ercim-news.ercim.eu

Number 119 October 2019

ERCIM NEWS

Special theme:

Smart Things Everywhere

KEYNOTE

JOINT ERCIM ACTIONS

SPECIAL THEME

The special theme “Smart Things Everywhere” has been coordinated by Margarete Hälker-Küsters (Fraunhofer AISEC) and Erwin Schoitsch (AIT)

RESEARCH AND INNOVATION

ANNOUNCEMENTS

Cognitive is the New Smart

The age of smart devices and applications is not only dawning, it is here. Every one of us interacts with connected devices multiple times a day, generating vast data to be analysed across a range of industries, applications, and scenarios. The number of objects equipped with a sensor and means for processing data is increasing exponentially. Correspondingly, the realisation of the "Internet of Things" is showing promising first results. For example, in telemedicine and patient monitoring, in the automotive industry, in building automation and in smart home applications, in logistics, to name only a few.

But we are not where we could be. Today's internet-based applications are still focused on classical paradigms, like collecting and exchanging data between peers and processing vast amounts of data in centralised hubs by "hyper scaler", which apply artificial intelligence (AI) to learn and to support decision making. However, essential for the future is an infrastructure that offers extended functions for knowledge generation. The internet of the future integrates technologies which emulate the cognitive abilities of humans – our perceptions as informed by all the senses; our awareness, imagination, and memory; our ability to plan, to orient ourselves, and to learn. It forms a network of cognitive technologies, and thereby is becoming a cognitive internet. The cognitive internet provides intercompany platforms to merge data from a wide variety of sources for it to be accessed in a controlled way. What's more, AI methods are integrated at the edge, e.g. right in the sensors. This has various benefits, such as allowing knowledge to be generated and used locally in real time, and GDPR compliant processes to be supported by reducing the amount of data that must be stored and processed outside the controlled edge devices.

Researchers from various disciplines have to interact in order to bring forward key technologies from sensors to intelligent learning methods in data processing and the cloud, for integrating these really "cognitive" abilities into the future internet. In order to respond to various industries' demands to combine applied specialist knowledge for comprehensive needs, the Fraunhofer-Gesellschaft set up a research cluster combining experts from different domains, to address these demands. The Cluster of Excellence Cognitive Internet Technologies CCIT [L1] launched its mission in 2018 with 13 Fraunhofer Institutes pooling their expertise in order to face the challenges of digitalisation and develop new solutions for a "Cognitive Internet". Cyber security and data protection play an important role to protect industrial knowhow and to support data sovereignty but also benefitting from the opportunities of digitalisation. Moreover, data processed by cognitive computing technologies must be trustworthy, otherwise the AI algorithms will produce, for example, inaccurate forecasts and wrong decisions based on manipulated or faked data.

The application scenarios for cognitive technologies are quite diverse, ranging from logistics to agile and mobile industrial manufacturing to autonomous driving. A self-organising production line, for example, has to be able to identify and locate components, machine parts, to be able to adapt its processing automatically to improve the production process. For applications such as autonomous driving, it is imperative that solutions based on ultraprecise cognitive models can detect traffic situations in real time and trigger an appropriate response. This requires new solutions that provide these cognitive abilities within cognitive sensors and cognitive edge components in vehicles and infrastructures. That way, these solutions will be able to respond very quickly to the given situation, plan proactively and take actions that are coordinated with the components in the immediate surroundings.

Prof. Claudia Eckert,
Director of the Cognitive Internet Technologies CCIT cluster of excellence and Head of the Fraunhofer Institute for Applied and Integrated Security AISEC.

Fraunhofer also develops speech-driven dialog systems with a special focus on domain-specific knowledge for application in various fields of business and industry. Using and combining state-of-the-art components for speech recognition, question/answering via knowledge graphs and speech synthesis, our technologies in particular address the concrete challenges and needs of industries, enterprises and B2B applications. Moreover, these technologies "made in Europe" ensure technological sovereignty, data can be stored and processed within secure data spaces and the methods of "informed machine learning" make sure that the systems can even be trained on small data sets.

Trustworthy electronics are also in our focus. They will become a key component especially for the implementation of products and solutions using artificial intelligence and machine learning. Integrity can only be achieved if all parts of these systems - especially the electronic hardware (CPUs, SoCs, sensors, memory) as the basis of all software components based on them - are trustworthy. Backdoors and Trojans must be excluded. However, owing to the strong internationalisation of various steps in electronics development, it is now difficult to guarantee this trustworthiness. Moreover, German and European companies today have hardly any alternatives to using electronic components from untrustworthy international sources if they want to defend or expand their world market leadership with innovative, digitised, AI-based products. Within the cluster CCIT activities have been started to support a new Fraunhofer initiative that aims to build Trusted Electronics Platforms (TrEP).

The ability to share real-time information from various sources in a controlled, i.e. data sovereign, and secure manner is a prerequisite for new business models and sustainable new value chains. Connected systems across company boundaries with cognitive capabilities and highly secure shared data spaces are key to maintain leading global market positions and to benefit from digitisation and global connectivity. United research disciplines, as displayed in the contributions of this journal, are correspondingly essential for the future.

[L1] https://www.cit.fraunhofer.de/

Simula Research Laboratory joins ERCIM

Simula is an ICT research laboratory that pursues important and fundamental problems of science and engineering in order to drive progress that is of genuine value to society. Closely integrated into the research is the education of future scientists and the development of commercial ventures.

Since its establishment in 2001, Simula has grown and now comprises three research units in Norway. Despite the geographic spread, the research foci remain concentrated on five areas: cybersecurity (in Bergen), communication systems and machine learning (in Oslo), and scientific computing and software engineering (in Fornebu).

Scientific computing underpins much of today's technological and societal advances. Simula develops new mathematical tools to enable multi-scale simulation models, from the cellular to the organ level, to interrogate fundamental biomedical questions. The large amount of data generated in this era offers enormous opportunities, Simula develops new algorithms for data-driven simulations and the efficient use of modern large-scale computers.

Simula pursues important and fundamental problems of science and engineering in order to drive progress that is of genuine value to society. Closely integrated into the research is the education of future scientists and the development of commercial ventures. Photos: Simula.

Software systems are becoming increasingly intelligent, dynamic, and unpredictable. Such complex systems require innovative validation and verification paradigms, in order to handle the unpredictability of future AI-enabled systems.

As modern society becomes increasingly digitized, it is of utmost importance that our communication infrastructures are extremely robust. Simula's researchers identify and investigate issues with society's telecommunications infrastructure and aim to provide new technology and solutions to address these issues. For instance, our networks must be both robust and flexible in order to support a wide range of applications, from automation to self-driving cars to the tactile Internet.

There is also a critical need today for new cryptographic solutions, not only to secure the processing of sensitive data on cloud computing platforms but also to protect against future attacks from quantum computers. Simula's research on cyber security has two primary goals: to acquire a deep understanding of how to protect sensitive information through the use of mathematical methods, and to design secure and reliable communications and storage solutions.

Machine learning and data science touch upon all of Simula's research fields, either by contributing to the development of these fields, or by actively using machine learning as part of the research. As with the other research foci, Simula's machine learning researchers conduct basic research in order to strengthen the link between theory and algorithms, but always with an eye towards linking algorithms to high-impact applications that are of particular relevance to society.

An integral part of conducting excellent research is educating and training tomorrow's scientists and technology experts. Simula provides educational opportunities from Master's through the postdoctoral levels and these activities are closely integrated with research and innovation. Simula partners with a large number of degree-awarding institutions, both within Norway and internationally, allowing Simula to offer a range of possibilities for research stays abroad.

Innovation activities are an inherent part of technology research, as the conceptual work is applied to real life. Simula provides both pre-seed support for technology entrepreneurs, with free office space and relevant support services, as well as real financial investments for promising start ups. By realizing the commercialization of research results, both via industrial partnerships and the creation of spin-off companies, Simula keeps its focus on solving problems that add value to society.

By joining ERCIM, Simula hopes to build a stronger European network within our scientific fields.

Please contact:
Kyrre Lekve, Deputy Managing Director Simula Research Laboratory, Norway
Kyrre.lekve@simula.no

2019 ERCIM Cor Baayen Young Researcher Award for Ninon Burgos and András Gilyén

The 2019 competition for the ERCIM Cor Baayen Young Researcher Award was highly competitive with 16 final nominees. A selection committee composed of members from ERCIM and Informatics Europe eventually awarded two outstanding young researchers in recognition of the outstanding scientific quality of their research and the impact on science and society they have already achieved in their short career: Ninon Burgos (CNRS) and András Gilyén (Caltech) with an honorary mention to Paris Carbone (RISE). The award will be presented at the European Computer Science Summit (ECSS) in Rome on Tuesday, 29 October 2019.

Ninon Burgos (left) and András Gilyén, the two winners of the 2019 Cor Baayen Young Researcher Award.

Ninon Burgos, nominated by Anne Canteaut (Inria) is a researcher in the ARAMIS Lab at the Brain and Spine Institute (ICM) of the French CNRS in Paris. She joined the ARAMIS Lab in 2017 when she was awarded a PRESTIGE Postdoctoral Research Fellowship (a European “Marie Sklodowska-Curie” fellowship programme). She completed her PhD in 2016 at University College London in the Centre for Medical Image Computing under the supervision of Sébastien Ourselin. In 2012 she obtained a MSc in Biomedical Engineering from Imperial College London and an Engineering degree from the French Graduate School in Electrical Engineering and Computer Science (ENSEA). Her research currently focuses on the development of computational imaging tools to improve the understanding and diagnosis of dementia. The main focus of her PhD was the development of image synthesis algorithms for MR-based attenuation correction in hybrid PET/MR scanners, and for detecting pathological abnormalities in the reconstructed PET data. During her first postdoctoral position at UCL, she decided to tackle the problem of automatic MR-based radiotherapy treatment planning in the pelvic region by developing a new framework combining segmentation and image synthesis. Ninon has already received high recognition from the research community. She got more than ten invitations to give presentations at external meetings, received three travel awards to attend major international conferences, a prize for her oral presentation at the PSMR 2015 conference in Italy, and a prize for the best publication in the European Journal of Medical Physics “Physica Medica” in 2017.

2019 COR BAAYEN AWARD

Winners (ex aequo):
- Ninon Burgos (CNRS), nominated by Inria
- András Gilyén (Caltech), nominated by CWI

Honorary mention:
Paris Carbone (RISE), nominated by RISE-SICS

Finalists:
- Mário Antunes (IT Aveiro), nominated by INESC-ID
- Constantinos Costa (Univ. of Pittsburgh), nominated by the Univ. of Cyprus
- Pavlos Fafalios (FORTH-ICS), nominated by FORTH-ICS
- Riccardo Guidotti (ISTI-CNR), nominated by ISTI-CNR
- Lucca Hirschi (Inria), nominated by Inria
- Csaba Kerepesi (SZTAKI), nominated by SZTAKI
- Sebastian Lapuschkin (Fraunhofer HHI), nominated by Fraunhofer Gesellschaft
- Elena Mocanu (Univ. of Twente), nominated by CWI
- Panagiotis Papadopoulos (Brave Software Inc.), nominated by FORTH-ICS
- Theofanis Raptis (IIT-CNR), nominated by IIT-CNR
- Christof Weiß (Univ. Erlangen-Nürnberg), nominated by Fraunhofer Gesellschaft
- Naomi Woods (Univ. Of Jyväskylä), nominated by VTT
- Marcin Wrochna (Univ. of Oxford) nominated by the University of Warsaw

https://www.ercim.eu/human-capital/cor-baayen-award

András Gilyén, nominated by Ronald de Wolf (CWI and University of Amsterdam), is a researcher at Caltech (previously at CWI until June 2019). András graduated on May 29, 2019, at the University of Amsterdam with the distinction “cum laude”. His thesis is about quantum algorithms, and more generally about the theoretical computer science (TCS) aspects of quantum computing, an interdisciplinary field combining physics, computer science, and mathematics. His thesis is based on a number of publications in the very best conferences: three papers in STOC and FOCS, which are the two most prestigious annual conferences in TCS, and are its analogues of Nature and Science; and one paper in SODA, which is the world’s premier algorithms conference. The thesis develops a new framework for quantum algorithms, called “quantum singular value transformation”. The work of András has already been quite influential in the field of quantum computing: other researchers have realized András’s new framework is extremely versatile and powerful, and are seeking out his help to apply it to their own problems. He recently became a postdoctoral fellow at the world-famous California Institute of Technology (Caltech), working with Profs. John Preskill, Fernando Brandão, and Thomas Vidick. He has also been accepted to participate in the very selective program on quantum computing in Spring 2020 at the Simons Institute in Berkeley.

Introduction to the Special Theme

Smart Things Everywhere

by Margarete Hälker-Küsters (Fraunhofer AISEC) and Erwin Schoitsch (AIT)

Market intelligence firm IDC predicts that by 2025 each of us will be interacting with a smart device several thousand times each day. This technological change will have considerable impacts on our lives and lifestyles, and will be accompanied by new challenges, opportunities and risks. In this edition we focus on smart applications in several domains, tackle a few technological aspects and deal with security and quality issues.

The basis is laid by the Internet of Things - IoT (interactions and communication between humans and smart devices) and the Industrial Internet of Things - IIoT (in industrial context, machine-to-machine communication). However, “smart things everywhere” is not just IoT or IIoT, or mobile phones – it means intelligence, cognitive systems and technology, machine learning and artificial intelligence, security, big data and cloud connectivity, involving many domains of everyday life and a main driver for the digital transformation of our world.

International standardisation is also reacting to these evolving developments and challenges, covering technical as well ethical and societal issues. Technical aspects are handled in the Joint Technical Committee of ISO/IEC JTC1, Subcommittee SC41 (“Internet of things and related technologies”) and SC42 (“Artificial intelligence”). For IoT, the main topics are: architecture; interoperability (a key issue that “massively deployed systems work”); applications and use cases; coordination involving liaison or study groups on IoT trustworthiness (a key issue for public acceptance and liability); IIoT; societal and human factors; and blockchain security under IoT restricted resources. In the area of AI, key activities focus on a framework for AI systems using machine learning and other foundational standards; big data; and particularly on trustworthiness of AI. This covers many important issues for the applicability of AI components and systems in critical applications, like risk management, interoperability, bias in AI systems and AI aided decision making, robustness of neural networks or algorithms, but also governance and ethical and societal concerns overviews.

The European Commission as well as large organisations, e.g. in context of automated driving regulations and recommendations, have set up documents and guidelines on ethical considerations and trustworthiness of AI and computerised decision making. The most important examples are the EC High Level Expert Group’s (HLEG) Report “Ethics Guidelines for Trustworthy AI” [L1], the report of Informatics Europe and ACM Europe “When Computers Decide: European Recommendations on Machine-Learned Automated Decision Making”[L2] , and the IEEE initiative on “Ethical aligned design” (even planning certification activities) [L3].

These considerations on international activities should complement the articles reporting about applications in various domains and on quality, safety, security and risk management in general of such widely deployed systems of “smart things”. The articles are grouped into five chapters:

- Smart Industrial Applications
- Smart Cities, Buildings and Homes
- Quality, Safety, Security and Risk Management
- Smart Things Networks and Platforms
- Other Applications of “Smart Things Everywhere”.

The keynote “Cognitive is the new Smart” by Prof. Claudia Eckert, Director of the Cognitive Internet Technologies CCIT cluster of excellence and Head of the Fraunhofer Institute for Applied and Integrated Security AISEC, shows us the way beyond the current conventional use of Internet, smart devices and IoT, towards a “Cognitive Internet”. This is a network of cognitive technologies for knowledge generation and sustainable value chains, while preserving trustworthiness, Intellectual Property Rights and GDPR (General Data Protection Regulation of the EC) compliance. This will allow to maintain Europe’s leading position in digital industry and business. CCIT, the Fraunhofer Cluster of Excellence »Cognitive Internet Technologies«, is an excellent example how to achieve these goals by joining forces in research and innovation.

On European level, particularly DG CONNECT, DG for Communications Networks, Content and Technology, who manages the EC “Digital Agenda”, the HORIZON Programmes, among them ECSEL, a Joint Undertaking and PPP (Public Private Partnerships between EC, national funding authorities and industrial resp. research partners) have taken up these challenges in their Work Programs, with research projects on Smart Manufacturing, Automated Driving, Smart Farming, Multi-Concern Assurance (Safety, Security, Performance and other Dependability properties), Smart Cities and Homes, and many traversal projects across domains and challenges. The pillars of these programs for “Digitalization of European Industry”, are IoT (physical meets digital), Big Data (value from knowledge) and AI and Autonomous Systems.

In the booklet “My agenda for Europe” of Ursula von der Leyen, the new President of the European Commission, a chapter is dedicated to “A Europe fit for the digital age”. It focuses on AI, IoT, 5G, and ethical and human implications of these technologies, empowering people through education and skills, and on protecting ourselves with respect to the risks of these technologies topics that are highlighted in the keynote of this issue of ERCIM News.

Highlights provided by the articles across domains, applications and challenges can be summarised as follows:

Industrial manufacturing processes place high demands on quality, efficiency and productivity to stay competitive. In order to reach this goal, machine learning (ML) and artificial intelligence (AI) methods are used for a wide range of data analysis, process control and production steering. ML usually needs a huge amount of data. In the production environment there are situations where only a small amount of high quality data is available, e.g. during the commissioning phase. In this case ML provides methods known as one-shot or few–shot learning. Data in industrial processes are provided by a variety of sensors, one example is the application of acoustic monitoring for quality assurance. The data from acoustic sensors can e.g. define the right time for preventive maintenance.

In today's digital world, large amount of data is collected from different sources. These data have to be analysed and visualised, and anomalies have to be detected within a short timeframe. Alarms and alerts have to be parsed in real time. Often self–learning algorithms are used to solve these kinds of problems. These new technologies will affect many areas, including manufacturing, financial services, healthcare, logistics, information security and also the operation of satellite constellations as well as spacecrafts.

In logistics one example is the exploitation of IoT technologies in supply chain traceability. From a wide range of devices, such as RFID, NFC, barcodes, Bluetooth, sensors, different types of data like velocity, temperature, humidity, location can be read. These data combined with ML algorithms can improve the whole supply chain for the benefit of all stakeholders. An efficient storage and pick up system based on wireless IoT technology for smart commissioning is another application field in the area of logistics.

Artificial intelligence technologies can also be applied to improve the safety of road traffic, to visualise noise pollution or to support farmers in their daily work.

Cameras monitor an intersection and send real-time and high-quality videos to a server where objects (pedestrians, cars, etc.) are identified and critical situations might be foreseen by machine learning technologies. This information will be sent to road users who can react accordingly. Based on a distributed acoustic monitoring system noise sources, noise levels and their contribution to the noise pollution at a specific location can be detected on a detailed level. Data from different weather stations are collected and analysed in a smart way to provide accurate weather forecasts (e.g. risk of freezing).

Knowledge–based medical diagnosis and therapy systems have been around since the 80s. Today, more advanced techniques enable further application areas, for example, remote diagnosis of allergic rhinitis is possible by analysing a short phrase uttered on a mobile phone. In an industrial working environment, cameras and sensors can provide information in a non-invasive way on user activities and states in order to detect health risks.

Last but not least are the smart applications we might have in the home in the future. The article "Transforming Everyday Life through Ambient Intelligence" gives insights into our future home life.

Many of the aforementioned applications are based on a wide range of devices which are connected via a network and exchange information. To design, configure, manage and maintain these kinds of networks is challenging in terms of interoperability, scalability, energy consumption, real-time data and security. A few articles are related to this topic.

With these new technologies and developments, further challenges and risks for safety, security and privacy (users) and data arise.

A key topic is security, which has different facets. First, security by design is desirable. "Security by design" refers to the systematic collection and assessment of security goals, threats and countermeasures. Based on the results, a secure system is designed and implemented.

Many of the embedded systems do not have the same level of security features that is common in standard operation systems. This is because devices are tailored to the specific application. A new approach, based on binary rewriting to retrofit already existing IoT systems, will be investigated in order to make the embedded systems more resilient against unauthorised access attempts. It is vital for smart systems to be protected against hacker attacks as effectively as possible.

The reliability of information in very different situations has to be ensured. This becomes crucial if, for example, personal data, contractual data, financial data and/or different stakeholders are involved. One answer is offered by blockchain technology together with the IoT and artificial intelligence. Another is the Arrowhead framework, which provides a chain of trust via mechanisms like certificates and secure on boarding. The International Data Space Association is conducting work in this area.

Systems used in real life (e.g. autonomous driving cars or medical diagnosis) have to guarantee safety, which means very high quality standards.

A characteristic of artificial and self-learning systems is that they may have an unsupervised unpredictable behaviour. Based on the criticality of the application, different quality assurance and test concepts must be developed in order to guarantee that the systems are reliable. Standards and certifications have to be evolved.

One approach to this is "cooperative risk management", where all smart devices share their information and act together. Another approach is to develop a hierarchical model for qualities for smart objects. Based on existing quality models, new dimensions are added, which take into account the capabilities that arise with smart objects.

Links:

[L1] https://kwz.me/hEH
[L2] https://kwz.me/hEE
[L3] https://kwz.me/hEK

Please contact:
Margarete Hälker-Küsters
Fraunhofer-AISEC, Germany
margarete.haelker@aisec.fraunhofer.de

Erwin Schoitsch
AIT Austrian Institute of Technology, Austria
Erwin.Schoitsch@ait.ac.at

Blockchain and AI – Cyber-Physical Production Systems

by Philipp Sprenger and Dominik Sparer (Fraunhofer IML)

A new joint project between Fraunhofer IML and TU Dortmund University combines blockchain, IoT and artificial intelligence (AI). The demonstrator for "cyber-physical production systems (CPPS)" integrates smart contracts and blockchain technology into a multi-agent system at shop-floor level. The system covers negotiations, financial transactions and agile self-organisation while employing only limited hardware resources in a near real-time environment.

Today's "cyber-physical production systems (CPPS)" are largely closed IoT applications with very limited contractual capabilities and interoperability. Integrating AI and IoT into production systems enables improvement of efficiency when it comes to finding optimal job routines in self-organised systems. The missing link here is the integration of business logic and management rules, which support contractual, administrative and financial processes to transfer the concept of CPPS to industrial operations. For this reason, researchers at Fraunhofer Institute for Material Flow and Logistics IML and TU Dortmund University are developing blockchain-based CPPS in the Innovationlab in Dortmund [L1].

The blockchain provides a distributed, tamper-proof dataspace for contractual as well as event-based data. AI-equipped software agents, which are researching and evaluating different organisational rules for better job-handling, serve as designated light nodes. This infrastructural element allows for distribution of new rules to all relevant shop-floor components via the blockchain. This feature also ensures reliability and trust for the generated transactions, which can serve as the basis for smart contract deployment [1]. The blockchain also keeps track of all agent-based decisions and distributes them within the network, so they can be adapted for different circumstances. At the same time, cryptotokens enable autonomous micro-transactions and support monetising actions on the shop floor at low transaction costs. This way, a pay-per-use model can be introduced, even if the software and hardware equipment is provided by different stakeholders.

Simultaneously, the blockchain supports an identity management concept for CPPS without any additional intermediaries [2]. This means that the integration of the CPPS into the blockchain's P2P-network provides every agent with a unique digital identity. This way, every transaction can easily be tracked to the device that produced it. Furthermore, agents can be introduced to or removed from the network without having to modify the system itself, greatly increasing agility and adaptability. The network itself provides the necessary legitimation for digital identities. These redundant servers, secure communication channels and digital identities assigned to all agent systems increase cyber-security and cyber-resilience [2].

In addition to this, the integration of blockchain-technology with AI-equipped agents supports decision-making processes by providing one uniform, transparent data pool and identical criteria for decision-making, effectively removing communication efforts (see Figure 1) [1]. By using smart contracts as enablers for autonomous process flows, the agent systems within the CPPS project can thus negotiate management actions and payments via micro-transactions more easily, while the blockchain-network itself is expandable more quickly and with reduced effort [1].

The project is a cooperation between Fraunhofer IML and TU Dortmund University, with all experiments being conducted in the Innovationlab in Dortmund. The researchers have different expertise in self-organising systems and want to investigate how the integration of administrative business processes and applications can be exe-

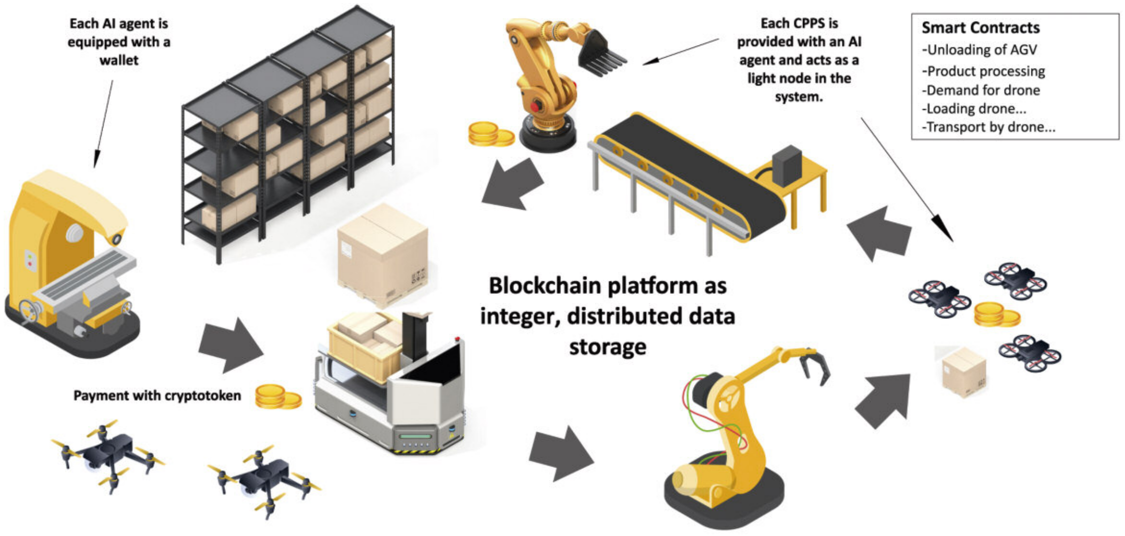
Each AI agent is equipped with a wallet
Each CPPS is provided with an AI agent and acts as a light node in the system.
Smart Contracts
-Unloading of AGV
-Product processing
-Demand for drone
-Loading drone...
-Transport by drone...
Blockchain platform as integer, distributed data storage
Payment with cryptotoken

Figure 1: CPPS-Demonstrator.

cuted successfully. The cooperation started in June 2019 and a proof of concept demonstration is planned for October 2019.

The CPPS-project is part of Fraunhofer IML's "silicon economy" vision. Within this scenario, logistics, as a major driving force of extensive digitisation, is situated at the centre of platform economies. Inspired by the B2C applications of Silicon Valley, the silicon economy is the B2B equivalent, connecting artificial intelligence, global interconnectedness, data-based business models, agile financial flows as well as smart contracting and distributed ledger technologies all in one digital habitat. Originally regarded as a mere service industry, the entire logistics discipline is expressible via algorithms and can thus serve as a powerful launch pad for distributed artificial intelligence and industrial applications of biological systems, for example for swarm intelligence. Moreover, because of its interface role, logistics also demonstrate ideal conditions for widespread implementation of those technologies across various industries and, because of its network structure, across the globe. The CPPS project is set to contribute to this vision of the silicon economy to enable fast and independent means of operation at shop floor level. Like its biological role model, self-organising swarms, the project aims to develop and propel equally skilled cyber-physical production systems for industrial applications and lay the groundwork for a large-scale introduction at a later stage.

Link:
[L1] https://www.innovationlab-logistics.com/

References:

[1] K. Salah, et al.: "Blockchain for AI: Review and Open Research Challenges", in: IEEE Access, vol. 7, pp. 10127-10149, 2019. https://doi.org/10.1109/ACCESS.2018.2890507

[2] D. W. Kravitz, J. Cooper: "Securing user identity and transactions symbiotically: IoT meets Blockchain", in: Global Internet of Things Summit (GIoTS), Geneva, Switzerland: IEEE, pp. 1-6, 2017. https://doi.org/10.1109/GIOTS.2017.8016280

Please contact:

Philipp Sprenger, Dominik Sparer
Fraunhofer Institute for Material Flow and Logistics, Germany
philipp.sprenger@iml.fraunhofer.de,
dominik.sparer@iml.fraunhofer.de

Product Quality Prediction in Batch Processes with Small Sample Sizes Using Siamese Networks

by Christian Kühnert (Fraunhofer IOSB), Johannes Sailer (Fraunhofer IOSB) and Patrick Weiß (Fraunhofer ICT)

Deep Learning algorithms usually need a large amount of data. Still, when analysing measurements from manufacturing processes, informative data in sufficient quantities is rather rare, making the task more complex. Therefore, the development of so-called few shot learning algorithms, focusing especially on the analysis of small data sets, is one of the current research topics of the Fraunhofer Cluster of Excellence.

Production processes, especially in manufacturing, must meet high standards for quality, efficiency and productivity in order to remain competitive. Owing to increasing process complexity and frequent plant conversions, however, production is often not performed efficiently. One way to increase plant efficiency is a systematic data evaluation, including data management and analysis of historical and streaming process data. For data analysis, machine learning (ML) and artificial intelligence (AI) have undergone rapid development and are currently at the centre of technical innovations. In applications with very large amounts of data available, e.g. in speech recognition or image analysis, deep neural networks are currently the state of the art.

In contrast to image analysis or speech recognition, the use of machine learning methods to analyse production data is a little more complex because even extensive production data often contains comparatively little information. For example, in a plant configured for serial production, process parameters are usually kept constant, leading to the point that always the same data set is delivered. On the other hand, data with a high information content but in small quantities, is generated during commissioning or after major modifications of the plant. An example is the use of new raw material where parameters or process control strategies have to be changed and adapted several times to achieve the desired results.

Hence, one current research topic in the Fraunhofer Cluster of Excellence Machine Learning is to make ML methods applicable for this type of data. ML methods, suitable for small data sets, are referred to as one-shot or few-shot learning and were first presented in 2006 [1]. Recent examples where one-shot learning showed good results are in the area of classifying images [2] and detecting the road profile for autonomous cars [3].

In a current use case, a method called Siamese Network [2] is used to predict the product quality of insulation panels produced on a foam moulding machine. A Siamese network, sketched in figure 1, is a neural network architecture in which two convolutional neural networks with the same weights work in parallel. Each network takes a different input vector, in this case the process data from two different productions, returning two comparable output vectors. From these two output vectors a distance measure is calculated, making Siamese networks feasible for classification tasks if one of the two input vectors is labelled. To validate the performance of Siamese networks on this type of production data, different recipes were carried out on the moulding machine and the quality of the resulting

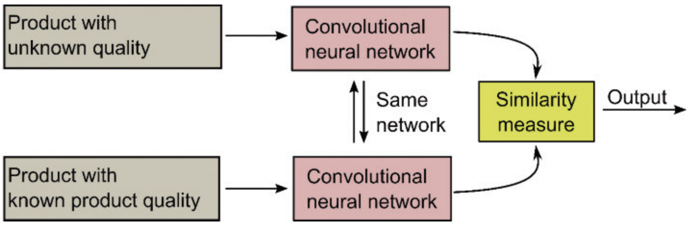
Product with unknown quality
Convolutional neural network
Same network
Similarity measure
Output
Product with known product quality
Convolutional neural network

Figure 1: rincipal design of a Siamese Network. Both convolutional neural networks share the same weights. Process data from a product with unknown quality is fed into one network and compared to a product with known quality. Depending on a defined distance measure (e.g. L2-norm) the similarity of the two products is calculated and a prediction of the quality can be made.

plates was measured. The resulting insulation plates were evaluated in terms of their welding degree, compressive strength and bending strength and separated into five classes: (1) Good quality, (2) brittle, (3) burned, (4) bent and (5) low compression. In summary, 130 experiments were conducted.

As a first step, the Siamese network was tested on the binary problem meaning to classify if a plate is of good or bad quality. For comparison and to validate the performance, a K-Nearest Neighbour algorithm (KNN) was taken as a baseline. Results showed that for the one-shot case, meaning only one good and one bad production was used for training, the Siamese network achieved on the average an accuracy of 83.2 % while the KNN ended up on 80.1 %, meaning the Siamese network outperformed the KNN by about 3 %. When taking the complete data set with a separation of 70 % training and 30 % test data, on average the network ended up with a 97.6 % classification accuracy, outperforming the KNN by about 2%.

For the classification of five classes, the network achieved a decent 52.3 % accuracy for the one-shot problem (keep in mind that 20 % chance would be a random guess) and finishing on an accuracy of 88.7 % for the whole data set. As for the binary case, KNN was outperformed by around 3 %.

In general, the results showed that even for small data sets ML methods are appropriate to analyse process data. Specifically, Siamese networks and the K-Nearest-Neighbour approach already show good results for one-shot learning in the classification of isolation plates for quality control. As expected, in a few-shot scenario, with more data available, the classification accuracy could be substantially increased. In all cases the Siamese network led to a better performance compared to the KNN.

References:

[1] L. Fei-Fei, et.al.: “One-shot learning of object categories”, IEEE transactions on pattern analysis and machine intelligence, 2006

[2] G. Koch, et.al.: “Siamese neural networks for one-shot image recognition.” ICML Deep Learning Workshop, 2015

[3] L. Huafeng et. al.: “Deep Representation Learning for Road Detection through Siamese Network”, Computer Vision and Pattern Recognition, 2019

Please contact:

Christian Kühnert
Fraunhofer IOSB, Germany
Christian.kuehnert@iosb.fraunhofer.de

Acoustic Quality Monitoring for Smart Manufacturing

by Sara Kepplinger (Fraunhofer IDMT)

“I describe this sound as ‘pling’ and the other as a ‘plong’…’”, the participant said, distinguishing between hockey pucks made from different materials. Similarly, it is possible to “listen” to acoustic quality in the context of industry. Based on a demonstrator using an air hockey table, we show an approach for acoustic quality monitoring, applicable to smart manufacturing processes.

Sensor technology already plays a major role in the industrial environment, and systems relying on machine learning (ML), sensor combinations, and data from all levels to improve production processes, are constantly evolving [L1]. However, the potential to use information gleaned from acoustic emissions has not yet been investigated in the context of industry. This is where our in-house demonstrator, an air hockey table, comes into play, enabling us to show vividly how we can recognise different material properties (in this case, different kinds of pucks) by their acoustic fingerprints. Players strike a small plastic disc (the puck), which moves on a small air cushion, from one side of the table to the other. With the help of measuring microphones and a supervised learning approach, it is possible to “listen” to which puck is being played (Figure 1).

The main advantage of our approach is the non-destructive quality control that does not interfere with the (production) process or demolish the inspected object. It combines various measurement and analysis steps: precise sound recording, pre-filtering of noise and useful sound, as well as intelligent signal analysis and evaluation using ML. Several steps are required prior to recording the data, including: determining how data will be acquired (i.e., type of microphone, position of sensors, description of the context, definition of recording length and method, sampling rate, resolution); context specific synchronisation with additional and other (sensor) data, and the definition of an annotation policy.

Once these requirements are defined, the actual recording of data considering different conditions for the later classification training takes place. In our case, the conditions include audible differentiation of three pucks of different weights (18/19/23 grams) and materials (3D printed (material: PA2200)/3D printed with rubber ring (material: PA2200)/hard plastic original), as well as different kinds of (simulated) background noises and player behaviour. We record continuously using an USB audio interface (into a mini PC in wave format using two directional microphones, one for each side of the air hockey field), with 44.1 kHz and 32 Bit. After the training the system based on various predefined conditions, the ad-hoc analysis and interpretation of data takes place. The classification starts during power-up of the air hockey table and is able to detect the different kinds of pucks on the field.

In an industrial environment, the installation of acoustic measurement technology is similar: a measurement microphone picks up the characteristic sounds of a given process. Depending on the application, we do the recording with several microphones or a microphone array. Algorithms of source separation (as one possible solution among others) filter out ambient noises, so that we separate background noise from useful noise. For this purpose, we use methods of source separation [1], originally developed to separate individual instruments from a music recording [2].

We use this approach, for example, in in-line quality assurance and predictive maintenance scenarios for industrial use cases. Here, the acoustic monitoring system can provide additional, and even more precise, information about the quality and condition of products or processes. This may be either integrated into existing measurement systems or used as a completely independent monitoring system.

The human auditory system is able to detect and distinguish individual sounds, even in noisy environments, and place them in specific contexts. In industrial manufacturing processes, it is often possible to make statements about the operating status of a component, engine or even an entire machine by listening attentively; experienced machine operators can hear problems or errors. However, so far it is quite difficult to detect these indications of faulty processes or products by any other means [3]. Our basic premise is that everything that is audible (and interpretable, e.g. recognisable as a difference) is also measurable as an indicator for quality.

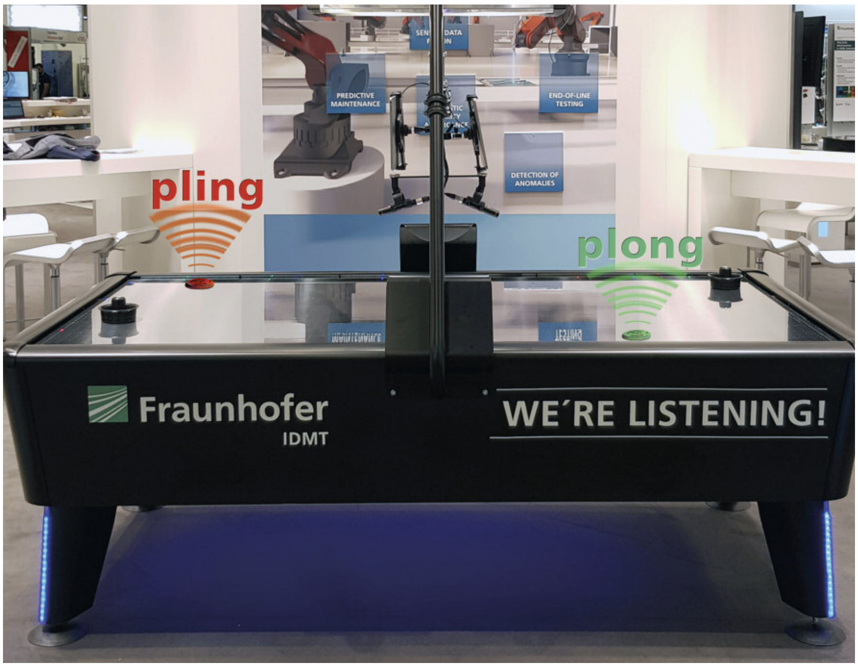

Figure 1: Air hockey table equipped with two measurement microphones recording the puck's sound (e.g., 'pling' and 'plong').

In order to reliably identify and automatically classify machine noises, the analysis steps mentioned above are indispensable. Until now, extensive training data has been required to train the respective system reliably. This is one of the challenges when it comes to optimising the recognition performance in practical use. One of the current challenges is the lack of available measurement data, which is why, as a first step, we have generated test data based on three application examples (electric engines, marble track, and bulk tubes) and made it available for testing purposes [L2]. At this stage, we have successfully tested our aforementioned method in practice, together with various industrial partners successfully, and reached the Technology Readiness Level (TRL) 6.

In the future, we would like to achieve a high recognition rate with less training data. Furthermore, we plan to develop a self-learning system that learns from acoustic measurement data to assess the quality of products and processes. Additionally, we are working to understand and interpret what acoustic information is heard by machine operators and inspectors, and how they perceive it. To this end, we are working on a parameterisation of the subjective factors.

Links:

[L1] J Walker, "Machine Learning in Manufacturing – Present and Future Use-Cases" https://emerj.com/ai-sector-overviews/machine-learning-in-manufacturing/.

[L2] Industrial Media Applications Datasets 2019. https://www.idmt.fraunhofer.de/datasets.

References:

[1] E., Cano, et al.: "Exploring Sound Source Separation for Acoustic Condition Monitoring in Industrial Scenarios", EUSIPCO2017.

[2] J., Abeßer, et al.: "Acoustic Scene Classification by Combining Autoencoder-Based Dimensionality Reduction and Convolutional Neural Networks", in: DCASE 2017.

[3] S., Grollmisch, et al.: "Sounding Industry: Challenges and Datasets for Industrial Sound Analysis (ISA)", in Proc. of EUSIPCO 2019.

Please contact:

Sara Kepplinger
Fraunhofer Institute for Digital Media Technology IDMT, Germany
sara.kepplinger@idmt.fraunhofer.de

Smart Pick-by-Light for Efficient Storage and Production Processes

by Hanna Herger (Fraunhofer IIS) and Thomas Windisch (Fraunhofer IIS)

The requirements for efficient production and manufacturing processes are becoming increasingly complex in global competition. Digitalisation should increase not only quality but also flexibility and the mobility of processes. The solution for meeting these requirements can be "wireless IoT technologies", which form a basis for networking and interaction between machines (M2M) as well as between information systems and the physical world. For IoT applications in the field of production and logistics, in particular, "wireless IoT technologies" are increasingly used to digitise processes and make them more efficient. This article shows how the field of commissioning is optimised by the wireless networking technology s-net®: Wireless pick-by-light systems enable more efficient and flexible order-picking processes and support interactions with people.

Wireless IoT technology for smart commissioning

The Fraunhofer Institute for Integrated Circuits (IIS) located in Nuremberg, concentrates on topics in the field of positioning and networking such as self-organising, wireless networking technology, wireless data transmission and the integration of sensor networks in production and logistics applications.

A technology developed by Fraunhofer IIS for extremely energy-efficient radio networking is s-net®. In recent years, it has formed the basis for the interaction between information systems and the physical world in increasing numbers of projects. In such cyber-physical systems (CPS), context-aware objects can record and transmit their application and environmental situation and interact independently and self-organising with several users, via s-net®. The integration of context-aware, networked objects is the key for a continuous digitisation of process-control and the Internet-of-Things (IoT).

Numerous IoT applications have already been implemented on the basis of s-net® networking technology. Especially in the context of Industry 4.0 and Smart Production, s-net opens up enormous potential. The digitised production benefits from s-net®, inter alia, in the areas of supported assembly and intralogistics, e.g. smart picking systems.

Commissioning is often the central function of warehouse logistics and has significant influence on other business areas, such as production and assembly. Digitisation is also changing the requirements for order picking: in addition to increasing efficiency and reducing the error rate, flexibility and mobility (for example, when changing product ranges or order fluctuations) are becoming increasingly important criteria. By using wireless pick-by-light systems, these requirements can be met.

Wireless pick-by-light: flexibility for mobile applications

Pick-by-light systems support the order picker with light signals for the intuitive and quick location of a storage compartment. Specific information such as the withdrawal quantity can be conveniently and quickly indicated on the shelf display - which results in a significantly higher picking performance. At the same time, a lower error rate is achieved and costly rework is reduced. However, common cable-connected pick-by-light systems only partially meet the requirements of flexible and low-effort picking. Fraunhofer IIS has therefore developed a wireless picking system based on the s-net networking technology, which has the advantages of previous pick-by-light systems and also meets the requirements for flexibility and scalability.

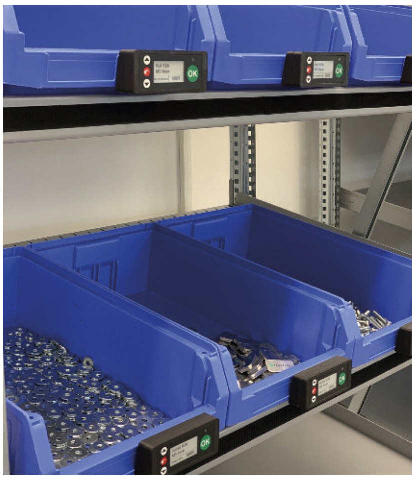

Figure 1: Wireless order handling solutions with the networking technology s-net® by Fraunhofer IIS.

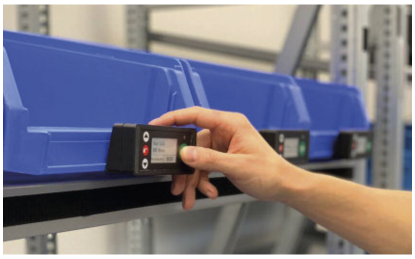

Figure 2: The self-organized radio communication between the individual compartment displays of the pick-by-light system allows an easy installation and is ideally suited for temporary storage structures and rapid remodelling of removal compartments and shelves.

With the aid of s-net technology, the pick-by-light system can be implemented wirelessly and mobile. Due to the extremely high energy efficiency, the shelf displays can last up to several years battery-powered. The self-organised communication between the individual compartment display nodes of the pick-by-light system makes it possible to easily install the compartment display on the shelves, and facilitates temporary storage structures and a rapid remodelling of removal compartments.

Optimisation of individualised assembly processes

A trend in the manufacturing industry, as in automobile production, is the increasing individualisation of products. Customer requirements need to be realised easily and cost-effectively up to batch size 1. For manual assembly processes, the complexity in terms of quality and efficiency increases, since different components have to be arranged in different orders for each product. This increases the demands on employees in terms of ensuring quality and efficiency of the assembly. The wireless pick-by-light system allows required objects, such as tools and correct components, to be made available in the right order directly at the assembly site.

TRILUM – Mobile Order Handling Solution

The wireless pick-by-light system [L3] based on s-net® technology is already available as a commercial product "TRILUM" from the cooperation partner AST-X GmbH. TRILUM guides employees to the correct storage position via LED light signal. Specific information such as the removal quantity can already be viewed on-site on an e-paper display in a convenient and time-saving manner. With the aid of an acknowledgment button, the removal of goods can be confirmed quickly and easily - the feedback to the warehouse management system is fully automatic and paperless.

TRILUM is mainly designed for use in logistics and production and especially suitable for flexible, temporary and mobile storage structures. The system is particularly easy to adapt, making it ideal for environments where take-up compartments and shelves are often redesigned or changed:

- Order picking processes: Light-guided order picking optimises warehousing processes in logistics and intralogistics - cost-effectively and paperless.
- Assembly operations: During assembly, required objects such as tools and correct components can be provided in the correct order.
- Distribution processes: Light-controlled distribution processes (put-to-light) ensure time-efficient, cost-efficient and error-free picking processes.

Links:
[L1] http://www.s-net-info.com/
[L2] https://trilum-solutions.com/
[L3] https://kwz.me/hEd

Reference:

[1] Hupp, Loidl & Windisch (2017), "Wireless IoT Für reinen Batteriebetrieb" in elektroniknet.de https://kwz.me/hEc

Please contact:

Hanna Herger, Thomas Windisch
Fraunhofer Institute for Integrated Circuits IIS, Germany
hanna.herger@iis.fraunhofer.de, thomas.windisch@iis.fraunhofer.de

Efficient Supply Chain Traceability Using IoT Technologies

by Alexandros Fragkiadakis (FORTH), Theoharis Moysiadis (Future Intelligence Ltd) and Nikolaos Zotos (Future Intelligence Ltd)

Traceability allows the identification and tracking of products as they travel through the supply chain, from the manufacturer to the consumer. Smart Product [L1] is a research project aiming to design and develop a secure IoT-based architecture for supply chain traceability, product origin verification and authenticity certification.

The Internet of Things (IoT) has revolutionised the technological means that have made feasible the interconnection and interaction between the physical and digital worlds. Billions of devices, worldwide, sense the physical world (ambient temperature and light, humidity, weather conditions, etc.), leading to the proliferation of numerous applications and initiatives including smart cities, e-health, precision agriculture, etc. Another area that IoT technologies can be successfully exploited is supply chain (SC) traceability, where products are recorded as they travel from the manufacturer to the consumer. An IoT system can read data from a plethora of devices such as smart tags (RFIDs, NFC, Barcodes, Bluetooth Low Energy), along with sensory data like ambient temperature and humidity, vehicle speed, geolocation, and if intelligently combined together, can effectively track the SC.

SC tracking offers numerous advantages to all parties involved (producers, retailers, consumers): food safety (quality deviation management, recalls in food crisis), perishable product protection, origin verification and brand certification, customer engagement and loyalty programs, monitoring, control, planning and optimisation of business processes remotely [1]. Furthermore, the utilisation of IoT technologies can offer low cost services, as the cost of sensors and other related equipment (tags, etc.) significantly decreases over time, owing mainly to the technological advances of the hardware manufacturing industry. At the same time, IoT technologies (communication protocols, interoperability standards, IoT/cloud architectures, security & privacy algorithms) have substantially matured. For these reasons, IoT-based SC tracking appears as an appealing and feasible solution.

The aim of the Smart Product project [L1] is the design, implementation and evaluation of an IoT-based platform for SC tracking, product authenticity certification and verification of origin, and consumer engagement. The platform will utilise IoT technologies at the device level (tags, sensors), as well as software and communication protocols, with emphasis on interoperability, energy-efficiency, security and privacy. This project aims to cover various weaknesses of existing SC tracking systems, such as: lack of monitoring beyond retail, lack of transparency, lack of usability, flexibility and immediacy, fragmentation of technologies and lack of interoperability.

At operational level, the Smart Product project will provide technological solutions by making the following feasible (Figure 1):

- SC tracking by recording product-related information such as storage conditions, geographical locations of production, storage and sales. This will make it possible for the producers to monitor the products, partially addressing the availability of products in areas where for various reasons it is not allowed (e.g. exporting

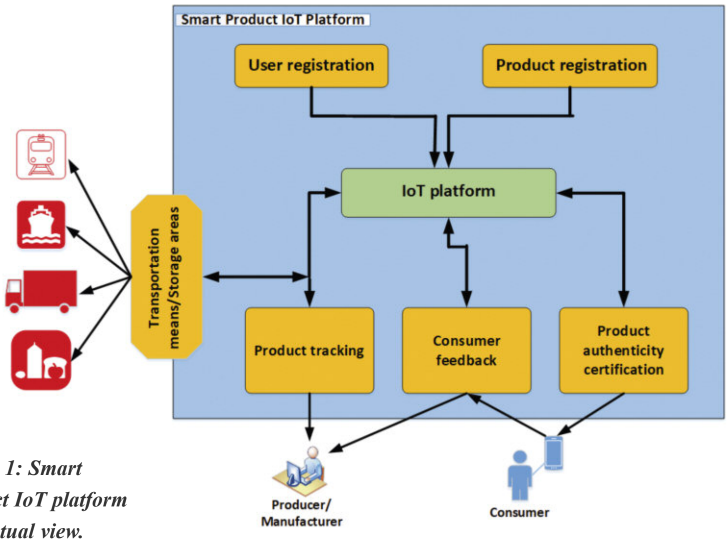

Figure 1: Smart Product IoT platform conceptual view.

smuggled products to other countries), but will also enable consumers to be informed about the origin of the products and their storage conditions during transportation. Prerequisites for successful monitoring and recording are: the existence of a smart label on the product or a package containing a number of identical products, the existence of equipment for reading the smart labels, the existence of equipment for measurement of storage conditions, such as ambient temperature and humidity sensors, and an appropriate database for data storing.

- Certification of products' authenticity and origin, limiting the distribution of counterfeit products and thereby increasing consumer confidence in genuine products. The prerequisites for making this possible are: the introduction of relevant product information into an appropriate database by certified users (producers, etc.), the existence of a smart label on the product that indicates with appropriate coding the identity of the product, and appropriate equipment for reading the smart label, such as a smart phone with the ability to read such labels.
- Stimulate consumer engagement with products through a smart phone application. Consumers, in addition to certifying the authenticity of the products, will be able to obtain further information on the product in question, such as its ingredients, how to use it, etc. By creating appropriate user interfaces on social networks, consumers will be able to further engage with producers/ marketers through appropriate marketing methods.

The platform will be demonstrated and evaluated through a concrete use-case (Figure 2) that employs producers/manufacturers and consumers, consisting of three steps:

- Step#1: A producer/manufacturer registers their product on the project platform, receiving a unique code that identifies it (e.g. Electronic Product Code) that is further stored in the platform's database. It will also be possible to state the monitoring parameters during product traceability (storage temperature, humidity, etc.). Product related information such as optimum usage patterns, ingredients, geographical location of production, expiry date and activation of feedback to producers/manufacturers will also be posted. In addition, it will be determined which of the above information will also be available to the consumers.
- Step#2: The products are monitored during their shipment from the place of manufacturing to that of disposal. Both in the storage and in the transport vehicles, there will be sensors that will record the monitoring parameters selected in Step#1. At the same time, the unique code of the transported products will be recorded and linked to the monitoring parameters while stored in the database.
- Step#3: The product is already on the shelf and steps # 1 and # 2 have been followed. The product bears a smart label that is read by a smart phone application that further communicates with the platform and collects all available data about the product and imprints it on the mobile screen for consumer information. Through the same application, consumers can send feedback to the producer/manufacturer, and can post on social networks their experience with the specific product.

The consortium consists of two partners, FORTH and the Future Intelligence Ltd that receive funding from the Operational Program of the Region of Epirus, Greece.

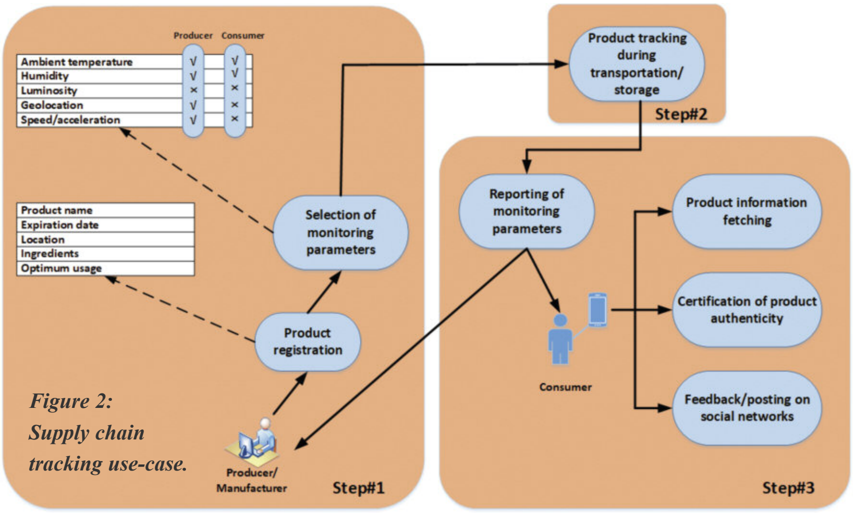

	Producer	Consumer
Ambient temperature	√	√
Humidity	√	√
Luminosity	×	×
Geolocation	√	×
Speed/acceleration	√	×

Product name
Expiration date
Location
Ingredients
Optimum usage

Selection of monitoring parameters
Product registration
Producer/ Manufacture
Step#1
Product tracking during transportation/ storage
Step#2
Reporting of monitoring parameters
Consumer
Product information fetching
Certification of product authenticity
Feedback/posting on social networks
Step#3

Figure 2: Supply chain tracking use-case.

Link:
[L1]: https://www.smartproduct.gr

Reference:

[1] C. Verdouw et al.: "Virtualization of food supply chains with the internet of things", Journal of Food Engineering, Elsevier, 2015.

Please contact:

Alexandros Fragkiadakis
ICS-FORTH, Greece,
alfrag@ics.forth.gr

Theoharis Moysiadis, Future Intelligence Ltd, Greece
tmoysiadis@f-in.gr

Nikolaos Zotos, Future Intelligence Ltd, Greece, nzotos@f-in.eu

Components and Tools for Large Scale, Complex Cyber-Physical Systems Based on Industrial Internet of Things Technologies

by Apostolos P. Fournaris and Christos Koulamas (ISI/Research Center ATHENA)

The application of Industrial Internet of Things (IIoT) technologies in large scale and complex cyber-physical systems of systems (CPSoS), such as those on large, tertiary sector buildings, energy grid and industrial production lines, still presents great challenges in the standardisation of the necessary mechanisms, procedures, components and tools for the deployment, configuration, commissioning and maintenance of their associated highly heterogeneous distributed subsystems. The problem is magnified if one also considers the typical, resource constrained nature of most of the networked embedded devices in a CPSoS, as well as the usually strict requirements in properties related to correct and safe operations, such as real-time, reliable and secure processing and net working [2].

In contrast to more restrained environments, e.g. modern smart home automation applications, in medium and large-scale installations, the costs of manual device configuration or substitution and reinstallation can easily reach unacceptable figures. Examples of such a scenario can be, the typical deployment procedures to be followed in a smart energy meter installation within a building where centralized IT infrastructure may not yet exist and for which different engineering teams may be responsible for different phases of the installation.

In I3T, “Innovative Application of Industrial Internet of Things (IIoT) in Smart Environments” project, several objectives and activities aim at the simplification of the aforementioned procedures by the design and implementation of the necessary components and tools for the deployment, configuration and management of the industrial wireless networked embedded systems existing at the base of the IIoT technologies ecosystem. We utilise the latest developments in IETF standardisation activities around the “thin-waist” of the IoT, being the IPv6-over-X/RPL/UDP/DTLS/CoAP related RFCs, and mainly the results of the 6TiSCH and ACE WGs as well as the latest development in embedded system on chip devices (SoC) that consist of embedded processors and on chip FPGA fabric. The activities involve the study of the highest achievable degree of the automatic initialisation and reconfiguration for different application scenarios, by identifying the different sets of parameters that need to be set off-line or online, and the limits between sets of parameters initialised statically at device implementation, dynamically during the deployment / commissioning time, or dynamically during the normal system operation and the whole system’s lifetime. Reconfigurability is also extended on the hardware level, where computationally demanding operations are implemented as hardware components that are dynamically

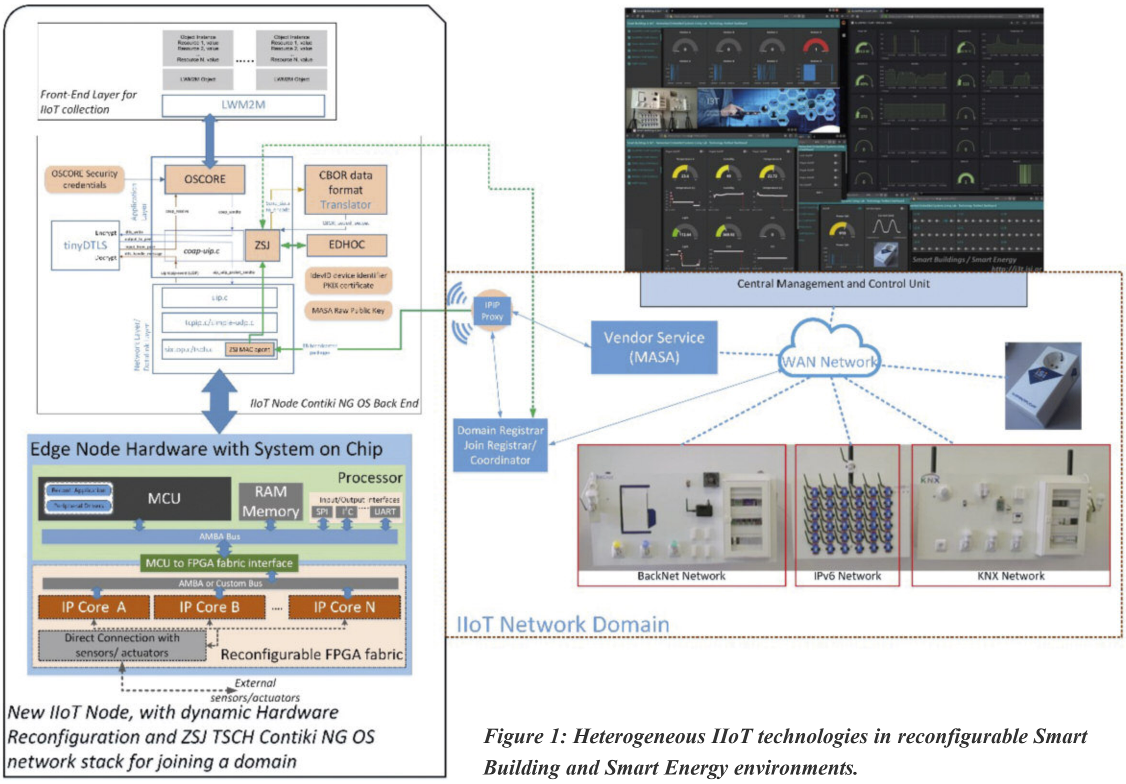
Object Instance
Resource 1, value
Resource 2, value
Resource N, value
LWM2M Object
Front-End Layer for IIoT collection
LWM2M
OSCORE Security credentials
OSCORE
CBOR data format Translator
Application Layer
tinyDTLS
Encrypt
Decrypt
coap-uip.c
ZSJ
EDHOC
IdevID device identifier PKIX certificate
MASA Raw Public Key
uip.c
tcpip.c/simple-udp.c
ZSJ MAC agent
Network Layer/ Datalink Layer
IIoT Node Contiki NG OS Back End
Edge Node Hardware with System on Chip
Processor
MCU
RAM Memory
Input/Output Interfaces
SPI
I²C
UART
AMBA Bus
MCU to FPGA fabric interface
AMBA or Custom Bus
IP Core A
IP Core B
IP Core N
Direct Connection with sensors/ actuators
Reconfigurable FPGA fabric
External sensors/actuators
New IIoT Node, with dynamic Hardware Reconfiguration and ZSJ TSCH Contiki NG OS network stack for joining a domain
Central Management and Control Unit
IPIP Proxy
Vendor Service (MASA)
WAN Network
Domain Registrar Join Registrar/ Coordinator
BackNet Network
IPv6 Network
KNX Network
IIoT Network Domain

Figure 1: Heterogeneous IIoT technologies in reconfigurable Smart Building and Smart Energy environments.

allocated on the FPGA fabric of modern cyber-physical system embedded processors.

The I3T architecture is associated with the network and end nodes of an IIoT infrastructure. We adopt industrial networks that use the TSCH operation of IEEE 802.15.4e and structure on top of it, a variation of the Zero-Touch Secure Join protocol [1] that enables the secure deployment of a new node in an IIoT network without user intervention. We further enhance this mechanism so that it can perform reconfiguration of the TSCH-6Top scheduling functions using secure CoAP messages. In parallel to that, we introduce a hardware/software co-design mechanism inside each IIoT end node SoC so that each device can support dynamic reconfiguration of its hardware resources. More specifically, the designer has available a series of hardware IP cores (with associated software drivers) that can be dynamically deployed after an application functionality analysis on the IIoT device. The analysis can highlight the computationally demanding operations that need to be accelerated by hardware in order to retain the IIoT real-time responsiveness and safety requirements. Target examples for hardware acceleration can be security/cryptographic primitive operations, and machine learning processing on fast and large time series data from the monitoring of electrical grid critical parameters or from machine and structural elements vibration sensors [3].

As an outcome of the above activities, we created an integrated, medium scale demonstrator for the smart energy and smart building environment, which includes widely used, wired and wireless, heterogeneous technologies of the relevant domains (e.g. BACNet, KNX, etc) which can also be attached to virtual environment simulators, in a hardware-in-the-loop fashion, capable of demonstrating the system operation at a larger scale. We have also designed and developed an IIoT node prototype that can support the I3T dynamic hardware accelerated reconfiguration and handle the TSCH software network stack (with all the proposed dynamic, zero touch secure I3T enhancements) [3].

Link:
[L1] https://i3t.isi.gr

References:

[1] M. Richardson: “6tisch Zero-Touch Secure Join protocol,” Internet Engineering Task Force Draft

[2] C. Koulamas et al.: “IoT components for secure smart building environments”, Springer, 2017

[3] A. Fournaris et al.: “Introducing Hardware-Based Intelligence and Reconfigurability on Industrial IoT Edge Nodes”, IEEE Design & Test, 2019

Please contact:

Apostolos Fournaris
Industrial Systems Institute / R.C. “Athena”, Greece
fournaris@isi.gr

Smart Municipality

by Jennifer Wolfgeher (FH Burgenland), Mario Zsilak (Forschung Burgenland) and Markus Tauber (FH Burgenland)

Digitalisation is already supporting individuals and smart cities in various ways. To increase automatization and digitisation, decisions must be based on trustworthy information. We are investigating the most common features of citizen participation and smart city platforms with the aim of determining the trustworthiness of the digital environment in this context.

Increased digitisation and automatization require applications that make it easy for citizens to report errors and be informed about possible incidents in their municipality. Such an incident management system can be citizen-based or automated via IoT (Internet of Things). An interaction between a citizen and the local authority may be to report an incident, e.g., an open manhole. An automated (IoT supported) version of this interaction would be the collection of relevant data from sensors. In any case, as actions are being triggered, both the citizen and the local authority needs information to be trustworthy.

Hence, on the municipality side, reliable and trustworthy data is essential, just as it is vital for citizens to know that they can rely on information they receive - such as disaster warnings. Trust is the most important requirement in emergency situations, in particular, but also in less extreme events like roadworks or construction areas. Existing smart city applications and citizen participation platforms provide features relevant for such scenarios.

In the literature many platforms and applications have been investigated [1],[2] on the basis of their features and potential to enhance sustainability, but the need for security and trustworthiness is rarely addressed. This is reflected in real-life citizen participation and smart city platforms (e.g. platforms from Austria and Luxemburg, smart city projects from Stockholm and Singapore [L1-L4]). In these smart cities, the applications designed to increase citizen participation and make residents lives easier, rarely guarantee the reliability of exchanged information. Whilst including a number of useful features (like interactive maps, online participation in political discussions, waste service information and smart parking), no existing system provides incident management as an application in a trustworthy IoT environment. To enable reliable incident reports from citizens, information must be communicated to the municipality in a trustworthy manner, allowing the local authority to take action in dangerous situations (e.g. open manholes). On the citizen’s side, only approved information about actual conditions from the municipality or sensors must be received, in order to ensure people’s safety. Thus, a trustworthy framework to enable secure communication and to link individual services from existing heterogeneous platforms is required.

Figure 1: Use case - incident report.

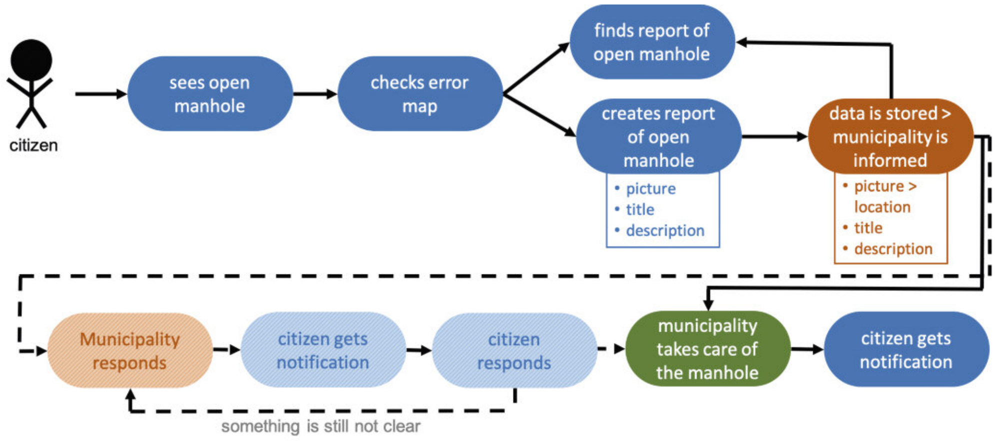
citizen
sees open manhole
checks error map
finds report of open manhole
creates report of open manhole
• picture
• title
• description
data is stored > municipality is informed
• picture > location
• title
• description
Municipality responds
citizen gets notification
citizen responds
municipality takes care of the manhole
citizen gets notification
something is still not clear

The Arrowhead framework [L5] is one example of a secure IoT framework that can integrate smartphones, sensors and external services in a smart city context. Sensors are already involved in the Arrowhead framework in different contexts, including industrial IoT (IIoT) and smart homes [L5]. This open-source framework provides many security functions by design, with the objective to facilitate security, reliability, real-time communication and safety in local cloud automation. With the chain of trust principle, it enables the usage of all services of the Arrowhead local cloud and the services of other clouds that are compliant with the Arrowhead framework, based on a secure on-boarding procedure.

In addition to the security advantages, Arrowhead supports a multi-cloud solution that would make it possible to stick together the “patchwork” of applications for smart city and citizen participation projects, e.g. the many applications of the smart nation of Singapore [L4] could be accessed via one service. The citizen could find the waste service as well as the parking service or tax service in one place. The applications could be in one cloud or distributed in separate, Arrowhead compliant, clouds, but the chain of trust would still be available from the citizen to each service.

The Arrowhead framework provides a chain of trust via various mechanisms, including certificates and secure on-boarding [3] to enable a trustworthy environment. Arrowhead is a fitting framework for the deployment of an incident report service in a trustworthy environment, capable of meeting the objectives of security, reliability, real-time communication and safety.

Further research, in the EFRE project “civis 4.0 patria” (FE07) will include the actual deployment of an incident management feature considering frameworks like Arrowhead to support the development of smart municipalities.

Links:

[L1] https://www.buergermeldungen.com
[L2] https://kwz.me/hEN
[L3] https://kwz.me/hEe
[L4] https://www.smartnation.sg
[L5] https://www.arrowhead.eu/

References:

[1] O. Gil, M. E. Cortes-Cediel, I. Cantador: “Citizen participation and the rise of digital media platforms in smart governance and smart cities”, Int. Journal of E-Planning Research (IJEPR), 8(1), 19–34, 2019.

[2] A. M. Pozdniakova, et al.: “Smart sustainable cities: The concept and approaches to measurement”, Acta Innovations, (22), 5–19, 2017.

[3] A. Bicaku, et al.: “Interacting with the arrowhead local cloud: On-boarding procedure”, in 2018 IEEE Industrial Cyber-Physical Systems (ICPS), pp. 743-748. IEEE, 2018.

Please contact:
Jennifer Wolfgeher
FH Burgenland, Austria
jennifer.wolfgeher@gmail.com

Teaching Sustainability and Energy Efficiency with the GAIA Project

by Georgios Mylonas (Computer Technology Institute & Press “Diophantus”) and Ioannis Chatzigiannakis (Sapienza University of Rome)

Today’s students are the citizens of tomorrow, and they should have the skills and tools to understand and respond to climate change. Green Awareness in Action (GAIA) has built an IoT infrastructure within 25 schools in Europe, in order to enable lectures that target sustainability and energy efficiency, based on data produced inside school buildings. The school community has reacted very positively to this approach and has reduced energy consumption as a consequence.

GAIA [L1], a Horizon2020 EC-funded project, has developed a large-scale IoT infrastructure in a number of school buildings in Europe. Its primary aim is to raise awareness about energy consumption and sustainability, based on real-world sensor data produced inside the school buildings where students and teachers live and work.

Overall, 25 educational building sites participated in GAIA, located in Sweden, Italy and Greece. The IoT infrastructure installed in these buildings monitors in real-time their power consumption, as well as several indoor and outdoor environmental parameters. However, this infrastructure would not be particularly useful without a set of tools to allow access to the data produced and provide functionality to sup-

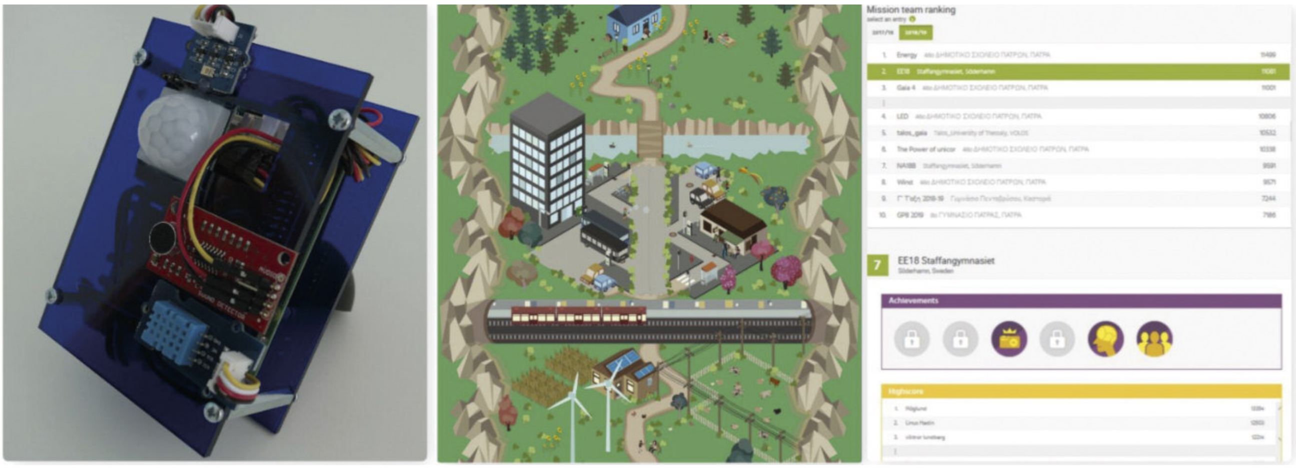

Figure 1: Some of the basic elements in the GAIA approach (from left to right): an IoT node installed inside classroom measuring environmental parameters, the world map of the GAIA Challenge on top of which students move between GAIA's "missions", and the ranking of different schools and classes participating in the Challenge.

port educational activities. The GAIA Challenge [L2] is a playful interactive platform aimed at students, designed to serve as an introduction to power consumption and energy saving. In addition, real-time data from sensors in the buildings and participatory sensing help to visualise the real-life impact of the students' behaviour and enable competitive gamification elements among different schools. The GAIA Building Manager online application offers visualisation of energy and environmental data.

The GAIA Challenge has been a success in GAIA's software portfolio, with over 3,750 registered students and teachers using it as an introduction to sustainability and energy efficiency concepts. The end-user audience of GAIA has largely comprised primary and junior high school students, but some high school and technical college students as well. This variability in students' age and origin has led to the use of a range of different approaches when applying GAIA's tools within educational activities.

The project completed its official activities in May 2019. One of its biggest achievements was the construction of an operational and reliable large-scale IoT infrastructure within 25 school buildings in Greece, Italy and Sweden. This infrastructure [2] includes over 1,200 sensing endpoints, which have been based on open-source components. By combining the tools and project methodology with data produced inside school buildings during related activities in the schools, we have energy savings of up to 15-20 %.

An important factor to encourage engagement is competition: students were intrigued by the prospect of competing with other schools and countries, and were further motivated to participate in GAIA's competitions for energy savings and related ideas. The project also held two official competitions between schools during educational years 2017-18 and 2018-19, which proved a valuable tool in motivating and engaging the schools participating in the project.

Following GAIA's completion, the network of schools built during the project, as well as the software tools developed, will continue to be active in the following school year (2019-2020). The consortium members will collaborate with other research teams and contribute to the community in this field via sharing datasets for experimental evaluation. There are also methodologies and educational activities available on the project website [L1], offering the opportunity to replicate the results of the project to any school/team wishing to do so.

This work has been supported by the EU research project "Green Awareness in Action" (GAIA), funded by the European Commission and the EASME under H2020 and contract number 696029.

Links:
[L1] http://gaia-project.eu/
[L2] https://gaia-challenge.com/

References:

[1] G. Mylonas, et al.: "Enabling Energy Efficiency in Schools based on IoT and Real-World Data", in IEEE Pervasive Computing, Vol. 17, Issue: 4, 2018, https://doi.org/10.1109/MPRV.2018.2873855

[2] D. Amaxilatis, et al.: "An IoT-based solution for monitoring a fleet of educational buildings focusing on energy efficiency", in MDPI Sensors, Special Issue in Advances in Sensors for Sustainable Smart Cities and Smart Buildings, 17(10): 2296, https://doi.org/10.3390/s17102296

[3] G. Mylonas, et al.: "An educational IoT lab kit and tools for energy awareness in European schools", Int. Journal of Child-Computer Interaction, Vol. 20, 2019, pp. 43-53, ISSN 2212-8689, https://doi.org/10.1016/j.ijcci.2019.03.003.

Please contact:

Georgios Mylonas
Computer Technology Institute & Press "Diophantus", Greece
mylonasg@cti.gr

Smart Intersections Improve Traffic Flow and Road Safety

by Martin Striegel (Fraunhofer AISEC) and Thomas Otto (Fraunhofer IVI)

Smart intersections help to address increasing traffic density and improve road safety. By leveraging data from infrastructure sensors, and combining and supplying those data to road users, their perception can be improved. This aids in protecting vulnerable road users (VRUs) and acts as a crucial building block for enabling automated and autonomous driving.

Increasing volumes of traffic are using municipal road infrastructure, with severe consequences for traffic efficiency and the safety of road users. Vulnerable roads users (VRUs), such as pedestrians or cyclists, are involved in 46 % of lethal accidents [1]. Exchanging information between road users increases their perception and is thus a critical building block to improve this situation.

The Smart Intersection, developed by Fraunhofer Institutes IVI, AISEC, HHI and IIS in the IoT-COMMS project [L1] from 2018 to 2019, installs cameras at traffic junctions to monitor traffic. Those cameras send real-time and high-quality video to a road-side-unit (RSU), which detects and classifies objects such as pedestrians or cars. Together with additional information such as position, speed and direction of movement, they are stored in a dynamic object map. The object map is then transferred from the RSU to cars via WLANp or LTE/5G-V2X.

This approach facilitates the cooperative perception of surroundings, enabling road users to utilise information from other sensors. Thus, they can recognise obstacles and other road users out of plain sight, which prevents accidents and allows more efficient traffic flows (Figure 1).

While modern cars can already utilise sensor-based object detection, parametrisation and categorisation of objects from within the moving car is challenging. Shifting those tasks to road infrastructure, on the contrary, allows reliable distinction between static and dynamic objects. Using spatially stationary cameras permits “learning” the area under surveillance, enabling early detection of critical situations and a quick reaction to them.

Previous approaches lack the bandwidth to transmit raw camera images between infrastructure components and thus transfer only detected objects. This limits the effectiveness of the object detection and impedes the comparison of detected objects, as each camera sees only a particular part of the big picture.

Figure 1: The smart intersection warns cars about the presence of pedestrians. This improves the safety of pedestrians and other vulnerable road users.

Design

The smart intersection, on the other hand, utilises high-speed mmWave transmission technology by Fraunhofer HHI to transmit raw video frames from multiple cameras to the RSU. There, object detection is performed using background subtraction algorithms developed by Fraunhofer IVI. Subsequently, objects are classified, for example “passenger car”, “cyclist” or “pedestrian” from multiview data. To do so, Fraunhofer HHI employs convolutional neural networks. In addition, classification accuracy is improved over time and trajectories can be calculated, permitting tracing of temporarily occluded objects.

The key benefit of this centralised detection and classification approach is that information from all cameras can be used. Thus, the same object can be seen in multiple perspectives, improving the detection quality significantly and increasing the area that can be perceived.

The object map is then written into a standardised collective perception message (CPM) and transferred from the RSU to cars via WLANp or LTE/5G-V2X. Thus, vehicles are enabled to expand the range of their own sensors significantly and safety-critical situations can be detected and controlled earlier or more safely at higher speeds.

Security and Privacy Concept

Another key component of the Smart Intersection is the security concept for protecting data against attackers. To provide security by design, the modular risk assessment (MoRA) method of

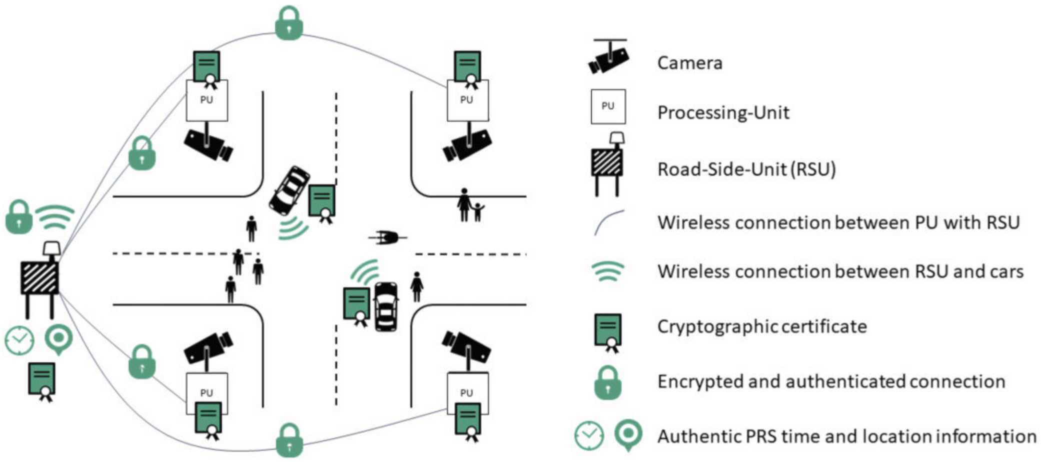
PU
PU
PU
PU
Camera
PU
Processing-Unit
Road-Side-Unit (RSU)
Wireless connection between PU with RSU
Wireless connection between RSU and cars
Cryptographic certificate
Encrypted and authenticated connection
Authentic PRS time and location information

Figure 2: Cameras capture the smart intersection and provide data to processing-units, which send the data over a secure link to the road-side-unit (RSU). The RSU sends a map of detected objects to cars, extending their perception range.

Fraunhofer AISEC was applied [2]. MoRA permits systematic collection and assessment of security goals, threats and countermeasures, resulting in a holistic and traceable security concept.

Human drivers and, in the near future, autonomous cars base decisions on information supplied by the smart intersection, relying on the authenticity of information. Thus, it is important to ensure that the smart intersection captures authentic information and that the data has not been altered throughout transmission and storage.

To ensure camera image authenticity, the RSU utilises a PRS Snapshot Sensor by Fraunhofer IIS, which captures tamper-proof Galileo PRS signals. This is used to apply unforgeable position- and time-stamps on the camera images and the generated object list. Public-key cryptography is used to sign and encrypt all messages during transmission and storage. Tamper-proof hardware prevents physical attacks. Figure 2 shows the smart intersection and the security concept.

To comply with data protection regulations, the centralised detection and classification of objects permits the localisation and anonymisation of number plates and faces of pedestrians in camera images. However, authorities desire to use the raw, i.e. not anonymised, video data captured by the smart intersection to investigate accidents. Confidentiality and authenticity of those images is ensured via cryptography as described above. Unlike standard camera surveillance systems, our solution stores all camera data encrypted in a trust-storage hosted in a high-security facility. Thus, these data can only be accessed after a judicial decision.

A proof-of-concept demonstration of the smart intersection has been deployed at Fraunhofer IVI in Dresden and will be used in further evaluation. We plan to enhance the intersection with further data-acquisition capabilities, transforming it into a primary building block of smart cities [L2].

Links:
[L1]: https://kwz.me/hEP
[L2]: https://kwz.me/hE1

References:

[1]: European commission, „2017 road safety statistics: What is behind the figures?”, 2018.
[2]: Jörn Eichler and Daniel Angermeier, “Modular risk assessment for the development of secure automotive systems”, published at 31. VDI/VW - Gemeinschaftstagung Automotive Security, 2015

Please contact:

Martin Striegel
Fraunhofer AISEC, Germany
martin.striegel@aisec.fraunhofer.de

Thomas Otto
Fraunhofer IVI, Germany
thomas.otto@ivi.fraunhofer.de

Smart Solutions to Cope with Urban Noise Pollution

by Jakob Abeßer and Sara Kepplinger (Fraunhofer IDMT)

Noise pollution, especially in urban environments, can have negative health impacts. Smart city applications for acoustic monitoring become essential to cope with the overall increasing noise pollution in residential areas. Based on measurement data from a distributed acoustic sensor network, a web-based application allowing for a real time visualisation of the citywide noise exposure was developed as part of the research project “Stadtlärm”.

As described by Berg and Nathanson in the encyclopaedia Britannica [L1] noise pollution is unwanted or excessive sound, which may derive from industrial facilities and other workplaces, highway, railway, and airplane traffic as well as from outdoor construction activities. The type of noise source influences which noise protection regulations are used. Therefore, one of our overall goals is to measure the exposure of citizens to noise in different parts of the city based on the German technical guidelines for noise reduction (TA Lärm).

Within the research project “StadtLärm” (German for “city noise”)

[L2] a distributed acoustic monitoring system [1] was developed whose sensor units are installed at the most relevant locations of both sound emission and sound immission (see Figure 1). Figure 1 gives an overview over the selected test site in the city of Jena, Germany. Acoustic sensors (orange squares) monitor the noise exposure at residential areas (purple) and nearby sound emission locations (red circles), such as open-air venues, a soccer stadium, as well as the most relevant transport routes, like streets (green) and train plus tram tracks (blue). Red arrows illustrate the most prominent sound propagation paths towards nearby residential areas. In addition to localised sound level measurements, the sensors automatically classify the most prominent sound events. This makes it possible to detect which noise sources contribute most to the noise pollution at a specific location. The system can provide valuable input data for systematic municipal planning in order to improve the residents' quality of life in the city.

The geographic location of the target area presents two main challenges. Owing to the valley-like elevation profile in Jena, most residential areas have higher elevation than the monitored sound emission locations. Furthermore, the prevailing westerly wind direction heavily influences the sound propagation at the test site. The sensor units compute a set of standardised noise level parameters in near real-time with a temporal resolution of 125 ms. In addition, a system for acoustic scene classification based on a deep neural network estimates the probability of nine different acoustic scene classes once every second. The network uses a cascade of convolutional layers with intermediate pooling in order to recognise temporal-spectral patterns in short-term spectrograms, which are characteristic of particular sound sources [2].

Connected by a communication system based on MQTT (Message Queue Telemetry Transport) [L3], these sensors communicate measurement data to a central server for data post-processing and storage. We implement the privacy-by-design approach by transmitting only time-localised probabilities of different acoustic scene classes.

The processing component (central server) publishes measurement results, which include noise level measurements as well as acoustic scene classification results. Furthermore, it offers request/response interfaces for retrieving historical measurement data. This can be useful, for example, to correlate noise level measurements with public events that have taken place, in retrospect.

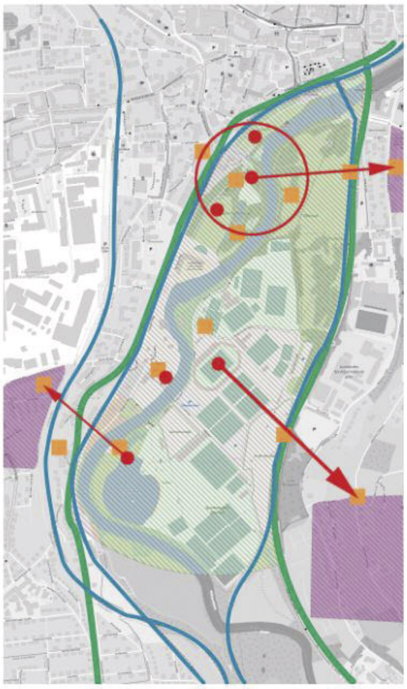

Figure 1: Targeted area in the city of Jena, Germany. Orange squares represent acoustic sensor units. Purple illustrates residential areas. Red circles show example sound emission locations. Streets are green and train and tram tracks are blue. Red arrows connect sound emission locations and nearby residential areas.

A web-based application receives and visualises live measurement data at different sensor locations on an interactive map. The map also provides additional layers of information about potential noise sources, such as road construction sites, locations of bottle banks, and playgrounds. It is possible to display the noise level values measured at the sensor units in the map application in real-time or aggregated over a configurable period. For example, the user can get a condensed overview of the past noise situation by eliminating the interferences caused by rush-hour traffic.

Finally, we will sketch several ideas for future improvements of the presented sensor network. First, displaying and communicating basic measurement results to the citizens via free and anonymous access to the application will contribute to the public acceptance of the system. In contrast, the city council can receive an extended account with further internal administrative information like conditions of the events' approvals and local residents' complaints. During the ongoing event, the regulatory agency of the city administration can be informed automatically via email if noise level limits are exceeded. Another approach to facilitate the administrative decision procedure is to predict future noise situations of planned events based on previously measured data.

We are extending the proposed framework for further application scenarios in an urban environment. For instance, the sensor units along the main streets allow traffic flow to be monitored and vehicle types to be identified. Furthermore, we will conduct long-term noise level analyses to identify quiet regions within the city. This information can be included in tourist recommendations, e.g., recovery area suggestions, or used for solutions addressing wellbeing and health [3].

Links:

[L1]: https://www.britannica.com/science/noise-pollution
[L2]: www.stadtlaerm.de
[L3]: https://mosquitto.org/

References:

[1] J. Abeßer, et al.: "A Distributed Sensor Network for Monitoring Noise Level and Noise Sources in Urban Environments," in: Proc. of FiCloud 2018.

[2] J. Abeßer, et al.: "Urban Noise Monitoring in the Stadtlärm Project – A Field Report", In: Proc. of DCASE 2019.

[3] S. Kepplinger, et al.: "Perspectives about Personalization for mHealth Solutions against Noise Pollution", in: Proc. of pHealth 2017.

Please contact:

Jakob Abeßer
Fraunhofer Institute for Digital Media Technology IDMT, Germany
jakob.abesser@idmt.fraunhofer.de

Transforming Everyday Life through Ambient Intelligence

by Constantine Stephanidis (FORTH-ICS)

The ICS-FORTH Ambient Intelligence (AmI) Programme is a long-term horizontal interdisciplinary RTD Programme aiming to develop pioneering human-centric intelligent technologies and environments which seamlessly support everyday human activities and enhance well-being through human-technology symbiosis.

ICS-FORTH has been investigating the potential of AmI technologies in domestic life [L1]. Inside the "Intelligent Home" simulation space located within its AmI Facility, everyday user activities are enhanced with the use of innovative interaction techniques, artificial intelligence, smart objects, ambient applications, sophisticated middleware, monitoring and decision-making mechanisms, and distributed micro-services. This tightly-coupled conglomeration of software and hardware components seamlessly interoperate so as to generate an ambient context-aware ecosystem that aims to: (i) improve the quality of life though appropriate monitoring of health-related variables (e.g. stress, sleep quality, nutrition); (ii) behave as an intelligent agent that communicates with the users in a natural manner and assists them in their daily activities; (iii) transform the environment into an ambient notification hub and personalised communication centre for the occupants (e.g. family members); (iv) implement a self-adaptive, energy-efficient and eco-friendly home control middleware; and (v) enhance leisure activities by providing entertainment experiences.

In particular, the hardware facilities of the Intelligent Home consist of: (i) an extensive grid of sensors and actuators that monitor and control various aspects (e.g. environmental conditions, energy and water consumption); (ii) smart commercial equipment (e.g. lights, speakers, locks, kitchen appliances, oil diffusers); (iii) smart devices and wearables (e.g. smart TVs, tablets, smartphones, smart watches, smart wristbands of various sorts); and (iv) technologically augmented custom-made artefacts. These components work in conjunction with the distributed computational framework of the Intelligent Home, named AmIHomeOS, so as to facilitate the realisation of comprehensive scenarios that accommodate various use cases (e.g. a family of four, an elderly couple, a disabled single adult) targeted to transform the domestic space into an all-inclusive environment that assists end users in an intelligent and personalised manner.

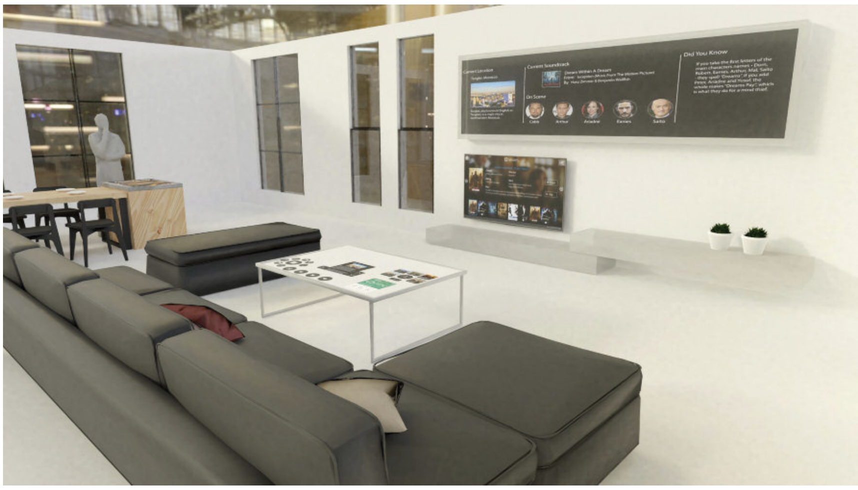

Figure 1: A photorealistic 3D rendering of the "Intelligent Livingroom".

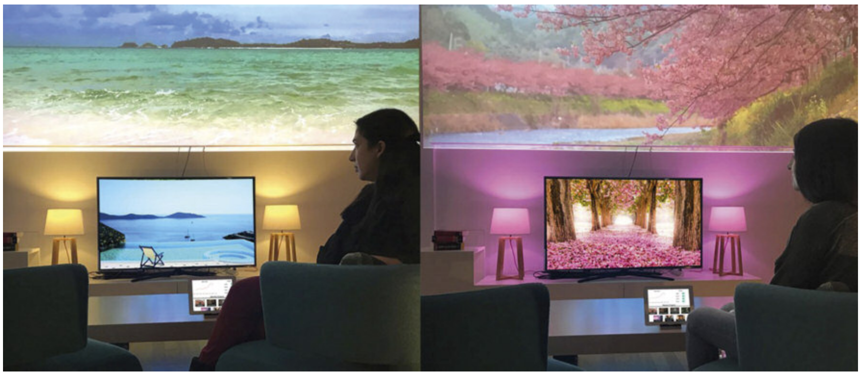

Figure 2: The CaLmi system in action.

The living room of the Intelligent Home has been transformed into a smart space that exploits ambient technologies (i.e. intelligent artefacts, technologically-augmented furniture) and services (e.g. video on demand, user profiling), as depicted in Figure 1. In this context, AugmenTable (a smart coffee table acting as a touch-enabled projection area, while being aware of the objects placed on it) and SmartSofa (a sofa equipped with sensors that detect user presence and posture and permit the execution of actions via mid-air gestures) interoperate with AmITV (a host of various interactive applications such as movie player, news reader, music player, etc.) and SurroundWall (an interactive wall-based display) in order to provide an enhanced viewing experience by transforming the overall space around the TV. Alongside, CaLmi aims to reduce the inhabitants' stress levels by real-time monitoring of various psychophysiological (e.g. electrodermal activity, heart rate) and contextual (e.g. workload, upcoming appointments, social life changes) variables and distributing into the environment appropriate relaxation programs that exploit the existing ambient facilities (Figure 2).

The bedroom is generally considered as the room of the house in which people retreat to unwind, relax, get ready and sleep. The intelligent bedroom supports these activities with the Hypnos framework, which aims to improve the quality of sleep by providing sleep hygiene recommendations. This relies on various sensors integrated under the bed, and the inhabitants' wearables to monitor their physical activity (e.g. movement, time in bed), physiological signals (e.g. respiration rate, heart rate, snoring) and various sleep-related parameters (e.g. time to fall asleep, time asleep, sleep cycles) while resting. Moreover, Smart Wardrobe (a custom-made closet equipped with various sensors and an embedded tablet) along with Smart Mirror (an interactive augmented reality system) aim to provide outfit recommendations based on contextual information (e.g. weather prediction, user's daily schedule, user preferences, clothes availability), allow users to be immersed in a "virtual mirror" where they can try them, and help them get dressed (i.e. helping them find what they are looking for).

Finally, the Intelligent Kitchen aims to integrate ambient technologies to support the inhabitants during the entire cooking process, ranging from the management of the inventory and shopping list, meal selection and preparation, to cleaning-up. Currently, in the kitchen, smart appliances and custom artefacts (e.g. smart cupboards, smart food containers) are accompanied by the AmICounterTop and AmIBacksplash systems, which constitute regular work surfaces augmented with technology so as to be transformed into interactive devices (e.g. primary or secondary displays, kitchen scales), assist the above tasks by communicating task-specific information to the user (e.g. highlight the ingredients needed for the current recipe step, recommend recipes based on goods about to expire, suggest buying more milk when the last bottle is about to be used).

Links:
[L1] http://ami.ics.forth.gr

Please contact:
Constantine Stephanidis
FORTH-ICS, Greece
+30 2810 391741
cs@ics.forth.gr

Managing the Trade-off between Security and Power Consumption for Smart CPS-IoT Networks

by Patrizia Sailer (Forschung Burgenland GmbH), Christoph Schmittner (AIT) and Markus Tauber (Fachhochschule Burgenland GmbH)

Making cyber-physical systems "smart" by managing the trade-off between security and resource usage is of utmost importance for building sustainable industrial systems. For example, addressing cyber security issues in such systems often require strong encryption. This may result in increased power consumption on devices that often depend on limited energy supply. In this work, we present an initial investigation into the usage of electrical power under different degrees of security in such situations to understand and quantify the level of reduction of power usage due to varying degrees of security.

The fourth industrial revolution is based on cyber-physical systems (CPS), where multiple components are interconnected over the Internet of Things (IoT). These components can communicate with each other as well as measure and change their physical environment. Such applications often require a high level of reliable and secure data acquisition, but often have limited energy resources. Thus, enabling policy-based adaptation to manage the trade-off between e.g. security, reliability and resource usage will help developing smart CPS. Supporting technologies may include self-*, deep-learning, or neural networks. In many applications it is crucial to maximise the lifetime of the components. Carrara et al. [1] have implemented an IoT-based management program to collect temperature and humidity data. Tauber et al. [2] have investigated energy efficiency and performance in a wireless-local area network (WLAN) to identify upper and lower bounds of energy efficiency due to different data flow characteristics. None of these have considered the impact of different security settings on power consumption, which is what we have addressed in this paper. This will help to understand the magnitude of power saving due to different levels of security by e.g. smart-/self-adaptation of the application.

In order to investigate the electrical power consumption under varying levels of security by different components of a CPS, we conducted measurements with a typical application, based on service-oriented architecture (SOA). The measurement setup consists of three Raspberry Pi v3 with Rasbian [L1] as the operating system. The setup emulates a network in which data is collected with a sensor and is sent via WLAN from client to server. The first Raspberry is equipped with a GrovePi+ [L2], which enables collecting data via a "DHT22 temperature and humidity sensor" [L3]. These data are saved in a list and sent to an HTTP server via a WLAN access point (AP). The second is the AP configured with HostAPD [L4], while the last one acts as the server. Each Raspberry Pi has its own Voltcraft Sem 6000 power plug (PP), which is responsible for measuring the power consumption. An "expect script" [L5] was used to gather the data from these plugs. To send the data from the client to the server, we implemented a "Spring Boot Rest Template" as a service that sends the collected data using a POST request. On the server we tracked the network communication via Wireshark and configured the following encryption suites (ES) to indicate the different power consumption:

- *ES_1: TLS_ECDHE_RSA_WITH_AES_128 _CBC_SHA (Weak)*
- *ES_2: TLS_ECDHE_RSA_WITH_AES_256 _CBC_SHA384 (Strong CBC)*
- *ES_3: TLS_ECDHE_RSA_WITH_AES_256 _GCM_SHA384 (Strong GCM)*

In our application we used TLS v1.2 as network protocol and ECDHE_RSA as key exchange algorithm, which is defined by a fixed ECDH key exchange signed with an RSA certificate [L6]. To illustrate the relationship between power consumption and encryption suites (ES), we used different key lengths for block encryption and a message authentication algorithm. In the weaker cipher suite (ES1), we used the Block Cipher Advanced Encryption Standard (AES) with independent block and key lengths of 128 bits, as opposed to the strong ones (ES2, ES3) with key lengths of 256 bits. Furthermore, the encryption options Cipher Block Chaining (CBC) and Galois/Counter Mode (GCM) vary in performance differences due to algorithm and security [3]. As algorithm for message authentication, we chose the Secure Hash Algorithm (SHA) with the different key lengths of 128 versus 384 bits.

With each encryption suite, we performed 21 test runs, each 25 minutes duration, with the collected data being sent from the client to the server every three seconds. To avoid latency or background jobs from the compiler, we performed a warm-up simulating a normal test run.

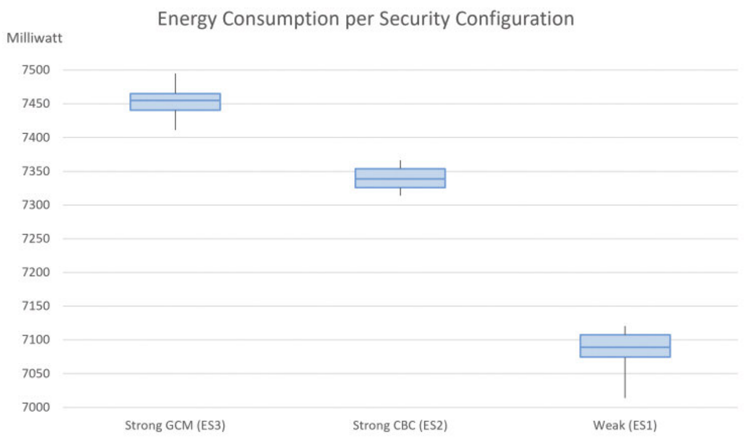

Figure 2: Consumption of electrical power in milliwatts at the client. This study, conducted as part of the EFRE project MiT4.0, shows that the choice of data protection measures via various cipher suites has a significant impact on power consumption: the stronger the level of security, the higher the power consumption. The issue of security needs to be further investigated and discussed in the IoT area, taking electrical power savings into account. Understanding the cooperation between these areas will make it possible to save more electrical power.

As shown in Figure 2, the results of the test runs show, that the stronger the cipher suite is, the more power is consumed by the client. The reason for that is that the client is responsible for generating the keys to communicate with the server via the selected encryption suite.

This study, conducted as part of the EFRE project MIT 4.0 (FE02), shows that the choice of data protection measures via various cipher suites has a significant impact on power consumption: the stronger the level of security, the higher the power consumption. Thus, it supports our initial idea that a significant power usage reduction can be achieved by reducing the strength of the cypher suite – if the situation allows for it. In future work we plan to extend those profiling activities as a basis for understanding the upper and lower bounds of power usage for smart CPS which will be able to manage the trade-off between security, reliability and resource usage.

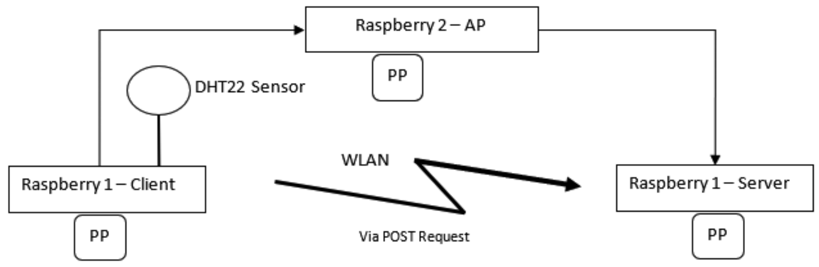

Figure 1: Architecture measurement setup. Three raspberries with power plugs (PP), which send data collected with DHT22 sensor from client via WLAN connected through access point to server. The power plugs measure the power consumption in milliwatts.

	KEY LENGTH AES	KEY LENGTH SHA	ENCRYPTION OPTION
WEAK (ES1)	256 bits	384 bits	CBC
STRONG CBC (ES2)	256 bits	384 bits	CBC
STRONG GCM (ES3)	256 bits	384 bits	GCM

Table 1: Overview differences used cipher suites.

Links:
[L1] https://kwz.me/hEL
[L2] https://kwz.me/hEl
[L3] https://kwz.me/hEp
[L4] https://kwz.me/hEg
[L5] https://kwz.me/hEr
[L6] https://kwz.me/hEY

References:

[1] M. Carrara, et al.: "An innovative system for vineyard management in Sicily", Journal of Agricultural Engineering, 41(1), pp. 13-18, 2010.

[2] M. Tauber, S. N. Bhatti, Y. Yu: "Application Level Energy and Performance Measurements in a Wireless LAN", IEEE/ACM GREENCOM 2011.

[3] Y. Hore, et al.: "Bitstream Encryption and Authentication using AES-GCM in Dynamically Reconfigurable Systems", Advances in Information and Computer Security, pp 261-278, 2008, ISBN: 978-3-540-89597-8.

Please contact:

Patrizia Sailer
Forschung Burgenland GmbH, AT
patrizia.sailer@forschung-burgenland.at

ISaFe - Injecting Security Features into Constrained Embedded Firmware

by Matthias Wenzl (Technikum Wien), Georg Merzdovnik (SBA Research) and Edgar Weippl (SBA Research)

The vast majority of the IoT is made up of computing devices that are highly specialised for their particular purpose. Owing to their specialisation and the resulting constraints, such as energy consumption and the deterministic fulfilment of deadlines (real-time requirements), these embedded systems tend not to share many security features in common with standard operating systems. We aim to provide automated approaches to implant security features into connected embedded systems to counter the lack of security features in the backbone of the IoT and improve their resilience against unauthorised access attempts.

The Internet of Things (IoT) is formed by an ever-growing number of interconnected embedded systems. An embedded system is a computing device designed for a specific purpose that operates under certain constraints (e.g., hardware and software tightly intertwined to monitor traffic conditions on a motorway). Use-cases like smart cities, Industry 4.0, Car2X communication and ambient assistive technologies promote networked embedded devices to become ubiquitous companions in our daily lives. Unfortunately, many of the computing devices in the IoT lack even the simplest form of security features, resulting in numerous successful attacks on IoT devices in recent years. Security features, specifically exploit mitigations, are procedures that reduce the likelihood of an unpatched software vulnerability being exploited by an adversary in such a way that it is possible to mislead the system to do unintended operations (e.g., address space layout randomisation (ASLR) that is built into every Microsoft Windows operating system derivate since Windows Vista (2007), as well as enabled by default in OpenBSD since 2003 and Linux since 2005). The lack of security features in embedded systems is primarily based on application specific constraints, hardware requirements (physical size, energy consumption, real-time requirements [1]), software diversity (full features OS like Linux, *BSD, Windows IoT Core, real-time OS like FreeRTOS, no OS at all) and security agnostic design decisions (e.g. unreflected usage of legacy code).

Well-known security features such as address space layout randomisation (ASLR), non-executable pages and stack layout transformation (SLX) are readily available for general-purpose operating systems. However, according to a study presented in 2017 [2], not all variants of general-purpose operating systems deployed to embedded computing platforms are configured to utilise security features. Additionally, many embedded operating systems lack these kinds of features altogether. This is primarily caused by the fact that many embedded operating systems aim to support a wide variety of target hardware, including even the smallest devices. For example, the widely used FreeRTOS supports target platforms ranging from Microchip's SAMD20 microcontroller with 16 kB non-volatile and 2 kB volatile memory to full featured IA-32 based systems. Consequently, support for inherent security features is omitted. Additionally, almost all security features introduced at an operating system level protect either every component of interest (e.g., function entry point) or none. This leads to an increase in resource consumption in terms of memory and runtime that might not be practicable (e.g., the implementation of

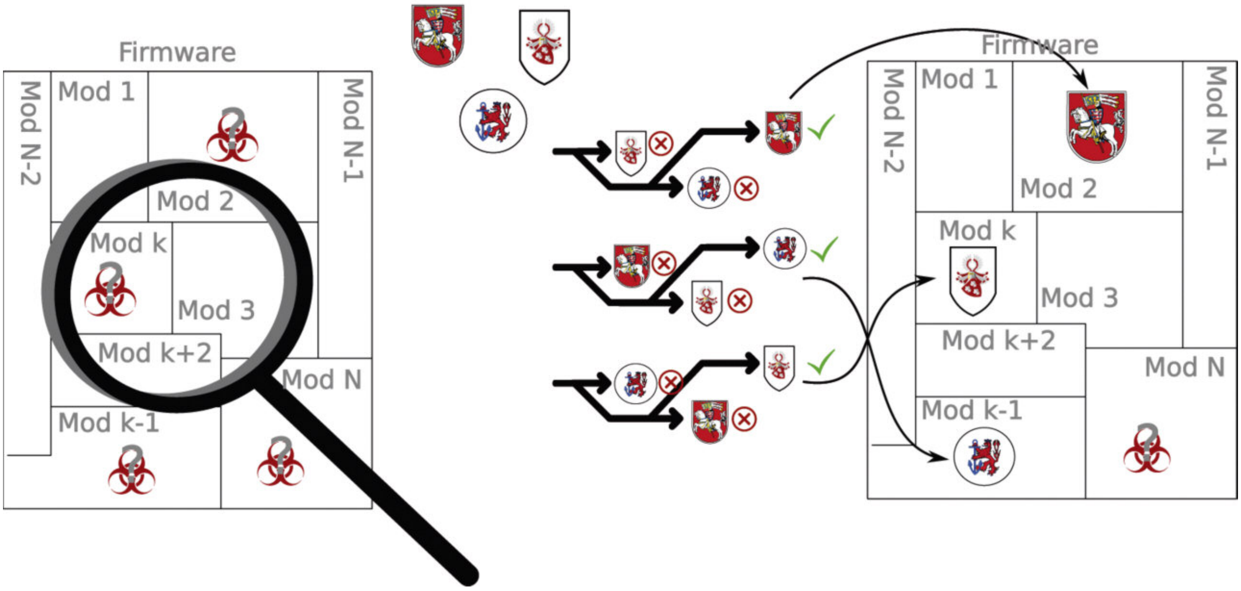

Figure 1:The core idea of project ISaFe read from left to right - Find locations within a firmware image that utilises data from external inputs in certain ways. Then choose and implant selected security features to mitigate possible attacks on the detected locations in such a way that as much as possible of the detected locations are protected and the firmware under investigations still adheres to its constraints (real-time, memory consumption, etc).

a shadow-stack at an operating system level doubles the amount of used stack memory in terms of return addresses). Nevertheless, patching all available embedded operating systems and already circulating legacy systems at source level is clearly illusive owing to their large heterogeneity.

Therefore, the aim of the FFG-funded project ISaFe [L1] is to provide automated approaches to implant security features into connected embedded systems to counter the lack of security features in the backbone of the IoT starting with September 2019. Together with the IoT startup Riddle & Code, whose aim is to provide an interface between embedded systems and blockchain technology, SBA-Research and the FH Technikum Wien, all situated in Vienna, Austria, pursue a novel approach based on binary rewriting to retrofit already existing IoT systems in order to make them more resilient against unauthorised access attempts.

Binary rewriting describes the alteration of a compiled and possibly (dynamically) linked program without having the source code at hand in such a way that the binary under investigation stays executable [3]. In the context of hardening this means to implant security features into the binary under investigation.

Unfortunately, constraints such as real-time requirements, or memory constraints are very common in embedded systems software. Thus, a straightforward approach like applying binary rewriting for hardening purposes on every point of interest (e.g., every function call, every branch) in a binary with constraints is not feasible due to the introduction of possible constraint violations. At the core, our approach requires us to solve an optimisation problem that maximises the number of protected vulnerable spots (locations within the firmware of interest that might be subject to attack, which we attempt to identify via taint tracking) with the most efficient security features, at the lowest possible costs - in terms of memory and run time overhead - while obeying the system’s constraints. The overall idea if the presented approach can be seen in Figure. 1.

Link:
[L1] https://kwz.me/hEB

References:

[1] H. Kopetz: “Real-Time Systems: Design Principles for Distributed Embedded Applications”, second edition, Springer, 2011.
[2] J. Wetzels: “Ghost in the Machine: Challenges in Embedded Binary Security”, Usenix Enigma, 2017. https://www.usenix.org/conference/enigma2017/conference-program/presentation/wetzels
[3] M. Wenzl, et al.: “From Hack to Elaborate Technique - A Survey on Binary Rewriting”, 2019, https://doi.org/10.1145/3316415

Please contact:

Matthias Wenzl
Technikum Wien, Austria
wenzl@technikum-wien.at

Georg Merzdovnik and Edgar Weippl
SBA Research, Austria
gmerzdovnik@sba-research.org,
eweippl@sba-research.org

Enabling Smart Safe Behaviour through Cooperative Risk Management

by Rasmus Adler (Fraunhofer IESE) and Patrik Feth (SICK AG)

Machines in an Industry 4.0 context need to behave safely but smartly. This means that these smart machines need to continue to precisely estimate the current risk and not shut down or degrade unnecessarily. To enable smart safe behaviour, Fraunhofer IESE is developing new safety assurance concepts. SICK supports related safety standards to implement these concepts in an industrial setting.

Smart machines need to behave safely in a smart way, predicting accidents and adapting their behaviour to avoid them, while efficiently fulfilling their original mission. The prediction of possible accidents is a challenging task. It requires perceiving the current situation and anticipating how it will evolve over time. For safety reasons, any uncertainty in this dynamic risk management must be compensated by a worst-case assumption. This limits the performance significantly. For instance, if a smart machine is not sure about the behaviour of another smart machine, it must assume the most critical behaviour and behave very cautiously even in situations where it is not necessary.

A promising approach for maximising performance while ensuring safety is cooperative risk management. If all smart machines, regardless of the manufacturer of the machine, shared their information about the current situation and its evolution, we could minimise the number of worst-case assumptions. However, the claims that each individual machine makes about the situation and its evolution could be wrong. Thus, Fraunhofer IESE proposes extending these claims to a machine-readable assurance case capturing the underlying reasoning and assumptions.

As shown in the left part of Figure 1, an assurance case is a reasoned and compelling argument, supported by evidence such as test results, that an entity achieves its objectives and constraints.

As shown in the right part of Figure 1, the argument generally consists of several reasoning steps. Each step states that a higher-level claim is true if some sub-claims are true and some assumptions hold.

The idea behind machine-readable assurance cases is to enable runtime evaluations of the assumptions used in the argumentation. A smart machine can evaluate if the assumptions used in the assurance case received from another smart machine hold in the current usage

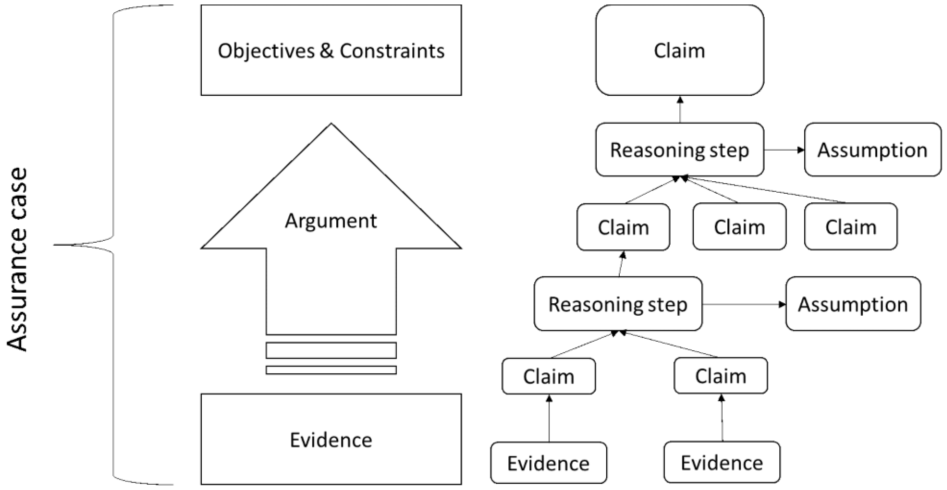

Figure 1: Illustration of an assurance case.

scenario. For instance, if the provided information is only correct if a temperature is within a certain range, then a smart machine receiving this information can check if the temperature is currently within this range. Furthermore, a smart machine may evaluate the assumptions with respect to the required integrity. If there are too many arguable assumptions, then the smart machine might decide that the integrity is insufficient.

Fraunhofer IESE is developing solutions for implementing this idea in the DEIS project [L1]. In this project, we use the term Digital Dependability Identity (DDI) for all the information that uniquely describes the dependability characteristics of a system or component [1]. We chose SACM [L2] as the basis for describing DDIs, because SACM is the meta-model of assurance case languages.

However, the concept of ensuring safety through machine-readable assurance cases is in conflict with traditional standards in the domain of industrial automation. In this domain, it is common to safeguard a critical application by using dedicated protective devices, mainly sensor devices with a fixed data evaluation function, that provide a certain safety function with a guaranteed level of integrity; e.g., a safety laser scanner detecting an intrusion into a preconfigured safety field. These devices and their provided functions are conformant to harmonised standards such as IEC 61496-3. Such standards describe for a specific technology – for example in the case of IEC 61496-3, active opto-electronic protective devices responsive to diffuse reflection – very precise requirements for achieving a particular integrity level.

In European countries, the use of such devices is heavily stimulated as the Machine Directive explicitly covers protective devices designed to detect the presence of persons. Along with the existence of specific harmonised standards, this has fostered the belief in the domain of industrial automation that safety equals conformance with such standards. This limits the possibilities for building safe systems, as it constrains the options for reducing the risk of a critical application to existing standardized protective devices. In particular, it hinders the implementation of the above-mentioned cooperative risk management.

To overcome this limitation, SICK supports the recently published technical specification IEC TS 62998-1. This technical specification gives guidance for the development of safety functions beyond existing sensor-specific standards such as IEC 61496-3. It facilitates the use of new sensor technologies (e.g., radar, ultrasonic sensors), new kinds of sensor functions (e.g., classification of objects, position of an object), combinations of different sensor technologies in a sensor system, and usage under new conditions (e.g., outdoor applications). To this end, IEC TS 62998-1 requires analysing the capabilities of a sensor used in the intended application. A system realising cooperative risk management can be considered as a safety-related sensor system with a respective safety function in this new technical specification.

To support future Industry 4.0 applications, sensor models could incorporate the information required by IEC TS 62998-1 and could become part of a machine-readable assurance case of the device. With this information stored in the assurance case, it would become possible to use the device with its provided sensor functions in critical applications unforeseen during the initial development of the device.

Links:

[L1] http://www.deis-project.eu/home/
[L2] https://kwz.me/hEb

Reference:

[1] E. Armengaud et al.: “DEIS: Dependability Engineering Innovation for Industrial CPS”, in: C. Zachäus, B. Müller, G. Meyer (eds): ‘Advanced Microsystems for Automotive Applications 2017’, Lecture Notes in Mobility. Springer, 2018.

Please contact:

Rasmus Adler
Fraunhofer Institute for Experimental Software Engineering IESE, Germany
rasmus.adler@iese.fhg.de

Patrik Feth
SICK AG, Germany
patrik.feth@sick.de

Assessing the Quality of Smart Objects

by Ivana Šenk (Inria and University of Novi Sad) and James Crowley (Inria)

With the encroachment of smart objects into our lives, it is increasingly important to define a principled approach to estimate the value of alternative technologies at design time, to define product specifications and to compare similar products. In mature domains, such approaches are based on properties that are referred to as "qualities". We are working to develop a hierarchical model for qualities for smart objects based on different modes of interactions.

What is a quality? Ideally, qualities are distinctive attributes that can be used to characterise things or to provide a standard for their comparison. The effective use of a quality requires an established way to assess it through an objective test. Such tests may be quantitative, providing a numerical value or qualitative, providing a symbolic label. In either case, the test must provide an objective reference for comparison.

The use and measurement of qualities has been studied in a number of related domains, including manufacturing, software engineering, communications, human-computer interaction, multimedia systems and Internet of Things. In general, the study of qualities in each of these domains has given rise to hierarchical conceptual frameworks, with the concept of "quality" being broken down into a variety of more detailed characteristics that can be evaluated with one or more tests.

In many domains, quality dimensions follow similar principles. For example, in manufacturing Garvin [1] proposed eight critical dimensions of quality that serve as a framework for analysis and product evaluation: performance, features, reliability, conformance, durability, serviceability, aesthetics, and perceived quality. Similarly, in software engineering the scientific and industrial communities have developed a variety of quality models, leading to the adoption of ISO/IEC standard 25010 [2], where quality is represented by two major aspects: product quality and quality in use. Product quality includes functional suitability, performance efficiency, compatibility, usability, reliability, security, maintainability, and portability. Quality in use includes effectiveness, efficiency, satisfaction, freedom from risk, and context coverage.

Similar patterns may be found in other domains, where we can find intrinsic qualities such as performance, functionality, conformance, correctness, reliability, efficiency, durability, reusability, maintainability, portability, compatibility, security and privacy; as well as user-related qualities such as usability, learnability, ease of use, naturalness of interaction, understandability, perceived quality, aesthetics, satisfaction, appropriateness, joy of use, usefulness, or trust.

Smart objects have emerged, as ordinary objects are increasingly augmented with digital technologies to provide embedded functions for perception, action, interaction and intelligence. To build a comprehensive quality model for smart objects, we propose to draw on the quality dimensions identified in related domains, and to identify new quality dimensions that take into account the capabilities that arise with smart technologies. Such new capabilities cover the features that make the objects smart, such as connectivity and cognitive abilities. These include abilities for sensing, actuating, interpretation, communication and networking, as well as abilities for easy, natural, and convenient interaction with people. A variety of qualities are possible, although the relative importance may vary with the type of smart objects.

We base our approach on different interaction dimensions of smart objects (Figure 1). We start with the intrinsic dimension that addresses the self-oriented qualities of the smart object, and add three extrinsic dimensions of interaction that the smart objects have with the environment, with other objects, and with people. Accordingly, we propose a hierarchy of detailed quality characteristics along the following main dimensions:

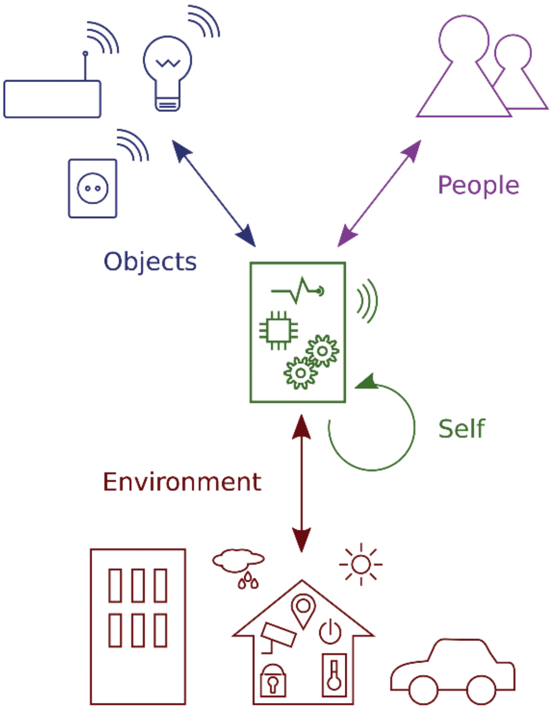

Figure 1: Interaction dimensions of smart objects.

1. Innate qualities, related to the core function of the smart object. These include functionality and conformance aspects, as well as qualities that relate to efficiency of the smart object, such as energy management, reliability, and availability, and qualities that take into consideration the

product lifecycle, such as durability, maintainability, portability, upgradeability, or recyclability.
2. Cognitive qualities, based on abilities to acquire, interpret and apply knowledge and skills. These include conceptual qualities that allow the smart object to behave in a appropriate manner, by perceiving, learning, understanding and reasoning about the environment, and taking relevant actions.
3. Social qualities, based on abilities for communication, networking, and inter-object relationships. These include aspects of social interaction between smart objects, from the types of connections that enable interaction to the possibilities that these connections provide. They cover composability through the diversity, capacity, compatibility and extendability of connections, availability of connections through discoverability, accessibility and transitivity, and sociability through the ways of forming relations, inter-object controllability and collaborative ability.
4. Suitability qualities, which include the capacity to act and interact with the users in a manner that is appropriate for the task and context. To measure suitability, we reflect upon the usability aspects such as learnability, ease of use, user control, understandability, or predictability, and acceptability aspects, such as privacy, security, trust, aesthetics, engagement, satisfaction, or usefulness.

We are currently working to establish adequate objective ways to test and measure these proposed quality aspects for smart objects. Some of these aspects can be tested under controlled laboratory conditions, while the others require evaluation under real world conditions. In order to assess our quality model we are applying the proposed tests to different types of existing smart objects. Our objective is to promote these ideas as an approach for defining, describing, measuring and comparing smart objects, with details to be refined and developed by the larger community. In addition, as the boundary between smart objects and smart services is increasingly fluid, we believe that any normative reference established for smart objects should also be useful for smart services including those provided by assemblies of smart objects.

References:

[1] D. Garvin: “Competing on the eight dimensions of quality,” Harv. Bus. Rev., 1987.

[2] “Systems and Software Engineering - Systems and Software Quality Requirements and Evaluation (SQuaRE) - System and Software Quality Models,” ISO/IEC 25010, 2011.

Please contact:

Ivana Šenk, Inria, France and
University of Novi Sad, Serbia
ivanas@uns.ac.rs

James Crowley
Inria, France
james.crowley@inria.fr

AutoHoney(I)IoT - Automated Device Independent Honeypot Generation of IoT and Industrial IoT Devices

by Christian Kudera, Georg Merzdovnik and Edgar Weippl (SBA Research)

The heterogeneous landscape of IoT devices poses new challenges to the deployment of honeypots. So far no generic honeypot framework exists that is capable of attracting attacks for the wide variety of hardware and software architectures. By combining real world device information and virtualisation techniques, we aim to build AutoHoney(I)IoT, a framework that automatically creates target device tailored honeypots for the (Industrial) Internet of Things, which are capable of convincing attackers that they are breaching a real device instead of a decoy.

The interconnection of physical devices, vehicles, household appliances and other objects with electronics, software, sensors and actuators has become an integral part of our modern lives. The industrial sector is also undergoing a change in device communication. Traditionally, automated factories and critical infrastructure were strictly separated from the Internet. However, since the advent of Industry 4.0, devices at control as well as supervisor level are frequently connected to the Internet to collect analytic data. The resulting network is called the “Internet of Things” (IoT) and “Industrial Internet of Things” (IIoT). Attackers seek to compromise such interconnected devices with malware campaigns [1, 2] to use them for spam distribution, Distributed Denial of Service (DDoS) attacks, cryptomining, or as an attack vector in Advanced Persistent Threat (APT) attacks. For this reason, interconnected devices are exposed to continuous threats and ongoing attacks. The large set of diverse hardware and software combined with the neglection of security best practices, such as the use of the same default credentials on all devices, the often non-existent update policies, and the lack of software hardening techniques render IoT and IIoT devices an ideal target for attackers. Many solutions have been proposed to monitor the Internet for malware infections. “Honeypots” are a common practice but owing to the heterogeneity of the devices they are substantially harder to implement in the IoT and IIoT domain than in the field of commodity systems (e.g., desktop computers, smartphones). The heterogeneous landscape of IoT and IIoT devices poses new challenges to the deployment of honeypots.

The goal of the research project AutoHoney(I)IoT [L1] is to provide a framework that automatically creates target device tailored honeypots for the (Industrial) Internet of Things, which are capable of convincing adversaries that they are breaching a real device instead of a decoy. The honeypots will be executed in an emulation environment that is able to interact with the outside world over common IoT and IIoT communication channels and allows fine-grained supervision techniques to be applied to monitor an adversary's behaviour throughout his entire attack. Figure 1 illustrates the overall concept of the AutoHoney(I)IoT framework. As input, the framework requires the firmware of the device to be virtualised as well as a database containing real world device information of common microcontrollers and system on a chips (SoCs). The framework's intended functionality is to analyse a firmware dump and find a Qemu [3] appliance capable of executing the firmware. Furthermore, external peripherals (e.g., network access) are attached as needed.

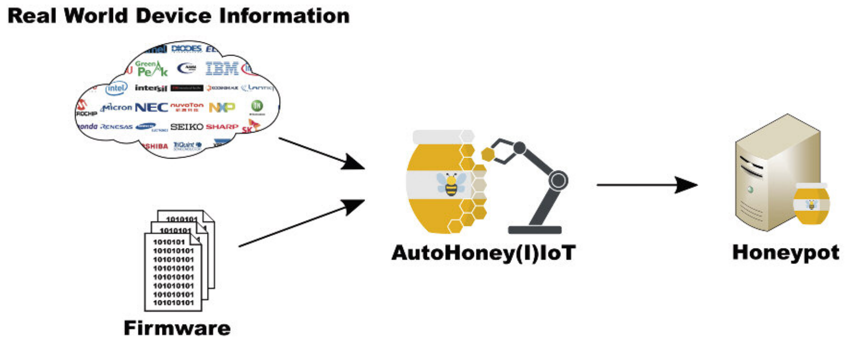

Figure 1: Overview of the AutoHoney(I)IoT framework.

The analysis is based on a two-step approach. The aim of the first step is to detect the architecture (e.g., ARM, MIPS, PowerPC, x86) of the firmware. The main idea of the second step is to apply fine-grained supervision techniques during the execution to identify the most appropriate processor. Therefore, the firmware is executed first on a generic processor corresponding to the identified architecture. Most likely this will result in an instruction mismatch, memory mismatch or internal peripheral mismatch, since there is a plethora of different embedded processors from various vendors with a wide variety of technical characteristics (e.g., central processing unit types, memory maps). During the emulation an algorithm attempts to identify an admissible embedded processor that is capable of executing the firmware dump. For this purpose, we create a database containing the real world technical characteristics of microcontrollers and SoCs from various vendors. The execution and fine-grained processor selection is repeated until a suitable processor is identified.

Since a convincing (I)IoT honeypot needs to expose authentic system behaviour and communicate with the outside world, the AutoHoney(I)IoT framework will provide: (i) a method to attach communication interfaces to the emulated embedded processor and bridge them to physical as well as simulated hardware, and (ii) a method to attach custom external peripheral device models to a communication interface whose behaviour can be scripted in a simple way.

The FFG funded project started in July 2019 and is expected to conclude in December 2021. It is jointly realised by SBA Research, FH Technikum Wien, TU Wien and Trustworks GmbH, all situated in Vienna, Austria.

Link:
[L1] https://kwz.me/hEt

References:

[1] K. Angrishi: "Turning Internet of Things(IoT) into Internet of Vulnerabilities (IoV): IoT Botnets", arXiv Prepr. ArXiv1702.03681, 2017.
[2] A. Spognardi et al.: "Analysis of DDoS-Capable IoT Malwares", in Proc. of INSERT, 2017.
[3] F. Bellard: "QEMU, a fast and portable dynamic translator", USENIX Annual Technical Conference, FREENIX Track. Vol. 41. 2005.

Please contact:

Christian Kudera, Georg Merzdovnik, Edgar Weippl, SBA Research, Austria
ckudera@sba-research.org, ,
gmerzdovnik@sba-research.org,
eweippl@sba-research.org

Yogurt: A Programming Language for the Internet of Things (IoT)

by Ivan H. Gorbanov, Jack Jansen and Steven Pemberton (CWI)

As IoT moves from the hands of professionals and academics into those of the general consumers, it becomes increasingly important to provide them with the appropriate tools for interaction. Yogurt is a domain-specific programming language for IoT, designed to tackle the disparity between powerful but complex languages and user-friendly environments with restricted capabilities.

"Yogurt" is an object-oriented declarative programming language for the Internet of Things (IoT). It allows the end-user to program their entire IoT ecosystem through one environment by leveraging the capabilities of Igor [L1], an architecture for unified access to IoT [1]. The language offers high expressiveness by incorporating mechanisms from traditional programming paradigms. Furthermore, the underlying programming model adopts a metaphor close to the users' real-life experience, thus reducing the learning effort required to adopt the language.

Traditionally IoT devices are programmed through a variety of high-level languages. This, however, requires an in-depth knowledge of computer programming which often takes years of training and practice to develop. In addition, the decentralised nature of the development of these technologies has led to the incorporation of many communication protocols and formal languages. Therefore, managing a complete system often requires knowledge of multiple programming environments. Both the consumer industry and academia have tried to address this problem by introducing systems which act as a central control point for the customer's IoT systems. Through these all of the user's devices can be programmed through one language. These programming facilities, however, are mostly optimised for usability in order to reduce the learning curve for inexperienced users. This, therefore, leaves a gap between highly expressive but difficult programming languages and user-friendly environments, which lack sufficient expressive power. Consequently, the main purpose of this work is twofold: to extend the work of Jansen and Pemberton [1] by providing a viable programming language for Igor, and to fill the abovementioned gap that currently exists in the programming facilities for IoT devices.

Yogurt's underlying programming model (Figure 1) achieves simplicity by providing abstractions analogous to the real world in order to reduce the gap between what the programmer is trying to achieve and how to achieve it. The "actor" abstraction represents all devices which would make up an IoT system. Each "actor" has a state, which represents what the device is doing in the real world and can perform "actions" which change that state. An "action" gets triggered by a change in the state of its own corresponding actor or that of another one. Conditions can also be added in the form of "guards" to allow the device's behaviour to vary based on its environment. This small set of general abstractions is at a high enough level so as to be easy to conceptualise. Furthermore, they are general enough to be applicable in the increasingly heterogenous collection of devices part of IoT.

The proposed model utilises several mechanisms from established programming paradigms to allow it to tackle the use cases that have come to be expected by users of this technology. The actor abstraction works as a class in object-oriented programming in order to allow the reuse of code, making writing Yogurt programs more efficient. Furthermore, it allows for the use of encapsulation and inheritance. This makes it easier to program more complex devices as a collection of simpler ones. In addition, by forcing the programmer to keep data and corresponding methods together, it is easier to keep track of dependencies and spot any conflicts that may arise from different methods trying to change the same data at the same time. Last but not least, the language is declarative, meaning that the user needs to specify what the result of the program needs to be rather than how exactly to achieve it. This feature brings two major advantages: it further reduces the learning curve as the programmer does not have to worry about concepts such as memory management and it makes the language context independent, meaning that a program which switches the lights on and off would remain the same regardless of the hardware devices used to implement the solution as the result will be the same.

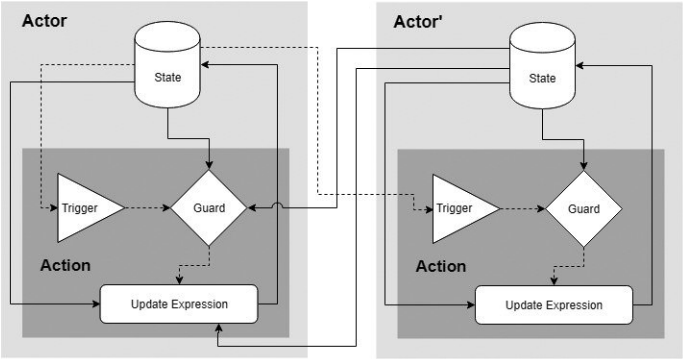

Figure 1: Yogurt programming model abstractions.

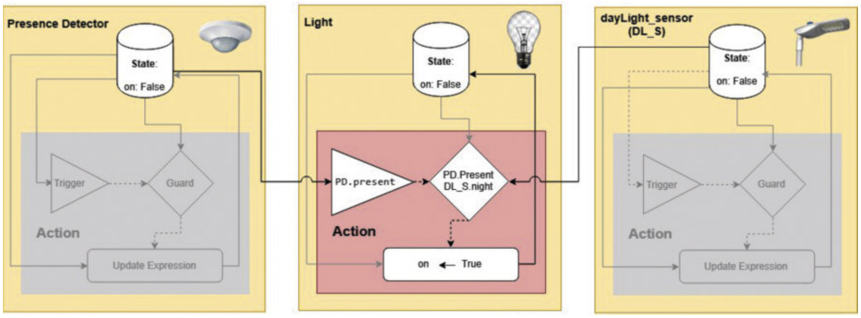

Figure 2: Light, presence sensor, daylight sensor example model.

The proposed textual representation of Yogurt uses human readable keywords to reduce the barrier to adoption and increase code readability. Here is a simple example of a light which turns on with a presence detector (PD) only when a day light sensor senses that it is dark outside:

```
on(PD.present):
        whenall(PD.present = True,
dayLight_sensor.night = True):
                on <- True
```

The current version has been tested using the Discount method for programming language evaluation [2]. Participants confirmed that the language was easy to pick and use because it allows them to think about the task in a way in which they would in the physical world, while suggesting changes that could be made to increase efficiency.

Links:

[L1] https://github.com/cwi-dis/igor

[L2] https://www.dis.cwi.nl/

References:

[1] J. Jansen, S. Pemberton: "An architecture for unified access to the internet of things", XML LONDON 2017 (2017).

[2] S. Kurtev, T. A. Christensen, B. Thomsen: "Discount method for programming language evaluation", in PLATEAU @ SPLASH. 1-8, 2016.

Please contact:

Jack Jansen
CWI, Netherlands
jack.jansen@cwi.nl

Smart End-to-end Massive IoT Interoperability, Connectivity and Security (SEMIoTICS)

by Nikolaos Petroulakis (FORTH), Konstantinos Fysarakis (Sphynx), Sotiris Ioannidis (FORTH), George Spanoudakis (Sphynx) and Vivek Kulkarni (Siemens)

Next generation networks, such as the Internet of Things (IoT), aim to create open and global networks for connecting smart objects, network elements, applications, web services and end-users. Research and industry attempt to integrate this evolving technology and the exponential growth of IoT by overcoming significant hurdles such as dynamicity, scalability, heterogeneity and end-to-end security and privacy. SEMIoTICS proposes the development of a pattern-driven framework, built upon existing IoT platforms, to enable and guarantee secure and dependable actuation and semi-autonomic behaviour in IoT/IIoT applications.

TThe introduction of digital technologies in economic and societal processes is key to addressing economic and societal challenges such as ageing of population, ensuring societal cohesion, and sustainable development. While the fifth generation (5G) mobile communications are already upon us, the next steps in their evolution will be key in supporting this societal transformation, while also leading to a fourth industrial revolution that will impact multiple sectors.

IoT appears to be an important pillar of 5G. Global networks like IoT create enormous potential for new generations of IoT applications, by leveraging synergies arising through the convergence of consumer, business and industrial internet, and creating open, global networks connecting people, data, and “things”. A series of innovations across the IoT landscape have converged to make IoT products, platforms and devices technically and economically feasible. However, despite these advances, significant business and technical hurdles must be overcome before the IoT’s potential can be realised.

Some important challenges and complexities include:

- Sustaining massively generated, ever-increasing, network traffic with heterogeneous requirements
- Adaptation of communication technologies for resource-constrained virtualised environments
- Provision of networking infrastructures featuring end-to-end connectivity, security and resource self-configuration
- Trusted information sharing between tenants and host systems.

Overcoming these challenges requires the implementation and deployment stack of IoT applications. The overall aim of SEMIoTICS [L1] is to develop a pattern-driven framework [1-2], built upon existing IoT platforms, to enable and guarantee secure and dependable actuation and semi-autonomic behaviour in IoT/IIoT applications. The SEMIoTICS framework supports cross-layer intelligent dynamic adaptation, including heterogeneous smart objects, networks and clouds. To address the complexity and scalability needs within horizontal and vertical domains, SEMIoTICS develops and integrates smart programmable networking and semantic interoperability mechanisms.

The SEMIoTICS architectural framework (Figure 1) has been envisaged and developed for efficient interconnectivity of smart objects. Each layer contains specific developed modules able to handle different aspects and guarantee different properties. More specifi-

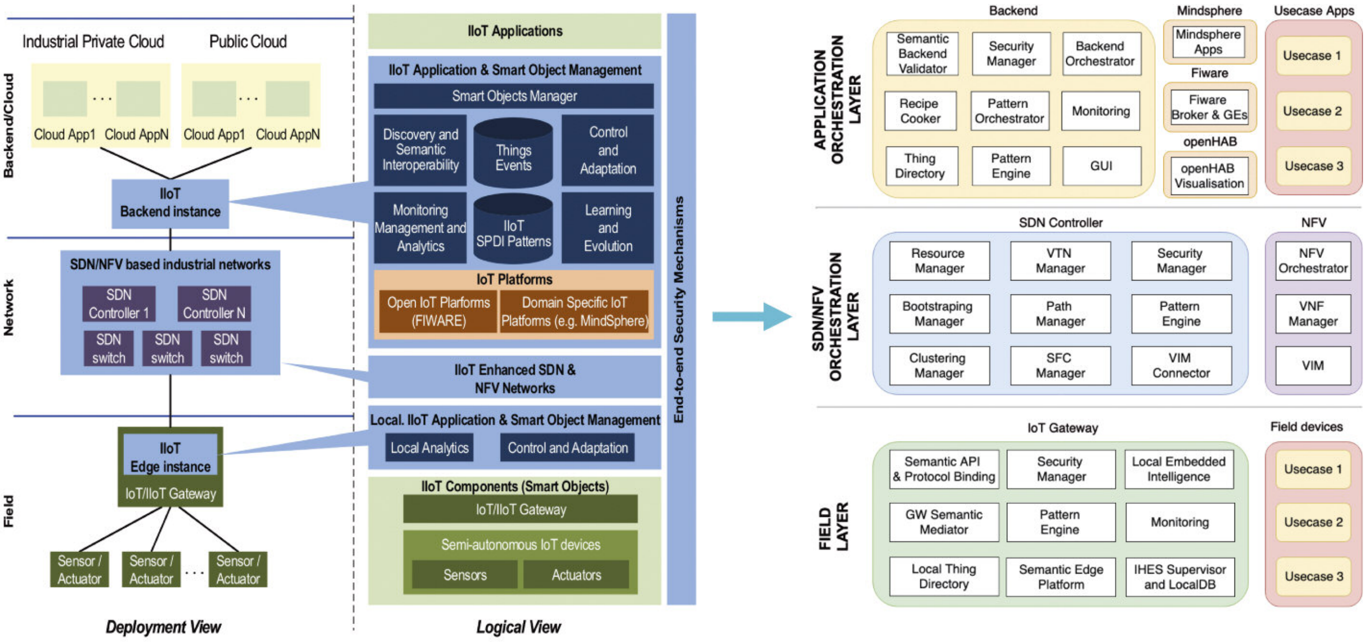

Figure 1: SEMIoTICS (i) envisaged architecture (ii) developed architecture.

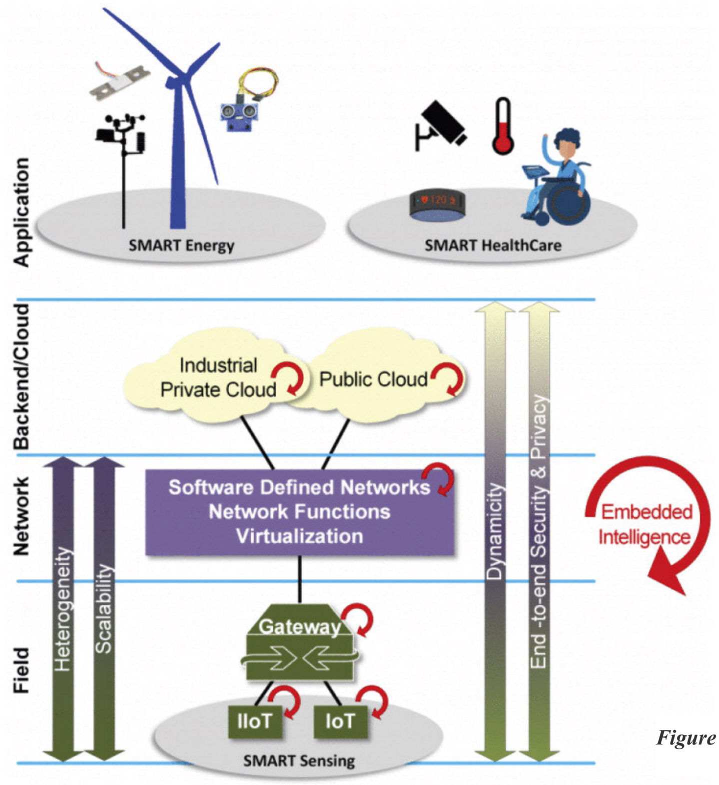

Figure 2: SEMIoTICS Use Cases.

cally, Software Defined Networking (SDN) Orchestration layer provides data and control plane decoupling resulting in a cloud computing approach that facilitates network management and enables programmatically efficient network configuration meeting different IoT application requirements related to security, bandwidth, latency and energy efficiency, using semantic information. Network Function Virtualization (NFV) Orchestration layer provides a flexible, programmable, dynamic and scalable networking paradigm, making it ideal for satisfying the QoS demands of SEMIoTICS use cases. Field layer is responsible for hosting all types of IoT devices such as sensors and actuators as well as IoT gateway which provides common way for communication and ensures enforcement of SPDI patterns in this layer. Finally, Application Orchestration layer consists of all applications receiving the communication from field layer.

The above is validated by industry, using three diverse usage scenarios in the areas of renewable energy, healthcare, and smart sensing (Figure 2) and will be offered through an open API.

- Use Case 1 - Smart Energy: This use case will showcase IIoT integration in Wind Park Control Network providing value added services such as local smart behaviour, predictive maintenance and monitoring etc. Current state of the art of Wind Turbine Controller in a Wind Park control network is typically an embedded or highly integrated operating system, which rigorously follows development and pre-qualification prior to deployment in the real world. Because of this slow process, new features, adding new sensors, actuators and related advancements require several months or even years to be fully matured and operational in the field.
- Use Case 2 – Smart Health Care: This use case employs the SEMIoTICS technologies to develop an Information and Communication Technology (ICT) solution aimed at sustained independence and preserved quality of life for elders with mild cognitive impairment or mild Alzheimer’s disease, with the overall goal of delaying institutionalisation: supporting both “aging in place” (individuals remain in the home of choice as long as possible) and “community care” (long-term care for people who are mentally ill, elderly, or disabled provided within the community rather than in hospitals or institutions).
- Use Case 3 - Smart Sensing: This use case offers an interesting specular approach to this scenario (influenced by “Edge Computing” or “Pervasive Computing”). The main assumption is that intelligent data processing shall take place at sensor level, and that distributed data classification and clustering is a key aspect for massive system scalability. Moreover, in this use case algorithms derived from AI techniques will be deployed at Gateway, down to MCU level, also allowing the system to online/self-learn from the environment. The latter is a quite challenging aspect in itself in the AI field.

SEMIoTICS is an IoT Security/Privacy Cluster European Union project funded under the H2020-IoT-03-2017 work programme, with Grant Agreement number 780315.

Link:
[L1]: http://www.semiotics-project.eu

References:

[1] N. E. Petroulakis, et al.: “SEMIoTICS Architectural Framework: End-to-end Security, Connectivity and Interoperability for Industrial IoT”, in 2019 IEEE Global IoT Summit (GIoTS), 2019.
[2] K. Fysarakis, et al.: “Architectural Patterns for Secure IoT Orchestrations”, in 2019 IEEE Global IoT Summit (GIoTS), 2019.

Please contact:

Nikolaos Petroulakis, ICS-FORTH, Greece, npetro@ics.forth.gr

Towards Improved Mobility Support in Wireless Sensor Networks

by Amel Achour, Lotfi Guedria and Christophe Ponsard (CETIC)

Sensor networks are developing at a fast pace and are facing new challenges such as sensor mobility resulting in topological changes within a network or in adjacent networks with roaming nodes. To address such situations, we are working to extend a reference open source implementation of a border router with extended smart bridge mode.

Today’s world is more connected than ever: both objects of production and everyday objects are connected and exchange information. This raises many challenges for wireless sensor networks (WSN) to ensure the proper quality of service under low power constraints, for example when the sensors are mobile. Surveys have shown that two thirds of applications involve some sort of mobility (e.g. cold chain management, environmental monitoring, vehicle tracking) [1]. Compared with telephony and internet, this problem has received relatively little attention in the literature, and although it has been the subject of ongoing surveys and technological studies like ours, no standard solutions yet exist [2].

Due to their limited resources, sensors are designed to communicate without any infrastructure and each sensor routes packets of his neighbour to reach the border router, which is the gateway between the WSN and the Internet. Modern sensors are smart and can communicate directly with external services resulting in a very large range of applications from smart cities and manufacturing to buildings and home automation. Two types of mobility can be distinguished. Micro mobility refers to sensor movement within the same sensor network domain with no change in the network address prefix. Macro mobility refers to sensor movement to a new sensor network domain with another prefix. Our current scope of work is micro mobility.

To exchange information in a WSN, the RPL protocol (Routing Protocol for Low-Power and Lossy Networks) can be considered [L1]. Since it is designed for lossy networks, it adapts the routing strategies to any topological changes dynamically, meaning it provides implicit mobility support. However, this management is not efficient because it can take a long time to update the topology and reach the new location of a mobile sensor. Our work aims to enhance the global mobility management in the context of RPL-based WSN and increasing the network resilience.

Our proposal is based on 6LBR (6Lowpan Border Router) reference implementation [3] provided in Open Source [L2]. It allows multiple border routers to aggregate WSNs in order to increase resilience. Among the available modes, we consider the smart bridge (SmartBridge) mode where each border router (BR 1 to 3) manages its WSN independently as depicted in Figure 1(a). When a mobile sensor moves from BR1 to BR2, BR2 updates its network configuration and creates new routes so the mobile sensor can communicate in this new location. As BR1 is not aware of this update and the node will stay in its routing table until a timer expiration is reached after several attempts to join it. Our proposal implements a synchronised smart bridge (SyncSmartBridge) mode based on the virtual root principle combined to a synchronisation mechanism. The goal of the virtual root is to provide a unified control plane view of all the concrete WSN within the domain. The synchronisation mechanism will ensure the broadcast of synchronisation triggers to

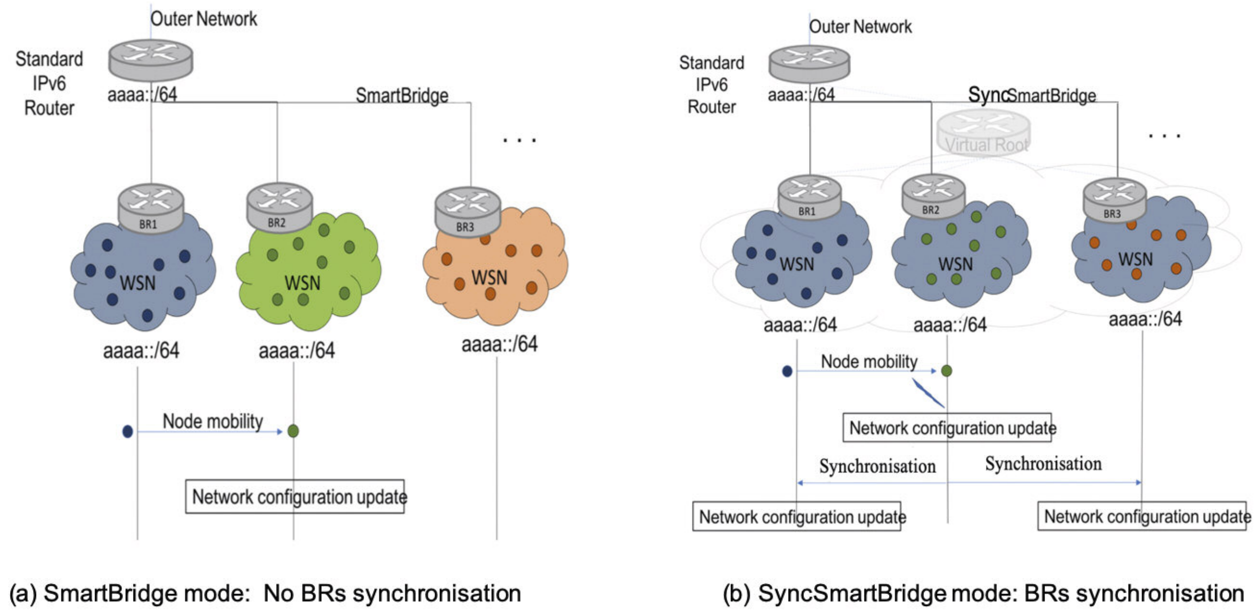

(a) SmartBridge mode: No BRs synchronisation

(b) SyncSmartBridge mode: BRs synchronisation

Figure 1: 6LBR based mobility management.

other BRs in order to update their topology and inform them about the new location of the mobile sensor. As shown in Figure 1(b), when a sensor moves from BR1 to BR2, it triggers a network configuration update in BR2. At the same time, BR2 broadcasts a synchronisation message to the other BRs so their topology and the sensor location is immediately updated.

Our protocol design is being validated using the Cooja simulator [L3]. The current work is focused on the scalability, performance assessment and robustness tests through a demonstrator that covers various baseline scenarios: one or multiple sensors moving in same/different directions, simultaneously or randomly. We also plan to consider a combination of more complex mobility scenarios to measure synchronisation delays in situations of high mobility and high network density.

Links:

[L1] https://kwz.me/hES
[L2] https://github.com/cetic/6lbr
[L3] https://anrg.usc.edu/contiki/

References:

[1] J. Wang, et al.: "A survey about routing protocols with mobile sink for wireless sensor network", Int. J. of Future Generation Communication and Networking, 2014;7(5).

[2] A. Achour, L. Deru, J-C. Deprez: "Mobility management for wireless sensor networks a state-of-the-art", Procedia Computer Science 52 (2015).

[3] L. Deru, et al.: "Redundant Border Routers for Mission-Critical 6LoWPAN Networks", Proc. of the Fifth Workshop on Real-World Wireless Sensor Networks, 2013.

Please contact:

Lofti Guedria, CETIC, Belgium
+32 472 56 62 70
lotfi.guedria@cetic.be

Designing IoT Architectures: Learning from Massive Spacecraft Telemetry Data Analytics

by Olivier Parisot, Philippe Pinheiro and Patrik Hitzelberger (Luxembourg Institute of Science and Technology)

"Decision Management System for Safer Spacecrafts" (DMSS) is a data analytics platform for the space domain that can detect anomalies in huge telemetry data acquired from numerous sensors in a complex IoT architecture.

Processing telemetry data is a challenge in any IoT architecture [1]. In the space sector, monitoring remote sensors is a common problem. The recent increase in downlink bandwidth and available processing power on spacecrafts have resulted in an increase in the number and the volume of telemetry measurements available to spacecraft controllers to monitor the health and safety of a spacecraft. In practice, the number of maintenance parameters is generally in the tens of thousands range. While providing a wealth of diagnostic information, these huge numbers overwhelm the ability of the human brain to oversee all telemetry measurements. Automatic checks and computer-aided data analysis can provide a daily overview for spacecraft operators and operations engineers.

In order to efficiently operate satellite constellations as well as spacecrafts, the Luxembourg Institute of Science and Technology [L1] and the Institute of Astronomy of KU Leuven [L2] have developed the "Decision Management System for Safer Spacecrafts" (DMSS) platform for the European Space Operations Centre (ESOC), the main mission control centre of the European Space Agency (ESA). DMSS offers a self-learning visual platform for anomaly detection in telemetry data coming from embedded sensors.

Based on complex data analytics algorithms [2], DMSS follows a visual analytics approach [3], providing interactive and visual representations of telemetry data and derived computational data. These representations come in the form of diagrams, such as Poincaré plots, KDE schemas, Heat maps, and time series plots (Figure 1), which are interactive and connected: choosing data in one diagram highlights corresponding data in another related diagram.

To help monitor the sensors, the main screen of DMSS presents a heat map displaying all telemetry anomaly scores for a specific day (Figure 2). The goal is to quickly identify which telemetry parameters potentially show an unusual behaviour:

- Blue means that no anomaly score was calculated for this date (due to missing or incomplete data – which is common with telemetry data from remote sensors).
- Grey means that the anomaly score is zero, thus the corresponding parameter has a standard behaviour for the selected date.
- Red means that the anomaly score is positive, this parameter has an uncommon behaviour; the user should click on it to see more details and investigate whether it is due to a real anomaly.
- Green means that the anomaly score is negative. Negative scores may occur for some score types when the anomaly score was previously positive and now the scoring system detects that the parameter behaviour is coming back to normal.

With this view, end-users can see if a group of telemetry parameters has a high anomaly score and see if one specific telemetry parameter has a high score compared to other similar telemetry parameters.

As a use-case, the data from two space missions operated by the European Space Agency were analysed: Mars Express (5,127 telemetry sensors i.e.

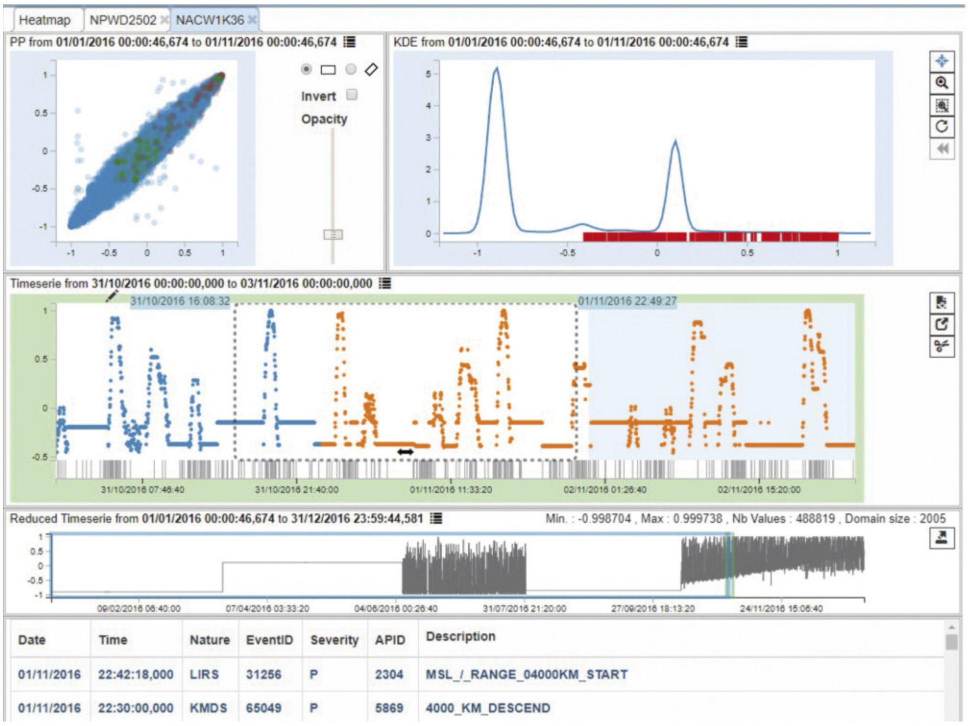

PP from 01/01/2016 00:00:46,674 to 01/11/2016 00:00:46,674

KDE from 01/01/2016 00:00:46,674 to 01/11/2016 00:00:46,674

Timeserie from 31/10/2016 00:00:00,000 to 03/11/2016 00:00:00,000

Reduced Timeserie from 01/01/2016 00:00:46,674 to 31/12/2016 23:59:44,581

Min. : -0.998704 , Max : 0.999738 , Nb Values : 488819 , Domain size : 2005

Date	Time	Nature	EventID	Severity	APID	Description
01/11/2016	22:42:18,000	LIRS	31256	P	2304	MSL_/_RANGE_04000KM_START
01/11/2016	22:30:00,000	KMDS	65049	P	5869	4000_KM_DESCEND

Figure 1: DMSS provides a dashboard to show telemetry data coming from remote sensors.

141 GB of data) and GAIA (28,209 telemetry sensors i.e. 1.46 TB of data for GAIA - including raw time series and statistical pre-calculations). For example, the number of points per time series is very variable, from a few samples up to 13 million samples. DMSS was able to display the potentially detected anomalies for the Mars Express mission. It also helped expert users to visually explore this important amount of data in a new way and try to find correlations between a potential detected anomaly and other parameters like events or commands sent to the spacecraft.

To conclude, DMSS helps to analyse large volumes of time series data coming from spacecraft sensors - which is a typical scenario for IoT applications. In this use case, operational risks are high, and efficient support of spacecraft operator engineers is of paramount importance. It is therefore crucial that software architecture and semi-automatic data analytics support perform efficiently and reliably. We anticipate that the approach could be applied to a large extent in other IoT scenarios.

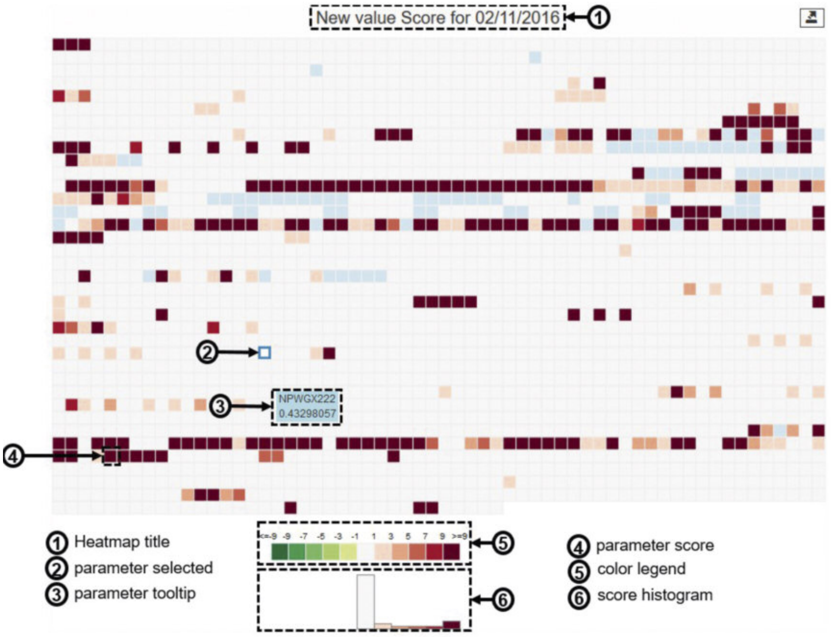

Figure 2: The heat map helps to the potential anomalies in telemetry data coming from these sensors.

Links:

[L1] https://www.list.lu/en/cooperations/lines-of-business/space/

[L2] https://fys.kuleuven.be/ster

References:

[1] Lazarescu et al.: “Design of a WSN Platform for Long-Term Environmental Monitoring for IoT Applications”, https://doi.org/10.1109/JETCAS.2013.2243032

[2] Royer et al.: “Data mining spacecraft telemetry: towards generic solutions to automatic health monitoring and status characterisation”, https://doi.org/10.1117/12.2231934

[3] Keim et al.: “Visual analytics: Definition, process, and challenges”, https://doi.org/10.1007/978-3-540-70956-5_7

Please contact:

Olivier Parisot, Philippe Pinheiro, Patrik Hitzelberger
Luxembourg Institute of Science and Technology
olivier.parisot@list.lu, philippe.pinheiro@list.lu, patrik.hitzelberger@list.lu

Autonomous Collaborative Wireless Weather Stations: A Helping Hand for Farmers

by Brandon Foubert and Nathalie Mitton (Inria)

In the context of smart farming, communications still pose a key challenge. Ubiquitous access to the internet is not available worldwide, and battery capacity is still a limitation. Inria and the Sencrop company are collaborating to develop an innovative solution for wireless weather stations, based on multi-technology communications, to enable smart weather stations deployment everywhere around the globe.

The French company Sencrop [L1] aims to help farmers in their hard daily work [1]. Sencrop manufactures and sells weather stations that autonomously collect data to help farmers accurately forecast the weather (e.g. risk of frost) and make sound decisions for the crop cycle (requirements for watering, fertiliser, etc). These stations collect accurate parcel-specific weather-related data, such as temperature, pluviometry, humidity, wind speed, etc. They rely on batteries that cannot be easily recharged or replaced regularly. Therefore, the Sigfox wireless technology is used to send the sensed data to the internet [L2]. Sigfox technology enables long range communications at a low energy expenditure. The Sigfox base stations to which the weather stations offload their data are deployed around the world to provide connectivity to customers.

But these advantages do not come without a price. Sigfox operates on the unlicensed ISM radio bands, which limits the time that a radio transceiver can emit by law. Sigfox technology thus suffers from low data rates and asymmetric wireless links. Furthermore, the Sigfox base stations are not yet deployed worldwide.

Sencrop weather stations require a worldwide ubiquitous network access to assist any farmer anywhere on Earth, but the coverage provided by Sigfox is insufficient. Furthermore, Sencrop needs to remotely upgrade the devices' firmware once a station is deployed; Sigfox can handle a small volume of regular monitoring data but is not equipped to support a whole firmware.

Indeed, the main issue faced by Sencrop is how to extend the current coverage and enable a bi-directional connectivity to the weather stations whilst limiting the energy consumption of devices. One possible solution is a weather station that integrates several wireless communication technologies, each providing different characteristics (delay, data rate, energy consumption, etc.), enabling the station to use the most appropriate technology for any given type of data to be sent.

As depicted in Figure 1, weather station "A" can reach Sigfox, as well as alternative technologies T1 and T2, and can autonomously pick the one that best suits its data requirements. T1 offers a large data rate and a bidirectional connectivity but with a high energy expenditure. It will be used for exceptionally large amounts of data, such as a firmware upgrades. T2 will only be used to send urgent data such as an alarm because it provides low delays at the cost of a medium energy consumption. Finally, Sigfox that features long delays and small data rates but benefits from a low power expense will be favoured for sending frequent and regular monitoring data. This solves the aforementioned issues, as station B

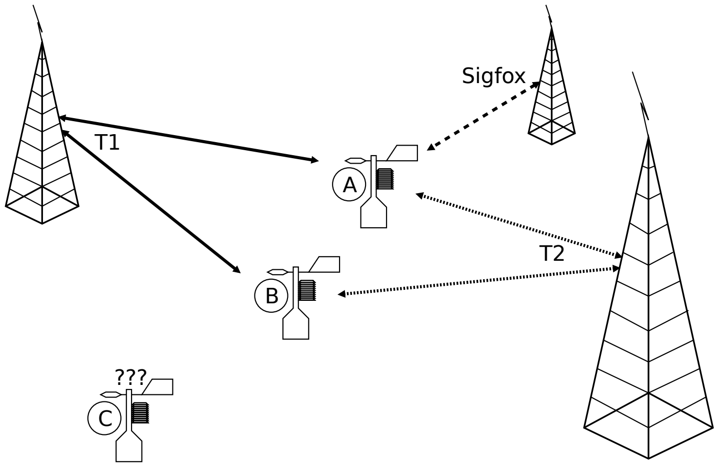

Figure 1: Autonomous multi-technology stations. A, B and C are weather stations. T1, T2 and Sigfox are different wireless techonologies.

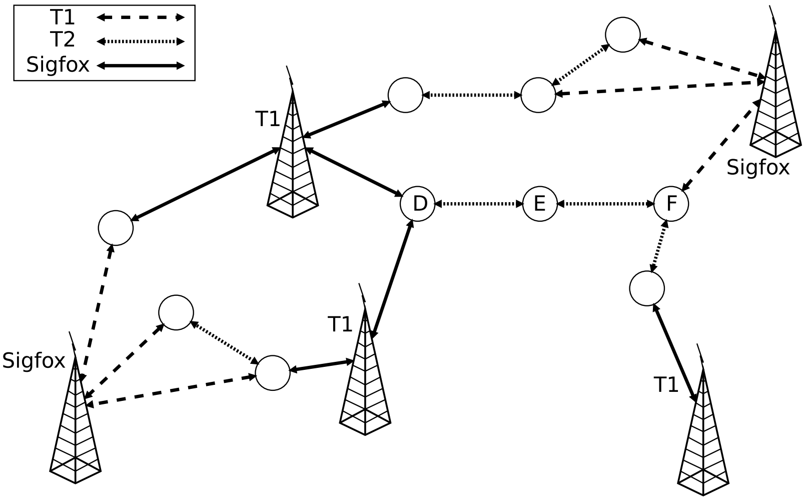

Figure 2: Multi-hop multi-technology network. D, E and F are weather stations. T1, T2 and Sigfox are different wireless technologies.

which is not in range of the Sigfox network can still use the other two technologies.

But this solution raises other issues. Adding wireless technologies and the ability to switch from one to another consumes energy and processing capabilities, resources that are limited on Sencrop stations. Moreover, stations are still dependent on a direct link to a base station. Areas deprived of any communication capability still exist on our planet in isolated rural environments. Inria and Sencrop jointly address this exciting challenge by designing a joint smart network interface selection, and a multihop multi-technology routing mechanism. It will dynamically adapt the data forwarding according to the local environment (remaining energy, connectivity, etc.) and the importance of the data by respecting environmental and hardware limitations of each station.

As depicted in Figure 1, if a station is isolated, such as station C, it may still be in the vicinity of another station, like B, through which it can reach a base station. This mechanism will not only extend the network coverage, but also save power and alleviate network traffic. Indeed, instead of using energy intensive long communication links, a data packet can follow a suite of more energy efficient smaller ones. For instance in Figure 2, station D can choose to offload its data through E and F (technology T2) and then to the base station (Sigfox) which might be less energy costly than using the direct link T1 to the base station. Stations will not only have to choose which wireless technology to use autonomously, but also to which neighbour they send their data to satisfy the data requirements.

This system is expected to improve the overall energy efficiency of Sencrop stations and enable a broader deployment to assist farmers with smart connected things everywhere around the globe.

This work was partially supported by a grant from CPER DATA, Sencrop, FEDER and I-SITE.

Links:
[L1] https://sencrop.com/en/
[L2] https://www.sigfox.com/en

Reference:

[1] A. Gregoire: “The mental health of farmers”, Occupational Medicine, Vol. 52, Issue 8, 2002, p. 471–476, https://doi.org/10.1093/occmed/52.8.471

Please contact:

Brandon Foubert
Inria Lille, France
+ 33 3 59 57 79 43
brandon.foubert@inria.fr

Automatic Detection of Allergic Rhinitis

by Gregory Stainhaouer, Stelios Bakamidis and Ioannis Dologlou (RC ATHENA)

The spectral characteristics of speech can be used to cluster individuals according to whether or not they suffer from an allergy. Based on the principles of adaptive modelling and fundamental frequency variations, as well as speech analysis by means of acoustic models, our technique achieves an efficient classification based on uttered speech over a mobile phone. The final decision is derived by combining the individual estimates, providing a tool for the automatic diagnosis of allergies.

The distinctive speech characteristics of people suffering from allergic rhinitis have the potential to enable remote diagnosis using speech analysis technology. We have developed a system for this purpose, consisting of two experts: the first provides a robust estimate of the jitter of the fundamental frequency and the second provides a confidence score derived from Hidden Markov Model-based acoustic modeling. These experts are calibrated using two sets of speech data, one from patients with allergies and a second from the same individuals after treatment. This supervised calibration allows the parameters of the two experts' algorithms to be fine-tuned and enables some thresholds to be derived, which are subsequently used to diagnosis others.

The notion of the rank of a signal, that is needed in this work, is defined by means of the rank of its covariance matrix [1,2]. Given the signal s and the nth order covariance matrix Cn, the rank of the signal s is equal to p, where p is the rank of the covariance matrix Cn and p < n. In case p = n for all n, the signal s is said to be full rank. Full rank signals are more complex to handle and are converted to limited rank by means of a singular value decomposition (SVD) based successive projections algorithm [3].

The fundamental frequency estimation algorithm is highly accurate and consists of two steps. During the first step the goal is to obtain from the original speech signal a new one where the fundamental frequency is the predominant frequency. This signal enters the second step of the algorithm that provides an approximate rank two signal of the fundamental frequency.

The first step is based on an iterative zero-phase filtering whose frequency response is a monotonically decreasing function. This guarantees that only the energy around the fundamental frequency remains as the number of iterations increases.

For the second step the SVD algorithm to reduce the rank of a signal by successive projections is used. Finally, the jitter of the fundamental frequency is computed as the variance of the distances between consecutive peaks of the rank two signal (Figure 1).

The second expert uses acoustic models for the derivation of confidence scores for the uttered phrases. These models provide estimates of the probability of features extracted from speech and give a string of phonemes. Each phoneme is represented by a Hidden Markov Model. For the Greek language, 32 phonemes are used to describe the various pronunciations. The acoustic modelling not only provides the phonetic transcription of the utterance but also

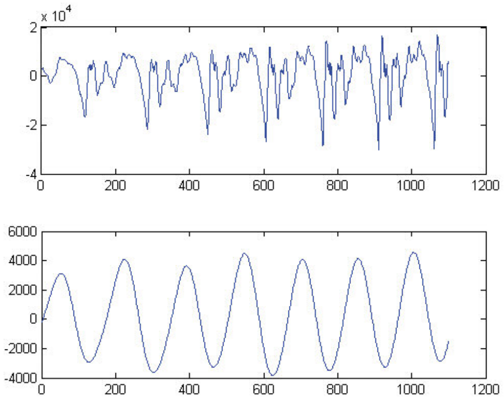

Figure 1: The top figure depicts a frame of voiced speech and the bottom figure depicts the corresponding rank two signal. Note that peaks are very clear and distances between peaks can easily be computed to ensure an accurate value for jitter.

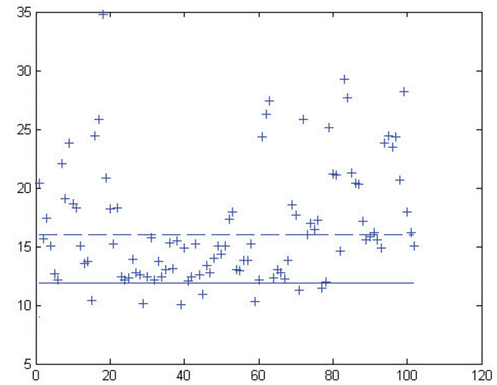

Figure 2: The solid line depicts mjc, the mean of healthy condition and the dashed line depicts mjs, the mean of allergic condition. The various + symbols depict unhealthy patients. Most of them lie beyond mjs and are safely classified. Many lie between the two means and decision becomes ambiguous and few are misclassified below mjc.

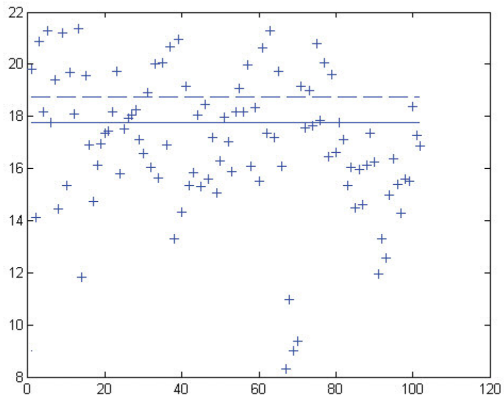

Figure 3: The dashed line depicts mpcg, the mean of healthy condition and the solid line depicts mpsg, the mean of allergic condition. The various + symbols depict unhealthy patients. Most of them lie below mpsg and are safely classified. Some lie between the two means and decision becomes ambiguous. There are also some misclassified cases lying above mpcg.

the probability of each phoneme which reflexes the confidence score.

The automatic diagnosis system was first trained and then tested using real life data collected from patients. For the training phase sixteen patients were used, who returned at a later stage when they were cured and uttered the same phrases. The two sets of experts that need to be trained are the jitter and the confidence score of the acoustic models.

In order to handle new patients, the mean values of the jitter for healthy and unhealthy subjects are computed, denoted by mjc and mjs respectively. These values are mjs=16 and mjc = 12. If jn stands for the jitter of the new patient then if jn > mjs the patient is unhealthy, whereas if jn < mjc the patient has no problem. For values of jn between mjs and mjc no accurate decision can be made (Figure 2).

The phrases are also processed by the acoustic models and a mean confidence score is produced. By denoting mps, the mean value that a patient gets before treatment, and mpc, the mean value that the same patient gets after treatment, one expects to find mpc > mps. A global mean value for all cured patients mpcg and a global mean for all unhealthy patients mpsg is computed to enable the classification of a new patient according to the confidence score criterion. These values are mpsg = 17.7 and mpcg = 18.7. If mpn stands for the mean confidence score of the new patient then if mpn < mpsg the patient is unhealthy, whereas if mpn > mpcg the patient has no allergy. For values of mpn between mpcg and mpsg no accurate decision can be made (Figure 3). A simple linear optimisation technique is used to combine the two experts in order to provide an overall decision.

The algorithms were developed in the Research Center ATHENA in Athens Greece. The work started in October 2018 and the first version was delivered in June 2019. The project “Patient Allergy Tracer” (project code: T1EΔK-02436) supported this work, which is implemented under the Action “Research, Create, Innovate”, funded by the Operational Program “Competitiveness, Entrepreneurship and Innovation” (NSRF 2014-2020) and co-financed by Greece and the European Union (European Regional Development Fund).

Future activities

Future plans focus on the performance of the algorithms both in terms of accuracy and speed. Improving the accuracy involves a better training process by including more patients in the experiment and also ameliorates the performance of the acoustic models. To improve the speed, fast SVD algorithms are needed for the implementation of the successive projections algorithm.

References:

[1] I. Dologlou, G. Carayannis: “Physical interpretation of signal reconstruction from reduced rank matrices”, IEEE ASSP, July 1991, pp. 1681-1682.

[2] I. Dologlou, S. Bakamidis, G. Carayannis: “Signal decomposition in terms of non-orthogonal sinusoidal bases”, Signal Processing, Vol. 51, June 1996.

[3] J.A. Cadzow: “Signal enhancement: A composite property mapping algorithm”, IEEE, ASSP, Vol. 36, No. 1, January 1998, 49-62.

Please contact:

Ioannis Dologlou
RC ATHENA, Greece
+302106875306, ydol@ilsp.gr

Supporting the Wellness at Work and Productivity of Ageing Employees in Industrial Environments: The sustAGE Project

by Maria Pateraki (FORTH-ICS), Manolis Lourakis (FORTH-ICS), Leonidas Kallipolitis (AEGIS IT Research),, Frank Werner (Software AG), Petros Patias (AUTH) and Christos Pikridas (AUTH)

Industrial environments can benefit from smart solutions developed on top of an infrastructure combining IoT and smart sensors that monitor workers in an non-invasive way, allowing the early detection and prevention of health risks. The sustAGE project aims to improve occupational safety and workforce productivity through personalised recommendations in two key industrial environments.

The use of health and well-being technologies in smart IoT ecosystems has been steadily increasing, and they are now found in environments such as smart homes, age-friendly workplaces and public spaces [1]. The deployment of smart sensors aims to support functional, physiological and behavioral monitoring, which can benefit both older adults facing gradual degradation of their motor and cognitive skills due to aging and workers in harsh environments performing arduous, stressful or health hazardous tasks [2]. Such smart systems can promote a reactive living and working environment, which provides appropriate and timely recommendations, acts preventively and mitigates health risks.

sustAGE [L1] is a H2020 EU project that aims to develop a person-centered smart solution fostering the concept of "sustainable work" for EU industries, thus supporting the well-being, wellness at work and productivity of ageing employees along three main dimensions. The first dimension is directed towards improving occupational safety and health based on workplace and person-centered health surveillance monitoring. The second aims to promote the well-being of employees via personalised recommendations for physical and mental health whilst the third supports decision making related to task/job role modifications aiming to optimise overall workforce productivity. The sustAGE solution is deployed in two challenging industry domains, specifically: (i) manufacturing, focusing on car assembly line workers; and (ii) transportation & logistics, focusing on dock workers involved in vessel loading/unloading operations.

The system functionalities build upon an IoT ecosystem, based on off-the-shelf sensors integrated into daily devices and in the work environment, considering both indoor (manufacturing) and outdoor (port) working conditions. The system gathers contextual information from the working environment and users' physiological signals, tasks, activities and behavioural patterns, in order to support user profiling and provide personalised recommenda-

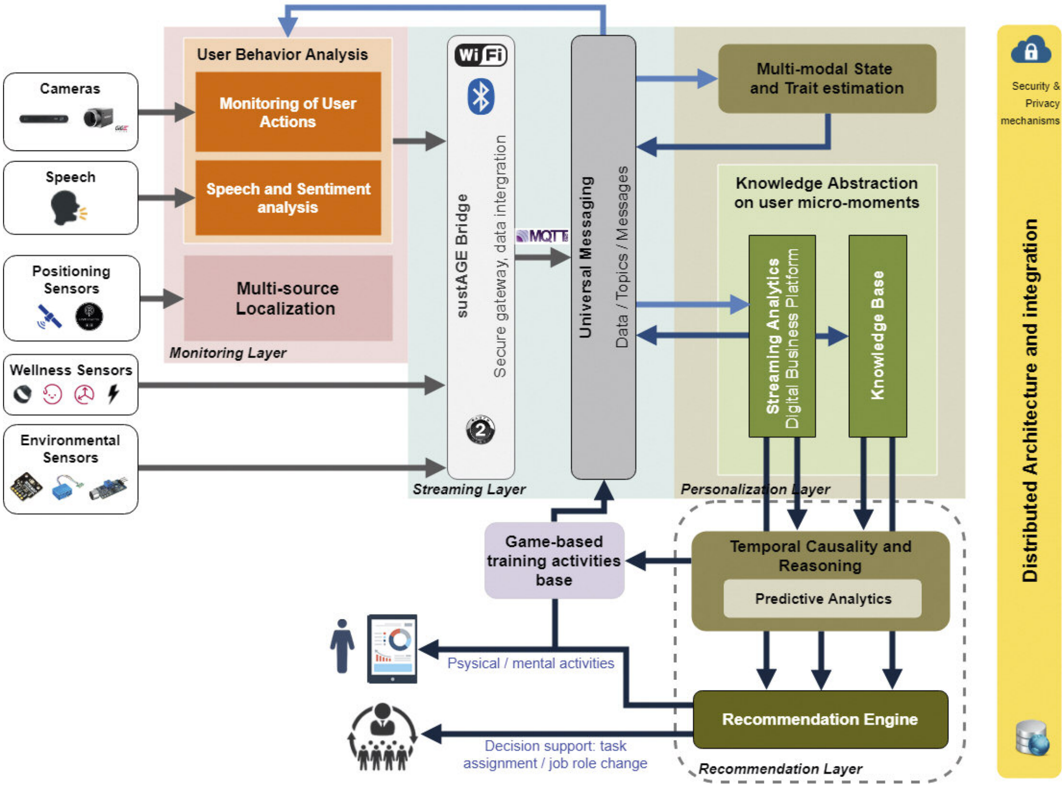
Cameras
Speech
Positioning Sensors
Wellness Sensors
Environmental Sensors
User Behavior Analysis
Monitoring of User Actions
Speech and Sentiment analysis
Multi-source Localization
Monitoring Layer
Wi Fi
sustAGE Bridge
Secure gateway, data intergration
MQTT
Universal Messaging
Data / Topics / Messages
Streaming Layer
Multi-modal State and Trait estimation
Knowledge Abstraction on user micro-moments
Streaming Analytics
Digital Business Platform
Knowledge Base
Personalization Layer
Security & Privacy mechanisms
Distributed Architecture and integration
Game-based training activities base
Temporal Causality and Reasoning
Predictive Analytics
Psysical / mental activities
Recommendation Engine
Decision support: task assignment / job role change
Recommendation Layer

Figure 1: sustAGE system architecture.

tions. Measurements collected from different devices and system modules support the definition of key micro-moments related to the user's daily schedule, work environment, workload, physical/emotional/mental state and social activities, whilst taking into account associated temporal variables. The IoT infrastructure comprises of: (i) Environmental sensors measuring air temperature, humidity, air quality, pressure, dust concentration and noise with custom, low cost Raspberry Pi/Arduino sensors; (ii) Passive cameras installed in key working locations, specifically stereo cameras for monitoring posture and repetitive user actions in the indoor manufacturing environment and mono cameras in the outdoor port environment for monitoring crane operators and dock workers involved in vessel loading/unloading; (iii) Beacons for localisation in indoor environments, achieving a precision of up to 10-20 cm within a range of 100 m; (iv) GNSS receivers built in smartphones for localisation in outdoor environments; (v) Wristwatch devices gathering physiological measurements, able to deliver notifications to users from the system; and (vi) Galileo-enabled Smartphone devices, offering centimetre accuracy and ability to communicate with the wristwatches.

The aforementioned devices/sensors collaboratively provide information on different user activities/actions (e.g. walking, bending, standing/sitting, pushing/pulling objects) and states (e.g. fatigue, discomfort), integrating temporal aspects and detecting specific events in the environment (e.g. user presence in specific areas, proximity to hazardous conditions). Moreover, the smartphone is the primary device for communication and multimodal interaction supporting natural language understanding and voice sentiment analysis. The adopted IoT configuration exhibits the advantages of unobtrusive user context interaction monitoring in a privacy-preserving way, since in private life, outside the working environment, only the wristwatch and the smartphone are used.

Figure 1 shows the system's architecture, which is conceptually divided into four layers, namely monitoring, streaming, personalisation and recommendation. The monitoring layer includes components that receive raw data from various sources, supporting processing near the end-devices that prevents potentially privacy-sensitive information from being sent to the system's upper layers. Cameras, wristwatches, environmental and location sensors as well as users' speech comprise the list of data sources that feed the system. The streaming layer includes the sustAGE "Bridge and Universal Messaging" modules that ingest the data through a secure gateway and prepare them for distribution to the rest of the system. Data streams of different protocols are buffered, homogenised and further sent to subsequent components via a common communication protocol that ensures speed and scalability in delivery of near real-time data. The Universal Messaging bus supports streaming of the real-time data, through the topics/messages across the IoT infrastructure and all other platform components. The personalisation layer includes personalisation mechanisms generating new information, through distilling the short-, medium-term states and long term traits, extracting knowledge abstractions and enabling user profiling to be updated regularly by exploring association with past episodes. The Apama streaming analytics engine [L2] performs analytics on the incoming data, referencing historical information where necessary, to identify previously occurred patterns. Generated knowledge is stored in the respective knowledge base to constantly improve analytics, reasoning and thus user recommendations. The recommendation layer consists of reasoning and recommendation modules which determine user recommendations, taking into account spatiotemporal constraints.

Environmental sensor measurements along with user state parameters are collected by the IoT platform of the streaming layer through dedicated gateways for monitoring and analytics. In sustAGE, the main requirements of the IoT management solution relate to: (i) interoperability, that is the ability to acquire data from devices from different vendors that use different protocols; (ii) scalability, i.e. the ability to handle increased amounts of data; and (iii) security and privacy of the communicated data. Apart from these challenges, efforts are focused on robust data analysis techniques to efficiently handle the data streams and support the determination of proper actions.

sustAGE is a multidisciplinary project involving ten European partners: Foundation for Research and Technology – Hellas (Greece, coordinator), Centro Ricerche Fiat, (Italy), Heraklion Port Authority (Greece), Software AG (Germany), Imaginary Srl. (Italy), AEGIS IT Research UG (Germany), Leibniz Research Centre for Working Environment and Human Factors (Germany), University of Augsburg (Germany), National Distance Education University (Spain) and Aristotle University of Thessaloniki (Greece).

Links:

[L1] https://www.sustage.eu

[L2] https://en.wikipedia.org/wiki/Apama_(software)

References:

[1] A. Kumar, H. Kim and G.P. Hancke: "Environmental monitoring systems: A review", IEEE Sensors Journal 13(4):1329-1339. 2013. https://www.doi.org/10.1109/JSEN.2012.2233469

[2] A. Almeida, et al: "A critical analysis of an IoT—aware AAL system for elderly monitoring", Future Generation Computer Systems, Vol.97: 598-619, 2019. https://www.doi.org/10.1016/j.future.2019.03.019.

Please contact:

Maria Pateraki, Manolis Lourakis
FORTH-ICS, Greece
pateraki@ics.forth.gr
lourakis@ics.forth.gr

Leonidas Kallipolitis
AEGIS IT Research, Greece
lkallipo@aegisresearch.eu

Frank Werner
Software AG, Germany
Frank.Werner@softwareag.com

Petros Patias, Christos Pikridas
Aristotle University of Thessaloniki (AUTH), Greece
patias@auth.gr, cpik@topo.auth.gr

High Performance Software Defined Storage for the Cloud

by Thomas Lorünser (AIT), Stephan Krenn (AIT) and Roland Kammerer (Linbit HA-Solutions GmbH)

Distributed Replicated Block Device (DRBD) is the de facto standard for redundant block storage. It is used more than 250,000 times world-wide and part of the official Linux kernel. DRBD4Cloud is a research project which aims at increasing the applicability and functionality of DRBD in order to enter new markets and to face future challenges in distributed storage.

Redundant data storage is a necessity for business continuity of virtually any cloud service. An intuitive approach is full data replication to multiple storage nodes, which is currently done by DRBD [L2]. However, due to large bandwidth requirements and storage overhead this is not feasible for large-scale deployments with many mirroring nodes, i.e., typical cloud settings. Other than DRBD, Ceph is the major backend block storage implementation. Ceph already offers integration with OpenStack. However, Ceph's performance characteristics prohibit its deployment in certain low-latency use cases, e.g., as backend for Oracle MySQL databases.

DRBD4Cloud [L1] will increase the performance and scalability (technical and organisational) of highly-available software defined block storage in dynamic large scale cloud deployments. It is based on DRBD, which is a high-performance low-latency low-level building block for block replication, offering key functionality of such systems. During the project DRBD will be integrated into OpenStack and DC/OS (the Distributed Cloud Operating System), the most prevalent tool suites to manage distributed computing resources. This will enable cloud backend providers to use DRBD technologies for replicated block storage, which is an inherently needed building block for highly reliable cloud storage offerings.

To ease the deployment and maintenance of DRBD a collection of key components will be offered as an easy-to-use software package. A component to orchestrate and monitor the storage environment consisting of a multitude of DRBD nodes will be developed. On top of this, web-based and API-based management and monitoring solutions will be developed in order to ease adoption of DRBD-based software-defined storage for cloud environments.

To guarantee availability, DRBD is currently storing up to 32 full data replicas on remote storage nodes. DRBD4Cloud will allow for the usage of erasure coding, which allows one to split data into a number of fragments (e.g., nine), such that only a subset (e.g., three) is needed to read the data. This will significantly reduce the required storage and upstream bandwidth (e.g., by 67 %), which is important, for instance, for geo-replication with high network latency. Additionally, specific schemes called secret sharing can even guarantee that the servers do not learn anything about the plain data,

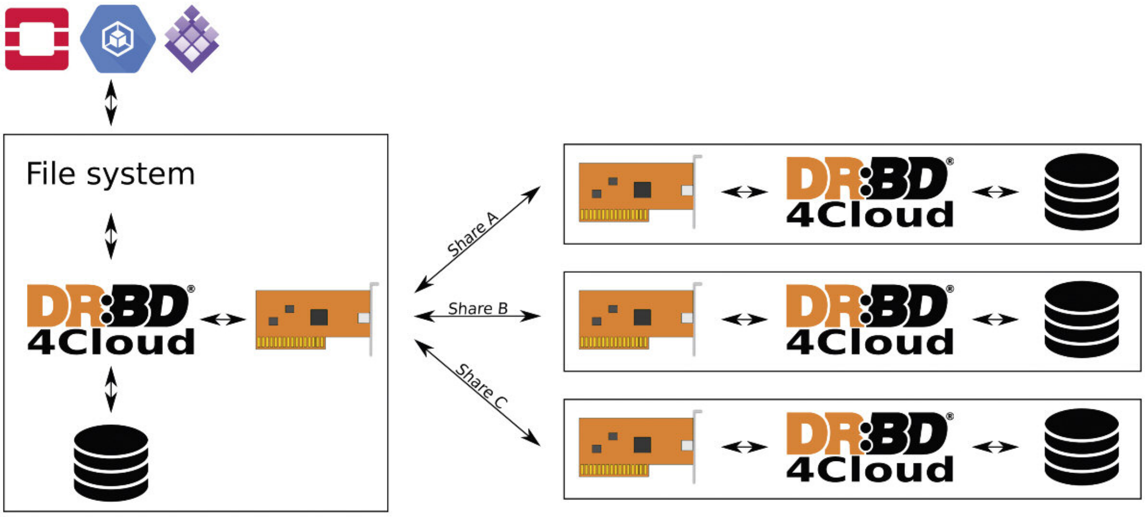

Figure 1: Depending on the configuration, no individual share contains any sensitive information about the replicated data, while the data can be recovered from any two of the shares to achieve high availability.

without requiring cryptographic keys [1]. This will allow for using public cloud storages without compromising confidentiality, making DRBD also usable to SMEs without own data centres but requiring highly available storage.

The feasibility of cloud integration of DRBD for small setups has already been demonstrated by means of a first proof-of-concept prototype which showed that the main challenge is to meet the scalability requirements of large scale deployments. Another challenge is the design and maintenance of a common shared driver core between the OpenStack and DC/OS integrations. Additionally, concerning erasure coding and secret sharing the main challenge is minimising the impact upon latency. Fortunately, in such schemes, read operations only require access to a subset of the storage nodes in order to retrieve the data, which positively affects the availability and latency. A first result on hardware acceleration of secret sharing [2] also shows the potential of low latency implementations on modern high-bandwidth network interface cards. As an optimisation, the monitoring solution could be utilised to optimise the global workload of the overall storage cluster and as results in [3] show, efficient auditing mechanisms can be used to verify the data integrity in the system although erasure coding and secret sharing are applied.

In summary, with DRBD4Cloud a new and highly optimised out-of-the-box solution for multi-node software defined storage in high-load cloud environments will be developed. The results will be delivered as software implementations, which similar to DRBD will (mostly) be published under an open source license. The extensions will allow DRBD – the de-facto standard for distributed replicated block devices – to cut the storage overhead by 50-80 % while guaranteeing practically equivalent levels of redundancy and will allow confidential data to be securely stored on semi-trusted cloud providers. Simplifying the deployment and enabling monitoring as well as integrating DRBD into cloud platforms will allow cloud providers to pick up the technology. At the end of DRBD4Cloud, all extensions will be delivered as internally tested prototypes. DRBD4Cloud is a joint effort of Linbit HA-Solutions GmbH, Pro-zeta A.s. and AIT Austrian Institute of Technology (AIT) which has received funding from the EUREKA Eurostars programme [L3].

Links:

[L1] https://kwz.me/hE3
[L2] https://www.linbit.com
[L3] https://www.eurostars-eureka.eu/project/id/11450

References:

[1] T. Lorünser, A. Happe, and D. Slamanig: "ARCHISTAR: Towards Secure and Robust Cloud Based Data Sharing", in CloudCom 2015, IEEE, https://doi.org/10.1109/CloudCom.2015.71

[2] J. Stangl, T. Lorünser, S.M. Dinakarrao: "A fast and resource efficient FPGA implementation of secret sharing for storage applications", DATE 2018, pp. 654–659, https://doi.org/10.23919/DATE.2018.8342091

[3] D. Demirel, et al.: "Efficient and Privacy Preserving Third Party Auditing for a Distributed Storage System", ARES 2016, pp. 88–97, https://doi.org/10.1109/ARES.2016.88

Please contact:

Stephan Krenn
AIT Austrian Institute of Technology GmbH
Stephan.Krenn@ait.ac.at

The Battle of the Video Codecs in Healthcare

by Andreas S. Panayides (Sigint Solutions and University of Cyprus), Marios S. Pattichis (University of New Mexico) and Constantinos S. Pattichis (University of Cyprus)

Video compression is the core technology in mobile (mHealth) and electronic (eHealth) health video streaming applications. With global video traffic projected to reach 82 % of all internet traffic by 2022, both industry and academia are struggling to develop efficient compression algorithms to match these unprecedented needs. To the best of our knowledge, this is the first performance comparison of emerging VVC and AV1 video codecs for medical applications.

Video codecs are currently experiencing unparalleled levels of growth driven by unprecedented video traffic demands that are projected to reach 82 % of all internet traffic by 2022 [1]. This growth is facilitated by advances in open-source video delivery protocols, such as Dynamic Adaptive Streaming over HTTP (MPEG-DASH) and WebRTC (Real-Time Communication (RTC)). Most importantly, industry initiatives such as Google's WebM project (leading VP8/ 9 video codecs development), and subsequent formation of the Alliance for Open Media (AOM) [L1] comprising key industry affiliates, have created a highly competitive environment, investing efforts towards creating the first ever royalty free video codec. In this context, AOM announced the code freeze of its debut encoder, AV1, in 2018, claiming the best encoding performance to date. At the same time, the Joint Video Experts Team (JVET) formed in October 2017 that has undertaken the development of H.265's successor, termed Versatile Video Coding (VVC), has just released its reference software abbreviated VTM (VVC Test Model) [L2].

In healthcare, video compression is a key enabling technology that is widely used for real-time medical video communications in a mobile-health setting [2], [3]. Application scenarios range from remote diagnosis and care to emergency incidence response, medical education, tele-robotics, and second opinion provision. A plethora of applications further extends to the home setting for assisted living applications. Moreover, video compression is also key in low-bandwidth applications such as remote diagnosis from isolated locations, developing countries, and disaster sites.

The objective of this study was to investigate the compression efficiency of well established (VP9 and H.265 (also termed high efficiency video coding (HEVC)), recently standardised (AV1), and emerging (VVC) video codecs and provide preliminary insights into their applicability in the healthcare domain.

Experimental Setup

The dataset used in this series of experiments consisted of 10 atherosclerotic plaque ultrasound videos with a video resolution of 560x448 at 40 frames per second, with a duration of 10 seconds, and yuv420p raw format (see Figure 1). Selected quantization parameters values for constant quality encoding, typical in video compression performance comparison studies, were {27, 35, 46, 55} for AV1, VP9 and {22, 27, 32, 37} for VVC, HM, and x265, to enable a fair comparison and a wide range of representative bandwidths given the video characteristics. Random access settings involved an intra update every 32 frames, aligned with the specified group of pictures (GOP) level. Default preset parameters available in every video coded were set at RandomAccess for VVC and HM, --good and –best for AV1 and VP9, respectively, and –placebo for x265.

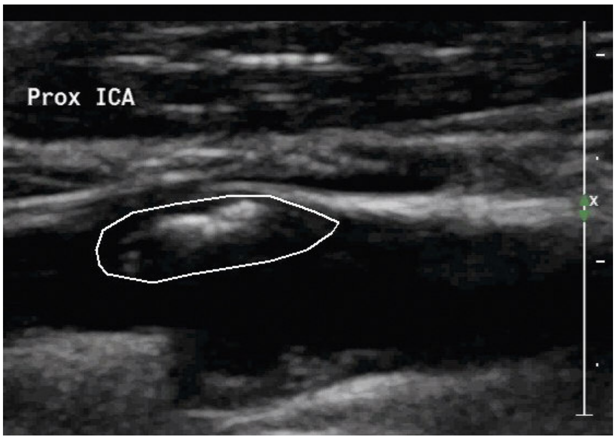

Figure 1: Original (uncompressed) ultrasound video image example of the internal carotid artery (ICA). The straight white line delineates the atherosclerotic plaque causing stenosis. Video resolution: 560 × 448, frame rate: 40 fps.

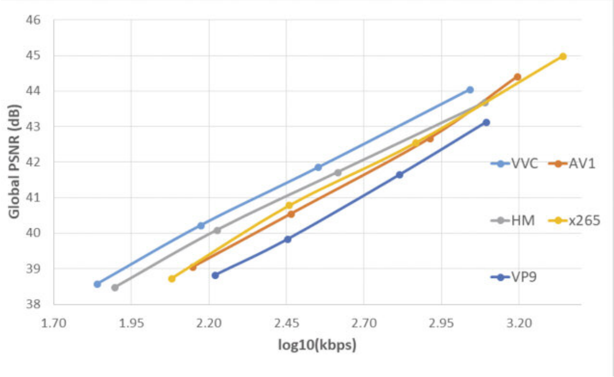

Figure 2: Video coding standards comparison. Rate-distortion curves depicting Peak Signal to Noise Ratio (PSNR) vs log (Bitrate) using mean values of the ten 560x416@40 frames/s ultrasound videos for all investigated QP values. VVC outperforms all rival video codecs.

	Bitrate savings relative to:			
Encoding	**HM (HEVC)**	**X265 (HEVC)**	**AVI**	**VP9**
VVC	23%	33%	33%	54%
HM (HEVC)		13%	18%	40%
X265 (HEVC)			4%	30%
AVI				25%

Table 1: Performance comparison of video coding standards in terms of overall bitrate gains. Savings were computed using the BD-Rate algorithm. VVC outperformed all other video codecs.

Performance Evaluation

Preliminary results depicted in Figure 2 and Table 1 show that VVC already appears to outperform AV1, despite VVC being in the early stages of the development process. For this particular experimental setup, bitrate demands reductions of 36 % were recorded for equivalent quality using the BD-RATE algorithm. Interestingly, both HEVC instantiations, namely HM (HEVC Test Model) and x265, demonstrated higher coding efficiency than AV1, with AV1 entailing additional bitrate requirements of 18 % and 4 %, respectively. It is important to note, however, that this finding is largely due to the limited video resolution of the investigated CCA videos. Ongoing experimentation shows that as video resolution increases, AV1 tends to outperform both HM and x265, which is consistent with the published literature. While not obsolete, VP9 is mostly used for benchmarking purposes in this study, significantly lacking in compression efficiency compared to all rival codecs. On a different note, besides x265 and VP9 that involve production level software optimization and hence can qualify for real-time performance, VVC, AV1, and HM incur long compression times, given that their principal intended usage is to validate their compression efficiency rather than for use in real-time video streaming applications.

Undoubtedly, more comprehensive experiments are needed to deduct the best video coding software, investigating different medical video modalities and higher video resolutions (high and ultra-high definition, and beyond). Studies on general-purpose videos in the literature revealed that encoding performance is largely affected by content and video size, and hence safe conclusions can only be drawn once experiments extend across the video resolution ladder and involve a sufficient number of medical video modalities. Such a study is already underway for emergency scenery and echocardiogram videos within the context of the Adaptive Video Control for Real-time Mobile Health Systems and Services (ACTRESS) project.

This work was co-funded by the European Regional Development Fund and the Republic of Cyprus through the Research and Innovation Foundation (Project: POST-DOC/0916/0023); Acronym: ACTRESS.

Links:
[L1] https://aomedia.org/
[L2] https://jvet.hhi.fraunhofer.de/

References:

[1] Cisco, V. N. I.: “Cisco Visual Networking Index: Forecast and Trends, 2017–2022”, White Paper, 2018.
[2] A. S. Panayides, M. S. Pattichis, C. P. Loizou, et al.: “An Effective Ultrasound Video Communication System Using Despeckle Filtering and HEVC”, in IEEE J Biomed Health Inform, vol. 19, no. 2, pp. 668-676, 2015.
[3] Z. C. Antoniou, A. S. Panayides, M. Pantzaris, et al.: “Real-Time Adaptation to Time-Varying Constraints for Medical Video Communications”, in IEEE J Biomed Health Inform, vol. 22, no. 4, pp. 1177-1188, July 2018.

Please contact:

Andreas S. Panayides
Sigint Solutions and University of Cyprus, Cyprus
+35722892757
a.panayides@sigintsolutions.com; panayides@cs.ucy.ac.cy

The VR4REHAB Interreg Project: Five Hackathons for Virtual and Augmented Rehabilitation

by Daniele Spoladore, Sonia Lorini and Marco Sacco (STIIMA-CNR, EUROVR)

Virtual and augmented reality technologies can support the process of clinical rehabilitation, making therapy more engaging, challenging and measurable for people with disabilities or injury, encouraging them to continue their rehabilitation outside of the clinical environment [1]. Although these technologies are used in some rehabilitation-related fields [2, 3], their extensive application in the clinical field is still limited.

The European Project VR4REHAB aims to foster the adoption of novel virtual reality (VR) applications for rehabilitation and health care. The project, a collaboration between seven European partners (Sint Maartenskliniek, European Association of Virtual Reality and Augmented Reality (EUROVR), St. Mauritius Therapieklinik, Teesside University, Royal Free London NHS Trust, Université de Lille 1 - Sciences et Technologies and Games for Health Europe), aims to create innovative VR technologies for rehabilitation clinics. Indeed, virtual and augmented reality can help patients recover at home, with or without therapists, by following virtual instructions while being monitored by a set of sensors. Moreover, the virtual and augmented application can calibrate the difficulty of tasks according to the patient’s ability, thus reducing boredom associated with too-easy tasks and the frustration that too-difficult tasks would generate. This can result in better engagement rates.

Figure 1: UK Hackathon (4 - 5 July 2018) & French Hackathon (14 - 15 June 2018). Participants of the event experiencing VR tools and pitching their ideas to an international audience.

In 2018, the VR4REHAB project ran five hackathons across Europe (in The Netherlands, Germany, France, United Kingdom and Belgium), gathering talented youngsters, entrepreneurs, patients, rehabilitation personnel and technical specialists, to develop new ideas about rehabilitation with virtual and augmented reality. Specifically, the project fostered the development of new solutions to tackle five clinical themes: (i) pain management; (ii) engagement, and immersion to promote treatment adherence, (iii) behavioural and cognitive training in children with brain injury, 4) lower limbs and mobility; (iv) training of upper limb movements.

Following the hackathons, the best ideas enter a development phase, the Game Jams, during which the selected concepts are developed in prototypes with the help of companies and ICT professionals; the developed prototypes will be tested in clinical trials throughout 2019. Finally, in the last phase of the project, the Challenges, the prototypes will become real applications: the developers, supported by professionals, will develop business plans to put their application on the market and to reach their target groups.

VR4REHAB's five hackathons proved that the hackathon idea was a success: more than 270 people between 18 and 45 years old, such as young innovators, developers, researchers, patients, and health care professionals, participated in the event.

The ideas generated during the hackathons will be stored and published into an online library: a digital platform in which the concepts will be available for further development. Moreover, the online library will be made available to health professionals, who can provide feedback or helping the developers in fine-tuning their applications. This platform will be integrated into EuroVRs' website by the end of the project.

This project was financed by the EUx InterReg NWE programme.

References:

[1] S. Arlati, et al.: "A Social Virtual Reality-Based Application for the Physical and Cognitive Training of the Elderly at Home", Sensors 19.2 (2019): 261.

[2] D. Spoladore, et al.: "Semantic and Virtual Reality-enhanced configuration of domestic environments: The Smart Home Simulator", Mobile Information Systems 2017 (2017).

[3] D. Baldassini, et al.: "Customization of domestic environment and physical training supported by virtual reality and semantic technologies: A use-case" 2017 IEEE 3rd International Forum on Research and Technologies for Society and Industry (RTSI). IEEE, 2017.

Please contact:

Sonia Lorini, European Association of Virtual and Augmented Reality, Belgium
communication@eurovr-association.org
http://www.nweurope.eu/VR4REHAB
https://www.facebook.com/VR4RehabProject

In Memory of Peter Inzelt

Dr. Péter Inzelt, former director of the Institute for Computer Science and Control (SZTAKI) passed away at age 75, on the 23rd of June, 2019.

Péter worked for the predecessor of SZTAKI from 1968 - first as researcher at Department of Control of Continuous Processes, then as head of the scientific department from 1981. He became chief financial officer of SZTAKI in 1987, and director in 1993. Péter led the Institute successfully for 21 years, up to age 70. During his long directorate period the Institute strengthened its professional reputation both at home and in the international scientific community.

His leadership style was always strongly characterised by managerial innovation, fairness in financial affairs, an anti-corruption attitude and social sensitivity. He introduced a new motivation system, launched a voluntary pension fund for the SZTAKI employees and their family members, purchased holiday apartments for the Institute, thus ensuring the summer recreation of many families and supported sporting activities within and outside the Institute. He also provided financial assistance to families suffering from health problems. And this list is far from complete.

Péter Inzelt (right) and Alain Bensoussan (left) signing SZTAKI's ERCIM membership in 1994.

Péter signed SZTAKI's ERCIM membership in November 1994 during the ERCIM meetings in Barcelona and was SZTAKI's representative until his retirement. He played an active role at the Board of Directors where he was highly respected. He always considered ERCIM as one of the most important organisations, if not the most important, that SZTAKI joined.

Last, I would like to cite his own words taken from his book published for the 50-year anniversary of SZTAKI:

"As for me, it was a pleasure to act as director here for 21 years, up to the limit of the legally possible age and my endeavour was always to serve the Institute in the best way. I hope some of my colleagues share this view. I wish much success to my successor, to the Institute, all its managers and employees."

SZTAKI considers Péter Inzelt one of our own. His memory will live on.

László Monostori, Director, SZTAKI

Call for Proposals

Dagstuhl Seminars and Perspectives Workshops

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Schloss Dagstuhl – Leibniz-Zentrum für Informatik is accepting proposals for scientific seminars/workshops in all areas of computer science, in particular also in connection with other fields.

If accepted the event will be hosted in the seclusion of Dagstuhl's well known, own, dedicated facilities in Wadern on the western fringe of Germany. Moreover, the Dagstuhl office will assume most of the organisational/ administrative work, and the Dagstuhl scientific staff will support the organizers in preparing, running, and documenting the event. Thanks to subsidies the costs are very low for participants.

Dagstuhl events are typically proposed by a group of three to four outstanding researchers of different affiliations. This organizer team should represent a range of research communities and reflect Dagstuhl's international orientation. More information, in particular, details about event form and setup as well as the proposal form and the proposing process can be found on

https://www.dagstuhl.de/dsproposal

Schloss Dagstuhl – Leibniz-Zentrum für Informatik is funded by the German federal and state government. It pursues a mission of furthering world class research in computer science by facilitating communication and interaction between researchers.

Important Dates

- Proposal submission: October 15 to November 1, 2019
- Notification: February 2020
- Seminar dates: Between September 2020 and August 2021 (tentative).

HORIZON 2020 Project Management

A European project can be a richly rewarding tool for pushing your research or innovation activities to the state-of-the-art and beyond. Through ERCIM, our member institutes have participated in more than 90 projects funded by the European Commission in the ICT domain, by carrying out joint research activities while the ERCIM Office successfully manages the complexity of the project administration, finances and outreach. The ERCIM Office has recognized expertise in a full range of services, including identification of funding opportunities, recruitment of project partners, proposal writing and project negotiation, contractual and consortium management, communications and systems support, organization of events, from team meetings to large-scale workshops and conferences, support for the dissemination of results.

How does it work in practice?

Contact the ERCIM Office to present your project idea and a panel of experts will review your idea and provide recommendations. If the ERCIM Office expresses its interest to participate, it will assist the project consortium either as project coordinator or project partner.

Please contact:
Peter Kunz, ERCIM Office, peter.kunz@ercim.eu

Editorial Information

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
+33 4 9238 5010, contact@ercim.eu
Director: Philipp Hoschka, ISSN 0926-4981

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

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

ERCIM News online edition: *ercim-news.ercim.eu/*

Next issue:
January 2020, Special theme: Educational Technology

Subscription
Subscribe to ERCIM News by sending an email to en-subscriptions@ercim.eu

Cover illustrtation: *source: Shutterstock.*

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

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

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

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

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

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

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

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

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

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

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

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

SBA Research gGmbH
Floragasse 7, 1040 Wien, Austria
www.sba-research.org/

SIMULA
PO Box 134
1325 Lysaker, Norway
www.simula.no

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

TNO
PO Box 96829
2509 JE DEN HAAG
www.tno.nl

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

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

VTT Technical Research Centre of Finland Ltd
PO Box 1000
FIN-02044 VTT, Finland
www.vttresearch.com

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