missing entrepreneurial skills, such as skills in innovation, marketing, sales, strategy, internationalization, leadership, human resource management, financial literacy and gaining financial capital. These areas were thus the focus of the pilot training course, which was supported by training manuals and case studies developed by the consortium partners, and which tested the Spin-Up training approach and course content in three countries – Finland, Portugal, and the Netherlands.

Notably, work was divided amongst the consortium members according to expertise and knowledge. Van Geenhuizen and Ye [2] produced the following spider diagrams illustrating the research findings on the entrepreneurial skills (Figure 1a-d – for the Netherlands, Finland, Poland and Portugal)

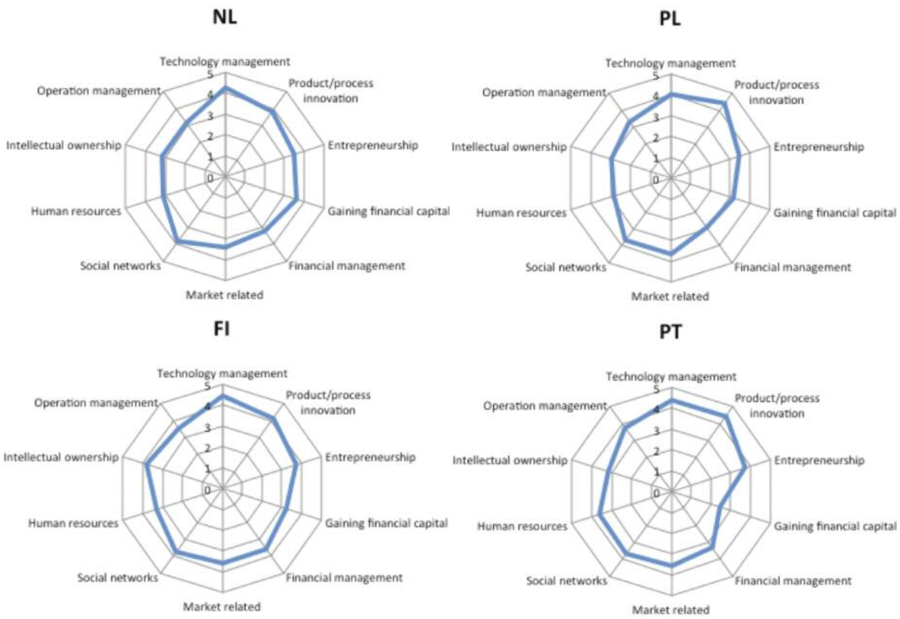

Figure 1a-d: Entrepreneurial skills map of spin-off firms in four countries (reproduced from [2])

Figure 1a-d show that trends differed among the four countries included in the research. Specific areas requiring improvement include: financial management (particularly in Poland and the Netherlands) and gaining financial capital (Portugal, and to a lesser extent, Poland and Finland); human resource management (Poland) and intellectual ownership (Netherlands, Poland and Portugal). For all countries involved, stronger areas included technology management and product/process innovation. Leadership skills were similarly strong in all four countries, though managers in the Netherlands tended to provide more opportunities to collaborators, while also giving more information to employees.

The pilot training program revealed that in addition to the need for basic entrepreneurial skills, entrepreneurs would benefit from customized, problem-oriented and interactive coaching or training. The involvement of specialists, such as a venture capitalist and a business angel, in the program, was greatly appreciated. In terms of organization, a modular structure was preferred, allowing USOs to participate per block, and participants expressed a preference for blocks to be scheduled outside of working hours.

The Spin-Up project enhanced the sharing of knowledge between project partners and with USOs, and longer term benefits will doubtless be observed as a consequence. Further work in this area is needed – in Europe as well as between Europe and other continents.

Link: http://www.spin-up.eu

References:

[1] M.A. Oliveira et al.: "Spin-Up: A Comprehensive Program Aimed to Accelerate University Spin-Off Growth", in proc. of ECIE 2013, vol. 1, pp.34-44, Brussels, Belgium
[2] M. Van Geenhuizen, Q. Ye: "Spin-Up Research Report", final version – July 2012.

Please contact:

Manuel Au-Yong-Oliveira, INESC TEC, University of Porto, and University of Aveiro, Portugal
E-mail: mao@ua.pt

A System for the Discovery and Selection of FLOSS Projects

by Francesca Arcelli Fontana, Riccardo Roveda and Marco Zanoni

Developing software systems by reusing components is a common practice. FLOSS (Free, Libre and Open Source Software) components represent a significant part of the reusable components available. The selection of suitable FLOSS components raises important issues both for software companies and research institutions. RepoFinder supports a keyword-based discovery process for FLOSS components, and applies well-known software metrics and analyses to compare the structural aspects of different components.

FLOSS (Free, Libre and Open Source Software) components are widely used in industry. They are managed on dedicated web platforms, called Code Forges, designed to collect, promote, and categorize FLOSS projects, and to support development teams. Some examples are Launchpad, Google Code, Sourceforge, Github, Bitbucket and Gnu Savannah. Code Forges provide tools and services for distributed development teams, eg, version control systems, bug tracking systems, web hosting space. Currently, Github hosts more than ten million software components, and is probably the largest. Sourceforge is one of the first Code Forges, created in 1999, with more than 425,000 components, all of them well categorized, while Google Code hosts about 250,000 components.

To find a component satisfying specific requirements, we can start from a simple query on a generic search engine or on each Code Forge. The risk is that we receive many useless answers, because FLOSS project names are often ambiguous and difficult to recover. Even retrieving the source code of a particular project can be difficult, because it may be managed by different portals, with different versioning technologies (e.g., Git, SVN, Mercurial (HG), CVS, GNU Bazaar).

RepoFinder is implemented in a web application and has been designed to retrieve and analyze the source code from open source software repositories. It performs these tasks through the following steps:

- Crawling of widely known Code Forges (Github, Sourceforge, Launchpad and Google Code) through their APIs or by parsing of HTML pages; during this step, projects are tagged with labels extracted from the web pages crawled and the other available metadata (e.g., tags, forks, stars).
- Analysis of the components, by integrating external software for the computation of metrics and the detection of code smells in Java code. We integrated the following tools: Checkstyle, JDepend, NICAD, PMD, Understand and another tool for metrics computation, called DFMC4J, developed at our Lab.
- Integration of data gathered by the code analyzers in a dataset, to be used for statistical analyses.
- Browsing of the data through a search engine based on simple google-like queries.

We have indexed more than 90,000 projects using RepoFinder, and found more than 2,000 labels. Analyses can only be applied to Java systems with a working Maven configuration. Maven ensures that all relevant sources are included in the analyses. It also retrieves all dependencies needed to compile the systems and to gather the required metrics correctly. In fact, Maven automatically downloads library dependencies, and depending on the configuration, can execute code generation tools (e.g., ANTLR, JOOQ), only selecting source code without test cases and, if required, can compile the system. RepoFinder can also visualize the results gathered for one system and compare the results of different systems.

The overall architecture of RepoFinder is represented in Figure 1.

We aim to apply and extend RepoFinder in the following directions:

- Analysis of software history. Software repositories contain important information about the development team and system maturity. Many analyses can be performed in this context; for example, to understand whether the quality of a software has been improved over time, and if new releases tend to be corrected quickly, evidencing a low maturity. This kind of data can be evaluated for an individual project or compared for different projects.
- Language support. Currently, RepoFinder analyzes Java systems configured to be built by Maven. Many other languages and build systems are available and widely exploited in open source development. Examples of other build systems for Java are Eclipse, Ant, Gradle, Ivy, Grape and Buildr. Analyzing additional languages could require the integration of further analysis tools. Other tools also exist that could retrieve useful information, like Cobertura (for systems with large test suites), or FindBugs and other code smell detection tools.

Other applications exist with similar functionalities, such as SonarQube and Ohloh. However, RepoFinder differs in that it combines discovery functionalities, more related to Ohloh, with analysis functionalities, which is the domain of SonarQube. There are a number of datasets that gather data coming from different Code Forges, which could be integrated in the crawling and analysis processes of RepoFinder. However, their availability is variable and, unfortunately, some of these data-collection projects are closed.

Please contact:
Francesca Arcelli Fontana and Marco Zanoni
University of Milan Bicocca, Italy
E-mail: {arcelli, zanoni}disco.unimib.it

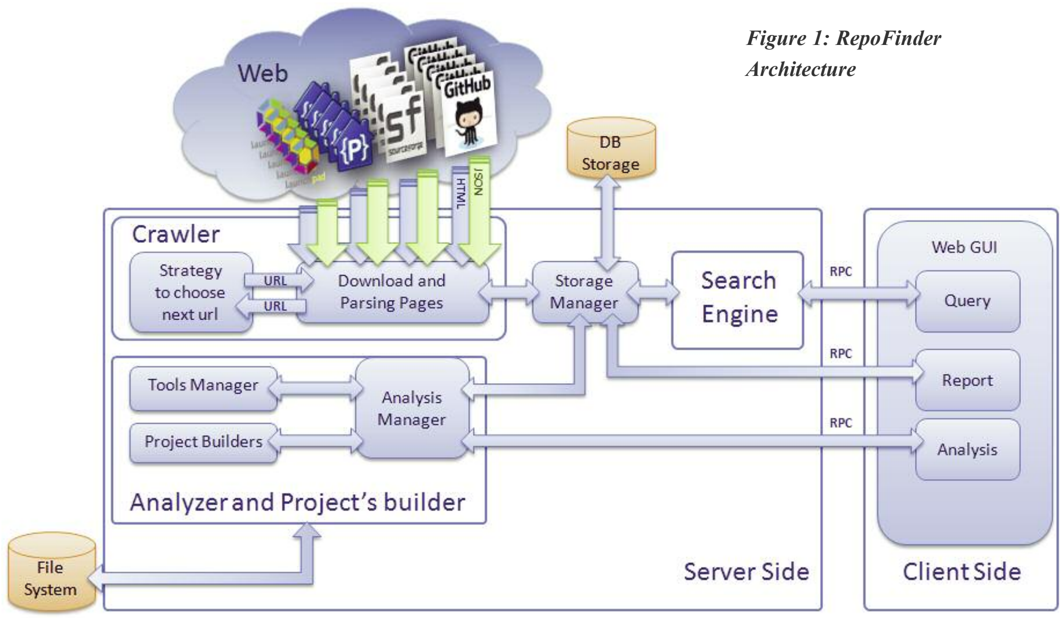

Figure 1: RepoFinder Architecture

VITRO – Vision-Testing for Robustness

by Oliver Zendel, Wolfgang Herzner and Markus Murschitz

VITRO is a methodology and tool set for testing computer vision solutions with measurable coverage of critical visual effects such as shadows, reflections, low contrast, or mutual object occlusion. Such "criticalities" can cause the computer vision (CV) solution to misinterpret observed scenes and thus reduce its robustness with respect to a given application

Usually, a large numbers of recorded images are used for testing in order to cover as many of these "criticalities" as possible. There are several problems with this approach:

- Even with a very large set of recorded images there is no guarantee that all relevant criticalities for the target application are covered. Hence, this approach is inappropriate for certification.
- Recording large numbers of images is expensive. Additionally, several situations cannot be arranged in reality due to reasons of safety or effort.
- For evaluating test outputs, the expected results (called "ground truth" or GT, see Figure 1) are needed. Usually these are generated manually, which is expensive and error prone.

A number of test data sets are currently publicly available, but mostly are not dedicated to specific applications. Their utility for assessing a computer's vision solution with respect to a target application is limited. Overall, the use of recorded test data does not provide a satisfactory solution.

Approach

To address this situation, AIT is developing VITRO, a model-based test-case generation methodology and tool set that produces test data with known coverage of application-specific typical and critical scenes. This development was started in the ARTEMIS project R3-COP (Resilient Reasoning Robotic Co-operating Systems). Figure 2 gives a schematic overview of all VITRO components.

Modelling: The "domain model" describes the objects (geometry, surface appearance etc.), which can occur in rendered scenes, together with relationships between and constraints on them, as given by the application (e.g. on their size and relative positioning). It also contains information about background, illumination, climatic conditions, and cameras, i.e. the used sensors. For certain families of objects, such as clouds, generators are available or can be developed on demand.

Criticalities: A catalogue comprising over one thousand entries has been created by applying the risk analysis method HAZOP (Hazard and Operability Study) to the realm of computer vision (CV). This process considered light sources, media (e.g. air phenomena), objects with their interactions and observer effects. The latter include artifacts of the optics (e.g. aberration), the electronics (e.g. noise), and the software. Only application-relevant entries of this catalogue are selected for the actual testing.

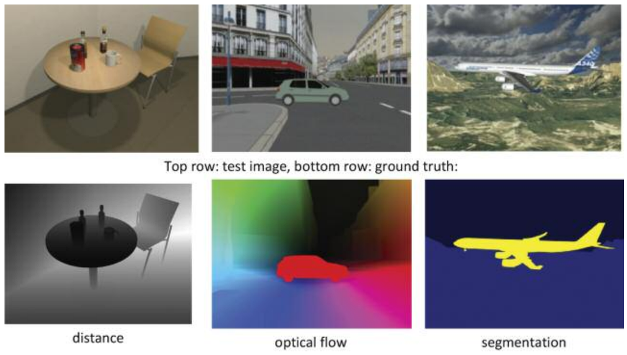

Figure 1: Examples of test images and ground truths generated by VITRO.

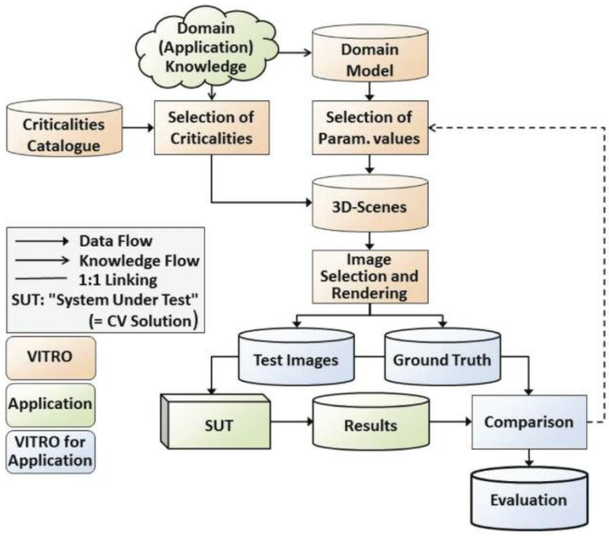

Figure 2: VITRO – schematic application process.

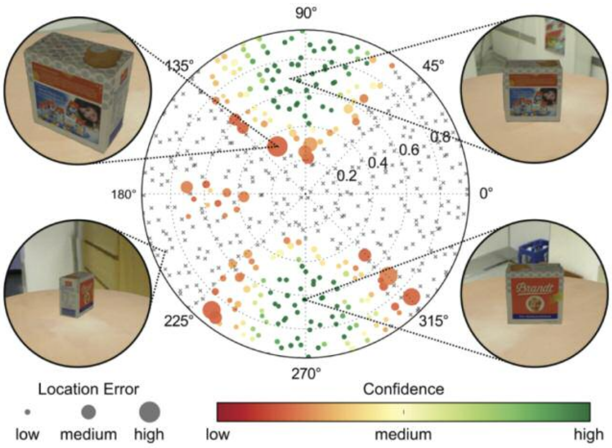

Figure 3: Example of test results visualization.

Scene generation: The parameters defined in the domain model, (e.g. object positions) establish a parameter space. This space is sampled with so-called "low discrepancy" which achieves an optimal coverage. Criticalities are included by means of further constraints or their occurrence is checked. Specific scenarios can be defined for the generation of characteristic curves or diagrams (see Figure 3).

Test data selection and generation: The previous step can yield redundant test images that are very similar to others. Therefore, test image candidates are characterized with properties derived from their original scenes, e.g. visual fraction of certain objects. These are used to group candidates and select representatives, which are rendered, and finally GT is generated for them.

Application: Generated test data can be applied to existing CV solutions for assessing their robustness with respect to the modelled application. But VITRO can already be used during development (e.g. test-driven development). In this case, simple scenes are generated first. Once these are processed robustly, increasingly challenging test data follow iteratively. Finally, adaptive and learning approaches can also be trained and tested.

An Application Example

Figure 3 shows just one of the possible visualizations of test results. The ability of a home service robot to locate a certain box depending on object distance and orientation was tested. Each dot represents a test case. Coloured dots denote successful detection, with colour corresponding to the algorithm's own confidence measure, and with size representing the location error against GT. Grey dots denote no detection. From this test run and its visualization we can derive that the confidence measure is good (no large green dots), but the object is detected only in few orientations.

Conclusion

VITRO provides test data sets perfectly tailored to specific applications with maximum expressive power about the robustness of the user's computer vision solution. Generated from models, the data are intrinsically consistent and precisely evaluable. Its use provides the following benefits:

- verifiable coverage of typical and critical situations
- automatic test data generation and test evaluation
- strongly reduced recording effort for real test data
- no manual definition of expected results
- testing of dangerous situations without risks
- applicable even during development
- support for adaptive/learning computer vision solutions
- results usable for future certification.

Link: http://www.ait.ac.at/vitro/

References:

[1] O. Zendel, W. Herzner, M. Murschitz: "VITRO – Model based vision testing for robustness." ISR 2013 - 44th International Symposium on Robotic (2013)

[2] W. Herzner, E. Schoitsch: "R3-COP – Resilient Reasoning Robotic Co-operating Systems"; DECS-Workshop at SAFECOMP 2013; in proc. of SAFECOMP 2013 Workshops, ISBN 2-907801-09-0, pp.169-172

[3] W. Herzner, M. Murschitz, O. Zendel; "Model-based Test-Case Generation for Testing Robustness of Vision Components of Robotic Systems"; in proc. of SAFECOMP 2013 Workshops, ISBN 2-907801-09-0, pp.209-212

Please contact:

Wolfgang Herzner; AIT Austrian Institute of Technology GmbH/ AARIT

E-mail: wolfgang.herzner@ait.ac.at

Building Systems of Aerial Drones

by Luca Mottola, Niklas Wirström and Thiemo Voigt

SICS researchers are investigating how to equip aerial drones with service-oriented interfaces and ways to facilitate the programming of a fleet of drones.

Aerial drones, such as the one in Figure 1, are formidable mobile computing platforms. Their ability to move in an almost unconstrained manner allows them to sense from, and act on, areas that cannot be reached by traditional devices. As the cost of aerial drones drops (see Links 1 and 2) and drone technology progressively moves into the mainstream (See Links 1,3), aerial drones are expected to add to the arsenal of technology that allows Cyber-Physical Systems to bridge the gap between the digital and physical worlds. This will enable applications that are currently unfeasible, such as pollution monitoring at altitude (see Links 4 and 5) and autonomous last-mile deliveries [3], while reducing costs compared with current practices.

Figure 1: Aerial drone

SICS researchers are currently investigating, for example, how to equip aerial drones with service-oriented interfaces, allowing existing business process and web mash-ups to transparently integrate the functionality provided by aerial drones. These devices are thus perceived as "mobile" web services providing sensory data that can be later input to standard online services at company back-ends or in the cloud. The ability to seamlessly blend aerial drones with existing IT practice will facilitate their large-scale adoption within industry and the public sector.

To complement the efforts above, SICS is researching ways to facilitate the programming of fleets of drones, unveiling the potential available when multiple such devices can cooperatively carry out a higher-level mission. This is achieved by means of a custom programming abstraction that gives programmers the illusion that the drone fleet can be programmed as a single device [1]. Such design hides the intricacies stemming from explicitly managing the communication and coordination among drones, greatly easing their programming.

In addition to designing and implementing the solutions above, SICS is also creating actual prototypes using commercially available drone technology. Based on a customized version of the widespread AR.Drone 2.0, SICS is applying the results of these efforts in multiple real deployments. For instance, in Aquileia (Italy) the programming systems above are helping archeologists at the “Domus dei putti danzanti” obtain aerial maps of the site. The 32,000 m2 area, partly shown in the picture, hosts the ruins of an ancient Roman house dating back to the fourth century BC. A video of one of the custom AR.Drone 2.0 in action in Aquileia can be viewed at youtube.com/watch?v=PPDGO-jc0Is.

Another area in which SICS is applying drone technology is in localization of wireless sensor networks (WSNs). Drones can be used as mobile anchor-points to localize stationary sensor nodes that are deployed over a given area. SICS has performed experiments by connecting a small IEEE 802.15.4 radio chip to a commercially available low-cost drone. The radio was used to perform range and signal-strength measurements to a set of stationary sensor nodes. After fusing the measurements, a median position accuracy of approximately 1.5 m was measured. This is considered appropriate for a wide range of WSN applications.

Links:
1. goo.gl/SPOIR
2. goo.gl/n292Bw
3. goo.gl/1qP8Jx
4. goo.gl/stvh8
5. goo.gl/3bjByr

Reference:
[1] L. Mottola et al.: “LiftOff: Spatio-temporal Programming of Multiple Coordinating Aerial Drones”, Technical report 2013.46 – Politecnico di Milano, Italy.

Please contact:
Luca Mottola
SICS, Sweden, and Politecnico di Milano, Italy
E-mail: luca@sics.se

MediaEval 2013 Evaluation Campaign

by Gareth J. F. Jones and Martha Larson

MediaEval is an international multimedia benchmarking initiative offering innovative new tasks to the multimedia research community. MediaEval 2013 featured tasks incorporating video analysis, speech search, music recommendation, analysis of affect and location placing of images.

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 image, video, speech, audio, tags, users, and context. 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 2013 was the fourth evaluation campaign in its current form, which follows on from the VideoCLEF track at CLEF 2008 and CLEF 2009.

MediaEval 2013 offered 6 Main Tasks and 6 Brave New Tasks coordinated in cooperation with various research groups in Europe and elsewhere. The following Main Tasks were offered in the 2013 season:

- Placing Task: Geo-coordinate Prediction for Social Multimedia. The Placing Task required participants to estimate the geographical coordinates (latitude and longitude) of images, as well as to predict how “placeable” a media item actually is. The Placing Task integrated all aspects of multimedia: textual meta-data, audio, image, video, location, time, users and context.

- Search and Hyperlinking of Television Content Task: This task required participants to find relevant video segments to an information need and to provide a list of useful hyperlinks for each of these segments. The task focused on television data provided by the BBC and real information needs created by volunteer home users.

- Spoken Web Search: Spoken Term Detection for Low Resource Languages The task involved searching for audio content within audio content using an audio content query. This task is particularly interesting for speech researchers in the area of spoken term detection or low-resource speech processing.

- Violent Scenes Detection in Film (Affect Task) This task required participants to automatically detect portions of movies depicting violence. P Violence is defined as "physical violence or accident resulting in human injury or pain". Participants were encouraged to deploy multimodal approaches (audio, visual, text) to solve the task.

MediaEval 2013 participants. Photo: John Brown.

- Social Event Detection Task: This task required participants to discover social events and organize the related media items in event-specific clusters. Social events of interest were planned by people, attended by people and the social media captured by people.

- Visual Privacy: Preserving Privacy in Surveillance Videos For this task, participants needed to propose methods for protection of privacy sensitive elements (i.e., obscuring identifying elements on people) in videos so that they were rendered unrecognizable in a manner that is suitable for computer vision tools and is perceived as appropriate to human viewers of the footage.

The MediaEval 2013 Brave New Tasks are incubators of potential Main Tasks for future years, and typically have smaller numbers of participants. The MediaEval 2013 Brave New Tasks were: Searching Similar Segments of Social Speech, Retrieving Diverse Social Images, Characterizing Emotion in Music, Crowdsourcing for Social Multimedia, Question Answering for the Spoken Web, and Soundtrack Selection for Commercials (MusiClef Task).

The MeviaEval 2013 campaign culminated in the annual workshop which was held at Reial Acadèmia de Bones Lletres in Barcelona, Spain from 18th-19th 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 2013 with a total of 95 papers appearing in the Working Notes, and 100 participants attending the workshop. In addition to organizer and participant presentations, the workshop featured invited presentations by Jana Eggink, BBC Research and Development, London and Frank Hopfgartner, Technische Universität Berlin, Germany. The Working Notes proceedings from the MediaEval 2013 workshop have again been published by CEUR workshop proceedings.

MediaEval 2013 received support from a number of EU and national projects and other organizations including: AXES, Chorus+, CUbRIK, COMMIT/, Promise, Social Sensor, Media Mixer, MUCKE, Quaero, FWF, CNGL, Technicolor and CMU. We are also happy to acknowledge the support of ACM SIGIR for providing a grant to support international student travel.

The MediaEval 2014 campaign is currently getting underway with task resgistration opening in March 2014. The tasks will run through the spring and summer, with participants invited to present results of their work at the MediaEval 2014 Workshop which will again be held in Barcelona, Spain in October.

Further details of the MediaEval campaigns are available from the MediaEval website.

Links:
MediaEval: http://www.multimediaeval.org
MediaEval 2013 online proceedings: http://ceur-ws.org/Vol-1043/

Please contact:
Martha Larson
TU Delft, The Netherlands
E-mail: M.A.Larson@tudelft.nl

Call for Participation

IEEE Cluster 2014

Madrid, 22-26 September 2014

Clusters have become the workhorse for computational science and engineering research, powering innovation and discovery that advance science and society. They are the base for building today's rapidly evolving cloud and HPC infrastructures, and are used to solve some of the most complex problems. In September, Madrid will be the meeting point of all these efforts thanks to the IEEE Cluster Conference.

The IEEE Cluster Conferences are one of the greatest initiatives of TFCC and serve as a major international forum for presenting and sharing recent accomplishments and technological developments in the field of cluster computing, as well as the use of cluster systems for scientific and commercial applications.

This year, the IEEE Cluster Conference will be held in Madrid from 22-26 September at the School of Industrial Engineering, Universidad Politécnica de Madrid. The organizing committee welcomes paper submissions on innovative work that describe original research and development efforts in the area of cluster computing.

Major topics of interest:
- system design, configuration and administrations
- system software, middleware and tools
- storage and file systems
- algorithms, applications and performance.

As the evolution of cluster technologies is expected to substantially impact emerging research areas like Data Science, the conference will offer specific tracks, tutorials and workshops on processing and storage within this key area. And for the first time, a "Student Mentoring Track" is planned.

Important dates:
- 24 April: Abstracts (required) due
- 2 May: Full papers due
- 19 May: Tutorial proposals due

More information:
http://www.cluster2014.org/

Call for Participation

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

Florence, Italy, 11-12 September 2014

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

Topics of interest include:
- design, specification, code generation and testing based on formal methods
- methods, techniques and tools to support automated analysis, certification, debugging, learning, optimization and transformation of complex, distributed, real-time systems and embedded systems
- verification and validation methods that address shortcomings of existing methods with respect to their industrial applicability (e.g., scalability and usability issues)
- tools for the development of formal design descriptions
- case studies and experience reports on industrial applications of formal methods, focusing on lessons learned or identification of new research directions
- impact of the adoption of formal methods on the development process and associated costs
- application of formal methods in standardization and industrial forums.

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

More information:
http://fmics2014.unifi.it/

Call for Presentations

NAFEMS European Conference: Multiphysics Simulation 2014

Manchester, U.K., 21-22 October 2014

The NAFEMS conference will provide an overview of state-of-the-art-methods for coupled and multiphysics simulations, mainly within the context of industrial applications and CAE.

The need for more realistic and accurate numerical simulations, coupled with the increasing availability of high-performance computing, continues to drive the demand for multiphysics simulations. Keeping up with this development is challenging both for the software vendors, who have to implement increasingly complex tools for evermore diverse applications, and for the users of the software, who are being asked to work with new physical models outside their main field of competence. It also poses a challenge to universities to include more interdisciplinary knowledge in the curriculum.

The conference brings together researchers, developers, teachers, and users of multiphysics simulation methods to present new results, exchange ideas and discuss the challenges. It is an excellent opportunity to meet other practitioners in the field of multiphysics and coupled simulations.

Topics include:
- industrial applications of multiphysics
- collaborations of industry with academia
- multiphysics and HPC
- optimization in multiphysics
- teaching multiphysics
- validation of multiphysics simulations

The conference is organized jointly by NAFEMS Multiphysics Working Group, MULTiPHYSICS, the International Society of Multiphysics and Fraunhofer SCAI.

More information:
http://www.nafems.org/events/nafems/2014/mp2014/

Call for Participation

W3C Workshop on the Web of Things

Berlin 25–26 June 2014

The Internet of Things is a hot topic in technical circles with many potential domains of applications, for instance, home automation, security, healthcare, manufacturing, smart cities and many more. The market potential is currently held back by fragmentation in the communication technologies and a lack of shared approaches to services.

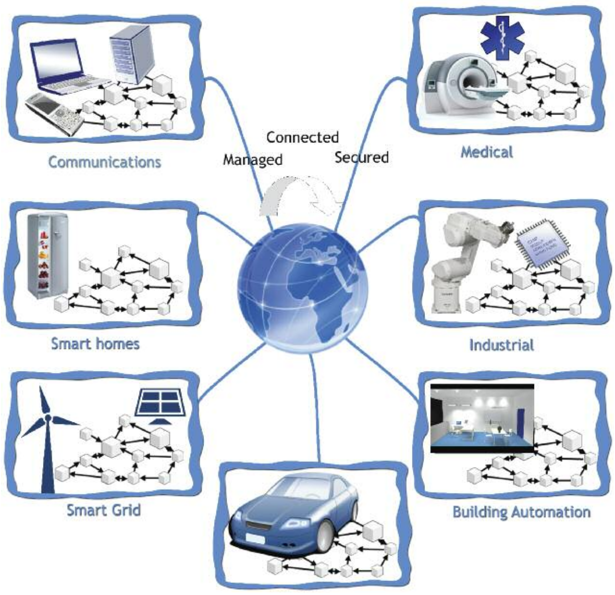

The Web of Things workshop topics.

W3C is holding a workshop in Berlin on 25-26 June, hosted by Siemens. We are hoping to encourage discussion on a wide range of topics, and to examine the potential for open standards as a basis for services, either between devices, at the network edge, e.g. in home hubs, or in the cloud. The combination of the Web of devices and the Web of data is expected to enable open markets of services with sweeping economic and societal impact.

The workshop is funded by the EU Compose (Collaborative Open Market to Place Objects at your Service) project. W3C/ERCIM is a partner of the project.

The workshop is free, but participants need to submit a brief expression of interest or a position paper.

More information:
http://www.w3.org/2014/02/wot/

CWI joins Sino-European Research Network LIAMA

Centrum Wiskunde & Informatica (CWI) has officially joined research network LIAMA. The Sino-European Laboratory in Computer Science, Automation and Applied Mathematics (LIAMA) is a research lab consisting of European and Chinese research institutes in the field of mathematics and computer science. LIAMA conducts research, training and transfer projects in these fields. CWI joined the lab as one of the founding members.

CWI will join two projects within LIAMA: CRYPT, aimed at studying secret-key and public-key cryptanalysis, and TEMPO, focusing on the development of safer and more reliable embedded systems.

Jos Baeten (general director CWI): “Joining LIAMA gives our researchers the opportunity to expand their international cooperation in Asia and work with fellow top researchers at Chinese universities. I believe that our membership will be mutually beneficial and I am looking forward to our future cooperation.”

LIAMA was founded in 1997 by Inria and the Chinese Academy of Sciences (CAS). In 2008 the lab was opened for other European research institutes. Another founding member is CNRS. Also in 2008, LIAMA became a National Centre for International Research of the Chinese Ministry of Science and Technology. It currently consists of 11 Chinese and 12 European research institutes.

More information:
LIAMA website: http://liama.ia.ac.cn/

MonetDB Solutions Newest CWI Spin-off

CWI’s 22nd spin-off company was recently launched: MonetDB Solutions. This consulting company is specialized in the development of database technologies for business analytics. It is established by the founders of MonetDB, the world's first and most widespread open-source column-store database management system, and led by Ying Zhang. Bundling the expertises of world-leading researchers and engineers, its core services involve technical consultancy, architectural review, performance assessment, migration and upgrade services and training.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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