



HAL
open science

The Internet of Things and The Web of Things

Emmanuel Baccelli, Dave Raggett

► **To cite this version:**

Emmanuel Baccelli, Dave Raggett (Dir.). The Internet of Things and The Web of Things. ERCIM, 101, 2015, ERCIM News Special Issue on The Internet of Things and The Web of Things. hal-01244735

HAL Id: hal-01244735

<https://inria.hal.science/hal-01244735>

Submitted on 16 Dec 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ERCIM



NEWS

www.ercim.eu

Special theme:

The Internet of Things and The Web of Things

Also in this issue:

Keynote:

by Cees Links, Founder & CEO
GreenPeak Technologies

Research and Innovation:

Mesh Joinery:
A Method for Building Fabricable
Structures

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 6,000 printed copies and is also available online.

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

Editorial Board:

Central editor:

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

Local Editors:

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

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

Cyprus: Ioannis Krikidis (krikidis.ioannis@ucy.ac.cy)

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

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

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

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

Artemios Voyiatzis (bogart@isi.gr)

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

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

Luxembourg: Thomas Tamsier (thomas.tamsier@list.lu)

Norway: Poul Heegaard (poul.heegaard@item.ntnu.no)

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

Portugal: Joaquim Jorge (jorgej@tecnico.ulisboa.pt)

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

Sweden: Kersti Hedman (kersti@sics.se)

Switzerland: Harry Rudin (hrudin@smile.ch)

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

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

Contributions

Contributions should be submitted to the local editor of your country

Copyright notice

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

Advertising

For current advertising rates and conditions, see

<http://ercim-news.ercim.eu/> or contact peter.kunz@ercim.eu

ERCIM News online edition

The online edition is published at

<http://ercim-news.ercim.eu/>

Subscription

Subscribe to ERCIM News by sending 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 2015, Special theme: Trustworthy Systems of Systems

The Internet of Things will Change our World

Everybody talks about the Internet of Things, the IoT... but how is the IoT actually going to change our lives?

The Internet of Things is creating a new world, a quantifiable and measurable world, where people and businesses can manage their assets in better informed ways, and can make more timely and better informed decisions about what they want or need to do. This new connected world brings with it fundamental changes to society and to consumers. By sensing our surrounding environment, the IoT will create many practical improvements in our world, increasing our convenience, health and safety, while at the same time improving energy efficiency and comfort. The IoT will be a new source of wealth creation.

IoT devices can be classified in three categories: (1) wearables, (2) smart home devices, and (3) M2M devices. The first two categories are the most important for consumers. 'Wearables' are the devices that people carry with them, which usually connect via Bluetooth to a smart phone, and from there to the Internet. This category includes devices such as smart watches, fitness bands and devices to help people to live more 'mindfully' – increasing the wearer's awareness of how well they sleep, how much they move around, monitoring their vital signs, etc.

Smart home devices are also part of the IoT and usually connect to the Internet via ZigBee low power wireless communication and the home router. These include all domestic devices, from lights and light switches to motion sensors, thermostats, door locks and automated curtains. Via its WiFi connection to the router, the smart phone also becomes an online dashboard and control device for Smart Home applications.

The third category, M2M (Machine to Machine) devices, comprises devices that are directly connected to the cellular network, such as cars that can report their location (in case of an accident or theft), or vending machines that can call in when their stocks are running low.

Many households and businesses have thermostats, weather stations, smart lighting, security and electronic door

locks, the majority of which are not currently interconnected. They are connected – but not to each other. The weather station does not provide information to the thermostat about the climate outside. The security system is not connected to the indoor motion sensors, nor to the electronic door locks (it does not automatically lock the forgotten back door when the inhabitants go out). In the future, all these systems will be interconnected, providing information to each other, and reacting accordingly. We are currently in an emerging state of the IoT, with individual vertical applications that operate as islands, and serve independent applications (such as security alarms, door locking, etc.). However, the real IoT will emerge when these applications cooperate, working together, and begin to use each other's 'awareness'. That is when the true IoT avalanche will start.

The key component of the IoT - whether wearables or smart home devices - is not the sensor, but the application. Connecting the sensors is difficult, but extracting information from data is the essence. Useful information extracted from the data can coach people by reaffirming when things go as planned or by alerting or taking action if something goes wrong; and data analytics can be used to compare situations, to coach and to provide feedback to help make improvements. This is slowly starting to dawn on manufacturers and service providers alike. People are interested in the IoT if it helps them to improve aspects of their lives. Improvements are not achieved by sensors alone: a completely different way of thinking is required, and it will take some time for the new paradigm to be fully embraced.

Privacy and security are key, together with data ownership. Note, these are not IoT issues, but general internet issues that are amplified by the growth of new applications. These issues already exist for the internet of people, and industry and government bodies are slowly starting to recognize them and take action.

The growth of the IoT can be compared with the growth of the automobile



Cees Links, Founder & CEO GreenPeak Technologies, www.greenpeak.com

industry. Picture the first cars hitting the road: there were no freeways, no road signs, no rules, no driving licences. Pedestrians did not know to get out of the way. Drivers did not know how to take turns at intersections. Neither drivers nor pedestrians understood the risks and liabilities, giving no consideration to liability and insurance.

We are currently at a similar stage with the IoT. Just as it took decades before all required infrastructure was in place around motor vehicles, it will take quite some time before it is in place around the Internet. Once I led the engineering team that successfully brought WiFi to the Mac laptop for Steve Jobs. After his team had achieved this, all others followed. Now WiFi is everywhere - but this took time and work. Similarly, in the context of the Internet and Internet of Things, there is a growing awareness that we need rules, training, legislation and enforcement. We are just starting to learn what might be needed.

The IoT will change the world in an even more profound way than has the Internet. If we ask our children today how the world existed before Internet, they are speechless. They have no comprehension of how people could communicate or even live their lives without the common place tools we have today. The same will happen with the IoT.

A decade from now, we will be dependent on the knowledge derived from the continuous stream of data from our wearables and our smart home devices, and we will have no idea how we managed the world and our lives before. We will be able to make better informed, more accurate and more timely decisions; and decisions that will improve our lives, save us money, and may even save our planet. The IoT will make the difference.

KEYNOTE

- 3 The Internet of Things Will Change our World**
by Cees Links, Founder & CEO GreenPeak Technologies

JOINT ERCIM ACTIONS

- 6 ERCIM “Alain Bensoussan” Fellowship Programme**
- 6 Cor Baayen Award 2015**
- 6 W3C Europe@20 Anniversary Event**
- 7 The Role of ERCIM EEIG**
by Antoine Petit

SPECIAL THEME

The special theme section “The Internet of Things and the Web of Things” has been coordinated by Dave Raggett, W3C and Emmanuel Baccelli, Inria.

Introduction to the Special Theme

- 8 The Promise of the Internet of Things and the Web of Things**
by Emmanuel Baccelli and Dave Raggett

IoT Operating Systems & Network Stacks

- 11 Open{WSN|Mote}: Open-Source Industrial IoT**
by Thomas Watteyne, Xavier Vilajosana, and Pere Tuset

- 12 RIOT and the Evolution of IoT Operating Systems and Applications**
by Emmanuel Baccelli, Oliver Hahm, Hauke Petersen and Kaspar Schleiser

- 14 FIT IoT-LAB: The Largest IoT Open Experimental Testbed**
by Eric Fleury, Nathalie Mitton, Thomas Noël and Cedric Adjih

- 15 OCARI: A Wireless Sensor Network for Industrial Environments**
by Tuan Dang, Pascale Minet and Erwan Livolant

IoT Security

- 17 Security and Privacy in the Internet of Things**
by Ludwig Seitz

- 18 Secure and Efficient Wireless Sensor Networks**
by Corinna Schmitt and Burkhard Stiller

Languages & Models

- 19 Creating Internet of Things Applications from Building Blocks**
by Frank Alexander Kraemer and Peter Herrmann

- 21 Interfaces to the Internet of Things with XForms**
by Steven Pemberton

- 22 Home Automation Devices Belong to the IoT World**
by Vittorio Miori and Dario Russo

Abstraction Layers & Programming Tools for IoT

- 23 Diopbase: Data Streaming Middleware for the Internet of Things**
by Benjamin Billet and Valérie Issarny

- 25 Programming Actors for the Internet of Things**
by Raphael Hiesgen, Dominik Charousset, Thomas C. Schmidt and Matthias Wählisch

RESEARCH AND INNOVATION

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

26 A Mobile End-User Development Environment for IoT Applications Exploiting the Puzzle Metaphor

by José Danado and Fabio Paternò

28 3D Web Visualization for Real-Time Maintenance of Smart Buildings

by Daniel Gaston, Christophe Joubert and Miguel Montesinos

30 COMPOSE: An Open Source Cloud-Based Scalable IoT Services Platform

by Dave Raggett

IoT Architecture

31 An Integration Gateway for Sensing Devices in Smart Environments

by Michele Girolami, Francesco Furfari and Stefano Chessa

33 Fighting Networking Heterogeneity in the Internet of Things

by Elias Z. Tragos, Vangelis Angelakis and Stefanos Papadakis

34 Tectons: Towards a Generalised Approach to Programming Systems of Systems

by Geoff Coulson, Andreas Mauthe and Markus Tauber

36 BETaaS: Building the Future Platform for Development and Execution of Machine-to-Machine Applications

by Carlo Vallati, Enzo Mingozzi and Giacomo Tanganelli

IoT Use Cases

37 Pro-IoT: Improving the Productivity of Companies by Using the IoT

by Heikki Ailisto, Nadine Pesonen, Pertti Raatikainen and Jonathan Ouoba

39 MusicBricks: Connecting Digital Creators to the Internet of Music Things

by Thomas Lidy, Alexander Schindler and Michela Magas

41 Deploying an NFC-based Inventory System: A Case Study

by Patrick Poullie, Thomas Bocek and Burkhard Stiller

42 TrakChain Estimates Costs for Track and Trace in the Internet of (many) Things

by Miguel L. Pardal and José Alves Marques

44 Mesh Joinery: A Method for Building Fabricable Structures

by Paolo Cignoni, Nico Pietroni, Luigi Malomo and Roberto Scopigno

46 Icing Detection and Protection for Small Unmanned Aircrafts

by Andrea Cristofaro and Tor Arne Johansen

47 LEARNMINER – Advanced Analytics for eLearning

by Christian Bauckhage, Markus Ketterl and Ulrich Nütten

48 Online Semantic Analysis over Big Network Data

by Chengchen Hu and Yuming Jiang

49 From Cooling Integrated Circuits to Efficient Electrical Power Generation

by Stephan Paredes, Brian Burg, Patrick Ruch, Ingmar Meijer and Bruno Michel

51 VAVID– Handling Big Data

EVENTS, IN BRIEF

Announcements

52 AITA 2015 - Advanced Infrared Technology and Applications

52 IC3K - 7th Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management

52 Information Security Conference - ISC 2015

53 ICEC 2015 - 14th International Conference on Entertainment

53 7th International Workshop on Software Engineering for Resilient Systems

54 PhD Positions Available in the BigStorage European Training Network

54 1955-2015: ITC Celebrates its 60th Anniversary!

54 SERENE 2015 - 7th International Workshop on Software Engineering for Resilient Systems

In Brief

55 25 Years of Python at CWI

55 MOOC on “Binaural Hearing for Robots”

55 Java Bug Fixed with Formal Methods

55 W3C Launches First HTML5 Course in New Partnership with edX



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. Fellowships are proposed according to the needs of the member institutes and the available funding.

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.

In order to encourage mobility:

- a member institute will not be eligible to host a candidate of the same nationality.
- a candidate cannot be hosted by a member institute, if by the start of the fellowship, he or she has already been working for this institute (including phd or post-doc studies) for a total of 6 months or more, during the last 3 years.

The fellows are appointed for 12 months either by a stipend (an agreement for a research training programme) or a working contract. The type of contract and the monthly allowance (for stipends) or salary (for working contracts) depend on the hosting institute.

Application deadlines

30 April and 30 September

More information and application form:

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

W3C Europe@20 Anniversary Event

W3C Europe will celebrate its twentieth anniversary with a symposium on Tuesday, 5 May in the Paris town hall.

The Web was born in Europe and W3C is celebrating the 20th anniversary of the European branch of W3C in Paris. Europe has played a fundamental role for the success of the Web. First of all, the founding of the World Wide Web Consortium (W3C) in 1994 became possible with the support by CERN, DARPA and the European Commission. Since 1995, when Inria hosted the European branch of W3C, the European Commission funded W3C activities in Europe, as for example the development of Web accessibility rules, enabling people with disabilities to participate equally on the Web. Since 2003, W3C Europe is hosted by ERCIM EEIG.

W3C will celebrate 20 years of work accomplished by European stakeholders for the benefit of the Web with a symposium on 5 May. Speakers will include, among others, representatives from Paris city council, the European Commission and the French government. The closing talk will be given by Tim Berners-Lee, Web Inventor and W3C Director. About 250 invited guests - global strategists, business leaders and developers - are expected to attend.

More information:

<http://www.w3.org/20/Europe/>

Call for Nominations

Cor Baayen Award 2015

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>

The Role of ERCIM EEIG

by Antoine Petit

Over the last 25 years, ERCIM has earned a reputation as one of Europe's leading expert groups in the areas of ICT and Applied Mathematics. In 2012, the ERCIM community was restructured and is now supported by two organizations, the ERCIM Association (AISBL) and the ERCIM European Economic Interest Group (EEIG). While the Association is carrying out and supervising all of ERCIM's scientific activities, the EEIG, composed of a subset of the Association's members, is responsible for managing the ERCIM Office and the European branch of the World Wide Web Consortium (W3C).

The ERCIM Office, hosted by Inria in its Sophia Antipolis Research Centre in Southern France, manages all administrative and financial aspects of the ERCIM community with a talented team of professionals.

Hosting and managing the European branch of W3C is a major task of the ERCIM EEIG. The change of the W3C European host from Inria to ERCIM in 2003 aimed at strengthening research relationships throughout Europe to better support Web technology development. ERCIM was a natural candidate as host since many of the European W3C Offices were already based at ERCIM member institutes. When Inria became the first W3C European host in 1995, W3C Europe had 50 members in four countries. Today, the W3C Membership in Europe exceeds 140 organizations, with representation in 20 countries in Europe. But more important, during the last twenty years we have witnessed the Web revolution with tremendous advances in Web technology, changing not only the world of research but also the whole of society. The evolution of Web standards is a challenging task, yet is crucial to maintain an open Web and to transform W3C into a more long-term international standardization organization. This year, W3C Europe will celebrate its twentieth anniversary at an event on Tuesday, 5 May in Paris, in conjunction with the ERCIM Spring meetings, and ERCIM representatives are cordially invited.

The ERCIM Office is also providing assistance and services to the whole ERCIM community. An important activity is the management of European projects involving ERCIM members or the W3C. Through ERCIM, our members have participated in more than 80 projects funded by the European Commission by carrying out joint research activities while the Office manages the complexity of the project administration, finances and outreach. The ERCIM Office also handles ERCIM AISBL financial matters and supports the whole ERCIM community in administrative tasks such as the management of the well established ERCIM Fellowship Programme, and in communications, for example by hosting and maintaining web sites and by producing ERCIM News.



© Inria / Photo C. Hebbly

Antoine Petit, Chairman and CEO of Inria, was nominated President of ERCIM EEIG by the ERCIM EEIG Board of Directors on 24 October 2014 in Pisa for a two year term from January 2015 to December 2016. Antoine Petit succeeds Michel Cosnard, who took office as ERCIM President in January 2011.

ERCIM has gained a leadership position for cooperating for excellence in research and is a recognized expert organization in Information and Communication Technologies in Europe. To maintain its position, and increase its impact on European strategy and policy, I'm convinced that ERCIM has to increase its membership by attracting the best research performing organizations and by developing new strategically important activities such as the "ERCIM white papers series" started last year with support of Domenico Laforenza, my fellow president of ERCIM AISBL. ERCIM groups of experts identify emerging grand challenges and strategic research topics in Information and Communication Science and Technology (ICST). In 2014, the first two white papers were published on the topics "Big Data Analytics" and "Cyber-Security and Privacy" with recommendations for future research. A white paper on Massively Open Online Courses (MOOCs), another current hot topic, is currently in work.

I also support Domenico's efforts in cooperating with leading European organisations, in order to play a role in shaping Europe's digital future through the platform of the European Societies in Information and Communication Sciences and Technologies (EFICST). I encourage ERCIM to strengthen relationships with organisations such as Informatics Europe and the ICT Lab, a Knowledge and Innovation Communities (KICs) of the European Institute for Innovation and Technology (EIT).

With our two organisations, ERCIM EEIG and ERCIM AISBL, I believe that we have the potential to strengthen our position as recognized leader in field of European research and innovation in ICST, to increase our impact on European strategy and policy, and to successfully continue our mission of "Cooperating for Excellence in Research".

Introduction to the Special Theme

The Promise of the Internet of Things and the Web of Things

by Emmanuel Baccelli and Dave Raggett

Improvements in electronics are enabling widespread deployment of connected sensors and actuators (the Internet of Things) with predictions of tens of billions of smart objects by 2020 (Cisco 2011). This raises huge challenges for security, privacy and data handling, along with great opportunities across many application domains, e.g., home and building automation, retail, healthcare, electrical grids, transport, logistics, manufacturing, and environmental monitoring (IERC 2014).

The Internet of Things started with work on radio frequency identity tags (RFID) and expanded to connected sensors and actuators, along with many communication technologies designed for different purposes (e.g., ZigBee, NFC, Bluetooth, ANT, DASH7, EnOcean). IPv6 has greatly expanded the address space compared to IPv4 and makes it feasible to give each device its own IP address. Considerable effort has been made on supporting IP all the way to constrained devices, e.g., the 6LoWPAN and RPL protocols for wireless (mesh) access, CoAP for lightweight message exchange and easy bridges to HTTP, and MQTT as a lightweight pub-sub protocol.

Today, however, product silos – a sign of market immaturity - are prevalent for the Internet of Things. This is where open standards and open platforms can play a major role in enabling the growth of rich open ecosystems that can realise the huge potential benefits for European countries. Can we repeat the runaway success of the World Wide Web and build a Web of Things? Turning that around, what can we learn from the Web in relation to encouraging an open market of services?

Much of the work on the Internet of Things (IoT) has focused on the technologies needed for constrained devices, long battery life, efficient use of wireless spectrum and so forth. Relatively little work has focused on applications and services.

One challenge is to design appropriate abstractions that hide details best left to platform developers. The first step in this direction is the use of the IPv6 / 6LoWPAN network stack on top of heterogeneous IoT hardware and link layers.

Another challenge is to provide end-to-end security. IoT devices are often very constrained, and this limits the options for handling security. A related problem is software upgradability: software upgrades are essential for addressing security flaws, as well as for updating devices to match changes in standards. However, this is a challenge on the most constrained IoT devices, which, to date, are not field-upgradable. This necessitates the use of better operating systems on most IoT devices, and/or gateways that can proxy and manage for other IoT devices that cannot afford a sophisticated OS.

With the increasing number of sensors, we all need to be confident that our privacy is safeguarded. This implies end-to-end encryption against eavesdroppers, strong mutual authentication and support for access control and data handling according to the data owner's policies. With the ability to combine data from different sources, it becomes necessary to track provenance so that the originating data owner's policies can be applied to derived data. This in turn motivates work on techniques for static analysis of service logic and dynamic enforcement of policies.

Given the anticipated very large numbers of sensors and actuators, it will be inevitable that some will fail, either through hardware faults, electrical noise or even botched upgrades. Services need to be designed to be resilient in the face of such failures. This will need to happen at multiple levels of abstraction. Resilience is also important for handling rapid changes in demand without overloading the platforms on which the services are running. Resilience is also key to handling cyber attacks. One approach to counter this is in depth defence with successive security zones and automatic trip wires for detecting intrusion and raising the alarm. Continuous monitoring can be combined with machine learning techniques for spotting signs of unusual behaviour.

Raw data often has limited value, only becoming valuable when it has been processed through multiple levels of interpretation that combines multiple sources of information, and provides results that are pertinent to a given context. This is where we can learn from nature by examining and mimicking the progressive processes involved in animal perception and recognition. The same is true for actuation, where high level intents can be progressively transformed into lower level control over different subsystems. What is needed to coordinate and synchronise distributed systems? As human beings, when we speak, our brains are able to coordinate the movements of many components, each of which have widely varying response times. The jaw bone is massive and needs to be set in motion well before the tip of our tongue, which can move much more quickly.

The Web places a strong emphasis on scripting, and the same will apply for the Web of Things. Scripts (e.g., Node.js using Javascript) could be used to define service logic for scalable cloud based platforms, for small scale platforms, e.g., as a

new breed of home hubs, and for device gateways that bridge the IoT and the Web.. However, further work is needed to determine which APIs are needed to support a broad range of use cases. For instance, the scalable cloud-based COMPOSE platform addresses use cases involving event streams. New work is needed to support continuously changing properties as opposed to discrete events, and to address the rich requirements for complex cyber-physical systems. These are likely to involve different requirements at different levels of abstraction, e.g., tight requirements on timing at a low level, and perhaps transactional robustness at a high level.

To enable open markets of services, we need a standard way to access the service descriptions, so that a search engine can cover services hosted in different clouds operated by different vendors. We then need a standard framework for representing descriptions along with standard vocabularies/ontologies. This needs to cover the purpose of a service, the interfaces it exposes or depends upon, its security and privacy related properties, and so forth. Interoperability depends upon having compatible semantics and data representations. What is needed to motivate the re-use of vocabularies? When existing vocabularies aren't a good fit to particular needs, what is needed to encourage the registration of a new vocabulary along with the assumptions that motivate it?

Where services have mismatching semantics or data formats, there is an opportunity for intermediaries to bridge the gaps. Search services can formulate plans for assembling services to fulfil the designated requirements. Such assemblies can be static or dynamic, e.g., all cars in this section of a road. Plans can also be applied to managing tasks, e.g., sequential tasks, concurrent tasks, and hierarchical arrangements of tasks, where tasks are associated with preconditions and postconditions. For the Web of Things, this corresponds to finding services that can perform these tasks, e.g., transforming data, identifying events that are implied by a combination of lower level events, or carrying out actions with a combination of actuators.

The 'things' in the Web of Things are virtual objects. They can represent real world objects such as sensors and actuators, people and locations, or even abstract ideas like periods of time (e.g., the 70's) or events (e.g., a football match, concert or play). The 'web' in the Web of Things refers to the idea that things' are accessible via Web technologies, e.g., HTTP at the protocol layer, or scripting APIs at the services layer. Where 'things' represent real-world objects such as people, the things can be considered as avatars that know about the person they represent, and can perform actions to achieve the goals of that person. This is related to the concept of the Social Web of Things in which things have relationships to you, to your 'friends' and to other things. The social network can provide a basis for routing notifications and for sharing services.

Google's 'Physical Web' (Google, 2014) is about beacons that broadcast web addresses to devices in their neighbourhood. This can be compared to walking down a crowded market street with all the store holders shouting out their wares and special offers. This calls for personal agents or avatars that are aware of your current interests and are able to recognise which beacons are relevant and which can be safely ignored. The agent could notify you directly or could perform tasks on your behalf.

Avatars are also related to the concept of personal zones as explored in the EU FP7 webinos project. Your personal zone is an abstraction that groups all of your personal devices and services. It provides an overlay model for secure messaging between your devices as well as offering zone wide services to trusted applications running on your devices. Your personal zone also acts as an avatar on your behalf offering services to your friends based upon the access control policies you set. For the Web of Things, personal zones offers a unified means for people to manage the data they own.

Another consideration is the lifecycle of services, their provisioning, and packaging as products for sale to consumers. As an example, consider a security camera purchased for use at home. The camera may be bundled with the software and services, or this could be purchased separately from a Web of Things marketplace. Either way, the user needs a really simple approach to installing the hardware and setting up the associated services. How is the camera 'discovered' and 'paired' with a service? How does the user provide additional metadata, e.g., giving it a name, describing its location, and setting its access control policy? The package could include an application that the user runs to set this up, and to install any required device drivers. The package could include other applications that enable the user to manage the device and service, and as a user interface for the service when it is up and running. In the context of a smart city, there may be a need to install and set up large numbers of sensors. This too should be as simple and painless as possible. The same should be true for managing software upgrades and for taking devices and services out of service as needed.

Suppliers and consumers of services need to reach agreements, and this can be formalised as contracts that cover payments, data handling policies and so forth. For open markets of services such contracts should be legally binding on the parties involved. Whilst data may be provided free, in other cases, some form of payment will be required, for instance, one off payments, per usage payments and subscription based payments. To enable open markets to operate across borders, there is a need for international standards around payments. Even if services are provided free of charge, they may be restricted to consumers in certain groups. Access control may be based on rules, e.g., as with the XACML XML access control language, or based upon the possession of tokens as with capability based systems.

Access control is related to identity management. Mutual authentication is needed to ensure that both suppliers and

consumers can be sure of the other party's identity. Identity verification is about linking identities to real world properties, e.g., the physical location of a sensor, or the full name and postal address of a human. There is a need for trusted agents that provide identity verification services. Trust is also important in relation to decisions about whether to use services: are they safe, do they come from a bone fide source, will they safeguard my privacy and so forth. This provides opportunities for agents that perform security and privacy audits of services. This can be complemented by crowd sourced reputations and reviews. Recommendation systems can further provide suggestions based upon what other people have looked at in similar contexts.

We have a lot to do. The papers in this special issue of ERCIM News provide a small sample of the research work that is helping to realise the potential for connected sensors and actuators, and the services that this enables.

Links:

IERC 2104: European Research Cluster on the Internet of Things, 2014: Internet of Things From Research and Innovation to Market Deployment:
http://www.internet-of-things-research.eu/pdf/IoT-From%20Research%20and%20Innovation%20to%20Market%20Deployment_IERC_Cluster_eBook_978-87-93102-95-8_P.pdf

COMPOSE: EU FP7 project 2012-2015,
<http://www.compose-project.eu>

Webinos: EU FP7 project 2010-2013,
<http://webinos.org>

Cisco, 2011: The Internet of Things — How the Next Evolution of the Internet Is Changing Everything:
http://www.cisco.com/web/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf

IoT2014: David Clark — Not making the same mistakes again:
<http://www.iot-conference.org/iot2014/keynote-speakers/>

Google, 2014: The Physical Web
<https://google.github.io/physical-web/>

Please contact:

Emmanuel Baccelli, Inria, France
 E-mail: Emmanuel.Baccelli@inria.fr

Dave Raggett, W3C
 E-mail: dsr@w3.org

Open{WSN|Mote}: Open-Source Industrial IoT

by Thomas Watteyne, Xavier Vilajosana, and Pere Tuset

The OpenWSN project is an open-source implementation of a fully standards-based protocol stack for capillary networks, rooted in the new IEEE802.15.4e Time Synchronized Channel Hopping standard. IEEE802.15.4e, coupled with Internet of Things standards, such as 6LoWPAN, 6TiSCH, RPL and CoAP, enables ultra-low-power and highly reliable mesh networks, which are fully integrated into the Internet. OpenMote is an open-hardware platform designed to facilitate the prototyping and technology adoption of IEEE802.15.4e TSCH networks, and is fully supported by OpenWSN.

Time Synchronized Channel Hopping (TSCH) was designed to allow IEEE802.15.4 devices (i.e., low power wireless mesh network devices) to support a wide range of applications including, but not limited to, industrial automation and process monitoring. It is based on a medium access technique which uses time synchronization to achieve ultra low-power operation and channel hopping to enable high reliability. Synchronization accuracy directly relates to power consumption, and can vary from microseconds to milliseconds, depending on the hardware and implementation.

The IEEE802.15.4e standard is the latest generation of ultra-low power reliable

networking solutions for Low-Power and Lossy Networks (LLNs). The memo RFC5673 (Industrial Routing Requirements in Low-Power and Lossy Networks) (Link 2) discusses industrial applications, and highlights the harsh operating conditions as well as the stringent reliability, availability, and security requirements for an LLN to operate in an industrial environment. In these scenarios, vast deployment areas with large (metallic) equipment cause multi-path fading and interference to thwart any attempt of a single-channel solution to be reliable; the channel agility of TSCH is the key to its ultra high reliability.

IEEE802.15.4e TSCH focuses on the MAC layer only. This clean layering

allows TSCH to fit under an IPv6-enabled protocol stack for LLNs, running IPv6 packet delivery in Low Power Wireless Personal Area Networks (6LoWPANs) (Link 3), IPv6 Routing Protocol for LLN (RPL) (Link 4) and the Constrained Application Protocol (CoAP) (Link 5). All of this complemented with the management interface and network operation specification currently being developed by the IETF 6TiSCH working group.

The OpenWSN project (Link 6) (see Figure 1) is an open-source implementation of the aforementioned protocol stack designed for capillary networks, rooted in the new IEEE802.15.4e TSCH standard, and providing IPv6 connectivity to ultra-reliable low power industrial mesh networks [1].

OpenWSN was founded in 2010 in the Berkeley Sensor & Actuator Center at UC Berkeley as an open-source project to implement, promote and contribute to the Industrial Internet of Things (IIoT), the next wave of innovation impacting the way the world connects. With a community of academic and industrial contributors, the OpenWSN protocol stack now runs on most popular low-power wireless platforms, and features a simulator and an ecosystem of cloud-based collaboration and continuous integration tools. OpenWSN has become the de-facto open-source implementation of IEEE802.15.4e-2012 Time Synchronized Channel Hopping (TSCH), the standard at the heart of the IIoT which enables ultra high reliability and low-power operation.

In 2013, work from the OpenWSN project contributed to the creation of IETF 6TiSCH (Link 7) Working Group, standardizing how to combine the ease-of-use of IPv6 with the performance of industrial low-power wireless technology. With over 290 members, IETF




Figure 1: OpenMote CC2538 main board. Composed by a ARM Cortex M3 SoC TI CC2538 Microcontroller. Open-Hardware design.




Figure 2: OpenWSN protocol stack diagram. At the MAC Layer features the new amendment of the IEEE 802.15.4 standard, namely the IEEE802.15.4e TSCH. On top the IETF standards suite provide IPv6 connectivity to low power mesh networked devices. OpenWSN is OpenSource and supports multiple commercial and open hardware platforms including OpenMote.

6TiSCH is spearheading the development of the Industrial IoT.

In 2014, members of the OpenWSN team founded OpenMote (Link 8), a startup company developing an ecosystem of hardware for the IoT, including the popular OpenMote (see Figure 2), the OpenBase/OpenUSB board to program it, and the OpenBattery to power it. Inheriting from previous ‘Berkeley Motes’, the OpenMote is widely considered the new generation of low-power wireless experimental platforms.

The OpenWSN ecosystem is devoted to accelerating the standardization and transfer of low power wireless technologies to the industry, enabling and promoting the adoption of low power wireless for industrial applications. OpenWSN is licensed under the BSD license; OpenMote Hardware is licensed under CERN OHL 1.2 hardware license. Both licences are neither

viral nor restrictive, enabling industrial users to take advantage of the platform without jeopardizing their developments or intellectual property. OpenWSN is in constant evolution, being updated with the latest standards in the field and becoming a central prototyping platform for future amendments and improvements to already existing standards and protocols.

Links:

1. <http://www.openmote.com/>
2. Industrial Routing Requirements in Low-Power and Lossy Networks: <https://tools.ietf.org/html/rfc5673>
3. Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks: <https://tools.ietf.org/html/rfc6282>
4. RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks: <https://tools.ietf.org/html/rfc6550>
5. The Constrained Application Protocol (CoAP): <https://tools.ietf.org/html/rfc7252>

6. OpenWSN Project: <http://openwsn.org/>
7. <http://tools.ietf.org/wg/6tisch/charters>
8. OpenMote Technologies <http://www.openmote.com/>

Reference:

- [1] T. Watteyne et al.: “OpenWSN: A Standards-Based Low-Power Wireless Development Environment”, Transactions on Emerging Telecommunications Technologies, vol. 23, issue 5, 480-493, 2012.

Please contact:

Thomas Watteyne
Inria, HiPERCOM2 team
E-mail: thomas.watteyne@inria.fr

Xavier Vilajosana
Universitat Oberta de Catalunya
E-mail: xvilajosana@uoc.edu

Pere Tuset
OpenMote Technologies
E-mail: peretuset@openmote.com

RIOT and the Evolution of IoT Operating Systems and Applications

by Emmanuel Baccelli, Oliver Hahm, Hauke Petersen and Kaspar Schleiser

The Internet of Things (IoT) is expected to be the next ‘big thing’. To date, however, there is no de facto standard software platform to program memory and energy-constrained IoT devices [1]. We expect an evolution of IoT software platforms that can be roughly compared to the recent evolution of software platforms powering smartphones.

Over the course of a few years, there has been an acceleration in the evolution of software platforms powering smart handheld devices - from a state where dozens of closed-source, slow-progress, rudimentary proprietary solutions were used, to a state where just two major players (iOS and Android) have imposed new de facto standards in terms of software platform interoperability, programmability and automatic updating.

The up side of this evolution is quicker-paced progress and the development of innumerable applications built on top of these platforms, which interact in innovative ways with one another, and with the cloud. It has become so easy to develop such applications (the infamous ‘Apps’) that hundreds of thousands of developers have been able to produce millions of Apps, which have been

installed/upgraded billions of times on smartphones and tablets. Arguably, the attribute of (re)programmability has had the most significant impact in this field, even more so than the evolution of handheld hardware.

On the dark side, companies pushing iOS and Android use such software platforms as tools to achieve questionable goals e.g., more or less covertly accumulating and monetizing personalized information. Furthermore, the true level of openness of iOS or Android is debatable, and their not being fully open facilitates the task of covertly accumulating personal information. The public is increasingly aware of this pitfall in terms of security and privacy, and it is therefore unlikely that people will accept IoT software platforms intrinsically tied to such monetization, or

lacking enough openness. In contrast, prominent examples have achieved high impact in the domain of computer networks, while remaining open, and therefore more trustworthy: for instance, community-driven efforts such as Linux, or the IETF [2], each fundamental to the emergence of the Internet as we know it today. While Linux demonstrated the power of open source software development and became the go-to operating system for Internet hosts, the IETF is an outstanding example of the positive impact of open standards, transparent standardization processes, and open discussion forums.

Until recently, IoT software platforms were in a comparable state to that of smartphone software platforms before Android and iOS. It is likely, however,




Figure 1: Typical device memory characteristics, and matching operating systems available on traditional Internet hosts vs IoT devices

that new standards will be set in the near future, in terms of IoT software platform openness, API, automated software updates and other characteristics that are necessary to enable a modern, large scale, secure ecosystem. Such an evolution will be game-changing for the IoT, and will fuel a new world of distributed applications developed by a large variety of actors, on top of the dominant open software platform(s). As yet it is unclear which IoT software platforms will emerge as dominant.

In practice, IoT software platforms face conflicting requirements: interoperability with the Internet, memory-constrained embedded programming, and portable, open-source code. Over the last decade, significant progress has been made in order to accommodate such requirements. Pioneer open source software platforms such as Contiki or TinyOS provided first attempts at accommodating these requirements, by exploiting non-standard coding paradigms, or by limiting code portability and the set of functionalities offered by the software platform (for instance, Arduino environment). Recently, more powerful but comparably low memory-footprint software platforms became available.

A prominent example is RIOT [3], an open source IoT operating system which enables programmers to develop applications on typical IoT devices,

with no learning curve (assuming prior experience with POSIX and Linux). On a wide variety of IoT hardware, RIOT enables the use of the standard coding languages (ANSI C and C++), well-known debugging tools (gdb, Valgrind etc.), and standard programming paradigms (such as full multi-threading), while being energy efficient and real-time capable. One way to gauge this effort is to make a parallel with IETF protocols such as 6LoWPAN or CoAP, which adapt IP to memory and energy-constrained IoT devices, without losing interoperability with the rest of the Internet. Similarly, a modern IoT operating system such as RIOT enables a full-fledged operating system on IoT devices with resources that are too limited for Linux to run on, without losing interoperability with state-of-the-art development tools and programmers.

RIOT is developed by an international community of open source developers (co-founded in 2013 by Inria, Freie Universität Berlin, and Hamburg University of Applied Sciences in the context of the SAFEST project), using transparent decision processes based on rough consensus, open discussion forums, open-source code, and open standards. The near-future goal of RIOT is to power a modern, large scale, evolutionary, and secure cyber-physical ecosystem, comprising heterogeneous IoT devices, distributed processes and applications (see Figure 1). Such appli-

cations are expected to leverage spontaneous wireless networks, Internet connectivity, the cloud and a dense, interconnected environment of communicating devices. In the mid and long term, such IoT applications are expected to not only enable optimization of existing processes, but also entirely new processes and the emergence of an enhanced reality, in which our interface with the Internet will no longer be predominantly a screen, but rather the objects of the cyber-physical system embodied by the Internet of Things itself.

Link:

The SAFEST Project, co-funded by ANR and BMBF, <http://safest.realmv6.org>

References:

- [1] C. Bormann et al.: "Terminology for Constrained node Networks", RFC 7228, Internet Engineering Task Force, 2014.
- [2] The Internet Engineering Task Force (IETF), <http://www.ietf.org>
- [3] RIOT: The Friendly Operating System for the Internet of Things. <http://www.riot-os.org>

Please contact:

Emmanuel Baccelli
Inria, France.
E-mail: emmanuel.baccelli@inria.fr

FIT IoT-LAB: The Largest IoT Open Experimental Testbed

by Eric Fleury, Nathalie Mitton, Thomas Noël and Cedric Adjih

The universal proliferation of intelligent objects is making Internet of Things (IoT) a reality; to operate on a large scale it will critically rely on new, seamless, forms of communications. But how can innovations be validated in a controlled environment, before being massively deployed into the real world? FIT IoT-LAB addresses this challenge by offering a unique open first class service to all IoT developers, researchers, integrators and developers: a large-scale experimental testbed allowing design, development, deployment and testing of innovative IoT applications, in order to test the future and make it safe.

IoT is not simply emerging as a major technology trend: it is already a reality with billions of existing devices. The Internet of Things represents a tremendous paradigm shift since Internet was designed; an evolution from pure end-to-end communication between an end-user device and a server in the Internet, to an Internet interconnecting physical objects that are freely able to communicate with each other and with humans. IoT builds on three pillars [1], related to the ability of smart objects to: (i) compute, (ii) communicate, and (iii) sense and interact with their environment.

Although IoT is already a reality, it is still maturing and waiting for its 'iPhone moment'. Several challenges remain, in particular relating to the standardization of efficient and universal protocols, and to the design and testing of IoT services and applications. Owing to their massively distributed nature, their design, implementation, and evaluation are inherently complex and tend to be daunting, time-consuming tasks. Required to overcome this hurdle is a representative, large scale, platform allowing researchers, IoT designers, developers and engineers to construct, benchmark and optimize their protocols, applications and services.

The FIT IoT-LAB testbed is our answer to these challenges. It offers a first class facility with thousands of wireless nodes to evaluate and experiment very large scale wireless IoT technologies ranging from low level protocols to advanced services integrated with Internet, accelerating the advent of groundbreaking networking technologies. IoT-LAB also provides dozen of robots to test and improve the impact of the mobility of IoT devices. FIT IoT-LAB's main and most important goal is to offer an accurate open access multi-user scientific tool to support design,

development, tuning, and experimentation related to IoT.

IoT-LAB is part of the FIT (Future Internet of Things) project which develops experimental facilities within a federated and competitive infrastructure with international visibility and a broad panel of customers. All facilities come with complementary components that enable experimentation on innovative services for academic and industrial users. FIT has received 5.8 million Euros in funding from the Equipex research grant program. The FIT consortium is coordinated by University Pierre et Marie Curie and composed of Inria, ICube laboratory from University of Strasbourg, Institut Mines-Télécom and CNRS. FIT is a platform federation. Such a federation of independent network experimental facilities is arguably the only meaningful way to achieve the required scale and representativeness for supporting Future Internet research. IoT-LAB testbeds are dispersed among six different locations across France offering access, for the first time, to 2728 wireless IoT fixed and mobile nodes equipped with various sensors (see Table 1).

Users can select and reserve the number and type of nodes they wish with

respect to different features, such as microcontroller (TI MSP430, ARM Cortex M3 and ARM Cortex A8), radio chip (2.4GHz or 800MHz) and additional functionalities (mobility / robots, accelerometer, magnetometer and gyrometer). Resources can be reserved on one or several sites at once. Services offered by IoT-LAB include:

- Total remote access to nodes reserved, e.g., allowing users to flash any firmware, without any condition or constraint. Any language or OS can be used to design, build, and compile applications;
- Direct access to a debugger server on each node so that all debugging can be performed remotely on the node (such as step by step code execution)
- Access to the serial ports of all nodes for a real-time interaction, with optional aggregation;
- Each node could be visible from Internet with end-to-end IP connection using IPv6 and 6LoWPAN for example;
- A very accurate power consumption monitoring of every node.
- Packet sniffer and analyzer on each node;
- A GPS module for some A8 nodes allowing a very precise end-to-end time synchronization, accurate moni-

	Grenoble	Lille	Paris	Rennes	Institut Télécom	Strasbourg	Total
WSN430 (800MhZ)	256					256	512
WSN430 (2.4GhZ)		256	120	256			632
M3	384	320	24		90	120	938
A8	256		200		70	24	550
Host Node	32	64					96
Total	928	640	344	256	160	400	2728

Table 1: IoT-LAB testbeds.




Figure 1: FIT-IoT LAB platform. Left: Strasbourg, centre: Grenoble, right: Lille.

toring and performance evaluation of communication protocols;

- A set of useful detailed tutorials, OS supports (Contiki, FreeRTOS, TinyOS, and RIOT) including full protocol stacks and communication libraries such as OpenWSN providing open-source implementations of IoT protocol standards;
- A unique fleet of mobile robots (WifiBot and TurtleBot);
- Strong extensibility through the availability of more than 100 empty slots on which users can physically

plug their own hardware devices, while benefiting from the IoT-LAB services and monitoring tools (a feature frequently required by both academic and industrial users).

Overall, FIT-IoT LAB testbed is a unique pioneer in the domain of IoT testbeds.

Links:

IoT-LAB: <https://www.iot-lab.info>
 FIT: <https://www.fit-equipex.fr>
 OpenWSN: <https://openwsn.atlassian.net/>

Reference:

[1] D. Miorandi, S. Sicari, F. D. Pellegrini, I. Chlamtac: “Internet of things: Vision, applications and research challenges”, *Ad Hoc Networks*, 10(7):1497 – 1516, 2012.

Please contact:

Eric Fleury
 ENS de Lyon / Inria, France
 Tel: +33 672 162 974
 E-mail: Eric.Fleury@inria.fr

OCARI: A Wireless Sensor Network for Industrial Environments

by Tuan Dang, Pascale Minet and Erwan Livolant

OCARI is a wireless sensor network designed to operate in industrial environments [1]. It is easy to deploy (i.e. ‘plug-and-play’), and is energy-efficient to support battery-operated nodes. OCARI nodes use commercial off-the shelf components. OCARI provides bounded medium access delays and the energy consumption of an OCARI network is predictable. In addition, the network is scalable (up to hundreds of sensor nodes) and able to support micro-mobility of nodes.

OCARI can be distinguished from WirelessHart and ZigBee by the following three characteristics:

- It relies on a mesh topology (see Figure 1), improving robustness and bandwidth use.
- It is self-configuring thanks to a dynamic multihop routing taking energy into account.
- It saves energy thanks to a deterministic medium access and sleeping periods for nodes.

Furthermore, OCARI has the advantage of being open source.

The OCARI stack comprises the Physical, MAC, Network and Application layers. It was first developed in the ANR OCARI project lead by EDF with Inria, LIMOS, LATTIS, LRI, Telit and DCNS. Today, it is actively developed by EDF, Inria, in collaboration with ADWAVE and other research labs. The network layer

includes OPERA (Optimized Energy efficient Routing and node Activity scheduling) that consists of:

- EOLSR, an energy-efficient routing including neighbourhood discovery and building of the routing tree rooted at the sink in charge of data gathering.
- OSERENA, [2] a colouring node protocol allowing a conflict-free schedule of node medium accesses to be built. Each node knows the time slots

in which it is allowed to transmit or it may receive from a neighbour. In the remaining slots it sleeps to save energy.

OCARI takes advantage of cross-layering [3] to optimize its performance. From the hardware point of view, an OCARI node consists of a radio frequency transceiver compliant with the IEEE 802.15.4 standard operating in

the 2.4GHz band and a 32 bits Cortex M3 microcontroller. OCARI is presently available on two platforms: Dresden deRFsam3-23T09-3 and ADWAVE ADWRF24-LRS (see Figure 2).

OCARI has been demonstrated many times, the most recent two demonstrations in November 2014:

- Demonstration during the Inria-Industry Telecom day: routing and colouring, fire detection on a DCNS ship construction site using ADWAVE hardware (see <http://www.inria.fr/en/innovation/calendard/telecoms-du-futur>).
- Demonstration of the steering committee of the Connexion Cluster project (<https://www.cluster-connexion.fr>): integration of sensors of various types in the OCARI network and interconnection of the OCARI network to the facility backbone by means of a OCARI/OPC-UA gateway built with Telecom ParisTech on a Raspberry Pi (see Figure 3).

Figure 1: OCARI topology.





Figure 2: OCARI on Dresden platform (left) and on ADWAVE platform (right).

Figure 3: Gateway OCARI/OPC-UA.



In the future, we plan to create an alliance, bringing together industrial end users (e.g., EDF, DCNS), sensor providers (e.g., DEF, Carmelec), research labs (e.g., Inria, Telecom ParisTech), integrators of network solutions (e.g., ADWAVE), SMEs and engineering service companies using wireless sensor networks (e.g., ATOS, PREDICT). The objectives of the OCARI Alliance are to create a sustainable ecosystem, gather feedback on real OCARI deployments and fund functional evolutions of OCARI.

Link: <http://www.ocari.org>

References:

- [1] T. Dang et al.: "Which wireless technology for industrial wireless sensor network? The development of OCARI technology", IEEE Transactions on Industrial Electronics, 2009.
- [2] I. Amdouni, P. Minet, C. Adjih "OSERENA: a coloring algorithm optimized for dense wireless networks", the International Journal of Networked and Distributed Computing (IJNDC), vol. 1, issue 1, November 2012.
- [3] K. Al Agha, et al.: "Cross-layering in an industrial wireless sensor network: case study of OCARI", JNW, Journal of NetWorks, special issue on Wireless sensor network: theory and practice, Academy Publisher, 2009.

Please contact:

Tuan Dang
EDF, STEP department, France
E-mail: tuan.dang@edf.fr

Pascale Minet, Erwan Livolant
Inria, France
E-mail: pascale.minet@inria.fr,
erwan.livolant@inria.fr

Security and Privacy in the Internet of Things

by Ludwig Seitz

The Internet of Things (IoT) has particular security and privacy problems. The Internet Engineering Task Force is designing authentication and authorization mechanisms for the most constrained devices which are part of the Internet of Things.

In many of the applications of Internet of Things (IoT), sensors measure variables such as speed, pressure, consumption, temperature or heart rate, and actuators control physical systems, such as brakes, valves, lights, power circuits, or automated drug dispensers.

What makes these scenarios interesting from a security and privacy perspective, is that they all affect the physical world, sometimes controlling critical infrastructure, and sometimes gathering very private information about individuals.

Clearly, there is a need for security and privacy protection in the IoT. Some of the devices used in the IoT have extremely limited memory, processing capacity and battery power; consequently, classical IT security mechanisms are often inadequate to cope with the unique security situations that arise. Many such devices operate on wireless networks, which only offer low bandwidth and which are prone to losing data-packets in transfer.

One would think that Moore's Law would fix these problems over time, by giving us more powerful processors, and cheaper memory modules. However, advances in this area go largely towards reducing the cost per unit, as well as power consumption, and not towards increasing performance.

We therefore need adapted security and privacy mechanisms that allow us to reap the potential benefits of the IoT, without endangering critical infrastructure or individual privacy.

Another problem is fragmentation, since security solutions are either not standardized, or are standardized only for one application area. This affects IT security in general, but IoT is particularly affected owing to the need for interoperability between devices produced by different vendors (otherwise we'd lose the 'Internet' from Internet of Things) and the rapid development of new technologies in that sector.



The author with a sensor node.

Privacy protection, on the other hand, depends largely on individual users to understand and configure security settings. This often requires a high level of IT security competence, and is therefore likely to fail more often than not. Addressing this issue is likely to greatly improve public acceptance of IoT consumer end products.

Currently the Internet Engineering Task Force (IETF) is working on various security topics related to the IoT. IETF is a large, international standardization organization, with a wide range of Internet related working groups. IETF has developed a number of protocols aimed at IoT applications, such as 6LoWPan, CoAP, and DTLS, and it currently has two working groups active in IoT security. The first, DICE, deals with adapting the DTLS [1] protocol to constrained environments, while the second, ACE, addresses authentication and authorization in constrained environments [2].

An example of the limitations that affect constrained devices is memory: it is assumed that the smallest devices capable of implementing some meaningful security measures would have around 10 KB of RAM memory and roughly 100 KB of persistent memory (e.g. flash-memory). As a comparison, the certificates that are sent over the network and processed as part of the HTTPS protocol (and the underlying

TLS protocol) have a size of at least 2 KB, and would therefore require a sizeable part of the small device's memory, making normal operation difficult.

While DICE's work will result in a protocol that allows these tiny devices to establish secure connections, the next security and privacy relevant question is what a device is allowed to do when it has successfully and securely connected to another. In traditional web services, these questions are answered by access control systems that usually query a database of authorization policies, defining who may access what under which circumstances.

It is obvious that the access control mechanisms designed for powerful servers cannot be applied for IoT without adapting them to the resource constraints described above. To this end, ACE is examining mechanisms where, for example, the burden of making access control decisions is shifted to a more powerful trusted third party and the device just has to enforce these access control decisions.

Thus the work on securing the IoT is ongoing, and the repercussions of the decisions that are made now will be affecting us for many years to come.

Links:

- (1) DICE WG:
<http://datatracker.ietf.org/wg/dice>
- (2) ACE WG:
<http://datatracker.ietf.org/wg/ace>

References:

- [1] E. Rescorla, N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012, <http://www.rfc-editor.org/info/rfc6347>.
- [2] L. Seitz, G. Selander, C. Gehrmann: "Authorization Framework for the Internet-of-Things", D-SPAN 2013.

Please contact:

Ludwig Seitz, SICS Swedish ICT
Tel: +46 703 49 9251
E-mail: ludwig@sics.se

Secure and Efficient Wireless Sensor Networks

by Corinna Schmitt and Burkhard Stiller

There exists a multitude of implemented, as well as envisioned, use cases for the Internet of Things (IoT) and Wireless Sensor Networks (WSN). Some of these use cases would benefit from the collected data being globally accessible to: (a) authorized users only; and (b) data processing units through the Internet. Much of the data collected, such as location or personal identifiers, are of a highly sensitive nature. Even seemingly innocuous data (e.g., energy consumption) can lead to potential infringements of user privacy.

The infrastructure of the Internet of Things (IoT) with its diversity of devices and applications, as well as the trend towards a separation of sensor network infrastructure and applications exacerbates security risks. A true end-to-end security solution is therefore required to achieve an adequate level of security for IoT. Protecting data once they leave the boundaries of a local network is not sufficient, especially when private and high-risk information are effected.

However, IoT is no longer limited to servers, routers, and computers with manifold resources. It also includes constrained devices – motes –, which are very limited in memory (approximately 10-50 KByte RAM and 100-256 KByte ROM), computational capacity, and power (supported by just a few AA batteries). These limited resources do not reduce the need to support end-to-end security and secure communications, but they make it much harder to meet these requirements. Depending on the specific

resources of these devices, the goal of secure WSNs is to support end-to-end security by a two-way authentication, an efficient data transport solution for the data, and a controlled data access, supporting the mobility of today’s users. Thus, different components were developed in the construction of SecureWSN and are illustrated in Figure 1. All components support hardware from different vendors with different resources. Different security solutions (TinySAM, TinyDTLS [2], or TinyTO [3]) were developed that are based on known algorithms from IP networks, like DTLS and BCK, and required adaptation (e.g., complexity) to fit these resources whilst supporting a heterogenous network structure.

End-to-end Security


Today, with so much personal information online, end-to-end security is essential in many situations. This represents the challenge for constrained devices usually used in WSNs.

SecureWSN tackles this challenge by developing three solutions for different types of resources.

TinyDTLS protects data from its source to the sink supporting confidentiality, integrity, and authenticity. RSA-capable (Rivest-Shamir-Adleman) devices are authenticated via X.509 certificates during the key-exchange in a two-way authentication handshake. Constrained devices perform a variant of the Transport Layer Security (TLS) pre-shared key algorithm. The data sink authenticates via certificate either directly with the mote or with an Access Control Server (ACS). ACS grants tickets to authenticated devices with sufficient rights. Motes request connection from their communication partner where key establishment is based on DTLS [2].

In comparison, TinyTO uses a Bellare-Canetti-Krawczyk (BCK) handshake with pre-shared keys. For the key generation, key exchange, signatures, and

Figure 1: Architecture and Components of SecureWSN.



encryption, the Elliptic Curve Cryptography (ECC) is used [3]. Thus, this solution saves resources and does not require a Certificate Authority (CA). It was shown that 192-bit ECC keys are as secure as 1024-bit RSA keys, which makes TinyTO a suitable alternative to TinyDTLS, supporting the same security functionality.

As sufficient resources are not always available to support the end-to-end security requirement, TinySAM was developed to support a one-way authentication. TinySAM uses the Otway-Rees key establishment protocol modified by the Abadi and Needham algorithm, where all nodes have an AES (Advanced Encryption Standard) key pair known by the key server. Two nodes build an individual session key pair for secure data exchange.

Owing to the diversity of applications and the amount of collected data, all these solutions also support aggregation in order to use the limited bandwidth (102 byte on MAC layer in IEEE 802.15.4) and energy as efficient as possible.

Data collected in WSNs consists of stable meta-information and sensor readings periodically measured and sent out in one message resulting in redundancy. Thus, the push-based Internet Protocol Flow Information Export (IPFIX) protocol serves optimization by dividing data into two small messages (template record and data record). The resulting TinyIPFIX protocol includes

special template records for motes and supports aggregation. Necessary header compression options were developed to reduce the overhead by required IPFIX headers [2].

Additionally, the SecureWSN approach includes the WSN configuration, management, and data handling of the WSN's owner by 'clicking buttons', all termed CoMaDa. CoMaDa works with a virtual representation of the real WSN network, displaying in real-time: (1) data collected and (2) the network status, as well as allowing for (3) mote updates (e.g., the degree of aggregation). The dedicated WebMaDa component [1] publishes the WSN data in the Internet and allows anyone who is authorized and has the appropriate credentials and rights, to view the WSNs.

Conclusions

SecureWSN consists of different modules supporting different end-to-end security modes, efficient data transport, aggregation, and controlled data access functionality. These solutions, which are currently available, are highly flexible, since each mechanism that is implemented can be selected depending on the requirements of the applications and hardware. As such, the approach of SecureWSN benefits any kind of IoT application that demands secure end-to-end support.

Continued work in this area will include further module developments and enhancements, such as pull requests and

ECC optimizations. Parts of SecureWSN were developed within EU projects SmartenIT and FLAMINGO and are part of the standardization process.

Links:

<http://www.csg.uzh.ch/research/SecureWSN.html>

<http://tools.ietf.org/html/draft-schmitt-ace-twowayauth-for-iot-01>

References:

- [1] M. Keller: "Design and Implementation of a Mobile App to Access and Manage Wireless Sensor Networks", Master Thesis, Universität Zürich, Communication Systems Group, Department of Informatics, Zürich, Switzerland, November 2014.
- [2] T. Kothmayr, et al.: "DTLS Based Security and Two-way Authentication for the Internet of Things", Ad Hoc Networks, Elsevier, Vol. 11, No. 8, November 2013, pp 2710-2723.
- [3] M. Noack: "Optimization of Two-way Authentication Protocol in Internet of Things", Master Thesis, Universität Zürich, Communication Systems Group, Department of Informatics, Zürich, Switzerland, August 2014.

Please contact:

Corinna Schmitt, Burkhard Stiller
Universität Zürich, Switzerland,
Tel: +41 44 635 7585,
+41 44 635 6710
E-mail: schmitt@ifi.uzh.ch,
stiller@ifi.uzh.ch

Creating Internet of Things Applications from Building Blocks

by Frank Alexander Kraemer and Peter Herrmann

Reactive Blocks is a toolkit to create, analyse and implement reactive applications. It helps developers to get concurrency issues right, reuse existing solutions, and brings model checking to the common programmer.

Internet of Things (IoT) applications connect hardware to communicate within a network, often with constrained resources. Therefore, even simple use cases quickly become very complex. In addition, IoT applications often combine several technologies, and only few programmers have the skill set to cover them all.

Model-driven development and model checking provide solutions to these problems. However, these solutions are barely used by programmers. Reasons for this are that model checking requires deep knowledge in formal methods to produce an appropriate input model that can be analysed. In addition, model-driven approaches

often fail to cover details in code in a suitable way.

With these concerns in mind, we developed the Reactive Blocks tool for event-driven and concurrent applications. It has its origins at the Department of Telematics at the Norwegian University of Science and Technology (NTNU),




Figure 1: The Reactive Blocks workflow. Blocks from libraries are combined to applications. Blocks are model-checked, and the application built via code generation.

and has been further developed by the spin-off company Bitreactive.

Building IoT Applications with Reactive Blocks

To appeal to the common programmer, the tool combines programming and modelling with an automatic compositional analysis, as shown in Figure 1: A developer creates a system from building blocks. Many of the blocks do not need to be developed from scratch but can be selected from libraries (1) The blocks are connected with each other (2) This is facilitated by behavioural contracts enabling the developer to understand the interface behaviour of a block without having to understand its inner details. Applications consist of a hierarchy of blocks. The applications are then analysed by automatic model checking. Once consistent, code is automatically generated and can be deployed as OSGi bundles or standalone Java applications (3).

To balance the benefits of generating code with that of manual programming, Reactive Blocks uses a combination of Unified Modelling Language (UML) with Java. UML describes the concurrent behaviour of building blocks with activity diagrams or state machines. State machines are also used to describe the contract of a block. The UML models only coordinate concurrent behaviour, and refer to Java methods for detailed operations on application programming interfaces (APIs) or other data. The Java methods are edited in the Java tools within Eclipse. Existing code can also be integrated, by encapsulating it into operations and building blocks. The code related to concurrency, i.e., code that decides when the operations are called, is generated from UML. This

is usually the kind of code that is cumbersome and error-prone to write manually.

Compositional Analysis

A model checker is built into the editor. The building blocks have a formal semantics and correspond to temporal logic formulas, so that the model checker can take the UML model as input. Developers can start model checking with a single click. The encapsulation of building blocks by contracts enables a compositional analysis, in which each block can be analysed separately. This minimises the problem of state space explosion. An error is shown to the developer as animation in the editor such that can be easily understood and fixed. Examples of errors are deadlocks or violations of protocols, which are detected because the UML model of the building blocks do not reach their final state, or because the external contract of a block is violated. In addition, the tool detects if an application does not honour the life cycle contract of a component when it runs in a framework like OSGi. Such errors are hard, if not impossible, to detect by testing. The mathematical basis of the building blocks in temporal logic formulas makes it possible to reason about other properties, such as reliability, security [1], real-time properties [2], or spatial constraints [3].

Fleet Monitoring Case Study

A case study was developed within the IMSIS project funded by the Norwegian Research Council. It consists of a system that monitors a fleet of trucks while they deliver goods. Their position relative to defined geo-fences is constantly monitored and compared to expected arrival times. The backend

system is informed about the arrival of freight. In case of deviations, alerts via email and SMS are sent. The system was created by programmers without specific knowledge in formal methods. Due to the built-in analysis, all interactions work correctly. Using building blocks, about 60% of the functionality could be taken from reusable libraries.

Link:

<http://bitreactive.com>

References:

- [1] L.A. Gunawan, P. Herrmann: "Compositional Verification of Application-Level Security Properties, in proc. of ESSoS 2013, LNCS 7781, February/March 2013.
- [2] Han F., P. Herrmann: "Modeling Real-Time System Performance with Respect to Scheduling Analysis", in proc. of UMEDIA 2013, Aizu-Wakamatsu, IEEE Computer Society Press, September 2013.
- [3] P. Herrmann, J.O. Blech, F. Han, H. Schmidt. A Model-based Toolchain to Verify Spatial Behavior of Cyber-Physical Systems. In 2014 Asia-Pacific Services Computing Conference (APSCC), IEEE Computer.

Please contact:

Frank Alexander Kraemer, Peter Herrmann
NTNU, Norway,
E-mail: kraemer@item.ntnu.no,
herrmann@item.ntnu.no

Interfaces to the Internet of Things with XForms

by Steven Pemberton

XForms is a language for describing interfaces to data, designed at W3C by researchers from industry and academia. It is a declarative language, meaning it describes what has to be done, but largely not how. The interface it describes does not have to run locally on the machine producing the data, but can be run remotely over the network. Since Internet of Things (IoT) computers typically have little memory and are low-powered, this makes XForms ideally suited for the task.

One of the unexpected successes of HTML was its adoption for controlling devices with embedded computers, such as home Wi-Fi routers. To make an adjustment to such a device, the user directs the browser to the IP address from which it is running and a small web server on the device serves up web pages that allow the user to fill in and submit values to change the working of the device.

However, the tiny embedded computers that form part of the IoT typically have memory in kilobytes, not megabytes, and lack the power to run a web server that can serve and interpret web pages. This calls for a different approach.

One approach is for the devices to serve up only the data of the parameters, so that those values can then be injected into an interface served from elsewhere. XForms [1], a standard that we have helped develop at W3C, is designed for exactly this type of scenario: although it is a technology originally designed for improving the handling of forms on the web, it has since been generalised to more general applications; version 2.0 is currently in preparation [2].

XForms has two essential parts: the first part is the model that specifies details of the data being collected, where it comes from, its structure, and constraints. It allows data from several sources to be combined, and data to be submitted to different places. It also ensures that as data is changed, that relationships between the data are kept up to date. The second part is the user interface that displays values and specifies controls for changing the values. These controls are specified in a device-independent manner that only describes what they are meant to achieve, not how to do it. This makes it easier to adapt the interface to different devices, screen sizes, etc., while still allowing the use of specific interfaces, such as radio buttons, via style sheets.

XForms has already been used for a number of years to control devices in this way at many petrol stations in the USA. Each device, storage tanks, petrol pumps, cash registers, and so on, contains a simple server that delivers its data as XML instances. XForms interfaces are then used to read and combine these values, and update control values (for instance the price of fuel being displayed on pumps).

As an example of how it could be used, Nest, a well-known producer of internet thermostats, has published the data-model interface to its devices. A simple interface to this could look like this:

```
<instance src="http://thermostat.local/">
<bind ref="ambient_temperature_c"
      type="decimal" readonly="true()">
<bind ref="target_temperature_c"
      type="decimal">
<bind ref="target_temperature_f"
      type="decimal" calculate=
      "../target_temperature_c*9/5+32">
<submission resource=
      "http://thermostat.local/data"
      method="put" replace="instance"/>
```

The following device-independent user interface control specifies that a single value is to be selected from the list of items, without specifying how that is to be achieved (using radio buttons, drop-downs or whatever):

```
<select1 ref="temperature_scale"
        label="Scale">
  <item value="f" label="°F"/>
  <item value="c" label="°C"/>
</select1>
```

Experience with XForms has shown that using it can reduce production time and costs by a factor of ten. As an example, one very large pilot project reduced production time from five years with thirty people to a single year with ten people. These remarkable savings are due largely to the declarative nature of XForms, which greatly reduces the

administrative programming needed. Several implementations of XForms are available, at least three of which are open-source. It is also part of the Open Document Format, implemented by Open Office and Libre Office. There is also a tutorial [3].

Link:

Nest API Reference:

<https://developer.nest.com/documentation/api-reference>

References:

- [1] J. M Boyer (ed.), XForms 1.1, W3C, 2009, <http://www.w3.org/TR/2009/REC-xforms-20091020/>
- [2] J. M. Boyer et al. (eds.), XForms 2.0, W3C, 2015, http://www.w3.org/MarkUp/Forms/wiki/XForms_2.0
- [3] Steven Pemberton, XForms for HTML Authors, W3C, 2010, <http://www.w3.org/MarkUp/Forms/2010/xforms11-for-html-authors>

Please contact:

Steven Pemberton
Chair of the W3C XForms WG, CWI,
The Netherlands
E-mail: steven.pemberton@cwi.nl

Home Automation Devices Belong to the IoT World

by Vittorio Miori and Dario Russo

We present a practical and scalable solution that aims to achieve the Internet of Things (IoT) paradigm in complex contexts, such as the home automation market, in which problems are caused by the presence of proprietary and closed systems with no compatibility with Internet protocols.

Home automation is an appealing context for the Internet of Things (IoT). We envisage future home environments with self configured embedded sensors and actuators (e.g., in consumer electronic products and systems) that can be controlled remotely through the Internet, enabling a variety of monitoring and control applications. Manufacturers will produce their own IP gateways so that proprietary domotic systems can be interfaced with an IPv4 enabled Ethernet socket.

By connecting the IP gateway directly to the Internet or through a home/residential gateway, the domotic system can be managed remotely using a PC, Smartphone or Tablet (Figure 1).

Unlike the IoT paradigm, a solution of this kind provides the home with a unique Internet access point (and a unique public IP address that can be assigned to the IP gateway or to the home/residential gateway in relation to the home network configuration) to control all the devices. In this scenario, the assigned public IP address does not identify a single device or function, but identifies the entire domotic network, and a customized software manager application able to locate devices and to activate their functions inside the domotic network is required.

One of the major problems in the home automation area is that different systems are neither interoperable nor interconnected. To tackle this issue, our laboratory has created a framework called DomoNet (Figure 2).

DomoNet [1] is an open source software released under the GPL license, following W3C recommended web standard technologies such as Web Services, SOA and XML. It constructs a unique view of the system that includes all the devices belonging to the different domotic systems available, through a set of modules that work as gateways. To implement interoperability, DomoNet defines a standard language, called



Figure 1: Home Remote Control.
Source: Home Care Reviews.

DomoML, for the semantic abstraction of heterogeneous systems in order to describe device functions, data types, messages and models of the interactions and communications among framework entities.

By exploiting the DomoML language, DomoNet has a uniform view of all the devices and their relative functions available on its network. In this context, all the differences between domotic systems are flattened out so that DomoNet can interact with each device using only this single high-level language. To physically execute operations on devices, DomoNet implements special entities named TechManagers. TechManagers are gateways that, on the one hand, interface DomoNet and the DomoML language, and, on the other, manage events on the specific domotic bus.

A software module has been developed to enable any home automation device to interface and interact via the IPv6 [2] network protocol. All home automation devices can thereby interact actively with the surrounding world, through their own IPv6 address that identifies them uniquely on the Internet. The system thus increases the user's ability to take full advantage of the benefits offered by the new IoT vision of the world.

When a device signals its entrance into the DomoNet network (Figure 3), its associated TechManager creates the corresponding DomoDevice. The formal representation of the DomoDevice is sent to the DomoNet server, which assigns it both a DomoNet and an IPv6 address, thereby providing a dual identification. In doing so, the DomoNet server compiles a bi-directional map in order to enable identification of address correspondences. The provision of this dual representation ensures backward compatibility with the applications and services previously created for the DomoNet ecosystem.

To process the request, the DomoNet server forwards it to the appropriate TechManager, which then translates the request into a formalism that the device can process and queries the device involved. The device's response is




Figure 2: DomoNet architecture.




Figure 3: The Test Architecture.

delivered to the TechManager and, after the DomoML translation process, is forwarded to the DomoNet server, which then sends it to the requester.

We are working on the development of a web server running inside the same computer that hosts the DomoNet server, in order to simulate a direct web interaction with the device identified by its IPv6 address. The Tomcat web server was used in the test architecture.

A future development will be to decentralize the entire DomoNet architecture in order to create virtual independent devices that act as agents capable of cooperating with each other according to the Digital Ecosystem architecture paradigm [3]. This approach would further the realization of the concepts underlying the IoT paradigm and promote a new model for thinking about the environment surrounding us and the objects contained therein.

Links:

Domotics Lab Home Page:
<http://www.isti.cnr.it/research/unit.php?unit=HA>
 DomoNet Home Page:
<http://sourceforge.net/projects/domonet/>

References:

- [1] V. Miori, D. Russo, M. Aliberti: "Domotic technologies incompatibility becomes user transparent", Communications of the ACM, vol. 53 (1) pp. 153 - 157. ACM, 2010.
- [2] M. Jung et al.: "Heterogeneous device interaction using an IPv6 enabled service-oriented architecture for building automation systems", in proc. of SAC'13, ACM, 1939-1941, 2013, DOI=10.1145/2480362.2480722
- [3] G. Briscoe, P. De Wilde: "Digital ecosystems: evolving service-orientated architectures", in proc. of BIONETICS'06, ACM, , Article 17, 2006. DOI=10.1145/1315843.1315864 <http://doi.acm.org/10.1145/1315843.1315864>

Please contact:

Vittorio Miori
 Domotics Lab, ISTI CNR, Italy
 E-mail: vittorio.miori@isti.cnr.it

Diopase: Data Streaming Middleware for the Internet of Things

by Benjamin Billet and Valérie Issarny

The Diopase middleware provides developers with new methods for writing distributed applications for the Internet of Things (IoT). Diopase leverages data streams as data model and continuous in-network processing as computation model, in order to deal with the challenging volume of data being continuously produced by the 'Things'.

The Internet of Things (IoT) is currently characterized by an ever-growing number of networked 'Things', i.e., devices that have their own identity together with advanced computation and networking capabilities: smartphones, smart watches, smart home appliances, etc. These Things are being equipped with increasing numbers of sensors and actuators that enable them to sense and act on their environment, linking the physical and virtual worlds. Specifically, the IoT raises many challenges related to its very large scale and high dynamicity, as well as the great heterogeneity of the data and systems involved (e.g., powerful versus resource-constrained

devices, mobile versus fixed devices, continuously-powered versus battery-powered devices, etc.). These challenges require new systems and techniques for developing applications that are able to: (i) collect and process data from the numerous data sources of the IoT and (ii) interact both with the environment using the actuators, and with the users using dedicated GUIs.

Solutions for the IoT currently rely heavily on third-party services and infrastructures, such as the cloud of Things where each device sends its measurements to centralized collection and computation points. Several prob-

lems emerge from these approaches, related to energy consumption (due to intense wireless communications) and device lifetime, network overload and privacy preservation. The challenge is then to allow Things to be much more autonomous and use third-party services only when required, such as in wireless sensor networks, but at a much larger scale.

To this end, the Diopase [1] solution aims at providing a common middleware layer that runs directly on the Things, enabling them to manage the huge volume of data continuously being produced (measurements, events, logs,

etc.) in a collaborative way. In this context, we consider data streams and continuous processing as the reference data and computation models for developing IoT applications. Continuous processing is indeed a very suitable paradigm for processing each piece of data one by one, without having to store the entire dataset.

Dioptase leverages a service-oriented architecture (SOA) revisited for supporting continuous processing. This SOA introduces the concept of stream services that produce and consume infinite data streams, in contrast to regular services that manage finite datasets. Stream services are classified into four families, depending on their roles:

- producers, which generate new data streams (e.g., from sensors);
- processors, which produce new data streams by continuously processing existing ones;
- storages, which store persistently or temporarily the data extracted from data streams and can serve them as new data streams when required;
- consumers, which acquire data streams for driving actuators or updating GUIs.

Once Dioptase is deployed onto a Thing, it enables developers to manage the Thing as an abstracted and homogeneous pool of resources that can execute stream services provided over time. These stream services can be developed by using the native languages of the device or a new lightweight stream processing language, called DiSPL, directly interpretable by Dioptase.

By composing these stream services, developers build their IoT applications as a set of tasks that are executed continuously by the hosting Things in a purely distributed manner. Developers are provided with dedicated tools for designing tasks graphically before injecting them into the network of Things at any time, using deployment Web services or through the Dioptase deployment server (DiDS). DiDS automatically computes where to deploy the stream services and then manages their execution over time. DiDS relies on a state-of-the-art task allocation algorithm [2], which guarantees the lifetime of each task, by minimizing the overall energy consumption or maximizing the fair allocation of tasks among the Things, according to the constraints

expressed by the developers and the characteristics of the available Things.

Dioptase features a customizable middleware architecture that is versatile enough to be deployed on a large class of Things that vary significantly in terms of resource availability (e.g.,

gate, adapt and split their own tasks according to their environment, their load, their available resources and their capabilities in a decentralized collaborative manner, taking the role of DiDS. As a benefit, the network of Things would adapt itself to minor and major changes over time.




Figure 1: Dioptase, a data streaming middleware for the IoT.

embedded systems, smartphones or plug computers), provided that these Things are able to communicate directly through the Internet infrastructure. As illustrated in Figure 1, Dioptase embeds various modules for interacting with existing components being used as part of the IoT infrastructure: discovery systems and registries, computation and offloading infrastructure, legacy sensor and actuator networks, etc.

This work was part of the larger CHOREOS project effort, which revisits the concept of choreography-centric service-oriented systems to introduce a dynamic development process and associated methods, tools, and middleware for the services in the Ultra Large Scale (ULS) Future Internet. Beyond CHOREOS, Dioptase represents a significant step towards an IoT where every Thing can be assigned tasks and complete them autonomously. As an ongoing area of research, Dioptase will be released soon as an open source project. For the near future, we plan to enable Things to automatically dele-

Link:

<https://mimove.inria.fr/dioptase>

References:

- [1] B. Billet, V. Issarny: "Dioptase: A distributed data streaming middleware for the future Web of Things", *Journal of Internet Services and Applications*, vol. 5, no. 13, 2014.
- [2] B. Billet, V. Issarny: "From task graphs to concrete actions: A new task mapping algorithm for the future Internet of Things", in *proc. of the 11th IEEE International Conference on Mobile Ad hoc and Sensor Systems (MASS)*, 2014.

Please contact:

Benjamin Billet, Valérie Issarny
Inria, France
E-mail: Benjamin.Billet@inria.fr,
valerie.issarny@inria.fr

Programming Actors for the Internet of Things

by Raphael Hiesgen, Dominik Charousset, Thomas C. Schmidt and Matthias Wählisch

The Internet of Things (IoT) enables a large number of devices to cooperate to achieve a common task. Each individual device is small and executes a confined software core. Collective intelligence is gained from distributed collaboration and Internet communication. Corresponding IoT solutions form large distributed software systems that pose professional requirements: scalability, reliability, security, portability and maintainability. The C++ Actor Framework CAF [1] contributes such a professional open source software layer for the IoT. Based on the actor model of Hewitt et al. [2], it aids programmers at a good level of abstraction without sacrificing performance.

The Internet of Things (IoT) is composed of many nodes, often with limited capabilities. Small software components that communicate via Internet protocol standards typically form a highly distributed work flow among machines with minimal human interaction. Traditional application scenarios include sensor networks that monitor data such as environmental or cyber-physical conditions. In addition to sensors, IoT networks include actuators that influence the environment. Complex applications are built from these nodes, for example: home automation, and tracking of health data. These systems enable machines to upload data to Internet servers, a task that originally required human interaction. Thus, they allow the tracking of data everywhere and anytime.

The main challenge for developers is to ensure an appropriate service quality even in unstable IoT environments. While working on a network of constrained machines, they need to supply synchronization primitives as well as mechanisms for error detection and propagation. To address these challenges, many developers fall back to low-level coding that requires specialized knowledge and is hard to maintain. As a result, code is barely portable, and often hand-crafted, which introduces high complexity, many sources of errors and little generality.

A distributed software layer as part of an operating system (OS) for embedded devices can resolve these challenges. While the OS eases portability and provides a standardized API for low-level functionality, the software layer uses these interfaces to provide a high abstraction for the application developer. For this purpose, we contribute the C++ Actor Framework (CAF) [1]. It is developed as an open source project at the Hamburg University of Applied Sciences and is open for contributions

from the community. CAF enables asynchronous communication between nodes by a reliable message passing facility. Following the 'shared nothing' paradigm of actors, it provides synchronization primitives as well as error handling capabilities at very low overhead. CAF runs on RIOT [3], the friendly operating system for the IoT.

As early as 1973, Hewitt et al. [2] proposed the (abstract) actor model to address the problems of concurrency and distribution. This model defines

aim at avoiding intermediate state in the network. Letting a runtime environment handle low-level functionality gives developers more time to focus on the application logic.

Our C++ Actor Framework has elaborated this model to allow for development of native software at a high abstraction layer without sacrificing performance. We provide exchangeable implementations of the runtime environment to make optimized use of




Figure 1: The CAF distributed software layer adds transactional message passing to IoT nodes.

entities called actors, which are concurrent, isolated and solely interact via network-transparent message passing based on unique identifiers. As a result, actors cannot corrupt the state of other actors. Furthermore, an actor can create new actors by an operation called spawn. This is often used to distribute workload, e.g., in a divide and conquer approach where actors divide problems and spawn new actors to handle the sub-problems concurrently. Furthermore, the actor model specifies error handling capabilities for distributed systems that allow for monitoring of subsystems and

system resources. This allows CAF to scale up to large performance-critical applications and down to small embedded devices. Its flexibility enables developers to test and verify their code on desktop machines before re-compiling and deploying the software on low-end devices. Hence, CAF provides a seamless development cycle, aiding developers to ship only well-tested components to the IoT. CAF can fill the gap between the high level of abstraction offered by the actor model and an efficient, native runtime environment.

IoT environments introduce specific constraints such as lossy networks, low-powered nodes, and communication capabilities limited to small packet sizes and intermittent connectivity. The original actor model was built on a strong coupling between components and cannot directly be transferred to the IoT. For example, the distributed error handling capabilities were not designed with constraints in mind and require adjustment. Furthermore, security mechanisms are not included and left to the runtime environment. However, IoT devices often carry private or critical data, require reliably and should remain resilient against node tampering. Hence, it is critical to provide encrypted communication as well as an authentication and authorization scheme for nodes.

We adapted CAF to the IoT environment by a new, loosely coupled communication layer as depicted in Figure 1. Part of its design is a transactional message-passing based on the request-response model offered by CoAP. This protocol offers the option of reliable message exchange as well as duplicate message detection. Each message exchange is independent and less vulnerable to connection failures than a coherent data stream. Our runtime environment can provide error propagation and mitigation in cases where messages

cannot be delivered after multiple retries.

A major focus in our adaption is a new network stack. The default implementation of CAF focuses on locally distributed hardware with many cores. As such, the network stack is built on TCP/IP in a straight-forward manner. In contrast, for the IoT, our transactional network stack targets IEEE 802.15.4 or Bluetooth LE. The IP layer deploys 6LoWPAN to sustain IPv6 compatibility, while UDP is used at the transport layer. With regard to security we rely on DTLS for encryption. Further, we are working on an authentication concept for IoT environments based on ID-based cryptography. Corresponding message exchanges will use the request-response model of CoAP.

As the number of devices connected to the IoT is expected to rise significantly in the coming years, we anticipate a severe demand for professionalizing application development in this domain. Consequently, a programming environment that offers a high level of abstraction on native code is needed to build robust and scalable applications of appropriate robustness and performance. We believe that CAF presents a promising advancement towards 'Programming the IoT' - conjointly with the embedded operating system RIOT.

Links:

<http://www.actor-framework.org>
<http://riot-os.org>

References:

- [1] D. Charousset, R. Hiesgen, T.C. Schmidt: "CAF - The C++ Actor Framework for Scalable and Resource-efficient Applications", in proc. of the 5th ACM SIGPLAN SPLASH '14, Workshop AGERE!, ACM, 2014.
- [2] C. Hewitt, P. Bishop, R. Steiger: "A Universal Modular ACTOR Formalism for Artificial Intelligence", in proc. of the 3rd IJCAI, San Francisco, CA, USA, 1973, Morgan Kaufmann Publishers Inc., pp. 235-245.
- [3] E. Baccelli, et al.: "RIOT OS: Towards an OS for the Internet of Things", in proc. of the 32nd IEEE INFOCOM, Poster (Piscataway, NJ, USA, 2013), IEEE Press.

Please contact:

Dominik Charousset, Raphael Hiesgen,
 Thomas C. Schmidt
 Hamburg University of Applied
 Sciences, Hamburg, Germany
 E-mail:
 {dominik.charousset,raphael.hiesgen,t.
 schmidt}@haw-hamburg.de

Matthias Wählisch
 Freie Universität Berlin, Germany
 E-mail: waelisch@ieee.org

A Mobile End-User Development Environment for IoT Applications Exploiting the Puzzle Metaphor

by José Danado and Fabio Paternò

'Puzzle', a mobile end user development framework, allows the end user to opportunistically create, modify and execute applications for interacting with smart things, phone functions and web services. The framework simplifies development of such applications through an intuitive jigsaw metaphor that allows easy composition on touch-based devices. Users immediately identified the utility of this feature, and found it easy to envisage using the framework in various scenarios in their daily lives.

The main purpose of this work is to allow users to do more with their existing devices and things within their homes or at work. This can be achieved by creating applications customized for individual needs on personal mobile devices. The Puzzle framework allows users to visualize the current status of intelligent objects, such as home appliances; how to operate them— for instance, through a voice command; apply the commands to

the object; and more generally, to prototype new Internet of Things (IoT) applications through a user-centred design approach in which users create and develop personalized applications.

In order to allow end users to create their own IoT applications and support their extensibility, the main challenges addressed in Puzzle [1] are twofold : how to create a User Interface (UI) that

seamlessly support IoT users to develop applications that meet individual needs, and how to create a flexible platform which supports customisation and interoperability of IoT applications. The mobile authoring environment has been designed taking into account how users can foresee the functions to compose and understand how the flow of their applications progresses, adopting a metaphor that is close to users' real life




Figure 1: Authoring and Execution of Puzzle Applications.

experiences, thus reducing their learning effort and increasing acceptance. On the other end, the platform needs to seamlessly control execution and communication through a plethora of objects and devices – e.g., light, heating, alarm systems, TVs, mobile devices. Consequently, it is important to be able to support interaction and communication between IoT devices and/or with the environment to exchange data and information ‘sensed’ about the environment, while reacting autonomously to events in the ‘real/physical world’, and influencing it by running processes that trigger actions and create services [2].

In the HIIS lab at CNR-ISTI, with support from an ERCIM post-doc fellowship, we have created Puzzle – a mobile end user development tool and a web-based platform to create applications including web services, phone functions and smart things. The mobile end-user development tool and related user interface (UI) considered current research in end-user development tools, limitations of mobile devices and focus on a user-centred design approach for IoT [3]. Puzzle is based on the metaphor conveyed by jigsaw pieces to stimulate end users to combine existing functions in ways that make sense to them. The decision to adopt that metaphor was based on its usage on other EUD environments, e.g., Scratch. In contrast to other visual languages, Puzzle adopts a higher level approach not just mimicking a traditional language through a graphical metaphor, but providing jigsaw pieces ready to be combined on the go, thus decreasing the learning curve and motivating users to explore and use it. Furthermore, jigsaw pieces were designed to facilitate users to combine them, and to solve errors and conflicts made during their combination.

Supported interaction techniques were developed in collaboration with end users, and their feedback was considered both to design and develop the UI and scenarios of use. The approach enabled us to avoid putting technology before the needs of the end user; consequently, allowing end-users to voice their requirements during the design and development of Puzzle, resulting in a user-friendly application able to interact with the physical environment.

At the framework level, end-users tinkering with IoT devices can stimulate creating and adding jigsaws as new building blocks to Puzzle, e.g. adding support to communicate with the lights in the house. New building blocks are created in Javascript; including their logic – i.e., inputs, outputs and functions, according to the Puzzle building block description. Once the building block is developed, it is added to a database storing its description to be combined and used in the authoring tool by end-users; thus enabling its execution management in Puzzle applications. During execution, building blocks follow top-down and left-to-right execution flow, where data outputted from the building block corresponding to a previous jigsaw is the input of a building block associated with the following jigsaw.

Use of IoT devices is triggered, during execution, inside a building block and through the use of web-services. Use of web-services provides a standard access to functionalities, hides platform dependent implementation and, consequently, provides the required interoperability in an IoT environment. Interoperability is also enforced through the Puzzle building block descriptions, which are used to check whether input and output jigsaw data

types match and manage the use of unpredicted values during application execution. This is important as the IoT is a heterogeneous environment; thus failure of one device needs to be mitigated by the environment.

Due to the heterogeneity of the environment, standard web-based technologies – e.g., HTML, CSS, Javascript and JSON – are used to support execution of Puzzle applications, wireless networks are used to support communication between IoT devices – e.g., Bluetooth or IEEE 802.15.4, and open hardware is used to foster contributions by a community of users – e.g., Arduino. Open hardware can be used to interface between proprietary devices and Puzzle, allow Puzzle to control IoT devices and be an incentive for a community of users – i.e., tinkering with current hardware - to increase hardware supported, and even create new solutions able to exploit Puzzle connection to existing hardware.

Links:

HIIS Lab: <http://giove.isti.cnr.it>
 Framework information at:
<http://giove.isti.cnr.it/tools/Puzzle/home>
 Documentation:
<http://hiis.isti.cnr.it:8080/AuthoringTool/v2/docs/>
 Video:
https://www.youtube.com/watch?feature=player_embedded&v=Asbqv7-QqqM

References:

- [1] J. Danado, F. Paternò: “Puzzle: A mobile application development environment using a jigsaw metaphor”, in *J. Vis. Lang. Comput.* 25, Elsevier, 2014, 297-315.
- [2] S. Holloway, C. Julien: “The case for end-user programming of ubiquitous computing environments”, in *proc. of FoSER '10, ACM, New York, NY, USA, 2010, 167-172*
- [3] I.P. Cvijikj, F. Michahelles: “The Toolkit Approach for End-User Participation in the Internet of Things”, in: D. Uckelmann et. al. (Eds.) *Architecting the Internet of Things*, Springer, 2011, 65-93.

Please contact:

Josè Danado, Fabio Paternò
 CNR-ISTI, Italy.
 E-mail: danado@isti.cnr.it,
fabio.paterno@isti.cnr.it

3D Web Visualization for Real-Time Maintenance of Smart Buildings

by Daniel Gaston, Christophe Joubert and Miguel Montesinos

Modern buildings are equipped with a variety of building automation technologies and sensors which are controlled, monitored and managed through heterogeneous information systems. To date, these systems show limited interoperability and tend to operate in isolation. The future will require a more comprehensive and efficient operation of buildings facilitated by way of integrated and smart cooperative management and automation systems. Web-based open-source technologies can create 3D virtual representations of the real-time geolocated activities within a building. Such a system can simplify both the maintenance and operation by facility managers and application services to building occupants.

The European project Building as a Service – BaaS (ITEA2 – 12011, November 2013-October 2016) aims to create a reference architecture and service framework for future building automation (BA) systems with extensive integration capabilities. Based on open IT standards and open-source technologies, the BaaS framework enables easy creation of new cross-domain BA services and the largely automated integration of existing BA components in a cost

effective and flexible manner. Figure 1 describes the integration of value-added visualization services on top of the reference BaaS architecture.

One of the biggest challenges for building automation is the collaboration of new components and systems with existing devices and infrastructures. We contributed to this technical challenge by creating 3D building smart objects with web capabilities out of the real-

time sensing BaaS platform. Such features facilitate complex user-created domain oriented applications that can be used both indoors and outdoors with a georeferenced environment and efficient XML interchange [1].

This feature has been specified as a functional requirement for the operation phase of the BaaS architecture. This requirement states that the architecture shall support the development of services that use building data models (building geometry, location and data of rooms, floors, sensors).

Based on this requirement, a set of virtual data points were defined in order to provide data to value-added services. Table 1 presents an overview of the different virtual data points that were implemented in the BaaS framework to give support for visualization web services.

Data points were defined based on a semantic meta-information that was specified for all phases of the BA service life-cycle in order to improve extensibility and reduce engineering effort of the BaaS platform. More specifically, the defined virtual data points used a location ontology describing the relation between building elements such Building, Floor, and Room. This ontology made use of the extended SSN ontology (for sensors and actuators) and OWL-S (for services) proposed in [2].

On top of the virtual data points, we provided a service layer implemented as REST services in order to enable real-time visualization of the building elements and sensor information through a web browser. The 3D visualization is served by an open-source WebGL Globe (Cesium) that enables the




Figure 1: BaaS framework instantiated with value-added 3D visualization services.



Figure 2: 3D graphical and data relation view of Patraix Building's 4th floor.

Table 1: Virtual BaaS data points for value-added visualization web services.

es.buildingLocation	Building location data point that provides location data of a building within a group of buildings. It provides a unique identifier, a building name, an icon URL, a tooltip text, a location (lat, lon, alt) in double precision, and a number of building floors.
es.floorModel	Floor model data point that provides a model data of a floor within a group of floors. It provides a unique identifier, a floor name, a building identification number, a floor level, and a geometry file identification number.
es.floorGeometry	Floor geometry data point that provides the geometry and additional values of a floor. It provides a unique identifier, a floor geometry file name, a floor geometry file size, geometry data, and the creation timestamp.

description of dynamic scenarios in virtual globes and maps. Basic architectural elements, as well as location and typology of sensors are represented in a geographical context. A specific JSON schema was described in the Cesium Language (CZML) to describe property values changing over time, such as status and information of sensors that can be provided as a stream to web clients to view dynamic values of building elements. Other values transmitted from the BaaS platform to the visualization component are the vertex positions in cartographic degrees of all floor rooms as well as the navigational model between building, floor and rooms based on the location ontology. Among the functionalities provided to the facility manager is the dynamic link between the status and information of the 3D scene with the information panel. Hence, navigation can be done either through the panel's elements, such as the list of sensors, or through the 3D scene billboards anywhere on the building.

The results to date are very promising: we integrated the previous visualization services with many partners in order to compose various demonstration scenarios on top of the Building as a Service platform, namely Smart Booking Room, Maintenance application, and Feedback application. Trials were carried out with end-users in a real environment, the Social Building of the Valencia City Council at Patraix. The facility manager, the room booking manager, teachers and citizens booking and using the rooms for their own needs (foreign language classes, expositions, movies, presentations, yoga, etc.) were given the chance to use the 3D web user interface to visualize directly the configuration and sensing information of a room (smart booking room and maintenance application), and interact with the facility manager if the room configuration fit or not their activities and preferences (feedback application). As a result, this work enabled to increase the usability of the smart building (increase room availability by 50%) and enhance the user experience for IoT services (with a move from manual management to a digital system by Internet on building actuators). Figure 2 illustrates the 3D graphical and data visualization interface part of the facility manager's maintenance application on the Patraix public building.

We plan to deploy and test the results of this project in the maritime sector [3] for global situational awareness in rescue, calamity and inspection operations in port infrastructures, and in the Smart Cities platform to scale the visualization capabilities to more public buildings.

In our research, we collaborated with several SMEs, Universities and Research Centres, including: Siemens AG; Materna; Kieback&Peter; TWT; Fraunhofer; TUM; TU Dortmund; Universität Rostock (Germany); Everis; UPV (Spain); KoçSistem; BOR; Defne; Smartsoft (Turkey); X.com; Masaryk University; and MDS (Czech Republic).

This research is also part of a horizontal task force with other ICT Future Internet projects - such as FIWARE, in particular FI-CONTENT 2 - that deals with building new innovative applications and services for every-day working and living environments, by making 'physical world information' easily available for smart services. Our work is partially supported by the Spanish MEC INNCORPORA-PTQ 2011, MiTYC TSI-020400-2011-29, and FEDER programs.

Link:

<http://www.baas-itea2.eu>

References:

- [1] M. Jung, J. Weidinger, W. Kastner, A. Olivieri: "Building Automation and Smart Cities: An Integration Approach Based on a Service-Oriented Architecture, pp. 1361-1367, IEEE Computer Society (2013).
- [2] J. Fons, D. Gaston, C. Joubert, M. Montesinos: "Internet of Things Applications for Neighbourhood Embedded Devices", ERCIM News 2014(98) (2014).
- [3] C. Joubert, M. Montesinos, J. Sanz: "A Comprehensive Port Operations Management System", ERCIM News 2014(97) (2014)

Please contact:

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

COMPOSE: An Open Source Cloud-Based Scalable IoT Services Platform

by Dave Raggett

Advances in electronics and communication technologies are stimulating the growth of low cost connected sensors and actuators. There are many potential application areas, for instance: home automation, security, healthcare, smart grids, integrated transport systems, and next generation manufacturing. To date, research in the area of the Internet of Things (IoT) has largely focused on the sensors, actuators and the communication technologies needed for long battery life and efficient use of spectrum etc. rather than on what's needed to encourage applications and services. Current IoT products exist in isolated silos. To fully realise the benefits, we will need scalable solutions for managing the vast amounts of data, and open standards that enable open markets of connected services.

The COMPOSE project is developing an open source cloud-based platform for IoT services. The run-time is highly scalable and implemented on top of CloudFoundry and Couchbase. This is complemented by the developer portal which supports discovery, registration, composition and deployment of services.


COMPOSE models services in terms of message streams, where each message corresponds to either a sensor reading or an actuator update, and is structured as one or more data channels. Sensors and actuators are virtualised in the COMPOSE platform as 'service objects'. These define the message streams

obtained from the sensors, or passed to the actuators associated with a given service object.

Services can define composite streams in terms of dynamic queries over other streams, for example, all temperature sensors within 1.5 Km of a given location. Services can also define new streams as a transformation from another message stream, for example, to transform the physical units for the data samples, or to smooth data samples. These mappings can be defined as simple expressions or more generally with scripts written in the JavaScript

programming language and executed using the Chrome V8 engine.

COMPOSE services are cloud-based. Each service object is associated with a 'web object' at the network edge. The web object provides an abstraction of the physical sensor or actuator, hiding the details of how these devices are connected. Web objects communicate with the COMPOSE platform using structured data represented with the JavaScript object notation (JSON). There is a choice between using HTTP or WebSockets. The latter is appropriate when the web object is behind a firewall that prevents the COMPOSE platform from opening HTTP connec-



tions with the web object. Here is an example illustrating how to push data to a service object:

```
PUT http://testbed.compose-
project.eu/thngs/<ServiceObjectID>/stream
s/<Strea...
{
  "lastUpdate": 194896802,
  "channels": [
    {
      "unit": "degrees",
      "type": "numeric",
      "name": "longitude",
      "current-value": 24.428239
    },
    {
      "unit": "degrees",
      "type": "numeric",
      "name": "latitude",
      "current-value": 1.3428239
    }
  ],
  "customFields": {}
}
```

Careful attention has been paid to security. COMPOSE uses encrypted sessions for protecting data exchanges. Data owners can set access control policies that limit who can access their data. Static analysis of service stream mappings is used to determine the provenance of derived data streams in order to comply

with the access control policies for the originating data owners. To make this static analysis practical, constraints are placed on the use of JavaScript language features. The analysis yields 'contracts' that are designed for use by the COMPOSE run-time system.

COMPOSE applications are able to access message streams via a RESTful interface over HTTP or WebSockets. This gives developers a choice of implementation technologies, e.g. HTML5 for the Open Web Platform, or as native applications on iOS or Android. OAuth2 provides the basis for application users to grant access to their data on the COMPOSE platform.

The COMPOSE developer portal supports a graphical authoring tool based upon an extended version of Node-RED. The portal enables developers to search for services and streams matching the query provided by the developer. Developers can also solicit recommendations based upon rankings provided by other developers.

The COMPOSE project is conducting a number of pilots to evaluate the utility of the framework in real world settings. The smart retail pilot features sensors within a supermarket that track the loca-

tion of shopping trolleys as customers move about the store, and which can later be combined with information on purchases. This data can be used to assist product positioning on supermarket shelves. The smart city pilot focuses on tracking free parking spaces at the Rovira i Virgili University in Catalonia. The smart territory pilot focuses on supporting skiers in the Trentino region of Italy. It combines sensor data with crowd sourced information covering the length of lift queues, the quality of ski centre facilities, and points of interest. Users can see their friends' locations and exchange notifications with friends, which facilitates social interaction.

Links:

<http://www.compose-project.eu/>
<http://www.cloudfoundry.com/>
<http://www.couchbase.com>
Chrome V8:
<https://developers.google.com/v8/intro>
JSON: <http://json.org/>
WebSockets protocol:
<https://tools.ietf.org/html/rfc6455>
OAuth2: <http://oauth.net/2/>
Node-RED: <http://nodered.org/>

Please contact:

Dave Raggett, W3C
E-mail: dsr@w3.org

An Integration Gateway for Sensing Devices in Smart Environments

by Michele Girolami, Francesco Furfari and Stefano Chessa

Smart Environments, and in particular Smart Homes, have recently attracted the attention of many researchers and industrial vendors. The proliferation of low-power sensing devices requires integration gateways hiding the complexity of heterogeneous technologies. We propose a ZigBee integration gateway to access and integrate low-power ZigBee devices.

Smart Environments, and in particular Smart Homes, have recently gained the attention of many researchers and hardware vendors. An increasing number of sensing devices, whose price is rapidly decreasing, are available on the market. Although such devices are becoming familiar in Smart Homes, user acceptance is limited by the fragmentation of the market. Heterogeneous technologies do not integrate seamlessly into a Smart Home; rather, each vendor offers its private vertical solution. ZigBee and OSGi play a predominant role in this scenario.

ZigBee offers a service-oriented approach for low-power devices, with the unique feature of defining a variety of profiles that standardize the functionalities of several classes of device (e.g., home automation, health care, smart energy). OSGi offers a component-based execution platform facilitating the deployment and management of software units.

The Wireless Network Laboratory at ISTI-CNR funds the ZB40 project [1], ZigBee API for OSGi Service

Platform. ZB40 faces the problem of providing an easy access to low-power sensing devices based on the ZigBee stack. It relies on the OSGi execution environment and it meets three basic requirements: (i) it provides a rich and flexible gateway for the ZigBee network, (ii) it extends the OSGi framework with an open mechanism to integrate ZigBee standard with a service-oriented approach and (iii) it defines an integration layer in order to access the ZigBee network by using other technologies.




Figure 1: The ZB40 architecture.

ZB40 is designed by considering a typical use-case of Smart Homes. A new ZigBee device is installed at home, for instance a standard smart plug for monitoring energy consumption. A user can discover the ZigBee device as soon as it is plugged in, without the installation of any specific custom driver; rather, all ZigBee equipment is recognized and integrated autonomously.

ZB40 is based on three layers, namely: the Access, Abstraction and Integration layers as shown in Figure 1. The Access layer directly communicates with the ZigBee network by means of a network

adapter (for example, ZigBee USB dongle). The Access layer interacts with the network in order to detect new nodes, removal of nodes, and to detect relevant events concerning changes in the network topology. The Access layer is designed to be vendor independent. In fact, it does not constrain the end-user to adopt a specific ZigBee network adapter to interact with the network; rather the Access Layer implements a general-purpose solution. The Abstraction layer is designed with the goal of adding more functionalities to the nodes detected. In particular, this layer refines the nodes discovered by the Access Layer with a profile-based

node. In this way ZB40 fully supports the profile approach followed by the ZigBee Alliance. The Integration layer exports the ZigBee functionalities to one or more target networks. Notable examples of exporters are the UPnP network, the REST paradigm or a successful experience of robotic ecologies such as the GiraffPlus robot [2].

WNLab has developed an integration layer for two core enabling and widely accepted technologies: UPnP and REST as well as the integration with two EU projects: the universAAL project and the GiraffPlus project. In Figure 2 we show the integration layer for the UPnP network. The Access Layer recognizes an OnOff Binary Light device installed at home. The Abstraction layer refines the device by adding the functionalities defined by the Home Automation Profile. Finally the UPnP Integration Layer maps the OnOff Light as a UPnP Binary Light. In this way, an UPnP client can discover and interact with the ZigBee Light from the UPnP network.

ZB40 has attracted the attention of several ICT companies as well as several research centres. ZB40 is an active open-source project with a growing community of users and developers.

Link:
<http://zb4osgi.aaloo.org>

- References:**
- [1] F. Furfari F et al.: “A Service-Oriented ZigBee Gateway for Smart Environments”, to appear in Journal of Ambient Intelligence and Smart Environments, JAISE Vol.6: 691-705, 2014.
 - [2] M. Girolami et al.: “The Integration of ZigBee with the GiraffPlus Robotic Framework”, Evolving Ambient Intelligence Vol.413:86-101, 2013.

Please contact:
 Francesco Furfari and Michele Girolami
 ISTI-CNR, Italy
 E-mail: {francesco.furfari, michele.girolami}@isti.cnr.it




Figure 2: The UPnP Exporter.

Fighting Networking Heterogeneity in the Internet of Things

by Elias Z. Tragos, Vangelis Angelakis and Stefanos Papadakis

The Internet of Things (IoT) aims to interconnect large numbers of heterogeneous devices to provide advanced applications that can improve our quality of life. The efficient interconnectivity of IoT devices can be achieved with a hybrid Cloud Radio Access Network (Cloud-RAN) and Software Defined Radio (SDR) framework that can overcome the heterogeneity of devices by seamlessly adapting to their communication technology.

The Internet of Things (IoT) presents itself as a promising set of technologies that enable the efficient and seamless interconnectivity of large numbers of devices, allowing them to be discovered and accessed by the services in an abstracted way. The IoT is considered to be one of the key technologies for 5G communications and a basic technological enabler towards the realization of ‘Smart Cities’. Smart Cities is a relatively new term for a new generation of cities that use new technology to drive competitiveness, sustainability, economic growth, energy efficiency and improving Quality of Life (QoL).

The successful deployment of Smart Cities calls for a unified ICT infrastructure to support the diverse set of applications for urban development. Up until now, most of these devices used standard Wireless Sensor and Actuator Networking (WSAN) technology based on IEEE 802.15.4. Recently, however, the particular demands of IoT applications have prompted the adoption of more advanced communication technologies, e.g., IEEE 802.11 or 4G/LTE which are capable of providing a higher quality of service. A networking architecture capable of allowing the efficient and secure interworking of a large number of heterogeneous devices, avoiding the interference between them is a quite challenging research topic. One has to consider the differences between the various communication technologies and protocols, in terms of carrier frequency, bandwidth, modulation and coding schemes, packet structures, packet sizes, etc. For example, a smart phone with a 3G network access can only exchange information with a ZigBee sensor via a third device that translates all necessary protocols and sends back the sensed data.

Furthermore, when large numbers of IoT devices are deployed in urban envi-

ronments where the ISM bands are already overcrowded the interference from external sources (other wireless networks) can severely impact network performance, since WSANs are very susceptible to other wireless transmissions. But how can we allow the devices

ized server units that manage a pool of base station resources, (ii) the distributed radio units (RUs) or SDR-based RUs (SRUs) that are located at remote sites and create wireless cells providing access to heterogeneous users and (iii) a high bandwidth low-latency link




Figure 1: Example of SDR-based heterogeneous C-RAN architecture.

to really ‘talk to each other’ (the basic motto of the IoT) without increased signalling in the network and long delays? How can we be sure that when an alarm is raised by an IoT device, this information will be prioritized and sent immediately to the respective target device without being lost due to collisions or interference?

A reliable architecture interconnecting heterogeneous IoT networks can be realized by integrating the benefits of Cloud-RAN [1] and Cognitive Radio [2] approaches. The Cloud-RAN architecture consists of (i) a cloud of central-

ized server units that manage a pool of base station resources, (ii) the distributed radio units (RUs) or SDR-based RUs (SRUs) that are located at remote sites and create wireless cells providing access to heterogeneous users and (iii) a high bandwidth low-latency link between the centralized unit and the remote sites. The advantage of this architecture compared with the standard C-RAN architectures that are proposed for LTE is the use of SDR-based radio units at remote sites that are able to handle simultaneous connections of various communication technologies. In this respect, only one SRU at each site has to be installed, avoiding the costs of installing multiple units for each communication technology. In such an architecture, the network management decisions (i.e., spectrum assignment, routing, scheduling, etc.) can be taken both locally at each SRU when they

relate with the users on each cell or centrally by the centralized server units when interaction with the neighbour cells are required or for optimizing local decisions. An example of this architecture is shown in Figure 1, where the different colours of the SRU cells denote different communication technologies and frequencies.

The software re-programmability of SRUs enables any updates regarding the communication protocols, installation of additional technologies, implementation of new networking standards, etc. to be done easily, saving significant costs associated with adding new hardware at each remote site. The ability of SDR to simultaneously handle different communication technologies can facilitate the virtualization of the remote units in a way that each unit will be seen as different Virtual Base Stations (VBS), each handling different communication technologies, e.g., VBS1 will handle IEEE 802.11, VBS2 will handle 3G, while VBS3 will handle IEEE 802.15.4 – based IoT devices.

The centralized server unit is able to perform an optimized management of the available network resources since it has a global view of the available resources at each SRU, and can reconfigure them easily at run-time to perform, for instance, traffic offloading when one unit is overloaded or to

change frequencies when there is increased interference. The architecture can also easily adapt to non-uniform traffic via the load balancing capability of the distributed pool of base stations. This pool can also share signalling, traffic and channel occupancy information of the active users in order to optimize the radio resource management decisions. Spectral efficiency is improved via the cognitive radio mechanisms for intelligent management of spectrum resources that can be applied on the SRUs together with joint processing and scheduling.

This architecture can have many advantages for IoT, owing to its ability to optimize spectrum access for the multiple heterogeneous devices, prioritizing access to network resources according to the service request in both a centralized and a distributed manner. For example, if two devices are in the same area, even if they use different communication technologies, they will exchange their data directly through the local remote unit reducing unnecessary signalling. Another advantage is the cost-efficient deployment of such an architecture, which only requires centralized management and operation, while the installation of new remote units can be done with a simple SDR device and the required software. On the road to 5G, where IoT is a basic technological pillar, such an architec-

ture could greatly help mitigate the heterogeneity of the involved devices and their efficient coexistence. Since the architecture is partly centralized, it can easily be integrated with existing IoT middleware platforms, playing the role of ‘communication manager’ - a functional entity proposed in many IoT architectures [3].

This work has received funding from the European Union’s Seventh Framework Programme (FP7/2007-2013) under grant agreements no 609094 and 612361.

References:

- [1] “C-RAN The Road Towards Green RAN”, China Mobile Research Institute White Paper Version 2.5, Oct 2011.
- [2] E.Z. Tragos, V. Angelakis: “Cognitive radio inspired M2M communications”, 16th International Symposium on Wireless Personal Multimedia Communications (WPMC), pp.1,5, 2013.
- [3] A. Bassi, et. al: “Enabling Things to Talk”, Springer, ISBN: 978-3-642-40402-3 (Print) 978-3-642-40403-0 (Online).

Please contact:

Elias Z. Tragos
FORTH-ICS, Greece
E-mail: etragos@ics.forth.gr

Tectons: Towards a Generalised Approach to Programming Systems of Systems

by Geoff Coulson, Andreas Mauthe and Markus Tauber

The world’s computing infrastructure is becoming increasingly differentiated into autonomous sub-systems (e.g., IoT installations, clouds, VANETs), and these are often composed to generate value-added functionality (systems of systems). But today, system composition is carried out in an ad-hoc, system-specific, manner, with many associated disadvantages. We need a generalized system-of-systems-oriented programming model that allows systems to be composed by application-domain experts, not just systems programmers. Furthermore, composition should occur in a principled way that generates well-understood compositional semantics and behaviour.

Today’s distributed computing environment is becoming increasingly diverse and differentiated in nature, to the extent that the world’s computing infrastructure is now far removed from the traditional picture of desktop PCs connected via fixed IP networks. The picture today includes: IoT infrastruc-

tures such as smart cities and buildings; environmental sensor and actuator networks using non-IP protocols; cloud systems based on clusters; ad-hoc networks such as MANETs and VANETs; and virtualized systems supported by network overlays. Furthermore, these various ‘systems’

increasingly need to interact with and respond to each other, typically in a dynamic on-demand manner. For example, WSNs need back-end clouds to process sense data, VANETs need to interact with smart cities when upon entry to the city, and overlay-based systems need to maintain resilience

properties when their underlying IP environment changes.

This is all leading to a world-view in which we need to be increasingly concerned with composing systems to build systems of systems [1]. Although this fact is becoming recognised in many research communities, surprisingly little work has been done on programming environments to facilitate system-of-systems composition. Instead, system composition is carried out using ad-hoc interfaces, and relies on detailed knowledge of the internals of the systems being composed. We argue that a generalised programming and modelling approach is urgently needed to enable the full fruition of the emerging system-of-systems world.

We are currently developing such a system-of-systems-oriented programming model. A key aspect of our approach is to assume that systems are self-contained, and that they interact and compose opportunistically. In this way, we see systems of systems emerging spontaneously as a result of opportunistically-arising, mutually-beneficial, time-bounded alliances between systems that dynamically discover potential partner systems in their environment. Examples of such opportunistic interactions/compositions include: i) exchange of traffic information between passing vehicles; ii) isolated islands of connectivity discovering a delay-tolerant overlay network through which they can interact; iii) opportunistic cross-layer optimization involving layers on different systems that happen to come in range of each other; and iv) a networked team of rescue workers dynamically interfacing with a local hospital's systems. The general pattern is one of loosely-coupled interaction between autonomous systems triggered by relevant events (e.g., proximity).

Our programming and modelling approach is based on a first-class programmatic abstraction of a 'system', which we refer to as a tecton. Essentially, a tecton is a distributed abstraction/representation of a self-contained, potentially-opportunistically-interacting, multi-node distributed system. It is intended that tectons are used uniformly to abstract the full range of 'systems', whether wired/wireless, fixed/mobile, large/small, static/

dynamic, low-level/high-level – e.g., user groups, MANETs, clusters, clouds, overlays, VPNs, sensor networks, or even (on a more 'micro' scale) devices that can plug-and-play with other devices.

The purpose of the tecton abstraction is to manage opportunistic interaction. This is achieved by providing tectons with programmatic condition-action rules that trigger when the tectons (or, more precisely, some of their nodes) come into contact, and determine under what circumstances and in what manner an interaction or composition should occur. This applies in both a horizontal and a vertical sense, involving (respectively) interactions between peer systems, and cross-layer interactions.

As an example, we might program a MANET tecton by (essentially) saying "IF any of the nodes in this MANET tecton come into contact with an 802.11 network tecton, THEN all its nodes should reconnect to their email servers via this node". By extrapolating from this example of vertical composition, it can be seen that condition-action rules enable the tecton programmer to specify, in a very general way, what should be done when one system encounters another. In summary, we envisage the whole process of developing, deploying and managing systems of systems as amounting to a process of defining horizontally- and vertically-composable tectons along with their associated condition-action rules.

In supporting the tecton abstraction, we have defined a layered architecture consisting of a Domain Specific Language/Ontology layer over a distributed tecton runtime layer. The upper layer is motivated by the fact that a single set of programming language concepts is unlikely to be acceptable, given the intended generality of the tecton approach. Instead, we look to support multiple programming tools that capture concepts of relevance to the domain of application. For example, a systems programmer working with network-level tectons might expect a DSL based on library calls embedded in C; whereas a building manager working with tectons that represent things like "all the windows on this floor" or "the group of today's visitors" may prefer a scripting or graphical approach. Similarly, the extreme variance in

domains of application dictate that the basic information concepts of interest to the programmer will vary according to the domain of application. Our approach here is to employ ontologies to structure these concepts.

The distributed tecton runtime offers an API that enables client code to carry out fundamental operations such as creating/destroying tectons, managing their node membership, setting the criteria for inter-tecton "contact", and coordinating the consequent action (e.g. interaction, composition, etc.). For example, setting the criteria for "contact" is achieved by providing the runtime with a predicate over ontology-derived <name, value> pairs. The runtime monitors these values in its distributed environment and triggers the corresponding action when the predicate "fires".

We are at the beginning stage of the development of the tecton concept, but are currently applying it in a CHISTERA project that investigates interoperable overlay networks and also in an IoT-based environment involving both WSNs and cloud infrastructures.

Link:

CHISTERA:

<http://www.chistera.eu/projects/dionasys>

Reference:

[1] M. Maier: "Architecting Principles for System of Systems", *Systems Engineering* 1 (4): pp 267–284, 1998.

Please contact:

Geoff Coulson, Andreas Mauthe
Lancaster University, UK
E-mail: geoff@comp.lancs.ac.uk,
a.mauthe@lancaster.ac.uk

Markus Tauber

AIT Austrian Institute of Technology,
Austria

E-mail: markus.tauber@ait.ac.at

BETaaS: Building the Future Platform for Development and Execution of Machine-to-Machine Applications

by Carlo Vallati, Enzo Mingozzi and Giacomo Tanganelli

Recently proposed IoT platforms are typically based on centralized and cloud-centric infrastructures. The BETaaS project aims at designing a platform for the execution of M2M applications in a distributed runtime environment running close to the physical environment where sensors and actuators are deployed. The main goal is to move the intelligence to the edge in order to allow localized, content-centric, and timely processing of smart objects data. The platform facilitates the integration of existing IoT systems, while providing software developers with a high-level, content-centric abstraction to access smart objects' resources.

Recent technology advances are evolving the objects we use in our daily lives into smart objects, i.e., regular objects empowered with communication and computational capabilities. Smart home automation solutions (e.g., smart thermostats, smart lights) and smart health products (e.g., wearable sensors) commercialized today are only a few examples representing the dawn of the forthcoming Internet of Things (IoT) revolution [1].

Interoperability among heterogeneous IoT systems is considered a key issue in this scenario. Current IoT solutions are typically vertical systems that are designed to serve a specific application acting in isolation with no or limited cooperation. Instead, a horizontal approach is needed to design future IoT platforms in order to facilitate the creation of a converged infrastructure providing diverse applications a seamless access to heterogeneous smart. This is the clear trend also followed by ongoing standardization activities related to IoT, e.g. oneM2M.

However, recently proposed horizontal solutions are typically based on centralized architectures where the intelligence to process and harmonize the data to/from heterogeneous IoT infrastructures is concentrated in the cloud. Although such solutions inherit the benefits of a cloud-based infrastructure, as a matter of fact they are not the best choice for many classes of IoT applications. Machine-to-Machine (M2M) applications, for example, are characterized by a limited scope in time and space (data needs to be processed only when and where generated) as they require simple and repetitive, closed-loop, but often highly time-sensitive interactions. Processing all data in the

cloud is generally inefficient in these use cases, and may be even unfeasible if the case requires stringent latency guarantees to process and react to collected sensor data [2].

BETaaS, Building the Environment for the Things-as-a-Service, is a European project co-funded by the European Commission under the 7th Framework Programme, which aims to overcome the limitations of cloud-centric IoT architectures. This is accomplished

through the design and development of a horizontal platform that leverages a distributed architecture to move the intelligence to the edge, i.e., close to the physical environment where sensors and actuators reside, in order to ensure timely application processing. The BETaaS platform, is designed to run on heterogeneous devices, called gateways, such as home routers or set-top boxes. The execution of M2M applications is then supported through a tight interaction among multiple gateways that form




Figure 1: BETaaS use case example.

a distributed runtime environment called a local cloud. On top of the local cloud, the platform offers M2M application developers a content-centric service-oriented interface, named Things-as-a-Service (TaaS), exposed to access smart objects connected to the platform regardless of their location and technology, cutting software development time and enabling code reusability.

The BETaaS platform, released as open-source software, is based on Java OSGi, a framework specifically designed to build modular platforms with a highly dynamic structure. The BETaaS architecture is based on a layered structure to ease the integration of existing systems and platform expandability. At the bottom, an Adaptation Layer allows integration of existing M2M systems through the implementation of plug-ins that provide a common interface to access the functionalities of the physical system. At the core, the TaaS layer implements the functionalities that allow gateways to cooperatively share resources, thus realizing the concept of local cloud. The result is the TaaS interface exposed to software developers to implement new services and external applications. To this aim, an additional Service Layer is included to host extended services, i.e., services that can run natively on the platform to implement custom control logic and can be installed at run-time exploiting the modular structure offered by OSGi. Eventually, the Service Layer exposes a unified RESTful interface, exploited by external applications running on external

devices to interact with smart objects and extended services. In addition to basic functionalities, the platform offers built-in support for several non-functional requirements aimed at supporting development of M2M applications and extended services: Context Awareness, Quality of Service, Security, Big Data Management and Virtualization. For an exhaustive description of these platform functionalities, the interested reader can refer to [3].

As an example of use case, consider a smart home environment with a set of already deployed vertical M2M systems: an alarm system equipped with presence sensors for surveillance, an environmental control system including temperature sensors for heating and cooling control as well as light switch actuators, and a garden watering system with humidity sensors. In this scenario, the BETaaS platform can be deployed on four gateways that cooperates to implement an all-in-one horizontal IoT system that enables the development of applications leveraging on smart objects working in isolation in the original design. For instance, an extended service that enhances the environmental control system can be installed to exploit humidity and presence information, e.g., turning down the air-conditioning when a window is opened as evidenced by the magnetic sensor part of the alarm system. External applications interacting with the platform and running on external devices can be deployed on a smartphone to expose a uniform control interface to end users.

The BETaaS platform achieves full separation between applications and the sensing and actuating infrastructure, allowing the development of applications by third-party developers. Generic applications ready to run in any BETaaS instance can be developed. This key feature can be exploited to trigger the creation of an M2M application market, from which end users can download and install applications on their private deployments.

Links:

BETaaS website:

<http://www.betaas.com/>

BETaaS project website:

<http://www.betaas.eu/>

BETaaS github repository:

<https://github.com/BETaaS/>

[BETaaS_Platform-Tools](#)

References:

[1] Atzori, Luigi, Antonio Iera, and Giacomo Morabito: “The internet of things: A survey”, *Computer networks*, 2010.

[2] Abdelwahab, at al.: “Enabling Smart Cloud Services Through Remote Sensing: An Internet of Everything Enabler”, *IEEE In-ternet of Things Journal*, 2014.

Please contact:

Carlo Vallati, Enzo Mingozi,
Giacomo Tanganelli
Department of Information
Engineering, University of Pisa
{c.vallati, e.mingozi,
g.tanganelli}@iet.unipi.it

Pro-IoT: Improving the Productivity of Companies by Using the IoT

by Heikki Ailisto, Nadine Pesonen, Pertti Raatikainen and Jonathan Ouoba

Internet of things (IoT) and the related Industrial Internet are recognized as one of the most significant technology-driven disruptions of the coming ten years. VTT Technical Research Centre of Finland has chosen IoT as a strategic area for research, development and innovation projects.

The Pro-IoT spearhead programme focuses on three technology areas, namely sensing technology, connectivity, and data-analysis with data security. These areas emanate from the needs expressed by European and especially Finnish companies, as well as on the strong competencies found in our

research centre. Each area is delineated in three domains that represent challenges for our societies: industrial asset management, connected health and digital society. In addition to the technological perspectives, the business aspects are also investigated, through a dedicated module, hence making it possible

to point out the relevance of the proposed solutions for companies. The complete structure of the programme is presented Figure 1.

In order to sustain living standards in Europe, an increase in the productivity of companies appears to be vital. As

such, the structure of the Pro-IoT programme is intended to provide the appropriate framework to initiate and maintain a productivity leap corresponding to the potential of IoT-based systems.

Research Highlights

The development of the Pro-IoT programme has resulted in many research projects. These projects have yielded significant achievements in various domains. A notable example, in the field of sensing, concerns wireless and battery-less devices, which represent a kind of Holy Grail. Such a paradigm would enable measurements in demanding conditions with low maintenance cost and without hindering production. VTT has developed a solution for measuring temperature, inclination, humidity, strain and other characteristics from machines and inside structures (e.g. walls) at distances up to 10 m, without batteries in the sensor. The principle is based on powering the sensor with radio frequency energy waves and using the return transmission to read out the resulting data (see Link 'Zero Power Sensor'). The IPR protected technology allows for several individually identifiable sensors to be operated in the same space.

Another relevant example concerns autonomous vehicles and machines with communication needs. A solution for communication in demanding conditions, namely a cluttered harbour environment, has been developed for and tested in Port of Singapore. The key issue was how to ensure reliable communication with predictable maximum delay between moving machines for their safe operation. The solution involved using two or more communication networks seamlessly to secure communication between piles of containers and other 'non-radio-friendly' objects.

At a more conceptual level, other research activities were intended to address the issues related to the multi-technology environment of mobile devices. Indeed, many of these devices (mobile phones, sensors, wearable objects) are endowed with multiple wireless technologies. This offers opportunities in terms of services and applications to deploy, provided that solutions are designed to effectively




Figure 1: VTT productivity with IOT program structure.

manage the multi-technology context. Therefore, a platform has been developed for energy-efficient communications between mobile devices that identifies the most appropriate combination of technologies for different exchanges [1].

Impact on Industry and Society

The disruption caused by Industrial Internet is comparable to earlier industrial revolutions initiated by steam, electricity and computers. Indeed, General Electric predicts extra growth globally worth 10 to 15 trillion US dollars during the next 15 years. This will depend largely on being able to leverage industrial internet technology to the full [2]. As such, the Pro-IoT programme contributes to this evolution by providing companies with appropriate tools to improve the performance of their industrial processes. As shown by the examples in the previous section, the emphasis is on realistic solutions adapted to a challenging environment, safe autonomous systems while targeting a reduction in the operating costs.

The Industrial Internet will not only increase productivity in industry and firms. By adapting its models to urban areas, it will also affect public services with the emergence and the consolidation of smart cities. The everyday lives of all of us as employees, consumers and citizens will benefit from it with more personalized and efficient end-user services.

Challenges Ahead

The Industrial Internet is not happening by itself; it requires coordinated and

intensified research effort involving the key stakeholders. We need to further address technological questions such as the need of more energy-efficient communication, the congestion of the Internet with the deployment of tens of billions of new nodes (sensors and smart objects). In addition, we have to find sustainable answers to the recurrent questions of information security and privacy in this new context. The safety aspect of the operations regarding autonomous machines, the issues of data ownership, as well as plausible and fair business models also represent major concerns.

Links:

The Pro-IoT spearhead project:
<http://www.vttresearch.com/impact/new-innovations-are-here/innovation-programmes/productivity-with-internet-of-things>
 Zero Power Sensor:
<http://knowledge.vtt.fi/ZeroPowerSensor>

References:

- [1] S. Chaumette and J. Ouoba: "A Multilevel Platform for Secure Communications in a Fleet of Mobile Phones", in proc. of MobiCASE 2014
- [2] P. Evans and M. Annunziata: "Industrial Internet: Pushing the Boundaries of Minds and Machines", 2012, http://www.ge.com/docs/chapters/Industrial_Internet.pdf

Please contact:

Ailisto Heikki
 VTT Technical Research Centre of Finland
 E-mail heikki.ailisto@vtt.fi

MusicBricks: Connecting Digital Creators to the Internet of Music Things

by Thomas Lidy, Alexander Schindler and Michela Magas

A new Horizon 2020 Innovation Action project fosters the uptake of music technology research and acts as an accelerator for bringing creative music tech ideas to market, while providing a mesh of technological 'bricks' to create an 'Internet of Music Things'.

Music is arguably the most pervasive of the performing arts. It has the power to enlighten us, cheer us up, bring us down and most importantly - bring us together. Substantial resources and personal energy have been invested into music technology research in recent decades. Despite the global success of European music tech companies such as Spotify, Deezer, Soundcloud and last.fm, the value chain from academic research to SMEs and large music industry corporations remains fragmented, preventing successful application of the results of innovative research. Yet, creative SMEs have been identified as the primary catalysts for innovation, and they can benefit vastly from innovative music technology research.

The Horizon 2020 Innovation Action 'MusicBricks' has been initiated with the aim of capitalizing on the current European leadership in successful music technology companies and secure a direct route from ICT academic research to the application and exploitation by SMEs and major industry players worldwide. The mission of the project is to foster an exchange between the music ICT experts from academic and national research organizations, and digital makers and content creators from creative industries SMEs.

Despite the remarkable progress in music information research, technologies provided by various projects are often difficult to apply owing to poor communication. Research results are typically scattered among institution Web pages and documentation is often neglected. Less user-friendly interfaces and steep learning curves hinder uptake by industry (even more so for SMEs and, in particular, one-person enterprises such as the typical music maker or app developer). In addition, corresponding technologies are frequently lacking proper engineering to facilitate interoperability with other technological components. Thus, complicated workflows

and little connectivity hinder the creation of larger systems that operate as a mesh of technological bricks connecting physical and tangible hardware and available music technology software. Yet this interoperability between tangible and wearable devices and the well-researched software components for music analysis and processing is heavily needed to make the 'Internet of Music Things' a reality [1].

The MusicBricks Project

MusicBricks responds to these problems, acting as a connector between research and industry, by identifying, wrapping and making existing tools available in easily usable formats. These tools, or 'bricks', shall require little learning, offer themselves for mash-ups, and are easily adopted and deployable by the creative SME digital makers and content creators. The project will foster interoperability by providing state-of-the-art connectors and interfaces to these technological building blocks. The resulting plethora of components - software endpoints, Web APIs, physical, tangible and wearable devices - are connected through and live on the Internet, realizing an 'Internet of Music Things', where many different kinds of analysis, processing, sensing, actuating or synthesis are taking place on different machines.

The consortium consists of the design research and innovation lab Stromatolite, the Institut de Recherche et Coordination Acoustique/Musique (IRCAM), the Music Technology Group of the Universitat Pompeu Fabra, the Fraunhofer IDMT Institute, the Vienna University of Technology, and is led by Sigma Orionis.

In the first step the project will provide the 'bricks' for creative makers: software tools and APIs, such as rhythm feature analysis and music similarity by Vienna University of Technology, key and tempo detection and melody and

bass transcription from Fraunhofer, melody curve extraction and API access to the freesound.org sample collection by Universitat Pompeu Fabra. On top of these APIs, graphical interfaces (GUIs) such as PlaySOM and Sonarflow offer semi-automated grouping and clustering of music while other visual frontends shall be easily used as well in conjunction with the tools and APIs. The third category of bricks is a new generation of Tangible User Interface (TUI) such as a wireless micro platform for motion sensing provided by IRCAM [2] This is complemented by portable computing platforms such as the Arduino, the Raspberry Pi and the Axoloti, a microcontroller specifically designed to create free-form digital audio instruments of a new kind.

The Internet of Music Things

All these technological 'bricks' form the foundation for a co-creational space where people - musicians, hackers, digital makers, creative SMEs - build entirely new applications. By interconnecting individual nodes - the provided API endpoints, network-enabled devices, sensors, microcomputers etc. - these bricks communicate with each other and start creating something bigger: a mesh of connected 'musical components' which eventually manifests as an 'Internet of Music Things'.

Events such as Hack days and Music Tech Fests fuel the creation of new mash-ups, hacks, audio tools, music instruments combined with unthought-of forms of application which has been demonstrated impressively in the past by innovations such as Siftables [3], compact devices with sensing, graphical display, and wireless communication that can be physically manipulated, recognize gestures and sense other nearby devices while able to interactively create audio; or performances such as the Brainwave Quartet in which music is performed live through the collective brainwaves of an ensemble



Figure 1: Music hackers interconnecting music technology to create interactive performances, new music interfaces or instruments.

wearing brain caps. Another great example is the “Wearable Axoloti Music Human Synthesizer”, a multi-user synthesizer where one user modulates the sound of another through touching, created with wearable electronics, conductive surfaces on T-shirts and Axoloti boards. It was awarded the ‘Internet of Music Things Hack Award’ at the Paris Music Tech Fest.

Events

MusicBricks both hosts and cooperates with a series of events, specifically the Music Tech Fest geared towards the creative developer community to actively test and utilize these tools, alongside others, to generate new ideas and create novel applications. The Music Tech Fest (Figure 1) has seen successful past events around the world and is a free, weekend-long event that provides an experimental and improvisational space where ideas are showcased, performed, demonstrated and discussed, while offering immediate assistance with the tools provided. It has a particular focus on interdisciplinarity: musicians meet hackers, researchers meet industry, artists and technologists come together.

Through this interdisciplinarity it enables a new dimension of creativity, leading to entirely novel ideas and applications. In past events creative hacker ideas have led to physical products and performances featuring new musical innovations such as a ‘Music Hat’ which converts brainwaves and head movements to music and various effects.

The most promising ideas and demonstrators generated at these events - whether a piece of software, a device, a new musical instrument, a performance or installation - will be supported by the MusicBricks Incubation Programme, which provides funding for residencies that enable creative makers to further develop their demonstrators to robust and market-ready prototypes and assists them through virtual and face-to-face collaboration. Feedback gathered during idea generation and incubation will assist in the consolidation and refinement of the available tools and interfaces. By exposing the resulting prototypes directly to big industry players and investors the project aims to reach the global market and gather

market feedback both for deployment of sustainable applications and to ensure long term impact on future research directions.

The next MusicBricks events will be Music Tech Fest Scandinavia May 29-31, 2015 in Umeå, Sweden, and a Music Hack Day at Sónar Festival Barcelona June 17-19, 2015. Until June 2016 MusicBricks will host a set of further events and support ideas through the incubation phase throughout this period. In order to ensure sustainability, the consortium is planning to set up a Music Tech Fund in conjunction with investment company Par Equity to further support and incubate the best prototypes and help take them to market.

Links:

Music Bricks Website:

<http://musictechfest.org/MusicBricks>

Music Tech Fest:

<http://musictechfest.org>

Results of previous Music Tech Fests:

<http://musictechfest.tumblr.com>

Videos from Music Tech Fests:

<https://www.youtube.com/user/MusicTechFest>

Reference:

[1] Music’s Internet of Things:

Heartbeats, Accelerometers...

Brainwaves?

<http://evolver.fm/2013/03/15/musics-internet-of-things-heartbeats-accelerometers-brainwaves/>

[2] Jules Françoise, Norbert Schnell,

Riccardo Borghesi, and Frédéric

Bevilacqua. Probabilistic Models for

Designing Motion and Sound

Relationships. In Proceedings of the

2014 International Conference on New

Interfaces for Musical Expression,

NIME’14, London, UK, 2014.

[3] D. Merrill, J. Kalanithi and P.

Maes. Siftables: Towards Sensor

Network User Interfaces. In the

Proceedings of the First International

Conference on Tangible and Embedded

Interaction (TEI’07). February 15-17 in

Baton Rouge, Louisiana, USA.

Please contact:

Thomas Lidy, Alexander Schindler

Vienna University of Technology

lidy@ifs.tuwien.ac.at,

schindler@ifs.tuwien.ac.at

Michela Magas

Stromatolite

michela@stromatolite.com

Deploying an NFC-based Inventory System: A Case Study

by Patrick Poullie, Thomas Bocek and Burkhard Stiller

NFC tags and transceivers are ubiquitous and well supported. Like many academic research groups, the Communication Systems Group CSG of the University of Zürich owns many physical devices, which are required for research and teaching. Traditionally, a printed, human-readable inventory and attached labels have been used to keep track of these devices. A new inventory approach was developed with the aim of simplifying the data acquisition using an NFC tag-based system supported by an Android application. This approach was implemented, tested, and used productively. However, due to technical difficulties - namely, a poor response of NFC tags on metallic material - the NFC-based inventory system did not simplify data acquisition and consequently, the NFC-based inventory system was changed back to the label-based system.

Near Field Communication (NFC) defines a communication channel for devices in close proximity (up to 10 cm) [1]. NFC has a wide range of applications as it can be used to transmit data between devices or between a device and an NFC tag. These tags are standardized – though some differences exist – and the data is stored in the standardized format NDEF (NFC Data Exchange Format) [2]. Depending on the type of tag, data from 48 Bytes (Ultralight) to 888 Bytes (NTAG216) can be stored. The very popular tag NTAG203 and its successor NTAG213 can store 144 Bytes of data. Most Android-based smartphones are equipped with an NFC transceiver, and the software support for NFC has been present since the Android version 2.3, which was released in December 2010. Apple also integrated the NFC technology into iPhone 6 and bases its mobile payment system on this technology, although not providing an open API. Thus, NFC reading and writing capable devices are widespread and ubiquitous.

Inventory System

The label-based Communication Systems Group CSG [3] inventory system for hardware acquisitions above a certain commercial value used a human readable, four character and four digit label. The label had to be printed with a specialized label printer and was attached to the device. The four digit number was used to identify a database entry, which contained information, such as device name, vendor, purchase date, and current holder. Retrieving inventory information about a device involved manually entering its number into a Web browser.

To avoid not only the need for a specialized printer but to also allow every group member to conveniently retrieve and change inventory information, a new tag-based inventory system was envisioned. Barcodes, QR-codes, and RFID tags were evaluated to label items. Bar- and QR-codes were not followed up for two reasons: first, creating human-readable tags would again imply the need for a specialized label printer, which would prohibit “decentralized”



Figure 1: App-based inventory system.

tagging of devices. Second, once these tags are created, information they store is static. This was a major drawback, because the new inventory system has to be able to operate offline, i.e., even without a database connection retrieving or updating inventory information of a device in physical reach had to be possible. Therefore, change of information on a tag and its synchronization later with the database was essential. Consequently, RFID tags had to be used.

Here, NFC tags qualified as the only option, as other RFID tags would require extra hardware for reading and writing tags. This would not only increase costs of the new approach but also reduce user convenience. Thus, since NFC is already supported by many smartphones and tablets, the NFC tags determined a valuable path.

Thus, the human readable inventory labels were replaced with NFC tags, to allow receipt and updating of inventory information upon scanning the tag with an NFC-enabled smartphone. To add a new device a database entry is created via the newly developed inventory smartphone app or via the inventory Web site. An NFC NTAG203 tag is placed on the device and an NFC-capable device used to store this entry's identifier and other essential information in those tags 144 Byte storage. Updating entries is simple: an NFC tag is scanned with the inventory app and all available information about this item is displayed immediately. This information (including location, room, and holder) can be updated within the app, and written to the tag and the database.

The smartphone inventory app was implemented in Java, while the backend was implemented in PHP communicating via HTTPS requests to the smartphone app. Authentication is done with LDAP, while all other data is stored in a MySQL database. Updating of device holders is the most frequently performed action, and the app's starting screen (cf. Figure 1) offers two buttons for this purpose. Upon selecting ‘Lend item’, the name of the new item holder is requested. Once entered, an NFC chip can be scanned, which, like the according data base entry, is updated

accordingly. The 'Return' button deletes the 'current holder' field of the subsequently scanned item (in the data base and on the NFC chip). 'Manipulate item' allows a database entry to be displayed and modified when the corresponding item is scanned. Thus, items can also be added or deleted here. With a 'More' click, additional information and a search function are displayed.

Evaluation

After the NFC-based inventory system was introduced, problems with using this approach in a practical environment soon became apparent. These problems were caused by the poor response of NFC tags, when any metal object was in their vicinity. As even some smartphones have metal cases, this resulted in a poor NFC communication reception. Thus, the evaluation recommends that smartphones with metal case are not used with an NFC-based inventory system. All following evaluations were performed with non-metallic smartphones.

The project identified a smartphone with a good NFC reception, a non-metallic case, and it was also expected to find suitable 'metal-compatible' NFC tags. In particular, explicit metal-shielded tags are expected to work on any surface. The tests utilized metal-

shielded tags ranging in price from a few cents to a few dollars. The cheap metal shielded tags performed as poorly as non-shielded tags, whilst the more expensive ones generally performed better than the low-cost tags. However, the evaluations have shown that even the most expensive and metal-shielded tags failed to work on a large set of metallic items.

Conclusions and Future Steps

Due to these problems with using the NFC system at metal objects, it was decided that an NFC tag-based approach was unacceptably unreliable, and the human-readable label system was reinstated. However, the smartphone app developed is still in use and determines a step forward in the entire inventory process, as it not only allows for an active identification of items – originally planned for by scanning their NFC tag –, but also allows for entering an identifier and all device-related information in a mobile manner, without the need for a Web-based inventory interface.

As very many metallic items were identified as difficult to read, even with current metal-shielded tags, the technology progress of new metal-shielded tags will be tracked in the near future for a possible return to the NFC technology.

Once metal-cased devices work well with a new type of metal-shielded tags, the inventory system will be switched back to an NFC-based inventory system. In conclusion, when planning an introduction of an NFC-based system, it is important to ensure that NFC tags on metallic items can be used by verifying NFC tags operations on metallic items.

References:

- [1] NFC and Contactless Technologies, NFC Forum, URL: <http://nfc-forum.org/what-is-nfc/about-the-technology/>, last access January 2015.
- [2] NFC Data Exchange Format (NDEF) Technical Specification, NFC Forum, URL: http://members.nfc-forum.org/specs/spec_list/, last access January 2015.
- [3] Communication Systems Group CSG, University of Zürich, Switzerland, URL: <http://www.csg.uzh.ch>, last access January 2015.

Please contact:

Patrick Poullie, Thomas Bocek, Burkhard Stiller
 Universität Zürich, Switzerland
 E-mail: poullie@ifi.uzh.ch,
bocek@ifi.uzh.ch, stiller@ifi.uzh.ch

TrakChain Estimates Costs for Track and Trace in the Internet of (many) Things

by Miguel L. Pardal and José Alves Marques

The TrakChain assessment tools take a description of a physical supply chain – relevant locations, how many goods are received, how often, etc. – and estimate the performance of track and trace queries in a modelled traceability system, providing predictions of how much processing and storage will be required for the working system. The tools were developed at Instituto Superior Técnico, Universidade de Lisboa, Portugal and were evaluated using a Pharmaceuticals supply chain case study.

The Internet of Things (IoT) promises benefits from a deeper connection between the virtual and physical worlds. One specific application area is logistics. The global economy depends on a wide range of supply chains that transfer goods from producers to consumers. The combined use of Enterprise Resources Planning (ERP) and Supply Chain Management (SCM) information systems has greatly improved the overall operational efficiency of supply chains.

However, to achieve further improvements, more up-to-date and precise information about the supply chain is required.

RFID is an IoT technology that allows detailed and automated data capture in the supply chain, as specified by the EPCglobal standards [2]. Tags are attached to the objects of interest and readers placed along the supply chain

locations generate event data, as illustrated in Figure 1.

A practical RFID traceability system should perform adequately for the large number of physical objects flowing in the supply chain [1]; and it should protect the sensitive business data from unauthorized access providing the desired data visibility [2]. The TrakChain project was proposed to evaluate both these aspects. It provides




Figure 1: Data about a physical object of interest, tagged with RFID, is captured along the supply chain. EPC IS is a standard for capturing and sharing event data.

tools to estimate and measure the computational and communicational costs; and also security tools. The project was developed at Instituto Superior Técnico, Universidade de Lisboa, Portugal; in collaboration with researchers from the University of Cambridge, UK; and the Massachusetts Institute of Technology, USA; and was aligned with normalization efforts by GS1 for 'RFID data discovery services' and 'event-based traceability'.

Traceability Cost Models

The cost models can compare traceability systems for different supply chains, so that the best architecture for a given setting can be found. It helps to find answers for questions such as:

Should the system be centralized or decentralized? Should data be copied to specific locations or referenced?

The data flow of the assessment tool is illustrated in Figure 2: given a supply chain characterization and system workflow specifications, the cost model can estimate processing times and data storage needs. The tool was validated with a case study in the Pharmaceuticals industry that compared solutions being proposed to ensure the authenticity of drugs [2].

Traceability Data Access Control

The supply chain participants need to trust that the traceability system will manage their data properly and enforce

data access control, otherwise they will not share their data [3].

TrakChain implemented visibility restriction mechanisms that can be used to define and enforce access control policies using RDF and SPARQL. The policies can be converted to a standard format, XACML, to reuse existing enforcement infrastructures with certified management and audit tools.

The expressiveness of the policies was evaluated against a set of requirements for a real-world pharmaceutical traceability system, and it was shown to be expressive enough to satisfy the business requirements. For example, it can specify dynamic conditions to allow the sharing of data with business partners downstream in the supply chain that are not known in advance.

Future Work

More data attributes, such as expiry dates and temperature readings, can also be made available and controlled by the traceability system. The increase in the safety and quality of products is a good example of how the Internet of Things can help change the world for the better.

Link:

<http://trakchain.net>

References:

- [1] M. L. Pardal, M. Harrison, and J. A. Marques: "Assessment of Visibility Restriction Mechanisms for Discovery Services", IEEE International Conference on RFID, 2012.
- [2] M. L. Pardal, M. Harrison, S. Sarma, and J. A. Marques: "Expressive RFID data access policies for the Pharmaceuticals supply chain", IEEE International Conference on RFID, 2013.
- [3] M. Eurich, N. Oertel, and R. Boutellier: "The impact of perceived privacy risks on organizations' willingness to share item-level event data across the supply chain", Journal of Electronic Commerce Research 10(3-4), 423-440, 2010. ISSN: 1572-9362.

Please contact:

Miguel L. Pardal
 Instituto Superior Técnico, Portugal
 Email:
Miguel.Pardal@tecnico.ulisboa.pt




Figure 2: Data flow of the TrakChain cost assessment tool.

European Research and Innovation

Mesh Joinery: A Method for Building Fabricable Structures

by Paolo Cignoni, Nico Pietroni, Luigi Malomo
and Roberto Scopigno

Mesh joinery is an innovative method to produce illustrative shape approximations suitable for fabrication. Mesh joinery is capable of producing complex fabricable structures in an efficient and visually pleasing manner. We represent an input geometry as a set of planar pieces arranged to compose a rigid structure by exploiting an efficient slit mechanism. Since slices are planar, a standard 2D cutting system is sufficient to fabricate them.

Over the last decade, a variety of rapid prototyping technologies have been developed to support the manufacturing process, especially for the fabrication of production-quality parts in relatively small numbers. However, while the printing resolution has improved substantially - and consequently, the accuracy in terms of reproduction has reached high standards - rapid prototyping is still perceived as being too expensive for the mass market. This is particularly true for large scale reproduction: only a few techniques can produce, even approximated, large scale actual copy within reasonable cost and time bounds.

We have introduced mesh joinery [1], a novel and practical approach to fabricate illustrative shape approximations made up of several interlocked planar pieces, called 'slices'. Slices can be easily fabricated even in large scale (both numbers and dimensions) using any 2D cutting device and then manually assembled through a sequence of simple well defined operations to obtain a rigid structure that resembles an input 3D model.

Our approach offers a truly low-cost solution owing to the simple cutting technologies employed and the relatively inexpensive material used (such as cardboard). Although the proposed slice structure approximates the original geometry and it cannot be considered as a 'exact physical copy', nevertheless, we believe that our method is attractive in many markets, such as in artistic or illustrative contexts, large scale approximate reproductions, free form scaffolding, and even in puzzles or toys, and where assembly is a key part of user experience.

Our approach is based on building interlocking arrangements composed of shallow flat ribbon-shaped pieces that follow a cross-field defined on the surface. We provide a novel formalism to automatically design a slice-to-slice interlocking system. This formalism provides enough degrees of freedom to follow complex cross-fields and, consequently, to efficiently approximate the global structure that characterizes the input shape. We also redesigned the traditional slit interlocking mechanism that is used to connect pieces in order to approximate generic 3D surfaces with greater flexibility. Additionally, we ensure a sufficient degree of physical sta-

bility of the final structure and provide the sequence of manual operations required for the assembly procedure.

Specifically, Mesh Joinery makes three major contributions that make it a significant advance in the field of large scale fabrication technologies.

First, we have extended the classical slit interlocking mechanism by providing additional structural degrees of freedom. In particular, we allow insertion movements that are not




Figure 1: Starting from a 3D model with a cross field the Mesh Joinery approach allows to automatically generate flat cut arrangements that well approximate the original shape.




Figure 2: Mesh Joinery is based on a novel slit based interlocking system that allows to connect flat pieces in a much more flexible way and that retains robustness through global relationships.

orthogonal to slices. In addition, we formulated non-orthogonal slice placement in a novel, structurally sound perspective. This new, more flexible, interlocking mechanism allows different slices to be connected without the traditional very limiting orthogonality constraints. These additional joinery degrees of freedom can be effectively exploited to represent complex models with few ribbon shaped slices that can run in the best way to approximate the original shape, far better arrangements to be generated than the current state of the art [2].

Second, we have introduced an efficient automatic strategy to approximate a surface with a set of slices. Slice placement is driven by an input cross-field [3] that represents in a compact, intuitive way the curvature flow over the surface. Our approach provides a set of appealing, uniformly distributed polylines lying on the surface of a mesh and agreeing with the given cross field. This curvature driven placement is very important to catch, with just a few pieces, the overall shape of the input object. In addition, the method takes into account slice insertion constraints and, while it does not theoretically guarantee that the mounting sequence is collision free, it yields arrangements that are practically assemblable and that exhibit a sufficiently robust slice structure. Our method may also take advantage of field symmetrization techniques for a better perception of the global structure of the object.

Finally, mesh joinery exploits an automatic procedure to ensure that the slice structure is physically achievable. First, it improves the final rigidity, acting upon the slit interlocking mechanism. Second, it ensures that the slice structure conforms to the physical constraints required by the manual assembling procedure. This procedure is specifically designed to deal with our extended slit mechanism.

The Mesh Joinery approach has been patented and it has been used in practice to build a number of real large models that have been shown at international conferences and workshops.

Link:

<http://vcg.isti.cnr.it/Publications/2014/CPMS14/>

References:

- [1] P. Cignoni, N. Pietroni, L. Malomo, R. Scopigno: "Field-aligned mesh joinery", ACM Trans. Graph, 33, 1, (February 2014).
- [2] Y. Schwartzburg, M. Pauly: "Design and optimization of orthogonally intersecting planar surfaces", Computational Design Modelling, 2012.
- [3] D. Panozzo, Y. Lipman, E. Puppo, D. Zorin: "Fields on symmetric surfaces", ACM Trans. Graph. 31, 4 (July), 2012.

Please contact:

Paolo Cignoni
ISTI-CNR, Italy
E-mail: paolo.cignoni@isti.cnr.it

Icing Detection and Protection for Small Unmanned Aircrafts

by Andrea Cristofaro and Tor Arne Johansen

In unmanned aerial vehicles, it is vital to be able to detect and accommodate ice adhesion on wings, control surfaces and airspeed sensors, since ice accretion modifies the shape of the aircraft and alters its measurements, thus changing the aerodynamic forces and reducing manoeuvring capability.

The use of unmanned aerial vehicles (UAVs) for civil applications – largely surveillance and monitoring - has increased markedly in recent years. UAVs are very often suited for harsh conditions that are unsafe for humans, such as those typically encountered in Arctic operations. Reliable, appropriate, and efficient UAV operations are needed in such conditions.

The phenomenon of ice accretion on aircraft wings and control surfaces is a well recognized problem in aerospace engineering: the modified shape of the leading edge due to ice, changes the lift, drag and pitch moment characteristics of the wing. For instance, an ice-covered airfoil may experience a 40% reduction in lift coefficient, while the drag may be increased by as much as 200%. A decrease in lift requires more engine power and implies a premature airfoil stall angle.

Large aeroplanes are commonly equipped with efficient anti-icing and de-icing devices; however, these are largely unsuitable for small aircraft, owing to their simple architecture and limited payload. Ice formations on aircraft surfaces during flight are typically caused by supercooled water droplets (SWD). When a water droplet cools, it does not freeze until it reaches a very low temperature; however, a droplet will freeze instantly, releasing latent and accreting ice, when it interacts with an external agent such as an aircraft. Both rate and amount of ice depend on the shape of the surface, its roughness, travelling body speed, temperature and droplet size.

A first challenge is the in-flight detection of icing, i.e., recognizing the accretion of ice layers using data provided by onboard sensors. This problem can be addressed by applying the results of a recent study by our research team [1],[2], in which we proposed a scheme for fault detection and isolation in overactuated systems based on Unknown Input Observers and Control Allocation methods. Unknown Input Observers are deterministic estimators that can be made insensitive to unmeasured inputs, while Control Allocation is very well suited to handle constraints and to readily reconfigure the control action in spite of actuator failures. Combining these two methods, a procedure has been designed to keep icing effects decoupled from mechanical faults and external disturbances such as wind gusts [3].

A second step towards achieving icing protection for UAVs, is the implementation of an efficient anti-icing and de-icing

mechanism. Very promising results come from carbon nanotechnology that can be used to make thin films of coating material to be painted on aircraft wings. Using the electricity onboard, the film can be heated to melt the ice. Such devices, however, are very energy demanding, thus it is vital to have an optimal scheme for the activation of the de-icing system to maintain aircraft stability. The combination of the aircraft dynamical model and the model describing the accretion of icing is the ‘Stefan problem’, which consists of a set of four partial differential equations: two heat equations, one mass balance equation and one energy balance equation.

Finally, in order to test and validate the efficiency and robustness of the icing detection algorithms as well as the icing protection module, field experiments in Svalbard are planned for the near future using the X8 UAV (Skywalker Technology Co, Ltd.).

The research at NTNU is a fruitful collaboration with other two ERCIM partner institutes, namely University of Porto (UPT) and University of Cyprus (UCY). In particular, we are currently working jointly with Prof. Pedro Aguiar (UPT) and Prof. Marios Polycarpou (UCY) who have made outstanding contributions to the project with their expertise in multiple model adaptive estimation and nonlinear fault detection, respectively.




Figure 1: Ice accretion effects on airfoils.

References:

- [1] A. Cristofaro, T. A. Johansen: “Fault-tolerant Control Allocation: An Unknown Input Observer based approach with constrained output fault directions”, in proc. of the 52nd IEEE Conf. on Decision and Control, Florence, Italy, 2013, pp. 3818-3824.
- [2] A. Cristofaro, T. A. Johansen: “Fault-tolerant control allocation using Unknown Input Observers”, *Automatica*, 50(7), 2014, pp. 1891-1897.
- [3] A. Cristofaro, T. A. Johansen: “An Unknown Input Observer approach to icing detection for unmanned aerial vehicles”, accepted for publication in proc. of IEEE American Control Conference 2015.

Please contact:

Andrea Cristofaro
ERCIM post-doc fellow at NTNU, Norway
E-mail: andrea.cristofaro@itk.ntnu.no

Tor Arne Johansen, NTNU, Norway
E-mail: tor.arne.johansen@itk.ntnu.no

LEARNMINER – Advanced Analytics for eLearning

by Christian Bauckhage, Markus Ketterl and Ulrich Nütten

LEARNMINER is an interactive social eLearning platform that integrates advanced analytics and linked data processing to assist users in structuring course material. Tools for analyzing crowd behaviour allow instructors to identify overly simple or difficult material and provide means for planning deadlines or tests. LEARNMINER allows for mobile access and provides interactive visualization of learning material and progress.

The interdisciplinary field of Web science studies technical and sociological aspects of the Web. Given that the Web is the largest human information construct in history whose rate of adoption and impact on society are unparalleled, Web science aims to understand the Web in order to engineer its future and to ensure its social benefit [1]. The research project SoFWiReD (Southampton Fraunhofer Web Science, Internet Research & Development), funded by the Fraunhofer ICON initiative, is a collaboration between the University of Southampton and Fraunhofer FOKUS and Fraunhofer IAIS. It involves web and data scientists with expertise in developing web platforms for data and knowledge driven processing that incorporate aspects of collective intelligence.


One of the platforms developed in SoFWiReD is LERANMINER, an eLearning solution that aims to simplify and optimize the process of learning and training on the job. It integrates social media features, linked Open Data processing, document analysis, and predictive analytics.

Confronted with large repositories of learning material, users may struggle to decide which item to work with next since this decision requires knowledge as to the content of the whole repository. LEARNMINER assists users by automatically determining document semantics and interrelations. It structures learning resources and displays them as a dynamic document graph which evolves with the progress of individual users as well as crowds of users. The knowledge graph thus helps users to prioritize study according to their own needs as well taking into account the experience of the overall population. In other words, the wisdom of the crowd is harnessed to facilitate the task of the individual. Analytics tools help instructors to track user progress or to predict when learning tasks will be accomplished. This in turn facilitates properly defined deadlines, homework, and tests.

LEARNMINER provides a semantic search engine for users to navigate through a repository of documents. Using semantic text mining [2], queries are matched to semantic categories extracted from a given repository. Search results are ranked according to their semantic similarity to a query and displayed together with the knowledge graph to allow users to contextualize material.

When displaying the knowledge graph, users have different options. In the basic view, documents are structured according semantic similarities and shown as interlinked nodes. Node colours reflect subject areas and the most general

Figure 1: Learn path visualization of a LEARNMINER knowledge graph.



document of a subject area is displayed as the central node of a subgraph. Node sizes are used to express document sizes (e.g., measured in text length); the bigger a node, the more time users are expected to devote to the corresponding document.

When displayed as a difficulty graph, the knowledge graph is organized according to how the overall user population performs on tests created for each document. The more difficult a document the farther away it is from the cluster centre. The time graph display organizes the knowledge graph such that it reflects the time that the user has spent working with a document. Documents close to the centre of the graph indicate material a user has (almost) completed, for example, by reading (almost) all pages of a document and taking corresponding tests. Documents farther from the centre indicate material which still requires more time to completion. Finally, the learn path display provides suggestions about how to continue studies by drawing a personalized path through the space of available resources. That is, this representation guides users based on their experience, background knowledge, and learning progress. While users are working LEARNMINER, their learn paths will adapt and individual and collective progress and success (e.g., measured in terms of grades obtained) are parameters used to compute suggestions for further reading.

LEARNMINER is currently employed in a corporate training environment, and users report that they especially appreciate its social features (e.g., support for study groups). By design, the platform is interoperable and can efficiently index and structure learning repositories for a wide range of subject areas.

References:

- [1] J. Hendler, et al.: “Web science: an interdisciplinary approach to understanding the Web,” *Communications of the ACM*, vol. 51, no. 7, pp. 60-69, 2008.
- [2] F. Metze, C. Bauckhage, and T. Alpcan: “Social and Expert Research in Online Communities”, in *Semantic Computing*, P. Sheu, et al. (eds), Wiley, 2010.

Please contact:

Ulrich Nütten, Fraunhofer IAIS, Germany
Tel: +49-2241-142915
E-mail: ulrich.nuetten@iais.fraunhofer.de

Online Semantic Analysis over Big Network Data

by Chengchen Hu and Yuming Jiang

Mobile applications, web services and social media generate huge data traffic from many sources. This is described as 'big network data' with the '4V-characteristics' - that is, variety, velocity, volume and veracity. Analysis on such big network data can discover problems, reveal opportunities and provide advice for the service and network providers for fine-grained traffic engineering with close to real-time adjustment or application acceleration [1].

Different applications tend to integrate several functionalities with various data formats. For example, the Twitter application produces network traffic such as tweeting, posting pictures, embedding video. Classical methods are limited to protocol or application identification [2]. We need to go beyond the packet and application analysis, when semantic information is the target [3]. Therefore, the first goal is for our method to exhibit fine-grained awareness, which analyzes user behaviour instead of traffic only related to a certain application. This implies that we need to use a general grammar to associate the unstructured data with user behaviour.

The second goal is to develop a flexible and uniform specification of the user semantic from network traffic. In previous work, heterogeneous and unstructured big network data in different formats are studied separately. It is a challenge to normalize the independent data structures and describe user behaviour in a unified framework to conduct a comprehensive analysis in a fine-grained semantic manner.

Finally, the third goal of our approach is that the method should operate at (close to) wire speed. Even when real-time analysis is not strictly needed, the off-line method is limited by storage because the analysis capacity and capabilities cannot keep up with the rate at which data is produced. Extracting and storing only useful information is a viable approach that needs to be explored further.

In order to achieve the above goals, we apply Deep Semantic Inspection (DSI), which contains a standard description to unify the various formats of different applications and finally obtain user semantics. Our basic idea is to extract a minimum but complete semantic for each user behaviour at wire speed, and then apply data analysis and data mining on the small sized semantic data instead of the raw traffic data. This process purifies the raw traffic and reduces the data volume by several orders of magnitude. Our preliminary experiment shows that the compression ratio between the raw traffic volume and our approach is in three orders of magnitude. As a result, the data volume for further user-defined high-level analyses can be significantly reduced to handle big, and increasing, network data.

We have designed and implemented a cross-platform system, named Semantic On-Line Intent Detection (SOLID) to realize our DSI approach. As shown in Figure 1, SOLID builds a processing pipeline in its kernel space, where a semantic parser translates the segments into the application semantic. A matching engine compares the application semantic with predefined specifications to output the user sketches. The kernel design allows simultaneous processing of multiple application specifications and multiple PDUs. Our implementation has demonstrated that on a real traffic trace, the SOLID-system achieves 17.6Gbps throughput with the cost of 709MB memory consumptions in x86 platform. To make the SOLID-system flexible so that it supports var-




Figure 1: An example and processing pipeline of DSI/SOLID.

ious analyses in the user space, we are now trying to improve the interface abstraction between the kernel and user space.

A few practical cases have been used to demonstrate the flexibility of SOLID, but more cases need to be investigated. This is necessary in order to acquire a thorough understanding of big network data with real traffic to analyze various factors, such as the application performance, user profiling, CP competition and application correlations.

The authors would like to appreciate Poul Heegaard of NTNU for his valuable suggestions on this article.

References:

- [1] T. Benson et al.: "Fine grained traffic engineering for data centers", in proc. of CoNEXT '11 (New York, NY, USA, 2011), ACM, pp. 8:1–8:12.
- [2] F. Yu et al.: "Fast and memory-efficient regular expression matching for deep packet inspection", in ANCS 2006. (Dec. 2006), pp. 93–102.
- [3] H. Li, C. Hu: "ROOM: Rule organized optimal matching for fine-grained traffic identification", in proc. of INFOCOM, 2013, IEEE (2013), pp. 65–69.

Please contact:

Chengchen Hu
ERCIM Fellow at NTNU, Norway
E-mail: chengchen.hu@item.ntnu.no, huc@ieee.org

Yuming Jiang
NTNU, Norway
E-mail: yjiang@item.ntnu.no

From Cooling Integrated Circuits to Efficient Electrical Power Generation

by Stephan Paredes, Brian Burg, Patrick Ruch, Ingmar Meijer and Bruno Michel

An earlier article in ERCIM News described a system for cooling integrated circuits with re-usable hot water. The result is a sharp increase in energy efficiency and a sharp reduction in the carbon footprint of data centers. The same technique is being applied to photovoltaic electrical energy production, again with a huge gain in system efficiency due to the parallel recovery of waste heat. In hot climates heat from datacenters as well as solar systems is converted into cooling by means of adsorption heat pumps.

The techniques described in the earlier article [1] involve using a network of microfluidic channels etched into the back side of the individual integrated circuit chips. The branched system is patterned after the cooling system of the mammalian brain that serves us all so well. Because water is so much more capable of heat transportation than air, the traditional air-conditioning system for a data center is unnecessary and the use of costly air-conditioning units and fans can be minimized. The integrated circuit chips are so designed that the "cooling" water can be so hot that the exiting hot water can be used for various purposes, heating a building or water desalination for examples, before being returned to the system for cooling.

One of the systems using a water-cooling infrastructure with heat recovery is the supercomputer SuperMUC at the Leibniz Rechenzentrum in Garching, Germany. The system has a peak performance of three Petaflops. At the time of its installation in May, 2012, it was the fastest computer in Europe and ranked number four in the TOP500 list of the world's fastest computers. The system runs successfully and problem-free since then and is currently being expanded with new hardware to double its performance to 6 Petaflops. This phase II expansion will become operative in June 2015.

The microfluidic cooling innovation has now also been applied to specially designed photovoltaic chips. The overall efficiency of this new approach to conversion of sunlight to electrical power was so high in an experimental system that a commercial system is being developed. High-efficiency multi-junction photovoltaic cells reach efficiencies for photovoltaic energy conversion in excess of 40%, but the sun's energy needs to be concentrated for their use to be cost-effective. This is being done through a system of parabolic mirrors which focus the solar energy more than 1500-fold on a dense array of multi-junction photovoltaic chips. At such concentration ratios, the heat flux at the photovoltaic cells is around 150 W/cm², which calls for high-performance cooling to avoid failure of the cells. Relying on the microfluidic cooling approach developed for the thermal management of microprocessors, the photovoltaic cells can be safely




Figure 1: Lab demonstration of a nine-cell, water-cooled photovoltaic receiver module.

operated under this extreme illumination and, at the same time, a large fraction of the incoming sunlight which is not converted to electricity can be captured as heat.

In the “Sunflower” implementation, 36 elliptic mirrors covering a 40 square-meter area achieve this focusing in a ten-meter high “sunflower”. The rated capacity of the system is 10 kW electrical power output at a solar-to-electrical efficiency of 25% and 22 kW thermal output at a solar-to-thermal efficiency of at least 55%. Overall this is an 80% usage of the sun’s energy. The hot water can be used to produce additional electrical power, heating (or even cooling) buildings or water desalination. In hot countries, the demand for heating is dwarfed by the demand for cooling. In the Middle East, for example, 90% of all electricity is consumed for air-conditioning in the summer months. The thermal output of the Sunflower system can be used to provide valuable cooling by means of thermally driven heat pumps, such as adsorption or absorption chillers [3]. The ability to drive a refrigeration cycle using heat from the concentrated photovoltaic system is a key value proposition in hot climates.

In fact, thermally driven heat pumps represent a promising technology to utilize even low-grade heat in scenarios where there is little or no demand for heating. Hot-water cooled datacenters in warm climates, for example, are still more effi-



Figure 2: Three energy-harvesting “sunflowers” in a field of conventional sunflowers.

cient than their air-cooled counterparts due to the minimization in air-conditioning and air-handling infrastructure. However, the recovered heat does not provide an added value and needs to be discarded. On the other hand, there is a continuous demand for cooling of power supplies, storage servers and other non-water-cooled datacenter infrastructure. Therefore, thermally driven heat pumps offer a solution to utilize low-grade heat recovered from hot-water cooled servers in order to provide the cooling for other air-cooled datacenter components. In this scenario, heat recovery combined with sorption cooling technology therefore enables further improvements in datacenter efficiency even in hot climates.

The synergies between the thermal management of servers and photovoltaic cells have led to the definition of high-efficiency systems in the fields of datacenters and energy generation. The design of the photovoltaic receiver modules with integrated cooling was done at the IBM Zurich Research Lab and the overall Sunflower is being built by a new company, Dsolar. Dsolar is part of Airlight Energy which specializes in solar power technology. Their contribution is an inexpensive implementation of the array of mirrors in each of which a reflecting plastic film is shaped by a pressure difference in a lightweight concrete casing. The sunflower has a tracking mechanism to follow the path of the sun in the course of a day. The innovative implementation is inexpensive to the point where the hope is that the power generation will be two or three times less expensive than standard technology.

Links:

<http://ercim-news.ercim.eu/en/79/special/using-waste-heat-from-data-centres-to-minimize-carbon-dioxide-emission>

Press release:

<http://www.zurich.ibm.com/news/14/dsolar.html>

TED talk:

<https://www.youtube.com/watch?v=j5Bh0AfuiUg&feature=youtu.be>

References:

[1] G. I. Meijer, T. Brunschwiler, S. Paredes, B. Michel, “Using Waste Heat from Data Centres to Minimize Carbon Dioxide Emission”, ERCIM News, No. 79, October, 2009.

[2] P. Ruch, et al.: “Toward 5-Dimensional Scaling: How Density Improves Efficiency in Future Computers”, IBM J. Res. Develop., vol. 55 (5), 15:1-15:13, October 2011.

[3] P. Ruch, et al.: “Heat-driven adsorption chiller systems for sustainable cooling applications”, in proc. of the 11th International Energy Agency Heat Pump Conference, Montréal, Canada, 2014.

Please contact:

Bruno Michel


IBM Zurich Research Lab, Switzerland

E-mail: bmi@zurich.ibm.com

VAVID– Handling Big Data

The German VAVID project develops improved techniques for data compression as well as new methods of data analysis, data management and interactive data visualization.

VAVID allows technology firms to get a better grip on the massive amount of data they need to handle. The partners in this project are developing methods to tackle the enormous volumes of data that accumulate at engineering departments. Examples of such data include simulation results and the sensor data received from machines and installations. VAVID works by using comparative analysis and data compression to reduce data to its relevant core. This saves on the costs of data storage and creates the transparency needed by engineers to optimize both production and products. VAVID is being coordinated by the Fraunhofer Institute for Algorithms and Scientific Computing (SCAI) and is receiving 2.2 million € in support from the German Ministry of Education and Research (BMBF) under its Big Data program.



The VAVID project is developing methods to tackle the enormous volumes of data that accumulate at engineering departments, such as data from simulation results.

In today's production environment, numerical simulation has become an indispensable part of product development. Before fabrication begins, it is essential to computationally analyze the product's characteristics in a way that mirrors reality as faithfully as possible. The computations and high-performance computer (HPC) systems required for this task are generating an ever growing mountain of data. An exponential rise in data volumes is also being seen due to the acquisition of sensor data during the operation of machines and plant. These measurement data allow engineers to draw important conclusions on how well control systems are working and how they can further optimize production.

The huge amount of data poses great challenges to technology companies. Moving large data packets around the company or to external partners, and archiving all this data, requires a powerful IT infrastructure, something that is often extremely cost-intensive. Moreover, important information

carried by data is frequently not recognized because the company does not have the necessary data extraction methods at its disposal. This is precisely where the VAVID project has set its sights. VAVID is the German acronym for "Vergleichende Analyse von ingenieurrelevanten Mess- und Simulationsdaten" or "Comparative Analysis of Engineering Measurements and Simulation Data."

The data being studied by the project are numerical simulation results data from the automotive and wind industries as well as measurement data taken from wind turbine monitoring systems. By performing joint and comparative analysis of data from different industries, the partners are in the first project phase developing a methodology for efficient data analysis. These methods and techniques are going into the creation of a high-performance data management system that will allow centralized data storage as well as efficient data access and retrieval.

The second phase will focus on the end user by developing innovative, universally applicable software components that can later be deployed in other industries. Examples of note here include the fields of aviation and mechanical and plant engineering (industrial plant).

The eight partners participating in the project reflect the cooperative effort being made to transfer academic knowledge to industrial practice and to orient research to the needs of industry: Fraunhofer SCAI, Bosch Rexroth Monitoring Systems GmbH in Dresden, GE Global Research in Garching, GNS mbH in Braunschweig, SCALE GmbH in Dresden, SIDACT GmbH in Sankt Augustin, Center for Information Services and High Performance Computing (ZIH) at TU Dresden and the Faculty of Computer Science, Databases department at TU Dresden. The project furthermore cooperates with associated partners AUDI AG in Ingolstadt, ParStream GmbH in Cologne and Volkswagen AG in Wolfsburg. The associated partners from the automotive industry contribute to the joint project by describing the real-world requirements of industry and by supplying relevant sample data for use in the project.

One aim of the VAVID project is to concentrate research and development efforts and point them toward Big Data technologies that can be used in many sectors of industry. A further objective is to generally strengthen the German information and communications technology industry and to boost the competitiveness of German companies. The VAVID project started in September 2014 and is set to run through August 2017.

Link:
<http://www.vavid.de>

Please contact:
Jochen Garcke, Fraunhofer SCAI
E-mail: jochen.garcke@scai.fraunhofer.de

Call for Papers

Information Security Conference - ISC 2015

Trondheim, Norway,
9-11 September 2015

The Information Security Conference (ISC), which started as a workshop (ISW) in 1997, is a well-established and highly reputable international conference that is held yearly. It has been held in five different continents. ISC 2015 is organized by the Department of Telematics at the Norwegian University of Science and Technology (NTNU).

Topics:

The conference seeks submissions on novel theoretical and practical results in: access control, accountability, anonymity and pseudonymity, applied cryptography, authentication, biometrics, computer forensics, critical infrastructure security, cryptographic protocols, database security, data protection, data/system integrity, digital right management, economics of security and privacy, electronic frauds, embedded security, formal methods in security, identity management, information dissemination control, information hiding & watermarking, intrusion detection, network security, peer-to-peer security, privacy, secure group communications, security in information flow, security for Internet of Things, security for mobile code, secure cloud computing, security in location services, security modeling & architectures, security and privacy in social networks, security and privacy in pervasive and ubiquitous computing, security of eCommerce, eBusiness and eGovernment, security models for ambient intelligence environments, trust models and trust policies.

Important dates:

- Paper submission deadline: April 27, 2015
- Acceptance notification: June 12, 2015
- Proceedings version: July 1, 2015

More information:

<http://isc2015.item.ntnu.no/index.php>

Call for Papers

IC3K - 7th Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management

Lisbon, 12-14 November 2015

The purpose of the IC3K is to bring together researchers, engineers and practitioners on the areas of Knowledge Discovery, Knowledge Engineering and Knowledge Management. IC3K is composed of three co-located conferences:

- KDIR - 7th International Conference on Knowledge Discovery and Information Retrieval
- KEOD - 7th International Conference on Knowledge Engineering and Ontology Development
- KMIS - 7th International Conference on Knowledge Management and Information Sharing.

Deadlines:

- Regular paper submission: 2 June 2015
- Regular paper authors notification: September 1, 2015
- Regular paper camera ready and registration: September 30, 2015.

More information:

<http://www.ic3k.org/Home.aspx>

Call for Participation

AITA 2015 - Advanced Infrared Technology and Applications

Pisa, Italy, 29 September - 2 October 2015

The 13th International Workshop on Advanced Infrared Technology and Applications will assess the state of the art of the technology in the Infrared bands and to present its most interesting applications. In the 13th AITA edition, special emphasis will be given to the following topics:

- Advanced technology and materials
- Smart and fiber-optic sensors
- Thermo-fluid dynamics
- Biomedical applications
- Environmental monitoring
- Aerospace and industrial applications
- Nanophotonics and Nanotechnologies
- Astronomy and Earth observation
- Non-destructive tests and evaluation
- Systems and applications for the cultural heritage
- Image processing and data analysis
- Near-, mid-, and far infrared systems.

Considering the importance of cultural heritage in Italy and the excellence achieved by Italian researchers and operators in this domain, a technical seminar entitled "Techniques for infrared diagnostics and documentation in the field of cultural heritage" will also be organized.

Important dates:

- 30 April 2015: Extended abstract submission
- 31 May 2015: Notification of acceptance
- 30 June 2015: Revised extended abstract submission.

The workshop is organized by the Fondazione "Giorgio Ronchi", the Istituto di Fisica Applicata "Nello Carrara" (CNR-IFAC), the Istituto per le Tecnologie delle Costruzioni (CNR-ITC), the Istituto di Scienza e Tecnologie dell'Informazione "Alessandro Faedo" (CNR-ISTI).

More information:

<http://ronchi.isti.cnr.it/AITA2015/>

Call for Participation

7th International Workshop on Software Engineering for Resilient Systems

Paris, 7-8 September 2015

The way software is developed is changing. It must take into account multifaceted constraints like unpredictable markets, evolving customer requirements, pressures of shorter time-to-market, etc. At the same time, software is controlling critical functionalities in several domains like transportation, health care, manufacturing, and IT infrastructures. As a result, modern software systems require on one side adding frequently (daily or weakly) new features, functionalities, or new versions of software artifacts according to changing contexts, business opportunities, or customer's feedbacks, on the other side ensuring their resilience – an ability of the system to persistently deliver its services in a dependable way even when facing changes, unforeseen failures and intrusions.

From this year SERENE is becoming part of EDCC, the European Dependable Computing Conference.

Topics

The SERENE 2015 workshop provides a forum for researchers and practitioners to exchange ideas on advances in all areas relevant to software engineering for resilient systems, including, but not limited to:

Development of resilient systems

- 3Incremental development processes for resilient systems;
- 3Requirements engineering & re-engineering for resilience;
- 3Frameworks, patterns and software architectures for resilience;
- 3Engineering of self-healing automatic systems;
- 3Design of trustworthy and intrusion-safe systems;
- 3Resilience at run-time (mechanisms, reasoning and adaptation).

Verification, validation and evaluation of resilience:

- 3Modelling and model based analysis of resilience properties;
- 3Formal and semi-formal techniques for verification and validation;
- 3Experimental evaluations of resilient systems;
- 3Quantitative approaches to ensuring resilience;
- 3Resilience prediction.

Case studies & applications

- 3Empirical studies in the domain of resilient systems;
- 3Methodologies adopted in industrial contexts;
- 3Cloud computing and resilient service provisioning;
- 3Resilient cyber-physical systems and infrastructures;
- 3Global aspects of resilience engineering: education, training and cooperation.

Contributions

We welcome relevant contributions in the following forms:

- 3Technical papers describing original theoretical or practical work;
- 3Experience/Industry papers describing practitioner experience or field study, addressing an application domain and the lessons learned;
- 3PhD Forum papers describing objectives, methodology, and results at an early stage in research;
- 3Project papers describing goals and results of ongoing projects;
- 3Tool papers presenting new tools or new versions of existing tools that support the development of resilient systems.

Important Dates

- Submission due: 24 April 2015
- 3Authors notification: June 19, 2015
- 3Camera ready papers: July 1, 2015.

The workshop is organised by the ERCIM Working Group SERENE.

More information:

<http://serene.disim.univaq.it/2015/>

Call for Participation

ICEC 2015 - 14th International Conference on Entertainment Computing

Trondheim, Norway, 30. September - 2 October 2015

The IFIP International Conference on Entertainment Computing is the primary forum for disseminating and showcasing research results relating to the creation, development and use of digital entertainment. The conference brings together practitioners, academics, artists and researchers interested in design, practice, implementation, application and theoretical foundations of digital entertainment.

Fun, joy and excitement have been strong driving forces for human development through history. Kids play and grow. Today we experience how tools, technology, methods and theory developed for the entertainment of people are taken into use for other purposes. A two way interaction is growing and we like to nurture it.

Authors with background in domains such as health, education, media, sport, are invited to contribute and participate, of course in addition to people working in the core areas of entertainment computing. We solicit paper, poster and demonstration submissions, as well as proposals for workshops and tutorials. IFIP International Conference on Entertainment Computing is the longest established and most prestigious conference in the field of entertainment computing. The conference is the primary forum for disseminating and showcasing research results relating to the creation, development and use of digital entertainment. The conference brings together practitioners, academics, artists and researchers interested in design, practice, implementation, application and theoretical foundations of digital entertainment.

More information:

<http://icec2015.idi.ntnu.no/>

1955-2015: ITC Celebrates its 60th Anniversary!

Ghent, Belgium, 8-10 September 2015

It is a pleasure to announce that the International Teletraffic Congress (ITC) will celebrate its 60th anniversary. Over the past 60 years, ITC has provided a forum for leading researchers from academia and industry to present and discuss key technological and methodological advances in the design, performance evaluation and control of communication networks, protocols and applications and in traffic measurement and management.

This year the 27th ITC congress takes place in Ghent, Belgium. The theme of ITC 2015 is “Traffic, Performance and Big Data” reflecting today’s networking challenges.

With the emergence of the Internet of Things, the number of devices being connected to the Internet is steadily growing and huge amounts of data are being generated worldwide. Big data brings new traffic and performance related challenges and calls for a deep revisit to the methodological tools that were traditionally used for performance evaluation and traffic engineering. New models and approaches are needed to investigate big data characteristics in terms of volume, velocity and variability and their impact on network performance; new solutions have to be designed to efficiently and securely manage information; new techniques are needed to support all phases of network planning, design and optimization.

ITC represents a wide and lively community of researchers and practitioners dedicated to push back the limits of knowledge in the area of networking. As such, ITC regularly organizes events like Congresses, Specialist Seminars and Workshops in order to discuss the latest changes in the modelling, design and performance of communication systems, networks and services. As a recognized forum, ITC events attract participants from all over the world.

Join the ITC community to celebrate this significant anniversary and prepare the future of ICT in the attractive city of Ghent, Belgium! <http://www.itc27.org/>

Call for Papers

SERENE 2015 - 7th International Workshop on Software Engineering for Resilient Systems

Paris, France, 7-8 September 2015

The SERENE 2015 workshop provides a forum for researchers and practitioners to exchange ideas on advances in all areas relevant to software engineering for resilient systems. It is co-located with the 11th European Dependable Computing Conference (EDCC) 2015.

Major topics of interest include, but are not limited to:

- Development of resilient systems (Incremental development processes for resilient systems; Requirements engineering & re-engineering for resilience; Frameworks, patterns and software architectures for resilience; Engineering of self-healing autonomous systems; Design of trustworthy and intrusion-safe systems; Resilience at run-time (mechanisms, reasoning and adaptation);

- Verification, validation and evaluation of resilience (Modelling and model based analysis of resilience properties; Formal and semi-formal techniques for verification and validation; Experimental evaluations of resilient systems; Quantitative approaches to ensuring resilience; Resilience prediction);
- Case studies & applications (Empirical studies in the domain of resilient systems; Methodologies adopted in industrial contexts; Cloud computing and resilient service provisioning; Resilient cyber-physical systems and infrastructures; Global aspects of resilience engineering: education, training and cooperation).

Important dates:

- Submission due: 24 April 2015
- Authors notification: 19 June 2015
- Camera ready papers: 1 July 2015

The SERENE Workshop is organised by the ERCIM Working on Software Engineering for Resilient Systems.

More information:

<http://serene.disim.univaq.it/2015>
https://www.linkedin.com/groups?home=&gid=4365850&trk=my_groups-tile-grp
<http://slideshare.net/SERENWorkshop>
<http://edcc2015.lip6.fr>



PhD Positions Available in the BigStorage European Training Network

BigStorage is an European Training Network (ETN) whose main goal is to train future data scientists in order to enable them and us to apply holistic and interdisciplinary approaches for taking advantage of a data-overwhelmed world, which requires HPC and Cloud infrastructures with a redefinition of storage architectures underpinning them – focusing on meeting highly ambitious performance and energy usage objectives.

The network is looking for 15 PhD students to fill exciting opportunities to gain in depth skills cross cutting the interesting topics of BigStorage. Positions are available based in both Academic and Industrial partner institutions in Spain, Germany, France, UK and Greece and they have a duration of 36 months. PhD students will take advantage of several training activities, such as Summer Schools, Seminars, Secondments and Internships.

More information:

A detailed description of the available positions can be found at <http://bigstorage.oeg-upm.net/jobs.html>

25 Years of Python at CWI



In February an interview with Guido van Rossum was published on the ACM website, celebrating 25 years of Python. This popular programming language originates from CWI in Amsterdam, where Van Rossum designed Python in December 1989 and finished the first working version in the first months of 1990. He named the language after his favourite British comedy series Monty Python's Flying Circus.

Science journalist Bennie Mols, author of the ACM article, writes: "For many years, Python has appeared on the list of the ten most-utilized computer languages in the world. Companies like Dropbox, Google, Mozilla, and Walt Disney Animation Studios are large users of Python, as are banks, universities and institutions such as NASA and CERN." On Python's success Van Rossum told Mols: "Python is easy to learn and easy to use. It's also a great tool for scientists who work interactively with data. Python has been available according to the open source philosophy, even before the term 'open source' was invented. Python users feel part of a community and are very active in making improvements."

More information:

<http://cacm.acm.org/news/183638-python-at-25/fulltext>

MOOC on "Binaural Hearing for Robots"

A Massive Open Online Course (MOOC) "Binaural Hearing for Robots" will be available on the platform France Université Numérique from 11 May to 12 June 2015. It will be taught by Radu Horaud, research director at Inria Grenoble Rhône-Alpes.

This course addresses fundamental issues in robot hearing and describes methodologies requiring two or more microphones embedded into a robot head thus enabling sound-source localization, sound-source separation, and fusion of auditory and visual information. The course is intended for students with good background in signal processing and machine learning. It is also valuable to PhD students, researchers and practitioners who work in signal and image processing, machine learning, robotics, or human-machine interaction and who wish to acquire novel competence in binaural hearing methodologies. The course material will allow the attendants to design and develop robot and machine hearing algorithms.

The course is provided by Inria through the project IDEFI uTOP (Open MultiPartner University of Technology).

More information:

<https://www.france-universite-numerique-mooc.fr/courses/inria/41004/session01/about>

Java Bug Fixed with Formal Methods

Researchers from the Formal Methods group at CWI fixed a bug in programming language Java in February 2015. They found an error in a broadly applied sorting algorithm, TimSort, which could crash programs and threaten security. The bug had already been known from 2013 but was never correctly resolved. When researcher Stijn de Gouw attempted to prove the correctness of TimSort, he encountered the bug. His bug report with an improved version has now been accepted. This version of TimSort is used by Android.

Java is broadly used because it provides a lot of support in the form of libraries. TimSort is part of the 'java.util.Arrays' and 'java.util.Collections' libraries. When a bug occurs there, it will appear on many places. Frank de Boer, head of the Formal Methods group says: "So far, it was one of the hardest correctness proofs ever of an existing Java library. It required more than two million rules of inference and thousands of manual steps. With such an important language like Java, it is important that software does not crash. This result illustrates the importance of formal methods for society." The study was co-funded by the EU project Envisage.

More information:

<http://www.cwi.nl/news/2015/java-bug-fixed-formal-methods-cwi>



W3C Launches First HTML5 Course in New Partnership with edX

In support of its on-going commitment to developer training programs, the W3C today announced a new member partnership with edX on its Massive Open Online Course (MOOC) platform under the name W3Cx. The W3Cx courses mean that thousands of developers around the world will have greater access to high quality training programs developed by W3C.

The first W3Cx course, HTML5-part 1, will focus on Web design fundamentals at an intermediate level. It is widely recognized that HTML5 is and will be the essential technology for organizations delivering applications across multiple platforms. HTML5-part 1 will be taught by Michel Buffa Professor at the University of Côte d'Azur (UFR Sciences). The W3Cx course HTML5-part 1 starts 1 June 2015.

More information:

<https://www.w3.org/2015/03/w3cx-launch.html.en>

<https://www.edx.org/course/learn-html5-w3c-w3cx-html5-1x>



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



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



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



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



SBA Research gGmbH
Favoritenstraße 16, 1040 Wien
<http://www.sba-research.org>



Centrum Wiskunde & Informatica

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/>



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/>



Fonds National de la
Recherche Luxembourg

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



Science & Technology
Facilities Council

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



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

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



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



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



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



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



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



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



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



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



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



University of Wrocław
Institute of Computer Science
Joliot-Curie 15, 50-383 Wrocław, 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/>