Number 109 April 2017

ERCIM NEWS

www.ercim.eu

Special theme:

Autonomous Vehicles

Also in this issue:

Research and Society:

Virtual Research Environments

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
Tel: +33 4 9238 5010, E-mail: 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 http://ercim-news.ercim.eu/ or contact peter.kunz@ercim.eu

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

Next issue
July 2017, Special theme: Blockchain Engineering

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

Contents

JOINT ERCIM ACTIONS

RESEARCH AND SOCIETY

This section on Virtual Research Environments has been coordintated by Keith G. Jeffery and Pierre Guisset (ERCIM)

SPECIAL THEME

The special theme section “Autonomous Vehicles” has been coordinated Erwin Schoitsch (AIT) and Fawzi Nashashibi (Inria)

Introduction to the Special Theme

Invited Article

RESEARCH AND INNOVATION

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

CALLS, IN BRIEF

Event

Calls

In Brief

ERCIM Membership

After having successfully grown to become one of the most recognized ICT Societies in Europe, ERCIM has opened membership to multiple member institutes per country. By joining ERCIM, your research institution or university can directly participate in ERCIM's activities and contribute to the ERCIM members' common objectives playing a leading role in Information and Communication Technology in Europe:

- Building a Europe-wide, open network of centres of excellence in ICT and Applied Mathematics;
- Excelling in research and acting as a bridge for ICT applications;
- Being internationally recognised both as a major representative organisation in its field and as a portal giving access to all relevant ICT research groups in Europe;
- Liaising with other international organisations in its field;
- Promoting cooperation in research, technology transfer, innovation and training.

About ERCIM

ERCIM – the European Research Consortium for Informatics and Mathematics – aims to foster collaborative work within the European research community and to increase cooperation with European industry. Founded in 1989, ERCIM currently includes 15 leading research establishments from 14 European countries. ERCIM is able to undertake consultancy, development and educational projects on any subject related to its field of activity.

ERCIM members are centres of excellence across Europe. ERCIM is internationally recognized as a major representative organization in its field. ERCIM provides access to all major Information Communication Technology research groups in Europe and has established an extensive program in the fields of science, strategy, human capital and outreach. ERCIM publishes ERCIM News, a quarterly high quality magazine and delivers annually the Cor Baayen Award to outstanding young researchers in computer science or applied mathematics. ERCIM also hosts the European branch of the World Wide Web Consortium (W3C).

> **"Through a long history of successful research collaborations in projects and working groups and a highly-selective mobility programme, ERCIM has managed to become the premier network of ICT research institutions in Europe. ERCIM has a consistent presence in EU funded research programmes conducting and promoting high-end research with European and global impact. It has a strong position in advising at the research policy level and contributes significantly to the shaping of EC framework programmes. ERCIM provides a unique pool of research resources within Europe fostering both the career development of young researchers and the synergies among established groups. Membership is a privilege."**
>
> ***Dimitris Plexousakis, ICS-FORTH, ERCIM AISBL Board***

Benefits of Membership

As members of ERCIM AISBL, institutions benefit from:

- International recognition as a leading centre for ICT R&D, as member of the ERCIM European-wide network of centres of excellence;
- More influence on European and national government R&D strategy in ICT. ERCIM members team up to speak with a common voice and produce strategic reports to shape the European research agenda;
- Privileged access to standardisation bodies, such as the W3C which is hosted by ERCIM, and to other bodies with which ERCIM has also established strategic cooperation. These include ETSI, the European Mathematical Society and Informatics Europe;
- Invitations to join projects of strategic importance;
- Establishing personal contacts with executives of leading European research institutes during the bi-annual ERCIM meetings;
- Invitations to join committees and boards developing ICT strategy nationally and internationally;
- Excellent networking possibilities with more than 10,000 research colleagues across Europe. ERCIM's mobility activities, such as the fellowship programme, leverage scientific cooperation and excellence;
- Professional development of staff including international recognition;
- Publicity through the ERCIM website and ERCIM News, the widely read quarterly magazine.

How to Become a Member

- Prospective members must be outstanding research institutions (including universities) within their country;
- Applicants should address a request to the ERCIM Office. The application should inlcude:
  - Name and address of the institution;
  - Short description of the institution's activities;
  - Staff (full time equivalent) relevant to ERCIM's fields of activity;
  - Number of European projects in which the institution is currently involved;
  - Name of the representative and a deputy.
- Membership applications will be reviewed by an internal board and may include an on-site visit;
- The decision on admission of new members is made by the General Assembly of the Association, in accordance with the procedure defined in the Bylaws (http://kwz.me/U7), and notified in writing by the Secretary to the applicant;
- Admission becomes effective upon payment of the appropriate membership fee in each year of membership;
- Membership is renewable as long as the criteria for excellence in research and an active participation in the ERCIM community, cooperating for excellence, are met.

Please contact the ERCIM Office: contact@ercim.eu

Call for Nominations

2017 ERCIM Cor Baayen Young Researcher Award

The Cor Baayen Young Researcher Award is awarded each year to a promising young researcher in computer science and applied mathematics. The award was created in 1995 to honour the first ERCIM President. The award consists of a cheque for € 5000 together with an award certificate.

Eligibility

Nominees must have carried out their work in one of the 'ERCIM countries': Austria, Cyprus, Finland, France, Germany, Greece, Hungary, Italy, Luxembourg, Norway, Poland, Portugal, Sweden and The Netherlands. Nominees must have been awarded their PhD (or equivalent) after 30 April 2014. A person can only be nominated once for the Cor Baayen Award.

The Cor Baayen award is a young researcher prize: The selection panel will therefore consider the quality of the PhD thesis and all the achievements done up to the nomination date.

The nominees must have performed their research during at least one year in an institution located in one of the countries mentioned above.

Submitting a nomination

Nominations have to be made by a staff member of an ERCIM member institute. Self nominations are not accepted. Each institute can forward two nominations to the final evaluation. If there are more than two nominations coming from an ERCIM member institute, it is the responsibility of this institute to select the maximum two nominees. Nominations must be submitted online.

Selection

The selection of the Cor Baayen Young Researcher Award winner is the responsibility of the ERCIM Human Capital Task Group, who might consult expert opinion in reaching their decision.

Deadline for nominations: 15 May 2017

Details and application form:
https://www.ercim.eu/human-capital/cor-baayen-award

The Cor Baayen Young Researcher Award is named after the first president of ERCIM and the ERCIM 'president d'honneur'. Cor Baayen played an important role in its foundation. Cor Baayen was scientific director of the Centrum voor Wiskunde en Informatica (CWI) in the Netherlands, from 1980 to 1994.

ERCIM "Alain Bensoussan" Fellowship Programme

ERCIM offers fellowships for PhD holders from all over the world. Topics cover most disciplines in Computer Science, Information Technology, and Applied Mathematics. Fellowships are of 12 months duration, spent in one ERCIM member institute. Fellowships are proposed according to the needs of the member institutes and the available funding.

Application deadlines for the next round: 30 April and 30 September 2017

More information: http://fellowship.ercim.eu/

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 80 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 attractive 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 as described above, either as project coordinator or project partner.

Please contact:
Philippe Rohou, ERCIM Project Group Manager
philippe.rohou@ercim.eu

Virtual Research Environments: How Researchers Really Collaborate

by Keith G. Jeffery and Pierre Guisset (ERCIM)

How researchers collaborate is per se a research topic! Many scientific problems are related across different domains: for example, the global climate changes involve knowledge from eco-system, ocean, atmosphere and earth as well as from energy science and human-related activity modelling. Scientists face great challenges in handling collaboration among different disciplines and in modelling and discovering knowledge in massively available data from a wide diversity of domains. Today, data-driven approaches are considered as good alternatives to drive scientific research activities.

Virtual Research Environments (VRE) are the online software systems enabling collaboration between researchers from different scientific domains. Creating efficient VRE is focusing significant interest from the research community. The main scientific challenges are:

- Data context issues (metadata)
- Data heterogeneity issues
- Fast-changing data issues
- Data quality issues
- Privacy issues
- User experience issues
- Software issues.

ERCIM is taking a leading role in Europe for driving research activities in VRE, as demonstrated by the contributions featured in this issue of ERCIM News.

We start with Keith Jeffery describing the heterogeneity of Research Infrastructures and their research communities and introducing VRE4EIC (www.vre4eic.eu), an H2020 research and innovation action led by ERCIM and targeting to overcome this heterogeneity.

Then, four contributions are highlighting critical aspects of efficient VRE:

- Yi Yin and Anneke Zuiderwijk (TU Delft, The Netherlands) are presenting the fundamental requirements for a VRE in the Big Data Era.
- Cesare Concordia and Carlo Meghini (CNR, Italy) are documenting a reference (software) architecture for enhanced VRE (e-VRE).
- Phil Archer (ERCIM/W3C) is addressing the issue related to metadata, by reporting on the The Smart Descriptions & Smarter Vocabularies Workshop (SDSVoc) hosted by CWI in Amsterdam end of 2016
- Leonardo Candela, Donatella Castelli and Pasquale Pagano (CNR, Italy) are proposing innovative methods for making the development and the deployment of VRE as easy and effective as possible.

And finally, three contributions are focusing on VRE usage and benefits in integrating Research Infrastructures:

- The European Plate Observing System (EPOS, https://www.epos-ip.org/) aims at creating a pan-European infrastructure for solid Earth science to support a safe and sustainable society, by Daniele Bailo and Manuela Sbarra (INGV, Italy).
- Zhiming Zhao, Paul Martin (both University of Amsterdam, The Netherlands) and Keith Jeffery are discussing the aim to design and develop a suite of standard solutions to those common problems based on the reference model of environmental Research Infrastructures (ENVRI-RM, see http://www.envriplus.eu/) and the e-VRE architecture proposed by VRE4EIC.
- Jorge dos Santos Oliveira, José Borbinha and Ana Teresa Freitas (Universidade de Lisboa) are describing the design and development of a VRE for supporting studies using metagenomic data applied to the oil and gas domain.

Globally, 70.000 researchers are working in related scientific fields and thus are potential end users of new VRE. It is thus of the utmost importance to provide them with the most efficient tools to support them in their findings. This is the mission of the VRE research community.

Please contact:
Keith G. Jeffery, Pierre Guisset, ERCIM,
keith.jeffery@ercim.eu, pierre.guisset@ercim.eu

VRE4EIC: A Europe-Wide Virtual Research Environment to Empower Multidisciplinary Research Communities and Accelerate Innovation and Collaboration

by Keith G Jeffery

Research Infrastructures and their research communities are heterogeneous. This is a barrier to one community (re-)using the assets of another. VRE4EIC aims to overcome this heterogeneity.

Research has increasingly become specialised into communities such as oceanography, ecology, geology, materials science. However, many phenomena can only be understood by bringing together the research activities of several communities. Examples include the relationship between shellfish pollution, algal blooms and agricultural use of nitrates or the relationship between health problems, climate and social conditions. Recently, many communities have developed pan-European research infrastructures (RIs) bringing together several national research teams and assets such as datasets, software, publications, expert staff, sensors and equipment. One way to assist and encourage interdisciplinary research is to bring together the communities and assets of the RIs.

However, this collaboration comes with complications. Each community has developed its own standards for research methods, data formats, software to be used, etc. This makes it difficult for an ecologist, for instance, to utilise oceanographic data. The heterogeneity is especially evident in digital representations of data, software, people, organisations, workflows and equipment. However, many of these assets are represented digitally by metadata providing a succinct description of the asset. The metadata standard chosen varies from community to community. On the other hand, there is a limited set of basic things (entities or objects) that are involved in research (for example, data, people, samples) and so the various metadata standards have some commonality in the things they represent – although they do so in different ways.

Thus, the 'line of attack' to provide multidisciplinarity for researchers is to try to harmonise the metadata and thus gain access to – and (re-)utilisation of – the assets. There are two basic approaches: the software broker approach provides mapping and conversion between pairs of metadata standards. This results in n(n-1) convertor pairs. The alternative approach is to choose a canonical superset metadata standard and convert each metadata standard to/from that. This results in n convertor pairs. This metadata-driven brokering is now regarded as the best approach [1]. However, again we have two choices; the canonical superset may be realised physically – so providing an 'umbrella' consistent metadata resource or catalog over all the participating RIs or the superset metadata may just be a reference syntax (structure) and semantics (meaning) and each RI provides its pair of convertors. The latter approach leads to an architecture with peer RI to RI communication, requiring quite some software at each RI to interact with the other RIs and generate appropriate workflows. The former leads to a system over the RIs – linked to them via Application Programming Interfaces (APIs) – commonly named a Virtual Research Environment (VRE), which has the advantage of a 'helicopter view' over the participating RIs and so can generate workflows optimally. Either way, the core of a VRE is the superset catalog (whether conceptual or physical).

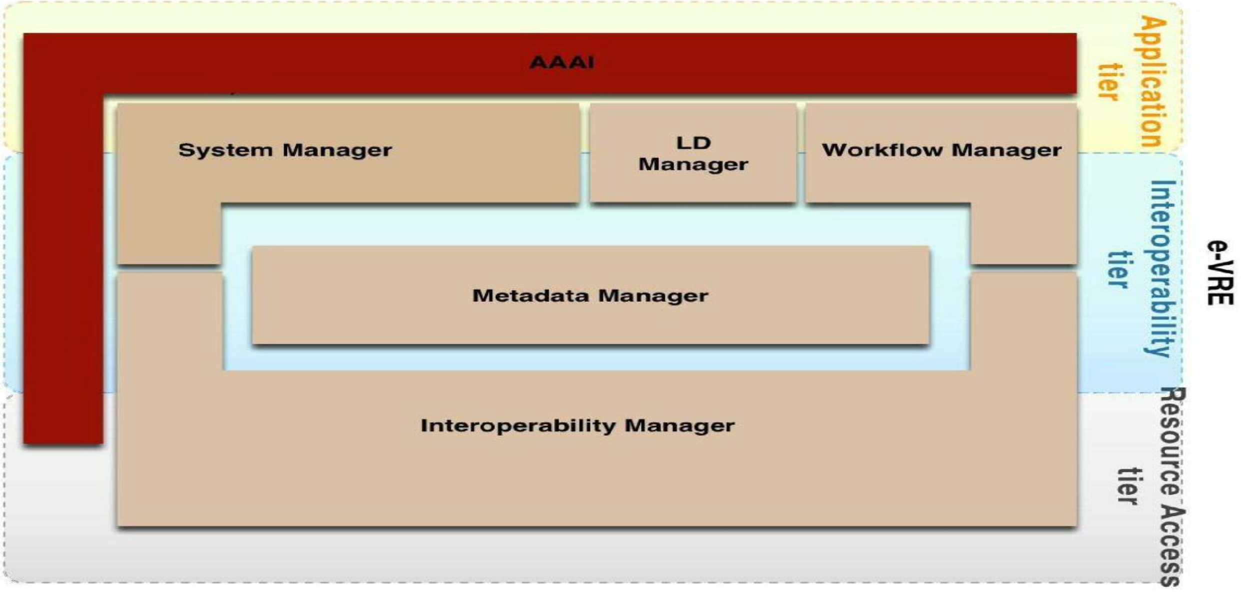

Moving from the Reference Architecture to the Technical Architecture

Figure 1: e-VRE architecture coneptual components.

VRE4EIC [L1] aims at providing a model for such VREs, which includes requirements, reference architecture and implementation on two use cases to demonstrate its feasibility and innovative impact. VRE4EIC has chosen CERIF (Common European Research Information Format: an EU recommendation to Member States) [L2] to denote the superset catalog.

In fact, a VRE provides more than access to the assets of RIs; it also provides researcher intercommunication through various means and software to generate workflows to harness the available analytics, visualisation and simulation capabilities of the RIs. Ideally the VRE workflow should be optimised to ensure co-location of data and software which means moving data to the software from the various RIs participating or – especially as datasets become larger – moving the software to the data. This has implications in terms of access rights, privacy and security and in finding an equitable method of 'payment' for use of the RI assets. The VRE may also use e-Is (e-Infrastructures) such as eternal curated storage or supercomputing services with the requirement to manage the deployment of (parts of) the workflow to these e-Is. The VRE should assist the researcher with research management; assisting in finding relevant research, assisting in research proposals, tracking research portfolio and cataloguing research outputs (such as scholarly publications, patents, datasets, software) since increasingly funding organisations utilise such information in planning future research programmes and in evaluating the quality of research proposals.

VRE4EIC has undertaken a considerable amount of requirements collection and analysis, and has characterised many RIs to understand their available interfaces. The architecture has been designed (Figure 1) and construction is underway. The prototype will be evaluated by the RIs that are in the project first, and then other RIs will be invited to evaluate the system.

In parallel, VRE4IC has been cooperating with other VRE projects, notably EVER-EST in Europe but also – via the VRE Interest Group of RDA (Research Data Alliance) – SGs (Science Gateways) in North America and VLs (Virtual Laboratories) in Australia. In parallel, the various metadata groups of RDA, coordinated by Metadata Interest Group (MIG), are working on a standard set of metadata elements – to be used to describe RI assets in catalogs -– which are not simple attributes with values but will have internal syntax and semantics [2].

Links:

[L1] vre4eic.eu
[L2] www.eurocris.org/cerif/main-features-cerif

References:

[1] Stefano Nativi, Keith G. Jeffery, Rebecca Koskela: "Brokering with Metadata", ERCIM News 100, http://kwz.me/W1

[2] Keith G Jeffery, Rebecca Koskela: "The Importance of Metadata", ERCIM News 100, http://kwz.me/W2

Please contact:

Keith G Jeffery
ERCIM Scientific Coordinator of VRE4EIC
keith.jeffery@ercim.eu

Scientists' Fundamental Requirements to Deal with their Research Data in the Big Data Era

by Yi Yin and Anneke Zuiderwijk (Delft University of Technology)

Collaboration among researchers from different disciplines is becoming an essential ingredient of scientific research. In order to solve increasingly complex scientific and social conundrums, research data needs to be shared among researchers from different disciplines. New technologies pave the way for unlimited potential for preserving, analysing and sharing research information. The methods used to leverage information technologies to deal with research data vary significantly among scientists, and likewise the requirements of individual scientists vary.

The need for data sharing

Science requires the collection and use of research data. The importance of data for science is equivalent to that of water for life. The proliferation of information communication technologies and other technological innovations has transformed how scientific research is conducted. The new trend in scientific research covers new research domains, funding sources and way of disseminating research results. More data is accumulated and scientists are more connected in the digital era, which is significantly changing all the phases of scientific research. Modern scientific research, which is increasingly data-intensive and complex, requires multidisciplinary collaboration as well as the support of large-scale research infrastructures and high-end experiment instruments, such as artificial intelligence, ecology science, health, biodiversity, culture and heritage research. In addition to obtaining funding from governments, research institutes or industries, crowdsourcing is also becoming a source of funding for scientific research: the widespread use of Kickstarter [L1] is a case in point. The way of communicating research results and output is also evolving. Open access and data sharing which allow everyone to access any research data, not limited to scientific publications and datasets, are increasingly favoured by society as a whole [1] [2]. Virtual research environments (VREs) can be used to support this more advanced type of research data sharing and to support collaboration among researchers.

The needs of researchers

In order to understand the fundamental needs of scientists to conduct various research activities in a VRE, ten interviews were conducted with scientists from different research domains. End-user requirements were collected for each of the activities in the lifecycle of scientific research. The requirements for developing a VRE fall into two categories: functional requirements, which describe what the system should do; and non-functional requirements, which include quality or performance attributes.

Building on the outcomes of the ENVRIplus project [L2] the elicited functional requirements were grouped into seven categories:

- *Data identification and citation*, which covers the identification of various types of research data and associated metadata and provides clear references for specific datasets in terms of citation. The use of persistent and unique identifiers for both data and metadata we found to be crucial for data identification and citation.
- *Data curation*, which includes all processes and activities to manage acquired datasets, for instance, detailed data management planning and workflows for data management.
- *Data Cataloguing*, which refers to the collection and cataloguing of information for various categories associated with research activities, for instance experiment equipment, data processing software, data products, publications, research individuals and organisations, research events and research objectives.
- *Data processing*, which covers the functionalities related to searching, processing and analysing data.
- *Data provenance*, which covers the functionalities to track the changes of datasets.
- *Collaboration, training and support*, which covers the functionalities related to research collaboration and training, for instance user interface configuration, establishment of research groups and the supervision of research progress.

Besides the functional requirements, the non-functional requirements specify quality attributes related to functional requirements for a VRE. According to the interviewed researchers, a quickly-accessible, reliable, easy-to-use, low-cost VRE is needed. Besides the performance-related requirements, VREs also need to consider ethical, legal and privacy and security perspectives according to the guidelines and principles defined by the 'European Charter for Access to Research Infrastructures' [L3]. The non-functional requirements are categorised into:

- *System performance* related requirements defined by FURPS+ and ISO 25010:2011, for instance performance efficiency, usability, reliability, maintainability, compatibility and portability of the information system.
- *Privacy, security, trust and legal* requirements, which specify that the whole development of the VRE should comply with all legislations, especially how the use of the VRE should be robust against cyber-attacks in terms of protected information privacy and security regulated by the new General Data Protection Regulation [L4]. Trust requirements specify the acceptable behaviour of the stakeholders in the VRE system, such as users, system developers and service providers.

From the survey and from the existing VRE-related projects, we elicited various requirements regarding the whole lifecycle of scientific research. However, scientific research is evolving all the time, and it is unrealistic for the designed VRE system to cover all the evolving needs of scientists. The designed VRE system therefore needs to be adaptive and flexible to connect to other existing research infrastructures to collectively serve scientists' needs.

Links:
[L1] www.kickstarter.com/
[L2] www.envriplus.eu/
[L3] ec.europa.eu/research/infrastructures/index_en.cfm?pg=access_ri
[L4] http://ec.europa.eu/justice/data-protection/reform/index_en.htm

References:

[1] Fecher, B., Friesike, S., & Hebing, M. (2015). What drives academic data sharing? Plos One, 10(2). doi:10.1371/journal.pone.0118053
[2] Zuiderwijk, A., & Janssen, M. (2014). Open data policies, their implementation and impact: A framework for comparison. Government Information Quarterly, 31(1), 17-29. doi:10.1016/j.giq.2013.04.003

Please contact:

Yi Yin, Anneke Zuiderwijk
Delft University of Technology, The Netherlands
Y.Yin@tudelft.nl, A.M.G.Zuiderwijk-vanEijk@tudelft.nl

A Reference Architecture for Enhanced Virtual Research Environments

by Cesare Concordia and Carlo Meghini (ISTI-CNR)

Virtual Research Environments (VREs) are rapidly becoming a popular technology for supporting scientists during their research work. In order to overcome the current heterogeneity, a reference architecture for VREs is being developed by the VRE4EIC Project.

The goal of a Virtual Research Environment (VRE) system is to decouple science from ICT complexity, by providing researchers with a facility that takes care of ICT, thus allowing them to focus on their work. In this sense, a VRE is a fundamental component of an e-Research Infrastructure (e-RI) as it makes the resources of the e-RIs easily accessible and reusable to the community of researchers that owns the e-RI. Here, by e-RI we mean 'facilities, resources and related services used by the scientific community to conduct top-level research in their respective fields' [L1] while 'resource' indicates any ICT entity that is of interest in an e-science community. Typically, a resource is owned by an e-RI that provides an identity for the resource and manages it, making it accessible and re-usable. Examples of resources include: datasets, workflows, algorithms, web services, computational or storage facilities, cloud endpoints etc. In general, a VRE is expected to:

- allow researchers to communicate with each other and to share and use the resources available in the community's e-RI
- allow researchers to advance the state of the art by building new resources as the result of processing existing resources with the available tools. Such processing may be

the application of an individual piece of software to a dataset, such as the extraction of certain knowledge from a single file; or, it may result from the execution of a complex workflow obtained by combining available services, including other workflows
- allow research managers to apply economy of scale models to access and manage resources that researchers or single organisations alone could not afford.

The most advanced e-RIs have developed their own VRE, showing awareness of the crucial role that a VRE can play for their researchers. Others are currently designing their VRE. However, the number of currently existing or designed VREs is very limited; more importantly, these VREs show a great heterogeneity in scope, features, underlying protocols and technologies, partially defeating the interoperability goal that lies at the very heart of a VRE. One of the major goals of the VRE4EIC project is to overcome this issue by proposing a reference architecture for an enhanced VRE (e-VRE). Based on a thorough analysis of the requirements of a VRE, and on the characterisation of an ample range of existing research infrastructures, the project has individuated three logical tiers in e-VRE:
- The Application tier, which provides functionalities to manage the system, to operate on it, and to expand it, by enabling administrators to plug new tools and services into the e-VRE.
- The Interoperability tier, which deals with interoperability aspects by providing functionalities for: i) enabling application components to discover, access and use e-VRE resources independently from their location, data model and interaction protocol; ii) publishing e-VRE functionalities via a Web Service API; and iii) enabling e-VRE applications to interact with each other.
- The Resource Access tier, which implements functionalities that enable e-VRE components to interact with eRIs' resources. It provides synchronous and asynchronous communication facilities.

In each tier, a set of basic functionalities has been grouped into six conceptual components:
- The e-VRE management is implemented in the system manager component. The system manager can be viewed as the component enabling users to use the core functionalities of the e-VRE: access, create and manage resource descriptions, query the e-VRE information space, configure the e-VRE, plug and deploy new tools in the e-VRE and more.
- The workflow manager enables users to create, execute and store business processes and scientific workflows.
- The linked data (LD) manager is the component that uses the LOD (Linked Open Data) paradigm, based on the resource description framework (RDF) data model, to publish the e-VRE information space – i.e., the metadata concerning the e-VRE and the e-RIs in a form suitable for end-user browsing in a semantic web (SM)-enabled ecosystem.
- The metadata manager (MM) is the component responsible for storing and managing resource catalogues, user profiles, provenance information, preservation metadata used by all the components using extended entity-relational conceptual and object-relational logical representation for efficiency.
- The interoperability manager provides functionalities to implement interactions with e-RIs' resources in a transparent way. It can be viewed as the interface of e-VRE towards e-RIs. It implements services and algorithms to enable e-VRE to: communicate synchronously or asynchronously with e-RIs' resources, query the e-RIs' catalogues and storages, map the data models.
- The authentication, authorisation, accounting infrastructure (AAAI) component is responsible for managing the security of the e-VRE system. It provides user authentication for the VRE and connected e-RIs, authorisation and accounting services, and data encryption layers for components that are accessible over potentially insecure networks.

Each conceptual component is further structured into one or more actual software components and possibly sub-components. The complete list of these components can be found in [1], which also provides the interfaces implemented by each component, and the usage relationships between such interfaces. Each interface is in turn articulated into a set of methods, whose signature is also provided.
For instance, the metadata manager component consists of four sub-components, each devoted to a specific metadata type: the user, resource, preservation and provenance catalogue. The project is now entering into its development phase, during which it will provide implementation for some components, retrofitting them to existing VREs, in order to enhance them.

Link:
[L1] ec.europa.eu/research/infrastructures/index_en.cfm?pg=what

References:
[1] Carlo Meghini (ed.): "Architecture Design", Deliverable D3.1 of the VRE3EIC Project. Available on demand by the project coordinator (see http://www.vre4eic.eu/).

Please contact:
Carlo Meghini, ISTI-CNR, Italy
Carlo.Meghini@isti.cnr.it

Smart Descriptions and Smarter Vocabularies

by Phil Archer (W3C)

Sharing data between researchers, whether openly or not, requires effort, particularly concerning its metadata. What is the minimum metadata needed to aid discovery? Once data has been discovered, what metadata is needed in order to be able to evaluate its usefulness? And, since it's not realistic to expect everyone to use the same metadata standard to describe data, how can different systems interoperate with the metadata that is commonly provided? These topics and more were discussed in Amsterdam in late 2016.

The Smart Descriptions & Smarter Vocabularies Workshop (SDSVoc) was organized by ERCIM/W3C under the EU-funded VRE4EIC project and hosted by CWI in Amsterdam. Of 106 registrations, it is estimated that 85-90 people attended. The event comprised a series of sessions in which thematically related presentations were followed by Q&A with the audience, which notably included representatives from both the scientific research and government open data communities.

Participants of the Smart Descriptions & Smarter Vocabularies Workshop.

The workshop began with a series of presentations of different approaches to dataset description, including the CERIF [L1] standard used by VRE4EIC, followed by a closely related set of experiences of using the W3C's Data Catalog Vocabulary, DCAT [L2]. It was very clear from these talks that DCAT needs to be extended to cover gaps that practitioners have found different ways to fill. High on the list is versioning and the relationships between datasets, but other factors such as descriptions of APIs and a link to a representative sample of the data are also missing.

High level descriptions of any dataset are likely to be very similar (title, licence, creator etc.) but to be useful to another user, the metadata will need to include domain-specific information. A high profile case is data related to specific locations in time and space, and one can expect this to be part of the general descriptive regime. On the other hand, highly specialized data, such as details of experiments conducted at CERN, will always need esoteric descriptions.

An important distinction needs to be made between data discovery – for which the very general approach taken by schema.org is appropriate – and dataset evaluation. In the latter case, an individual needs to know details of things like provenance and structure before they can evaluate its suitability for a given task.

A descriptive vocabulary is only a beginning. In order to achieve real world interoperability, the way a vocabulary is used must also be specified. Application Profiles define cardinality constraints, enumerated lists of allowed values of given properties and so on, and it is the use of these that allows data to be validated and shared with confidence. Several speakers talked about their validation tools and it's clear that a variety of techniques are used. Validation techniques as such were out of scope for the workshop, although there were many references to the emerging SHACL [L3] standard. Definitely in scope however was how clients and servers might exchange data according to a declared profile. That is, apply the concept of content negotiation not just to content type (CSV, JSON, RDF or whatever) but also the profile used. The demand for this kind of functionality has been made for many years and proposals were made to meet that demand in future standardization work.

The workshop concluded that a new W3C Working Group should be formed to:

- revise and extend DCAT;
- provide guidance and exemplars of its use;
- standardize, or support the standardization elsewhere, of content negotiation by profile.

A full report on the event is published by W3C [L4] along with the agenda that links to all papers and slides, and a complete list of attendees. The W3C membership is being consulted on the formation of the new working group, expected to begin its work in May 2017.

Links:

[L1] www.eurocris.org/cerif/main-features-cerif
[L2] www.w3.org/TR/vocab-dcat/
[L3] www.w3.org/TR/shacl/
[L4] www.w3.org/2016/11/sdsvoc/report

Please contact:
Phil Archer, W3C
phil@philarcher.org

Making the Development and Deployment of Virtual Research Environments Easy and Effective

by Leonardo Candela, Donatella Castelli and Pasquale Pagano (ISTI-CNR)

Virtual research environments are emerging as an invaluable tool for scientists, enabling professionals in different fields to collaboratively and seamlessly access and use resources (computing, datasets, services) spread across several providers. This solution is particularly relevant in long-tail science contexts, i.e., when researchers and practitioner communities lack dedicated resources to perform their research. Implementing such a solution requires an approach that is open, flexible, and can easily evolve.

Virtual Research Environments (VREs) are web-based, community-oriented, collaborative, user-friendly, open-science-compliant working environments for scientists and practitioners working together on a research task [1]. They share commonalities with Science Gateways (SG) and Virtual Laboratories (VL). The overall goal is to provide scientists and researchers with integrated and user friendly access to data, computing and other services that are usually spread across multiple diverse data and computing infrastructures. Furthermore, they are designed to enact and promote collaboration among their members.

To develop and operate this type of working environment, it is necessary to: (a) develop a set of interoperability solutions that can interface the VRE services with the subset of existing resources that are relevant to a particular research project and offered by 'third-party' providers, (b) develop a set of basic user-friendly services promoting collaborative and open-science-friendly interaction among the VRE members, (c) consider short and long term provisioning of both the working environment and the products resulting from it.

Owing to these characteristics, the development and operation of such environments is incompatible with 'from scratch' and 'isolated' approaches. Developing and maintaining a 'mediator' to interface with a given e-infrastructure requires significant investment in time and effort that can be better afforded by applying economies of scale and scope (namely, by using it in the context of many VREs).

To meet this need, we have designed and developed a technology (actually, a software system) named gCube [2] capable of creating and operating an e-infrastructure offering VREs with the 'as-a-service' paradigm.

gCube has been progressively endowed with: (i) a rich array of 'mediators' for interfacing with existing 'systems' and their enabling technologies including distributed computing infrastructures (e.g., EGI) and data providers (e.g., by relying on standards like OAI-PMH, SDMX, OGC W*S) as well as for making it possible for third-party service providers to easily exploit gCube facilities (e.g., OAuth, OGC W*S, REST APIs); (ii) a set of basic services including a shared workspace where the objects used and resulting from VRE activity (beyond simple files) can be stored, organised and accessed as if they were in a 'standard' file-manager; a social networking area where members of each VRE can have discussions, share news and other material of interest, rate each item of a discussion, classify the discussion items by hashtags, refer to people or groups thus to call for actions from them, etc.; a user management area where authorised people are allowed to manage VRE membership, to create groups, assign members to groups, assign roles to members, invite new members, etc.; an open, customisable and extensible set of facilities made available for the needs of the specific community. These include a project management and issue-tracking system with a wiki, a rich and extensible data analytics platform, a flexible 'products' catalogue where any (research) artefact produced in the VRE that is worth being published can be easily made available by equipping it with rich metadata including licence and provenance, a rich array of domain data management facilities. VREs are created by using a wizard-based approach where a VRE designer is simply requested to select (among the existing ones) the facilities and resources he/she is willing to have in the VRE, and then upon approval the VRE is automatically provisioned and made available by a web-based portal.

This technology is currently enacting the D4Science e-Infrastructure] and exploited to create and operate more than 70 diverse VREs [L1]. Overall, these VREs are serving more than 3,100 (returning) scientists in 44 countries across a rich array of diverse communities usually associated with international initiatives and projects, e.g., i-Marine (fisheries and marine biodiversity scientists), BlueBRIDGE (fisheries and aquaculture scientists, educators & SMEs), SoBigData.eu (social mining scientists), ENVRI+ (environmental scientists), AGINFRA+ (agriculture scientists), PARTHENOS (cultural heritage practitioners), EGIP (geothermal scientists), OpenAIRE-Connect (multidisciplinary community dealing with scholarly communication and open science, EDISON (data science educators).

Link:
[L1] services.d4science.org/explore

References:

[1] L. Candela, D. Castelli, P. Pagano: "Virtual Research Environments: An Overview and a Research Agenda" Data Science Journal. 12, pp. GRDI75–GRDI81, 2013. DOI: http://doi.org/10.2481/dsj.GRDI-013
[2] M. Assante, et al.: "Virtual research environments as-a-service by gCube", PeerJ Preprints 4:e2511v1 2016, https://doi.org/10.7287/peerj.preprints.2511v1

Please contact:

Leonardo Candela
ISTI-CNR, Italy
leonardo.candela@isti.cnr.it

EPOS – European Plate Observing System: Applying the VRE4EIC Virtual Research Environment Model in the Solid Earth Science Domain

by Daniele Bailo and Manuela Sbarra

The European Plate Observing System (EPOS) aims at creating a pan-European infrastructure for solid Earth science to support a safe and sustainable society. In accordance with this scientific vision, the mission of EPOS is to integrate the diverse and advanced European Research Infrastructures for solid Earth science relying on new e-science opportunities to monitor and unravel the dynamic and complex Earth system.

The EPOS architecture [L1] is composed of three connected technical and organisational layers (Figure 1):

- National research infrastructures (NRI)
- Thematic core services (TCS)
- Integrated core services (ICS).

Multidisciplinary data and data products are organised and governed by the thematic core services (TCS) and are driven by various scientific communities encompassing a wide spectrum of Earth science disciplines (e.g., EIDA/ORFEUS in seismology [L3]). Such communities rely in turn on data provided at a national level by national research infrastructures (NRIs).

TCS data, data products, services and software will be integrated into integrated core services (ICS), a platform that will ensure their interoperability and access to these services by the scientific community and other users.

From the above organisational architecture, it is quite clear that EPOS integrated core services are an example of a virtual research environment (VRE). EPOS ICS is indeed currently one of the use cases that will take advantage of architecture paradigm developed by the VRE4EIC project [L2] at the implementation stage.

An architectural model was developed in VRE4EIC, following a rigorous method that started from the use cases, elicited the requirements, listed the system functionalities and finally developed an abstract architecture on the basis of the functionalities. The architecture is based on a modular paradigm, and the functionalities are implemented by a number of core components. In order to validate the VRE4EIC components, EPOS has been selected as one of the use cases to demonstrate the robustness and scalability of the VRE4EIC model.

In the framework of EPOS ICS implementation, one of the main components is the metadata manager. As described elsewhere [1], metadata is the main concept over which the whole EPOS VRE – in agreement with the VRE4EIC model – builds upon. The metadata manager, implemented as a metadata catalog, contains all the information that the system needs to manage in order to satisfy a user request. It contains the descriptions in a digital representation of the EPOS e-Infrastructure, including data, software, users and resources such as computers, detectors or laboratory equipment. As recommended by VRE4EIC, it uses CERIF [L4] as a tool to harmonise information on research. CERIF describes datasets but also software, services, users and resources such as computers, datastores, laboratory equipment and instruments and has already been used in the context of environmental sciences [2].

Another key component of EPOS implementation is the interoperability layer (Figure 1). This layer contains all the sub-components (mappers, convertors etc.) devoted to the communication with the national research infrastructure. One of its main functions is metadata exchange: metadata from data providers should be harvested by the ICS main system (the VRE). It requires the construction of mappers, to map the community specific metadata to CERIF, and convertors. When the metadata cannot be harvested, as in the case of provision of service based systems (e.g., web mapping services that provide access to geological maps), the web services need to be mapped in order to enable the VRE to broker the access to data. This layer also includes all the semantic mapping, which is currently one of the greatest challenges to tackle. Examples of mapping from most known metadata ISO standards have been studied and implemented in the EPOS VRE [3] and in other contexts [2]. Mappings from one metadata standard to another are currently under development in VRE4EIC in the framework of WP4 with a dedicated tool, namely the '3M' (mapping memory manager) [L6] developed by ICS-FORTH.

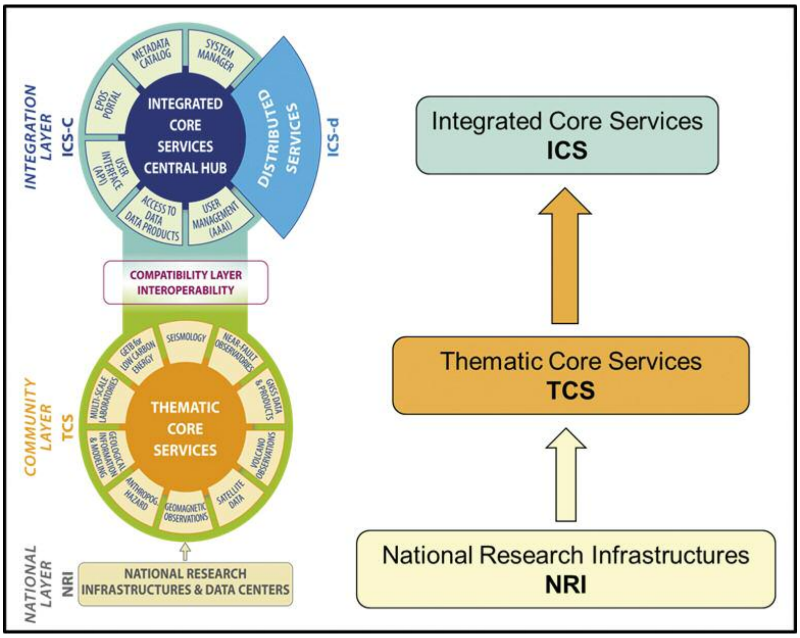

Figure 1: Key elements of the EPOS Functional Architecture

Other components in the EPOS ICS system include: (a) an AAAI (Authentication, Authorisation, Accounting Infrastructure) module [L5], dedicated to the secure authentication and authorisation of users also with different authentication systems (SAML, OAuth, OpenID, X.509 and related products such as EduGAIN, Shibboleth, Kerberos and others); (b) a system manager and a workflow engine, that manage the user request, establish the list of actions to be executed, execute the workflow and assemble the output.

In this context, VRE4EIC is providing the above architectural framework, together with state-of-the-art methods and best practices, to enable any research infrastructure to get to the level of a virtual research environment. Virtual research environments enhance the functionalities of the research infrastructures, that usually provide access to domain specific data, by widening the level of integrated information available to the user. In the case of EPOS, information about specific data, together with usage computational services and tutorial, training and other types of documentation that can enable any citizen scientist to have knowledge about solid Earth processes, will be available.

The innovation impact of the VRE4EIC architecture – based on metadata brokered paradigms – is enormous and not yet fully exploited, especially in the environmental data domain, where the integration of different types of data (e.g., seismic, geological, geodetical) together with computational tool and remote collaborative tools, can help to monitor and unravel the dynamics and the complexity of the Earth system.

Importantly, the use and re-use of data and the harmonised access to services that the creation of VREs can enable, will foster the creation of new scientific data products, which can be made accessible through the same EPOS VRE platform. The potential impact of this new scientific delivery system for environmental sciences and for the society is still unexplored.

Links:

[L1] https://epos-ip.org/
[L2] http://www.vre4eic.eu
[L3] www.orfeus-eu.org/
[L4] www.eurocris.org/cerif/main-features-cerif
[L5] epos-ip.org/glossary/aaai
[L6] 139.91.183.3/3M/

References:

[1] K. Jeffery, D. Bailo: “EPOS: Using Metadata in Geoscience. Metadata and Semantics Research, 170–184, 204, http://link.springer.com/chapter/10.1007/978-3-319-13674-5_17

[2] E. Boldrini, et al.: “Integrating CERIF Entities in a Multidisciplinary e-infrastructure for Environmental Research Data”, Procedia Computer Science 2014 vol: 33 pp: 183-190.

[3] D. Bailo, et al.: “Mapping solid earth Data and Research Infrastructures to CERIF”, 0, 9–11, 2016, http://dspacecris.eurocris.org/handle/11366/537

Please contact:

Daniele Bailo
EPOS Management Office ICT – INGV, Italy
+39 06-51860728, daniele.bailo@ingv.it

VRE in the Data for Science Approach to Common Challenges in ENVRIPLUS

by Zhiming Zhao, Paul Martin (University of Amsterdam) and Keith G. Jeffery (ERCIM)

Environmental Research Infrastructures (RIs) face common challenges regarding data management and how best to support the activities of scientists at all stages of the research data and experimentation lifecycle. The ENVRIPLUS ‘Data for Science’ theme aims to design and develop a suite of standard solutions to those common problems based on the reference model of research infrastructures (ENVRI-RM) and the e-VRE architecture proposed by VRE4EIC.

Environmental science research is increasingly dependent on the collection and analysis of large volumes of data gathered via wide-scale deployments of sensors and other observation sources. To study the development of earthquakes or volcanoes for example, one needs continuous observation of the surrounding geographic regions and their underlying strata in order to obtain the data necessary to model various seismological processes and their interactions. Depending on the problem scale and geographical focus, these observations can only be provided by sources distributed across different countries, institutions and data centres. Moreover, such research activities also often require advanced computing and storage infrastructure in order to analyse, process and model the data, and to perform simulations. Advanced research support environments (i.e., specialised infrastructure to support research) are clearly needed to better enable researchers to access data, software tools and services from different sources, and to integrate them into cohesive experimental investigations with well-defined, replicable workflows for processing data and recording the provenance of results for peer review.

A recent publication [1] identified several kinds of support environment that must be made to work together to support data-centric research:

- computing, storage and network infrastructures, e.g., provided via EGI [L1], EUDAT [L2] and GEANT [L3], also called e-Infrastructures (e-Is);
- services for accessing, searching and processing research data within different scientific domains, called Research Infrastructures (RIs), e.g., ICOS [L4], EPOS [L5] and EURO-ARGO [L6] for the atmospheric, earth and marine sciences respectively; and
- environments for providing user-centred support for discovering and selecting data and software services from different sources, and composing and executing application workflows based on them, called Virtual Research Environments (VREs) [L7] or Science Gateways (SGs) [2]. These different types of supporting environments often overlap with each other, as shown in Figure 1.

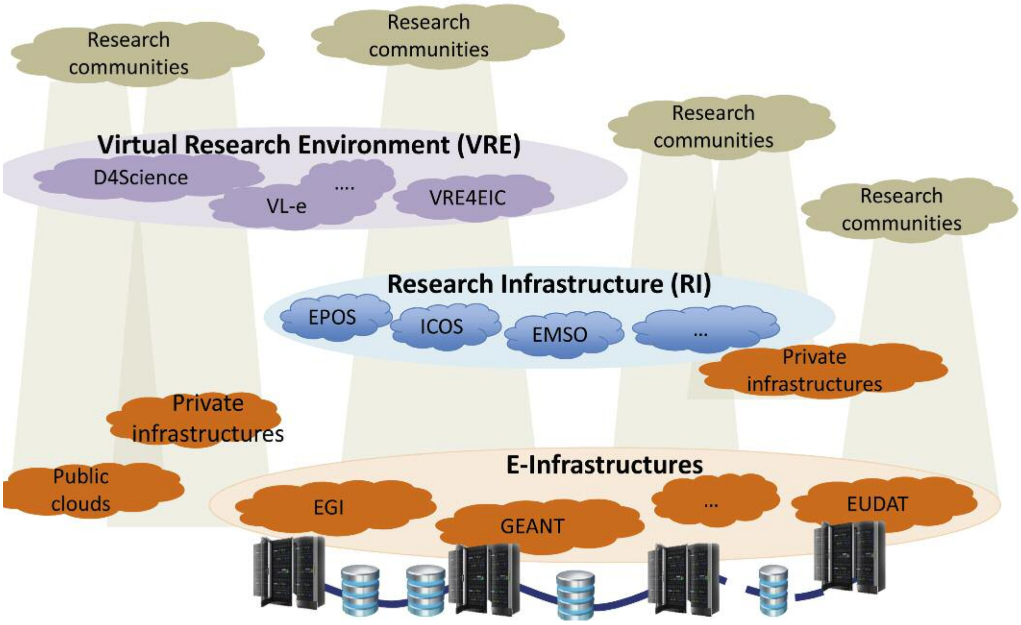

Figure 1: Different research support environments and the role they play in ICT activities initiated by user communities.

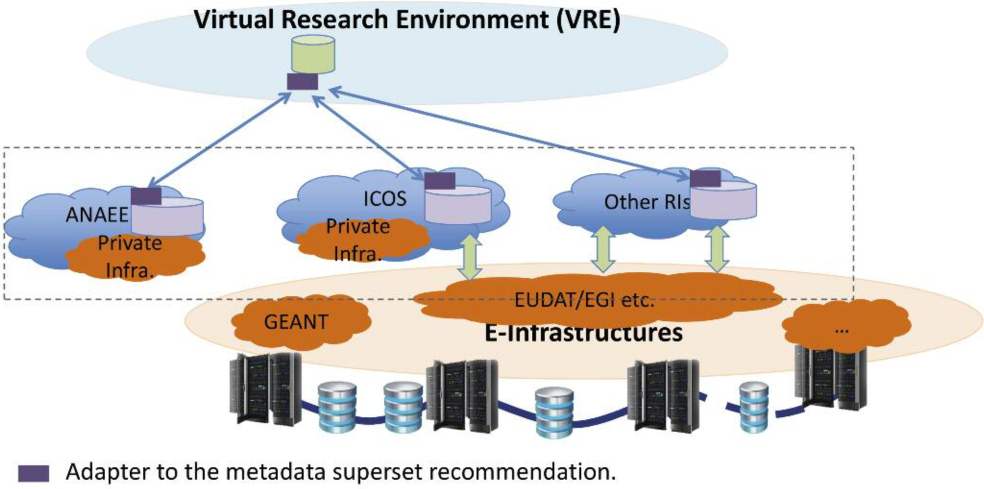

Figure 2: Metadata superset recommendation in ENVRIPLUS to enable future interface to overarching enhanced virtual research environments (e-VRE).

Within the EU Horizon 2020 project ENVRIPLUS [L8], the ‘Data for Science’ theme investigates and develops interoperable solutions to common problems that environmental RIs face for managing data and supporting the activities of scientists throughout the research data and experimentation lifecycle, and encompasses work on the development of common data services within a common semantic framework. Problems being addressed include:

- how to identify and cite data from different sites or infrastructures;
- how to control the quality of nearly real-time data from sensors and annotate them;
- how to catalogue the data and to allow users to search and access data from different sites or infrastructures;
- how to support scientists to perform experiments using data, software tools and resources from different remote infrastructures;
- how to effectively manage the infrastructure resources in the scientific experiments and allow scientists to achieve their goals more quickly; and
- how to effectively record the events and results generated during experiments so that scientists can reproduce them independently. Sharing solutions to those common problems will not only reduce development costs but also promote interoperability between different infrastructures.

Besides being important pillars for user communities in their respective domains, environmental RIs are also intended to support interdisciplinary research as well as contribute directly to cross-domain initiatives such as Copernicus [L9] (contributing to GEOSS [L10]). This requires standard policies, models and e-infrastructure to improve technology reuse and ensure coordination, harmonisation, integration and interoperability of data, applications and other services. The Data for Science theme follows a ‘Reference Model guided’ approach. It builds upon abstract concepts derived from the analysis of common operations of RIs and subsequently defines an architectural reference model for environmental RIs in general. Early results from specific RIs in construction have been carefully reviewed in order to identify good technology candidates for the realisation of the various common services needed, and a number of interactions have been carried out at various levels with computational e-infrastructures (such as EGI), data infrastructures (such as

EUDAT), and other initiatives (such as VRE4EIC [L11]) that work on related issues. The e-VRE reference architecture in the VRE4EIC project, for example, is being used to guide the development of interfaces to access data and software resources from ENVRIPLUS RIs. Figure 2 shows the basic idea.

ENVRIPLUS, a four year project, is approaching the end of its second year. Version 2 of ENVRI RM is available; it has been used to guide the design of new identification/citation, processing, optimisation, curation and cataloguing services. CERIF and CKAN have been recommended for prototyping a cross-RI catalogue service. The Open Information Linking for Environmental RIs (OIL-E) framework developed in ENVRIPLUS has also been aligned with CERIF in collaboration with the metadata team in VRE4EIC. Furthermore, a recommendation for how to use metadata catalogues as a basis for constructing federated services at VRE-level for interacting with individual RIs and underlying e-infrastructure has been produced and introduced to the ENVRIPLUS RI community (see Figure 2).

Links:
[L1] www.egi.eu
[L2] www.eudat.eu/
[L3] www.geant.org/
[L4] www.icos-infrastructure.eu/
[L5] www.eposeu.org/
[L6] www.euro-argo.eu/
[L7] www.jisc.ac.uk/rd/projects/virtual-research-environments
[L8] www.envriplus.eu
[L9] www.copernicus.eu/
[L10] www.geoportal.org/
[L11] www.vre4eic.eu

References:

[1] Zhiming Zhao, et al.: "Time critical requirements and technical considerations for advanced support environments for data-intensive research", Proc. IT4RIS workshop Porto 29 November-2 December 2016.
[2] Mark A. Miller, Wayne Pfeiffer, Terri Schwartz: "The CIPRES science gateway: enabling high-impact science for phylogenetics researchers with limited resources", in Proc. of XSEDE '12, ACM, New York, NY, USA, 2012.

Please contact:

Zhiming Zhao
University of Amsterdam, The Netherlands
z.zhao@uva.nl

A Virtual Research Environment for the Oil and Gas Domain

by Jorge dos Santos Oliveira, José Borbinha and Ana Teresa Freitas (INESC-ID/IST, Universidade de Lisboa)

Microbial enhanced oil recovery gained a new meaning with the development of metagenomics. This genomic method involves identifying the collective genome of a microbial community, including those that cannot be cultivated in controlled conditions. However, this method generates petabytes of genomic data, presenting specific computational challenges, demanding the development of new collaborative research platforms. The main goal of this project is to develop a virtual research environment (VRE) that can support studies using metagenomic data with application to the oil and gas field.

In 2010, UK's Joint Information Systems Committee (JISC) supported a landscape study on VREs. The research area with the largest number of identified examples of VREs was arts and humanities, followed by engineering and physical sciences and computer science, with biotechnology and biological sciences in fourth place [1]. The prominence of biotechnology and biological sciences VREs is interesting, given that there are usually ethical concerns with sharing data in these domains.

The most popular examples of VREs with an impact in biotechnology and biological sciences, are: (i) myExperiment [L1] – the main goal is to share and discuss workflows by providing networking tools to researchers; (ii) BioVel [L2] – a virtual e-laboratory which provides a graphical environment where researchers can design and construct analysis pipelines, before they are deployed through the BioVeL Portal and myExperiment; (iii) DECIDE [L3] – a VRE that helps the medical community with patient examination and diagnosis; and (iv) Genomics Virtual Lab [L4] that promises to take the IT out of bioinformatics by providing a cloud-based suite of genomics analysis tools.

Several factors can be seen as a barrier for the adoption of VREs in biotechnology and biological sciences, for instance, the complexity and consequent learning curve of such systems, often requiring IT knowledge; the 'cultural' habit of the (error-prone) physical notepad; and the broader approach to most of the research topics.

A ready-to-use application that has organisational and time-saving advantages over a physical notepad is required in order to increase the adoption of VREs in these domains. In this project, this particular question was addressed by acquiring feedback from laboratory teams during the development stages and by offering domain-specific information that is difficult to find just by browsing the web.

The new oil and gas VRE

This new oil and gas platform [L1] is a responsive web-based working environment, providing: (ii) workflow sharing, both bioinformatics and bench; (ii) access to algorithms and tools;

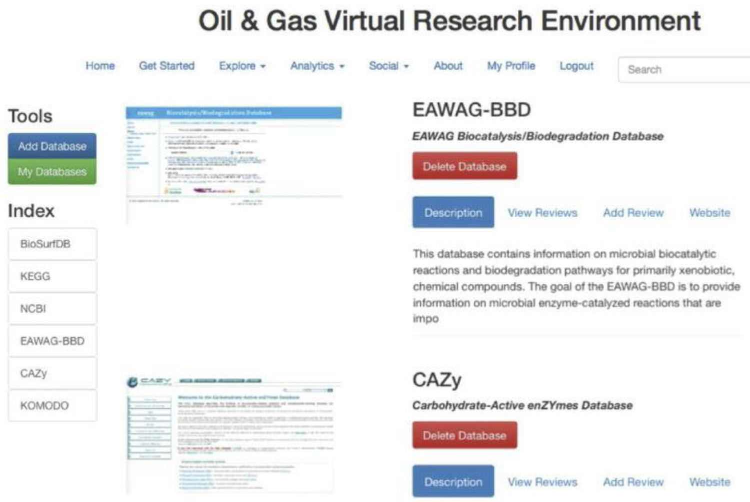

This database contains information on microbial biocatalytic reactions and biodegradation pathways for primarily xenobiotic, chemical compounds. The goal of the EAWAG-BBD is to provide information on microbial enzyme-catalyzed reactions that are impo

Figure 1: Database page, an example of the gallery-style sharing platform, where users can view, comment and review tools, or even submit their own.

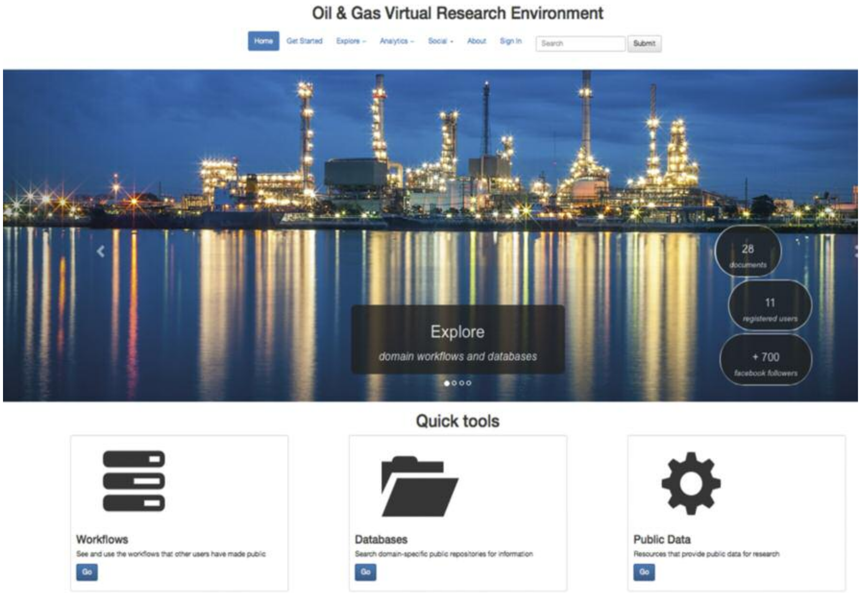

See and use the workflows that other users have made public

Search domain-specific public repositories for information

Resources that provide public data for research

Figure 2: Home page of the oil and gas VRE.

(iii) access to databases and datasets related with biodegradation and biosurfactants; (iv) access to processing resources; (v) access to storage resources; and (vi) access to social tools.

In particular, the sharing of tailor-made computational and bench workflows is very useful in metagenomics research due to the novelty, heterogeneity and size of the data that is generated.

Another main problem faced by VREs is their sustainability. Each VRE tends to create all services from scratch, using resources that require high maintenance and large budgets. Consequently, activity within most VREs declines with time and they become dead resources.

In order to overcome this important issue, we adopted the strategy of relying on existing technology and resources as much as possible, using e-infrastructure standard APIs and formats, and following an API-first philosophy. For example, social share and discussion resources are Facebook, myExperiment and BitBucket; user authentication is based on the Google Sign In platform; and Google Analytics tool was integrated into the platform in order to provide real-time and long-term statistics of the VRE usage. All this is build in the flexible, cloud-friendly and JavaScript-centric MEAN Stack. This VRE also takes advantage of other existing specialised services, for example for the BioSurfBD database [2].

The project team

This project started in 2015 as part of a larger initiative aiming to develop a multidisciplinary approach to identify and characterise patterns of taxonomic and functional diversity of microbial communities present in reservoirs [3].

The project brought together more than 20 researchers, from three research institutes: (i) the Information and Decision Support Systems Lab (IDSS) research group from INESC-ID, Portugal, that has the knowledge to develop the computational biology algorithms, databases and the VRE; (ii) the Bioinformatics Laboratory of LNCC, Petropolis, Brazil, that is responsible for providing workflows and metagenomics sequencing data; and (iii) the Department of Biochemistry, Universidade Federal do Rio Grande do Norte, Brazil, that is responsible for the laboratory experiments and feeding the VRE with data.

Due to its flexibility, this VRE can be used in other domains. For example, a similar solution for the area of precision medicine has already been proposed.

Links:

[L1] www.myexperiment.org
[L2] www.biovel.eu
[L3] applications.eu-decide.eu
[L4] nectar.org.au/
[L5] aleph.inesc-id.pt/vre/, www.facebook.com/ogvre/

References:

[1] A. Carusi and T. Reimer: “Virtual research environment collaborative landscape study,” UK’s Joint Information Systems Committee, Tech. Rep., 2010.

[2] J.S. Oliveira et al.:“BioSurfDB: knowledge and algorithms to support biosurfactants and biodegradation studies,” Database 1-8, 2015.

[3] J.S. Oliveira et al.: “Biogeographical distribution analysis of hydrocarbon degrading and biosurfactant producing genes suggests that near-equatorial biomes have higher abundance of genes with potential for bioremediation,” submitted to BMC Microbiology, 2017.

Please contact:

Jorge dos Santos Oliveira, University of Lisbon, Portugal
+351.213100300, jorge.oliveira@tecnico.ulisboa.pt

Introduction to the Special Theme

Autonomous Vehicles -

by Erwin Schoitsch (AIT)

Highly automated and autonomous systems are currently a key issue in many application domains: automotive, transport in general (railways, metro-lines, aircraft, space, ships), industrial automation, health care, cooperating mobile robots and related machines (e.g., fork lifts, off-road construction engines, smart farming, mining, all kind of drones and robots for surveillance, rescue, emergency services, maintenance). Highly automated and autonomous systems play an important role in the ‘digital transformation’ - the strategic and disruptive evolution towards a ‘digital society’, which is the key focus of European Research in Horizon 2020. This is depicted in Figure 1, which is taken from an official presentation of DG CONNECT at various events.

At a European level, research initiatives in Horizon 2020 include smart mobility and related applications for highly automated vehicles, such as those used in smart farming, marine, construction, logistics, manufacturing, smart cities and even health care (elderly care homes, hospital environments). This comprises general ICT research as well as domain-related programmes.
Some relevant programs are:

- ECSEL (Electronic Components and Systems for European Leadership, a joint undertaking between the EC, national member states and private partners from the industrial associations ARTEMIS, AENEAS and EPoSS),
- SPARC (robotics research, a PPP between EC and euRobotics) and the
- ‘Large Scale Pilot’ Programmes (e.g., in Smart Farming and Mobility), which are closely connected to the ‘Internet of Things’ Innovation Initiative (AIoTI and the EC).

At a national level in many European countries, substantial research is being conducted on autonomous vehicles and in particular (highly) automated driving. These efforts are not restricted to large countries like Germany and France - for example, the Austrian Federal Ministry for Transport, Innovation, and Technology (BMVIT) has launched a call to set up and run a public test region for automated vehicles, the ‘Austrian Light-vehicle Proving Ground’ (ALP.Lab) starting in 2017. For another example see ‘RECAR: Hungarian REsearch Center for Autonomous Road vehicles is on the way’.
There are many challenges to consider:

- Safety and security, privacy, dependability in general (see articles under ‘Generic Challenges’)
- Sensors and actuators
- Software development, life cycle issues
- System integration
- Connected vehicles, V2X connectivity
- Cooperative driving and transport systems, systems-of-systems aspects
- New mobility (multi-modality enabled by highly automated/autonomous vehicles)

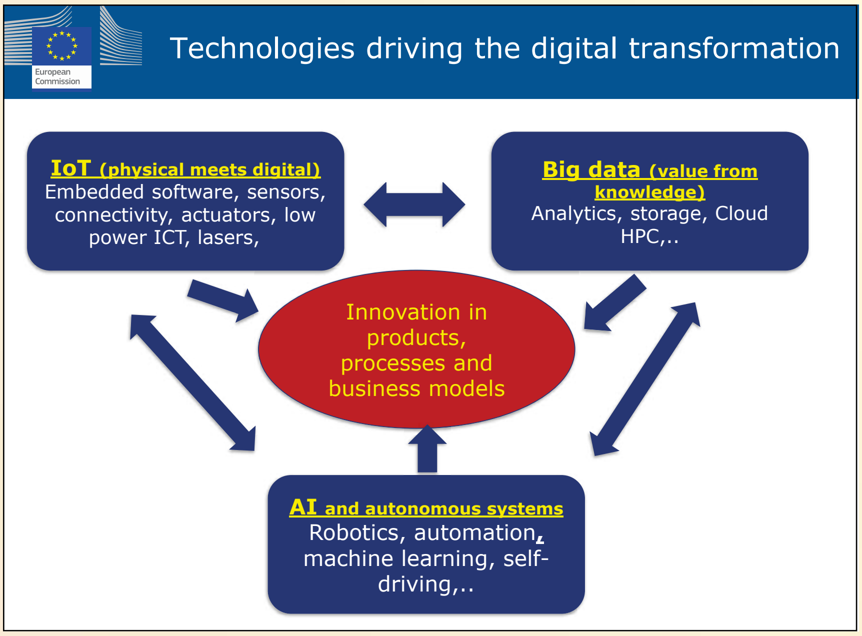

Technologies driving the digital transformation

IoT (physical meets digital)
Embedded software, sensors, connectivity, actuators, low power ICT, lasers,

Big data (value from knowledge)
Analytics, storage, Cloud HPC,..

Innovation in products, processes and business models

AI and autonomous systems
Robotics, automation, machine learning, self-driving,..

Figure 1: Technologies driving the digital transformation (source: European Commission, DG CONNECT, W. Steinhögl).

- Imulation and control
- Verification and validation
- Standardisation
- Situation understanding, cognition, decision making
- Path planning, (precision) maps, localisation and navigation
- Environmental awareness, self-learning,
- Human interaction and (public) acceptance, and
- Societal, ethical and legal aspects.

This includes impact from other sciences such as big data, IoT, artificial intelligence, communications and cloud, mechatronics and semiconductors, which contribute to meet these challenges. Connected cooperative autonomous vehicles are adaptive systems-of-systems. In this context, we have to consider several levels of system autonomy:

- the vehicle (robot) as such (level 1, local autonomy, self-dependence),
- the fleet/swarm/ad-hoc group of connected vehicles (level 2, increased amount and chances for information and adaptation of control), and
- the regional/global level 3 (throughput, environmental friendly operation, saving of resources), which needs to be considered for traffic or logistics optimisation or multi-modal transport, for instance.

There is a big difference between development and use in specialised fields of application, where trained operators and/or structured environments are involved (like construction, manufacturing, on-site operations, railways/metros, aircraft and space) and where the general public and public spaces set the requirements (road transport, smart cities/buildings/homes and care). ‘Mixed traffic’ of autonomous and traditional vehicles is the most demanding scenario, and in urban environments the ‘vulnerable road users’ (people, bicycles etc.) will still remain as partners (see also ‘Hello human, can you read my mind?’). Therefore, the Roadmaps for automated driving foresee five levels of ‘take over’ from the driver, the highest one being urban traffic. Similar levels are defined for other transport systems like railways and aircraft (see article on ‘Cross-Domain Fertilisation in the Evolution towards Autonomous Vehicles’ and the key note).

For many businesses, ‘Digital Transformation’ will be disruptive – some will vanish, some will change, but also new challenges and chances will arise and roles change. One example may be that for fully autonomous cars, insurance and liability will become the OEM/manufacturer’s responsibility and no longer be with the driver, the driver’s licence will become a vehicle licence. With reduced individual car ownership, OEMs may shift from pure selling to fleet management and maintenance of autonomous car fleets, because vehicles are then used mainly on demand (see also ‘How the Digital Business Model can Transform and Boost the Car Industry’).

This issue of ERCIM News covers many diverse aspects of the Special Theme, without being able to claim completeness.

The first group of contributions discusses rather generic challenges, like dependability of autonomous controls, safety and security co-engineering (addressing connected, intelligent automated vehicles) and the chances for disruptive innovations in advanced robotics by adaptive autonomy.

The largest block of articles is about automotive topics, ranging from the importance of new digital business models for the car industry to particular applications, specific development and technology paradigms and national research initiatives, including considerations for cooperative systems, connected vehicles and human factors and models for traffic safety and optimisation.

Several articles address cross-domain or very particular challenges and technologies, including agriculture, railways, ships and underwater robots, UAVs, machines, issues of swarm and fleet management in context of off-road and road vehicles, to name just a few.

References:

[1] ARTEMIS Strategic Research Agenda 2016, ARTEMIS Industrial Association, March 2016
(SRA 2016: https://artemis-ia.eu/publications.html)

[2] MASRIA 2017 – Multi-Annual Strategic Research and Innovation Agenda for ECSEL Joint Undertaking (AENEAS, ARTEMIS, EPoSS) https://artemis-ia.eu/publication/download/masria2017.pdf

[3] Position Paper Safetrans Working Group “Highly automated Systems: Test, Safety, and Development Processes” (2016) (cited in: D. Watzenig, M. Horn (Eds.): “Automated Driving”, Springer 2016, ISBN 978-3-319-31895-0)

Please contact:
Erwin Schoitsch, AIT Austrian Institute of Technology
Erwin.Schoitsch@ait.ac.at

Invited Article

Automated Driving

by Daniel Watzenig (Virtual Vehicle Research Center, Graz)

Automated vehicle technology has the potential to be a game changer on the roads, altering the face of driving as we know it. Many benefits are expected, ranging from improved safety, lower stress for car occupants, social inclusion, reduced congestion, lower emissions, and more efficient use of roads due to optimal integration of private and public transport.

Over the last three years, public authorities from many countries have presented action and innovation plans to facilitate the development and stepwise introduction of automated vehicles. Those plans cover actions for a multitude of technical and non-technical aspects that need to be taken into account. Several forecasts predict a limited availability of automated driving functions by 2020 (partial and conditional automation) and a wide availability by 2040 including high and full automation [1], [2], [3].

In particular, substantially improved vehicle and road safety is one of the expected major achievements. However, the ultimate safety test for automated vehicles will have to determine how well they can replicate the crash-free performance of human drivers, especially at the level of partial and conditional automation within mixed traffic. In order to be accepted by drivers and other stakeholders, automated vehicles must be reliable and significantly safer than today's driving baseline.

Automated vehicles by nature rely on sensing, planning, reasoning, and acting (or re-acting). A suite of vehicle sensors based on different sensing modalities (such as radar, lidar and camera) along with external sources (V2X) and detailed digital maps gather raw data of the vehicle's environment, driving situation, and ambient conditions. Sophisticated algorithms interpret the data, process it, and convert it to commands for the actuators (steering, braking). Most of the state-of-practice vehicles and prototypes rely on in-vehicle sensors and require little digital infrastructure communication while a greater connectivity between vehicles and their infrastructure is known to be beneficial. This entails the development of common communication protocols, security standards, and investment in new types of infrastructure or upgrade of existing ones. Reproducible testing of such highly connected, cooperative, bi-directionally interacting, and automated systems of systems has become one of the biggest challenges.

Figure 1 highlights the complex architecture that has to be designed, implemented, validated, and finally qualified for automotive purposes (ISO 26262 and ISO PAS 21448 SOTIF). Reliability and safety (for occupants and other road users) have to be guaranteed at any time, in any weather conditions, and in any traffic situation.

As the technology for automated driving develops, the research and development focus is shifting from basic research towards demonstration of technological readiness and 24/7 availability of functions and systems. The shift of automated driving to the next level of maturity requires several technological advances and fields of action including (but not limited to):

- Fail-operational behaviour (multiple redundancy at low cost)
- Assessing safety, reliability, and robustness (independent and reproducible test)
- Harmonised driving scenarios (set of agreed test scenarios)
- Self-learning, self-adapting, self-configuring architectures (qualifiable and certifiable)
- Demonstrating cybersecurity and privacy
- Dependable power computing (1 GByte/s of data has to be processed in real-time)
- Human factor (especially for SAE J3016 level 3)

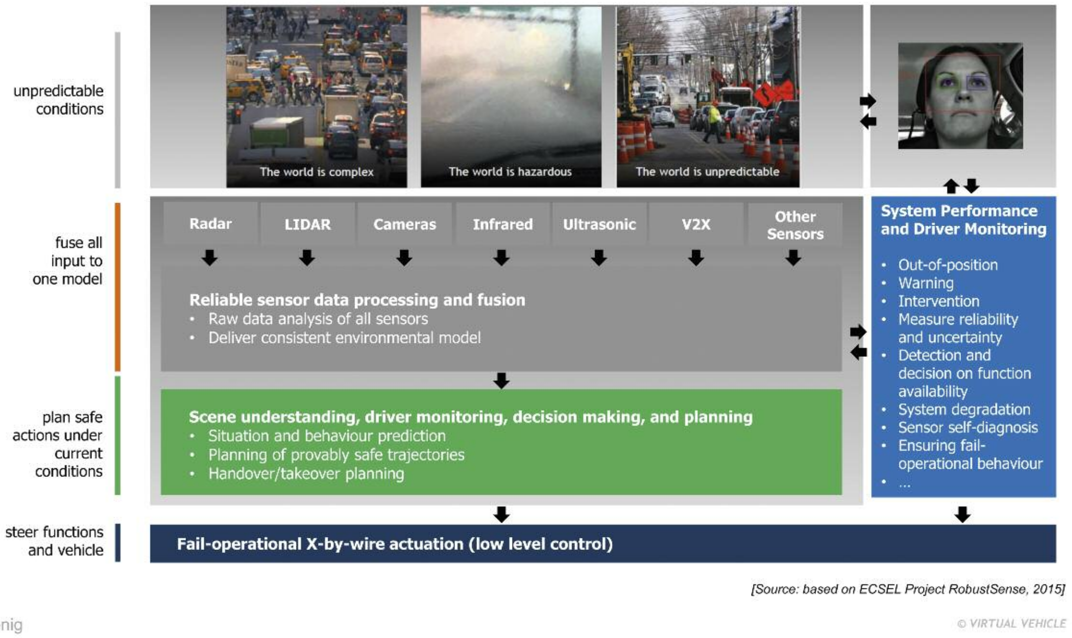

unpredictable conditions

The world is complex

The world is hazardous

The world is unpredictable

fuse all input to one model

Radar | LIDAR | Cameras | Infrared | Ultrasonic | V2X | Other Sensors

Reliable sensor data processing and fusion
- Raw data analysis of all sensors
- Deliver consistent environmental model

plan safe actions under current conditions

Scene understanding, driver monitoring, decision making, and planning
- Situation and behaviour prediction
- Planning of provably safe trajectories
- Handover/takeover planning

System Performance and Driver Monitoring
- Out-of-position
- Warning
- Intervention
- Measure reliability and uncertainty
- Detection and decision on function availability
- System degradation
- Sensor self-diagnosis
- Ensuring fail-operational behaviour
- ...

steer functions and vehicle

Fail-operational X-by-wire actuation (low level control)

[Source: based on ECSEL Project RobustSense, 2015]

Figure 1: Automated driving architecture.

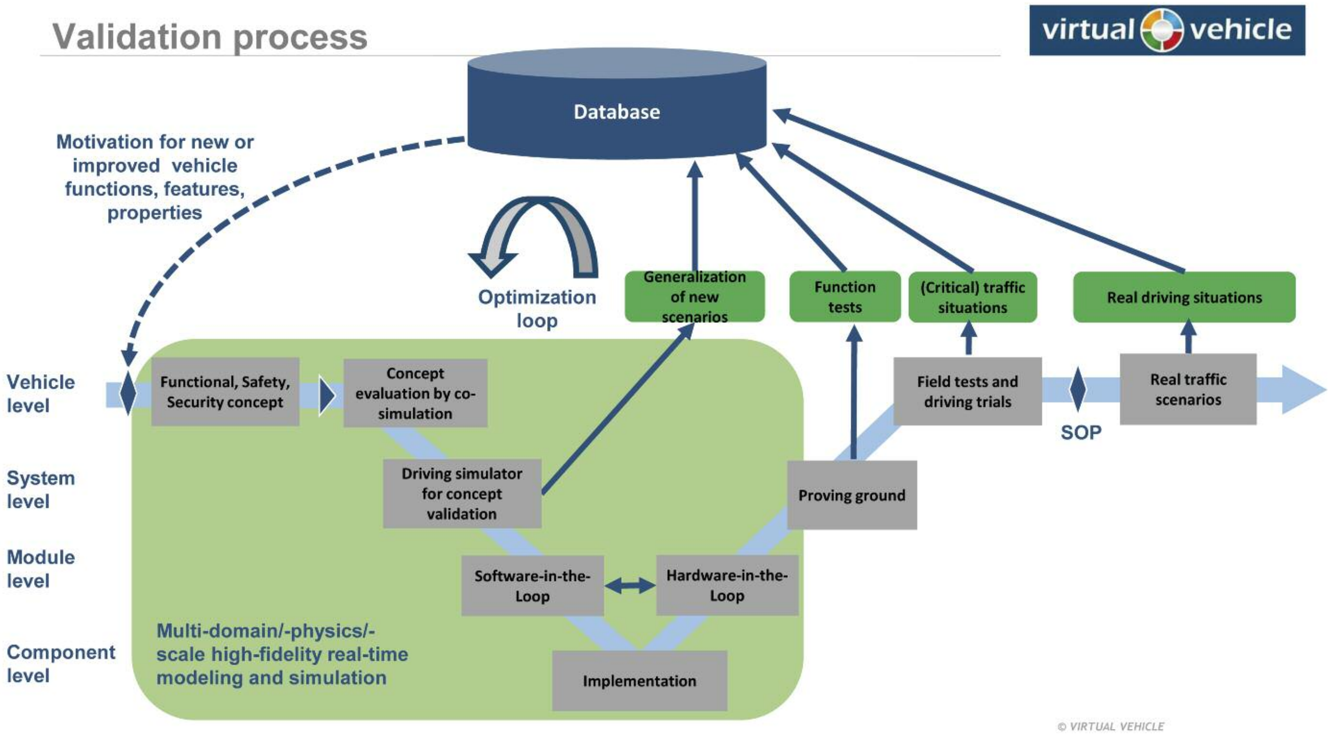
Validation process
virtual vehicle
Database
Motivation for new or improved vehicle functions, features, properties
Optimization loop
Generalization of new scenarios
Function tests
(Critical) traffic situations
Real driving situations
Vehicle level
System level
Module level
Component level
Functional, Safety, Security concept
Concept evaluation by co-simulation
Driving simulator for concept validation
Software-in-the-Loop
Hardware-in-the-Loop
Implementation
Proving ground
Field tests and driving trials
SOP
Real traffic scenarios
Multi-domain/-physics/-scale high-fidelity real-time modeling and simulation
© VIRTUAL VEHICLE

Figure 2: The use of real-world data to develop and test safety-critical automated driving functions.

- Environment modelling and perception in real-time
- Advanced data fusion (low, intermediate, high level fusion) and robust control
- Digital infrastructure embedding
- Real-world data-driven engineering feedback.

Despite tremendous improvements in sensor technology over the last couple of years, pattern recognition techniques, robust signal processing, control system design, computational power (multi-core and many-core technology), V2X, and other system technology areas, market introduction of a fully automated vehicle that is capable of unsupervised driving in an unstructured environment still remains a long-term goal. Even for structured environments, further research is needed to exploit the full potential of road transport automation.

One of the keys to succeed is the tight engagement of real-world data and virtual development methods to design new or to improve existing vehicle functions and features as well as methods, processes, and tools for development and validation. Figure 2 highlights the ‘closed engineering loop’ in order to improve methods, tools, and processes for virtual approval.

Seamlessly connecting models, data, and processes with openness to a certain extent along the entire vehicle life-cycle will significantly contribute to a decrease in road driving test kilometres, to lower development costs, and ultimately to improved reliability and safety. Austria has already reacted in order to strengthen its position within Europe. The Austrian Federal Ministry for Transport, Innovation, and Technology (BMVIT) has launched a call to set up and run a public test region for automated vehicles. All relevant industrial and academic partners from the automotive and infrastructure domains have joined forces and are on the way to establishing the Austrian Lightvehicle Proving ground (ALP.Lab) starting in 2017 which will cover motorways, federal highways, urban areas, and border crossings to Slovenia.

All of the above mentioned fields of activity need to be further investigated to exploit the full potential in order to ensure a safe, reliable, acceptable, and secure behaviour of automated vehicles embedded in their intelligent environment at all times. Along with the technological advances, progress in legislation, liability, and insurance is essential and urgently needed.

Meanwhile many collaborative European research projects focusing on different aspects of automated driving have been recently launched (e.g., ENABLE-S3, Inframix, TrustVehicle and AutoDrive) and several European initiatives (like ECSEL JU, EARPA, ERTRAC and SafeTRANS) are driving this topic forward, finally reflecting the strong European movement in the field of automated driving.

Recommended reading:

D. Watzenig, M. Horn, Automated driving – safer and more efficient future driving, Springer, ISBN 978-3-319-31895-0, http://www.springer.com/de/book/9783319318936, 2016.

References:

[1] Automated driving roadmap, European road transport research advisory council (ERTRAC), 2015.

[2] Roadmap on smart systems for automated driving, European technology platform on smart systems integration (EPoSS), 2015.

[3] ECSEL Austria, Austrian research, development & innovation roadmap for automated vehicles, http://www.ecsel-austria.net/newsfull/items/automated-driving-roadmap.html, March 2016.

Please contact:

Daniel Watzenig, Virtual Vehicle Research Center, Graz, Austria
Daniel.Watzenig@v2c2.at

Dependability for Autonomous Control with a Probability Approach

by Lan Anh Trinh, Baran Cürüklü and Mikael Ekström (Mälardalen University)

For the last decade, dependability – the ability to offer a service that can be trusted – has been the focus of much research, and is of particular interest when designing and building systems. We are developing a dependable framework for an autonomous system and its control.

A shift from automatic to autonomous control has emerged in the development of robots. Autonomous control allows the robot to have freedom of movement as well the ability to directly interact with humans as well as other robots in a collaborative environment. Having a dependable platform for autonomous control becomes absolutely crucial when building such a system.

The concept of dependability originally derives from software development and can be defined as 'the ability to deliver service that can justifiably be trusted' (Avizienis et al. [1]). The dependability of a system is evaluated by one, several or all of the following attributes: availability, reliability, safety, integrity, and maintainability. The implementation of dependability starts with an understanding of the threats to the system's dependability, which may include failures, errors, and faults. Therefore, four means have been developed to protect system dependability: fault prevention, fault removal, fault forecasting, and fault tolerance.

Mälardalen University is hosting a long term project, DPAC – Dependable Platforms for Autonomous system and Control [L1], with the main aim of implementing dependability for different platforms for autonomous systems and their control. Within the profile, three main fields are covered that include hardware with heterogeneous system architecture (HAS), software design and autonomous control. In this project, the dependability of autonomous control system plays a vital role.

A fault is the root of every failure occurring inside or outside of the system. Nevertheless, the most pressing challenge is how to predict the frequency of faults and at what moment a fault occurs, thus fault analysis is presented to minimise the probability of a fault and to estimate when faults will happen in the system. Thereafter, other means are developed to protect the dependability with respect to the analysis of faults.

Various approaches to fault analysis, such as Petri Net (PN), fault tree analysis (FTA), failure modes effects and criticality analysis (FMECA), and hazard operability (HAZOP) are introduced in the work of Bernardi et al. [2]. However, our research focusses on PN, since the PN framework provides not only a probability approach for fault analysis and fault prevention in both development and operational stages of designing a system but also for mitigation of the implementation progress. For instance, as an extension of PN, a stochastic Petri net can be combined with Markovian models to evaluate the probability of current state and the probability of future fault events for a fault prognosis process. In our study [3], a coloured time PN is utilised for fault tolerance analysis of multi-agents in a complex and collaborative context.

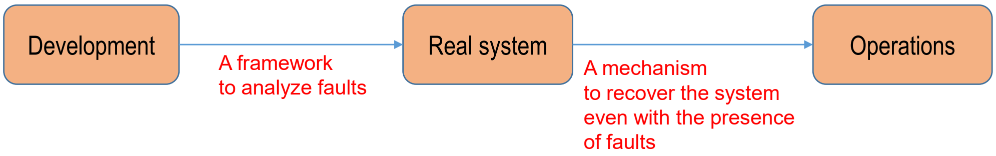
Development
A framework to analyze faults
Real system
A mechanism to recover the system even with the presence of faults
Operations

Figure 1: The development and operational stages in building a dependable control system.

In our current research, three different stages are considered to be implemented, which relate to three means: fault prediction, fault prevention and fault tolerance for autonomous control. For fault prediction, the probability model is established to estimate the time fails of the system. The time that a fail could happen is modelled by an exponential distribution. The design of the autonomous control system can be presented by a network of nodes like a Petri net and the fails are propagated from one node to others. Each type of failure is computed by separated variables and the hierarchy structure may be taken into account for a complicated agent interaction. The probability of a successful task depends on the parameters of the system. Note that there are two types of parameters: those from the environment that cannot be changed i.e., non-configurable, and the others that are configurable. Correspondingly, the probability of failures could be estimated based on the defined values of those parameters and the architecture of autonomous systems.

After designing a mathematics model of fault prediction with respect to configurable parameters, expectation maximisation (EM) is applied to minimise failures while taking into account hidden parameters from the environment. This stage is actually the implementation of fault prevention as we attempt to find the best model for the system to avoid further failures in future.

Finally, the fault tolerance allows the system to continue to work even as fails happen. To do so, there are different algorithms and modules running in parallel in the system. Each module will use the fault prevention design as described above. When one algorithm or module fails, it is immediately replaced by another. It is important to note that the decision must be done in advance i.e., before the fails happen, otherwise it could be too late. Again the fault analysis combined with an artificial intelligence algorithm are used for decision making.

Overall, this architecture allows all necessary means to be implemented to preserve the dependability of the system. Using a graphical probabilistic model, the means could be implemented in a real robot to facilitate dependable autonomous control.

Link:
[L1] www.es.mdh.se/projects/414-DPAC

References:

[1] A. Avizienis, et al.: “Basic concepts and taxonomy of dependable and secure computing”, IEEE Transactions on Dependable and Secure Computing, 2004.

[2] S. Bernardi, J. Merseguer, D.C. Petriu: “Model-Driven Dependability Assessment of Software Systems”, Springer, 2013.

[3] T. Lan Anh, C. Baran, E. Mikael: “Fault Tolerance Analysis for Dependable Autonomous Agents using Colored Time Petri Net”, 9th International Conference on Agents and Artificial Intelligence, ICAART 2017.

Please contact:

Lan Anh Trinh
Mälardalen University, Sweden
anh.lan@mdh.se

Safety and Security Co-engineering of Connected, Intelligent, and Automated Vehicles

by Christoph Schmittner, Zhendong Ma, Thomas Gruber and Erwin Schoitsch (AIT)

Connected, intelligent, and autonomous vehicles pose new safety and security challenges. A systematic and holistic safety and security approach is a key to addressing these challenges. Safety and security co-engineering in the automotive domain considers the coordination and interaction of the lifecycles, methodologies, and techniques of the two disciplines, as well as the development of corresponding standards.

Connected, intelligent, and autonomous vehicles transform traditionally mechanical and electrical cars into ‘networked computers on wheels’. Along with the many technology breakthroughs and benefits, challenges of safety and security become imminent and real. The electrical and electronic systems that control an automated vehicle are no longer immune to cyberattacks commonly seen in IT systems. A combined safety and security approach is necessary to address the challenges that have arisen in recent years, including co-engineering activities, methodologies, techniques, and a coherent approach in relevant standards.

The correct identification of safety and security goals is the first step in the development lifecycle of a system. The identification of hazards and assets reveals potentially vulnerable parts of a system. Subsequently, a first concept architecture for these parts is defined which can then be analysed to identify potential weaknesses, e.g., if a failure or an attack could trigger an intolerable risk. In both cases, requirements are defined which aim at preventing such risks. Different methods should be used to address various issues during the development lifecycle with different levels of detail. Figure 1 displays the most common methods in the respective phases of the V-model.

A useful security technique is threat modelling, which defines a theoretical model of perceived threats to a system. We developed a systematic approach to apply threat modelling to automotive security analysis and combined it with the Failure Mode and Effect Analysis to FMVEA, Failure Modes, Vulnerabilities and Effects Analysis [1]. Threat modelling should be performed in all phases of the development lifecycle. Different levels of detail can be used along the lifecycle with different objectives in each phase. In the concept phase, mod-

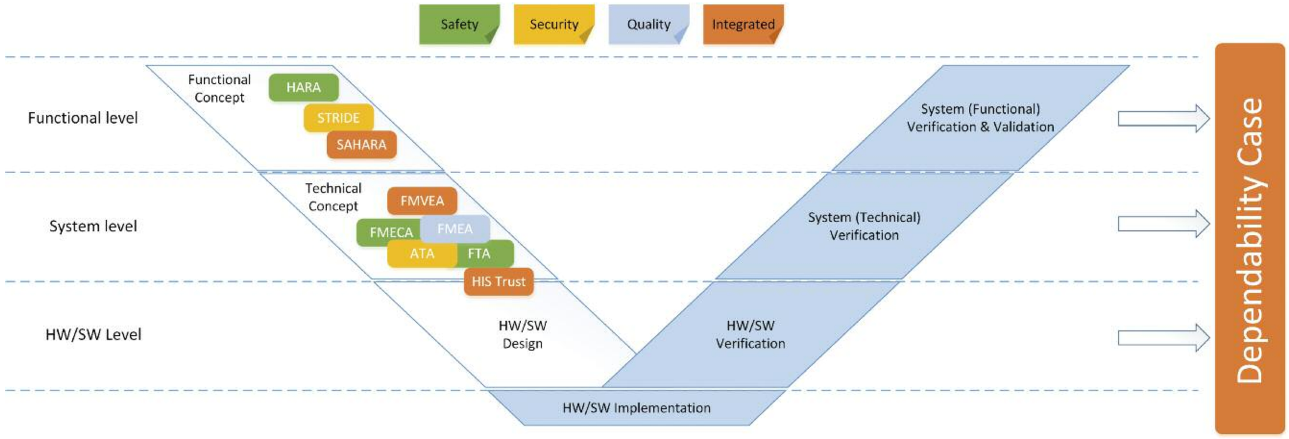

Figure 1: Dependability engineering in the development lifecycle.

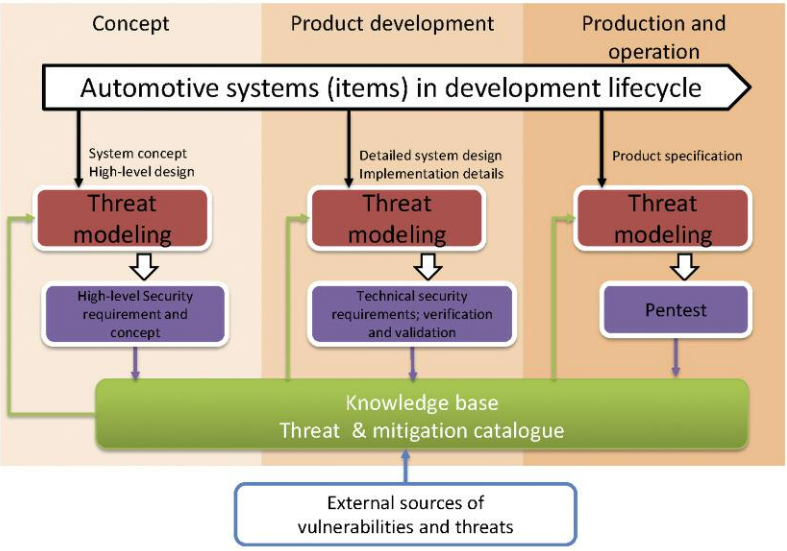
Concept
Product development
Production and operation
Automotive systems (items) in development lifecycle
System concept
High-level design
Detailed system design
Implementation details
Product specification
Threat modeling
Threat modeling
Threat modeling
High-level Security requirement and concept
Technical security requirements; verification and validation
Pentest
Knowledge base
Threat & mitigation catalogue
External sources of vulnerabilities and threats

Figure 2: Iterative threat modelling and mitigation during the development lifecycle.

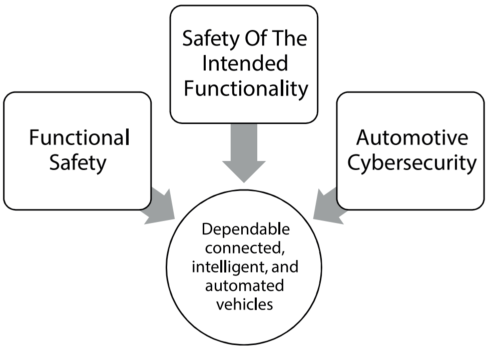
Safety Of The Intended Functionality
Functional Safety
Automotive Cybersecurity
Dependable connected, intelligent, and automated vehicles

Figure 3: SotIF approach combined with safety and cybersecurity co-engineering.

elling results in high-level security and safety requirements and security concepts. In the product development phase, it can define technical security and safety requirements for functional, security and safety design. It can also be used to discover design vulnerability and flaws and to specify comprehensive requirements that can be verified and validated in unit and integration testing in an iterative way in parallel to system design and implementation. In the production and operation phase, it prioritises risks and prepares penetration testing on completed automotive components and systems. A knowledge base is continuously enriched by the output from threat and failure modelling activities, enabling the reuse of artefacts across different projects. Further, related vulnerabilities and threats from external sources are promptly incorporated into the threat and mitigation catalogue.

A new version of the automotive functional safety standard ISO 26262 is currently under development. Although not suited for completely autonomous cars, automated functions are considered to a significantly higher degree than in the first edition. This is supported by the new standard development SotIF (Safety of the Intended Functionality). SotIF describes nominal performance metrics for sensor systems and automated functions. This regulates the area where a system may cause a hazard without a failure in the traditional understanding. The processing algorithm made, based on the received understanding of the environment, a hazardous decision without a fault in the system. This could be caused by a limitation in the sensor algorithm or signal noise or insufficient performance of a sensor. A new ISO/SAE automotive security standard completes these activities. All three standards need to consider the increased interaction and co-engineering between system-, safety- and security- engineers. AIT is involved in the development of all three standards and is a member of the cybersecurity and safety task group that developed the Annex. The goal of interaction and communication points is to lay the groundwork for a workflow with shared phases [2].

AIT has further developed automotive safety and security co-engineering in the Artemis project EMC2 and demonstrated it for a Hybrid Electric Powertrain control system of AVL. In the SCRIPT project, we applied the newly published SAE J3061 standard for conducting TARA in the development of a secure communication gateway for autonomous off-road vehicles [3]. We are also working towards an efficient and model-based approach to multi-concern assurance including safety; security, reliability, and availability in the scope of ECSEL project AMASS. AIT will take the next step towards safe, secure and cost-efficient automated driving in the ECSEL project AUTODRIVE starting in 2017. The interaction point approach of ISO 26262 Edition 2 will be the object of research in the ECSEL project AQUAS also starting in 2017.

Links:

[L1] www.ait.ac.at/en/about-the-ait/center/center-for-digital-safety-security/

[L2] www.emc2-project.eu/

[L3] www.tttech.com/company/research-projects/austrian/script/

[L4] www.amass-ecsel.eu/

References:

[1] C. Schmittner, Z. Ma, E. Schoitsch, T. Gruber: "A Case Study of FMVEA and CHASSIS as Safety and Security Co-Analysis Method for Automotive Cyber-physical Systems," in 1st ACM Workshop on Cyber-Physical System Security, Apr. 2015, ACM, pp. 69-80.

[2] E. Schoitsch, C. Schmittner, Z. Ma, T. Gruber: "The Need for Safety & Cyber-Security Co-engineering and Standardization for Highly Automated Automotive Vehicles", AMAA 2015, Berlin, Germany, July 2015.

[3] C. Schmittner, et al.: "Using SAE J3061 for Automotive Security Requirement Engineering", in International Conference on Computer Safety, Reliability, and Security, pp. 157-170. Springer, 2016.

Please contact:

Christoph Schmittner, Zhendong Ma, Thomas Gruber, Erwin Schoitsch
AIT, Austria
christoph.schmittner@ait.ac.at,
zhendong.ma@ait.ac.at,
thomas.gruber@ait.ac.at,
erwin.schoitsch@ait.ac.at

Adaptive Autonomy Paves the Way for Disruptive Innovations in Advanced Robotics

by Baran Cürüklü (Mälardalen University), José-Fernán Martínez-Ortega (Universidad Politécnica de Madrid) and Roberto Fresco (CREA)

Machines are going to become smarter with autonomy as the new standard. Consequently, new services will emerge. For this to happen, however, we need a new approach to autonomy, which assumes different levels – and perhaps more importantly, machines that can handle these different levels. Adaptive autonomy is a means to achieving this goal.

One of the major challenges in future robotics is to design systems that can collaborate with each other. This is radically different from classical robot automation in which industrial robots, protected by fences, carry out repetitive tasks. The societal impact of these advances is not hard to imagine. In this context, precision agriculture is a highly relevant application domain that is going to be subject to disruptive innovations. The effects of climate change are already influencing access to arable land and world population growth needs to be taken into account. Thus, there is a real risk of food shortage in the future if we do not develop autonomous systems that can collaborate to solve real-world complex problems. There is also a societal dimension to agriculture. If we look at agriculture and its short-term challenges, we see that in many industrial countries the depopulation of rural areas makes it harder to find skilled labour. This, despite that fact that only 2-5% of the total work force work is in this sector. With depopulation comes reduced investment in these regions – in infrastructure, for example. A lively countryside is only a fiction if there are neither jobs nor services in these areas.

Adaptive autonomy in combination with advanced human-robot interaction (HRI) could be part of the solution to this problem. The former assumes that the level of autonomy of a system, e.g., a farm vehicle, is subject to change based on its interactions with other autonomous systems or humans. Developing generic agent architectures to allow this to happen is a challenge. Advances in dependable hardware and software systems are also needed, since they allow autonomous systems to be safe and reliable. For an agricultural application, the overall system architecture is very complex, because it needs to be configured for both the vehicle navigation system (i.e., prescription maps) and for the farming task itself (i.e., seeding, crop establishment, plant care and selective harvesting). Thus, a huge quantity of information, site conditions (i.e., soil conditions) and data must be sensed and processed (involving sensors on board, actuators).

HRI is the other essential technology in this context, since these technologies do not always allow seamless user-friendly interaction that is also highly regarded by expert users. Lowering the cognitive load is important for minimising hazards. As mentioned above, adaptive autonomy agent architectures for robots can only create complex interactions through simple concepts such as ‘willingness to help’, and ‘willingness to ask for assistance’ [1]. The approach here assumes a high degree of interaction and complex relations (see Figure 1), instead of complex agent architectures. Mälardalen University, Sweden, is involved in two projects in which adaptive autonomy plays an important role: DPAC – Dependable Platforms for Autonomous Systems [L1], a national project, and ECSEL JU funded SWARMs – Smart and Networking Underwater Robots in Cooperation Meshes [L2]. The latter is coordinated by the second author of this paper.

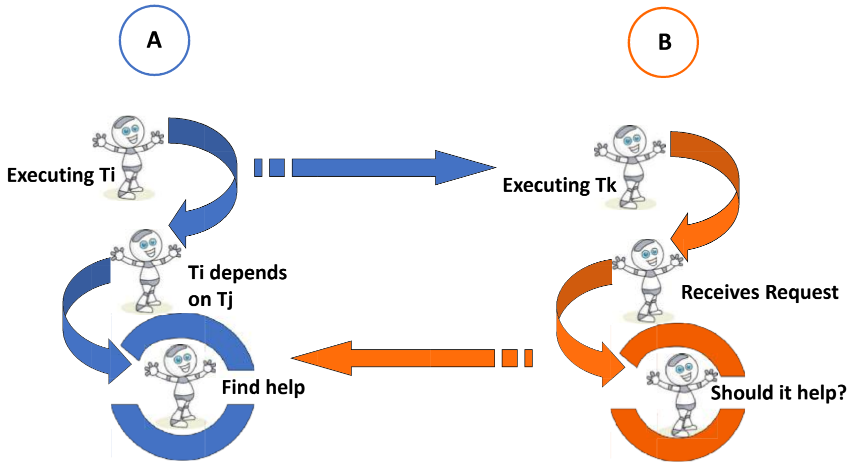

Figure 1: A high degree of interaction and complex relations.

When adaptive autonomy is extended to advanced robot-robot interaction and combined with HRI, we have a system for precision control of a group of autonomous systems by only a few operators. Therefore the role of a farmer changes and he becomes a specialised and innovative operator that can orchestrate a heterogeneous set of robots. Note that heterogeneity is an important assumption in this context. Different types, such as vehicles, robots, airborne solutions, etc. will be used for finding the right blend for a specific challenge. This will inevitably lead to mixed human-robot groups. For these groups to function seamlessly the autonomous systems need to be ethical [2].

Another research direction is underwater robots, which is associated with hazardous operations. In this application domain, in the SWARMs project, the first and second author have assumed that orchestration of the robots through high-level planning is needed. Despite the complexity of the missions, the operator is only monitoring the mission, including generation and execution of the plan. In both application domains, as well as other similar ones, the assumption is to minimise the intervention of the human operator. This is possible to achieve using autonomous systems. In any case, adaptive autonomy contributes to advanced robot-robot interaction schemes. Therefore, collaborative robots will get closer to humans working in teams – and even outperform them in certain tasks.

Links:
[L1] www.es.mdh.se/projects/414-DPAC
[L2] swarms.eu/

References:

[1] M. Frasheri et al., Towards Collaborative Adaptive Autonomous Agents, Conf. of ICAART, 2017.
[2] G. Dodig-Crnkovic, B. Çürüklü, Robots – Ethical by Design. Ethics and Information Technology, Springer, 2012.

Please contact:

Baran Cürüklü
Mälardalen University, Sweden
+46 (0)73-9607453
baran.curuklu@mdh.se

How the Digital Business Model can Transform and Boost the Car Industry

by Ioannis Chrysakis (ICS-FORTH)

Moving from the traditional to the digital business model gives the opportunity to car manufacturers and new non-traditional players of the car industry to create the autonomous cars of the future.

For many years, car manufacturers (or OEMs) have managed to increase their sales by applying traditional business models to the industry, and by continuously providing better customer experiences. In fact, the automotive industry has created competitive advantages through advances in the fields of mechanical and electrical engineering. More precisely, improvements in value depend on being able to create more cost-effective engines, provide extra safety packs that may include auto braking, chassis control etc. and fancy multimedia systems with a lot of connectivity functionalities (see Figure 1). However, these features are fairly easy to copy. For example, although multimedia screens that connect to Android or iOS devices were initially included by only some car makers, they seem to have become a standard feature for the upcoming new models of almost all manufacturers.

The digital disruption can potentially transform the car industry as it provides new opportunities for smart and autonomous cars which can now communicate, socialise and collaborate with other things, including other vehicles, traffic lights, mechanics, parking lots and dealers, thus enabling them to participate in a broad ‘system of systems’ [1]. New business models, though, are required in order to make manufacturers defend their value by redesigning customer engagement and expectations. The concept of the digital business model [2] can bring new capabilities to customers and the desired defending value for OEMs. Some of the digital services that can be encapsulated in cars may include insurance services according to driving style, travel suggestions, infotainment, monitoring car diagnostics, driver health services etc. The autonomous cars of the future could run any kind of software in their existing multimedia screens offering a great in-vehicle experience for drivers.

We will now present the digital business model focusing on two dimensions of exploiting software in the car industry that are closely related to autonomous driving behaviour. The first dimension has to do with location-based services. These may exploit maps and provide drivers with guidelines and information (gas stations, sights, traffic, shops etc.) according to their destination and activate the auto-driving functionality for preferable routes where it is applicable. Dedicated software could be developed by auto manufacturers or created through collaboration with non-traditional players like Google, Apple, Facebook or other start-up companies by integrating their services and their collected data. The partnerships between OEMs, suppliers, and service providers can benefit from sharing the costs of autonomous vehicle technology. On the other hand, OEMs in particular need to maintain control over their individual value creation and their potential success in the emerging ecosystems. The second dimension has to do with remote-analytics that are based on the live data analysis of the car. The exploitation of analytics could result in more cost-effective and safe driving by providing service reports and faults, active suggestions for fuel, type of driving, safety warnings etc.

In both cases, the added value of guidance could be delivered through software that requires access to satellite or mobile networks through the screens of each car. This software could boost car sales or alternatively it can be applied by using a subscription model to a centralised platform that shares specialised guidelines per driver or/and per car. Finally, this model defends the value by producing more autonomous cars that encapsulate the above features, are self-driven and can prevent failures. (See Figure 2).

In fact, a variety of companies have already entered the ecosystem, seeking to capitalise on opportunities created by digital disruption and the application of digital business models, for example, Automatic.com [L1] and moj.io [L2] which perform live car monitoring and

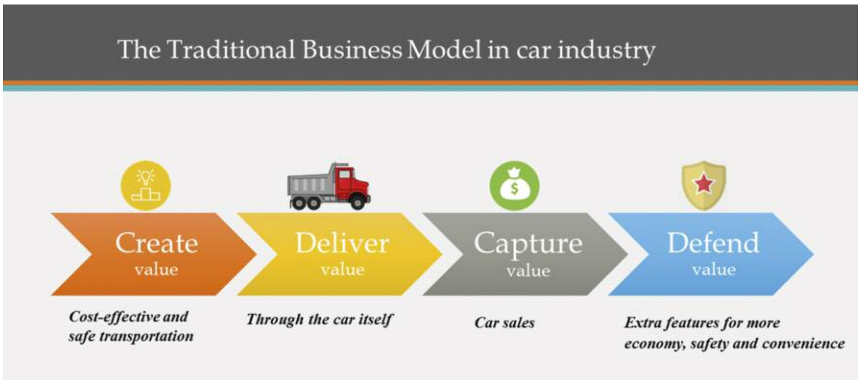
The Traditional Business Model in car industry
Create value
Deliver value
Capture value
Defend value
Cost-effective and safe transportation
Through the car itself
Car sales
Extra features for more economy, safety and convenience

Figure 1: The traditional business model in the car industry.

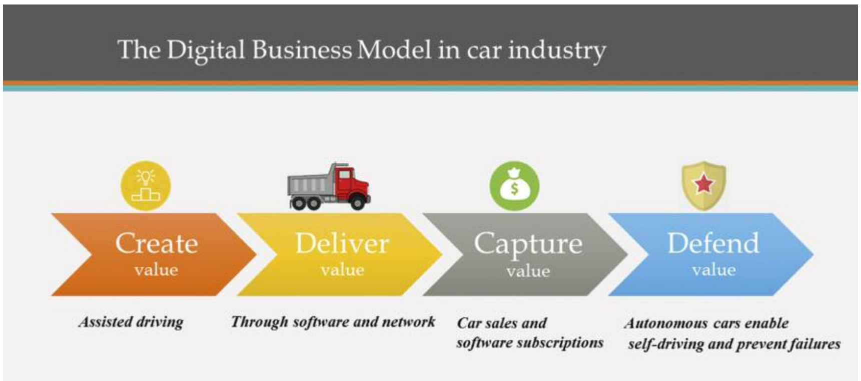
The Digital Business Model in car industry
Create value
Deliver value
Capture value
Defend value
Assisted driving
Through software and network
Car sales and software subscriptions
Autonomous cars enable self-driving and prevent failures

Figure 2: The digital business model in the car industry.

offer a lot of functionalities to connected cars. OEMs can collaborate with such companies or develop their own solutions by investing in new resources and technologies. Recently, Volvo [L3] announced that it uses machine learning algorithms to expose predictive analytics. Furthermore, the idea of autonomous driverless cars gives great opportunities for less cars, more parking spaces and an eco-friendly means of regional transport to customers that could not be necessary the actual car owners. Thus, Ford, GM, Audi, BMW, Tesla, Nissan, Mercedes-Benz, and others have promised to deliver self-driving cars in the next five years or so [L4].

Finally, since the disruptive technology-driven trends create new opportunities for autonomous smart cars, OEMs should align their strategic priorities to differentiate their products and services, in order to reshape the value proposition and finally create the desired competitive advantage.

Links:

[L1] www.automatic.com/how-automatic-works/
[L2] www.moj.io/
[L3] globalsmt.net/industry_news/amazing-ways-volvo-uses-big-data-machine-learning-predictive-analytics/
[L4] www.wired.com/2017/01/self-driving-cars-approach-auto-industry-races-rebuild

References:

[1] D. Wollschläger, et al.: "Digital disruption and the future of the automotive industry". IBM Center for Applied Insights, September 2015.
[2] O. A. El Sawy, et al.: "Business Modelling in the Dynamic Digital Space: An Ecosystem Approach", pp 13-20, Springer, 2012.

Please contact:

Ioannis Chrysakis,
ICS-FORTH, Greece,
+30 2810 391638
hrysakis@ics.forth.gr
http://www.ics.forth.gr/~hrysakis

RECAR: Hungarian Research Centre for Autonomous Road Vehicles is on the Way

by Zsolt Szalay, Domokos Esztergár-Kiss, Tamás Tettamanti (Budapest University of Technology and Economics) and Péter Gáspár, István Varga (SZTAKI)

RECAR is a Hungarian research centre for autonomous vehicle technology providing internationally unique education and research with close cooperation with industrial partners. The research centre has technologically advanced laboratories to facilitate high quality R+D+I activities. RECAR is also expected to contribute to the increase in qualified workforce in the automotive industry in Hungary.

Info-communication technology is creating big changes in road transportation (vehicles, infrastructure, and passengers) as well as society in general [1]. In the automotive industry, continuous expansion of automation is occurring [2]. In Hungary to date, research projects in the field of autonomous vehicles have been performed separately and in parallel, with little collaboration between institutes.

RECAR (REsearch Centre for Autonomous Road vehicles) is a comprehensive organisation that helps its members to combine their expertise to improve research outcomes. RECAR's aim is to connect academic and industrial expertise, as well as education and research, in order to facilitate the training of highly qualified professionals and to strengthen research competences. As a first step in establishing RECAR, in January 2016 the leaders of Budapest University of Technology and Economics (BME), Eötvös Loránd University (ELTE) and Hungarian

Academy of Sciences, Institute for Computer Science and Control (MTA SZTAKI) signed the RECAR program's collaboration agreement.

Based on international trends and incentives of national companies in the automotive industry, the partners initiated two new master programs (MSc): Autonomous Vehicle Control Engineering hosted by BME and Computer Science for Autonomous Driving hosted by ELTE. The foundations for the master programs have already been laid down and the detailed education plans are on the way together with the elaboration of course materials. Importantly, the master programs were jointly defined by both academic and industrial partners in order to address real industry needs. The programs will be launched by 2018.

Laboratories that go beyond the state of the art are required to perform education and research in the area of autonomous vehicle technology. The partners have defined a testing and validation structure of five layers [3], including laboratories that serve both educational and research needs (Figure 1). The laboratory tests are composed of technology research labs, component analysis labs, system integration labs, and last but not least a vehicle-in-the-loop lab. The labs will have unique and high quality equipment with most of the functionalities required for autonomous vehicle development and testing.

Several new research projects, which take advantage of the technology offered by these innovative laboratories, will begin in the near future (Figure 2). The research topics were defined based on current international trends as well as initiatives of the industrial partners. The research topics also reflect the EU's recent Gear 2030 body activities that aim to formulate a harmonised and competitive European vision for the 'Connected Car and Automated Driving'. Accordingly, the work will be realised in the form of ongoing industrial projects and research collaboration by seven main research groups (with examples):

1. Development of autonomous vehicle demonstration platform:
   - autonomous vehicle prototype development
   - demonstration of autonomous functions;

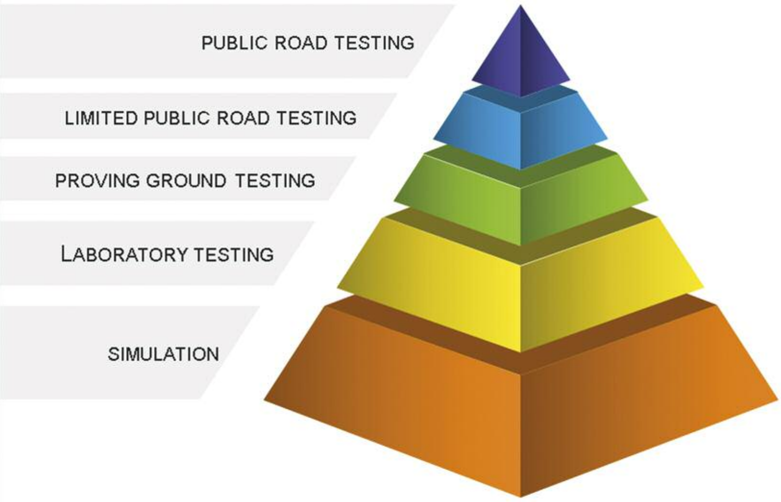

Figure 1: Autonomous vehicle testing and validation layers [3].

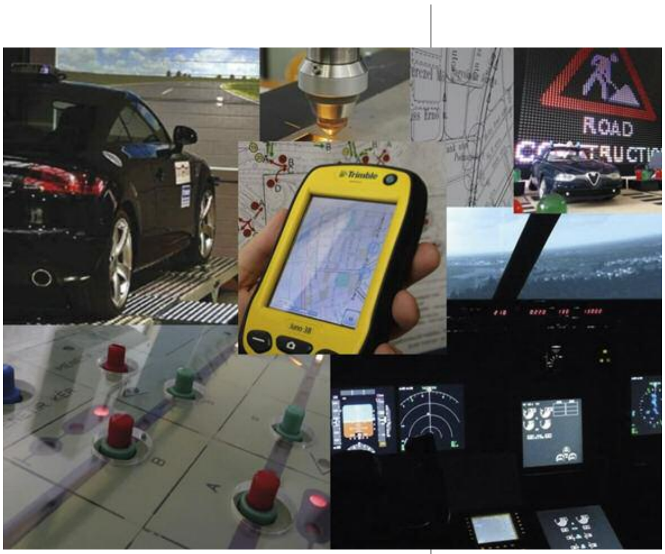

Figure 2: Example research topics.

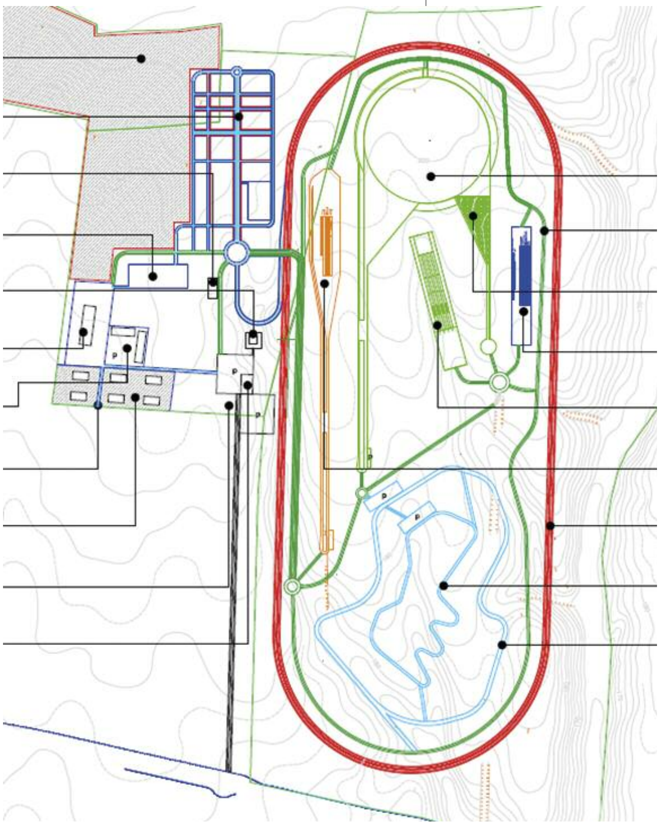

Figure 3: Automotive proving ground dedicated to autonomous vehicle testing and validation at Zalaegerszeg.

2. Control of autonomous vehicles:
   - cooperative vehicle control
   - optimal trajectory planning
   - cooperative navigation
   - vehicle motion measurement and estimation;
3. Communication systems within and among vehicles:
   - wireless communication of vehicles (V2X)
   - distributed measurement and information processing cyber security of autonomous vehicles;
4. Environment sensing of autonomous vehicles:
   - optimal sensor architecture
   - environment sensing based on multi-sensor fusion
   - radar-sensor based target classification;
5. Intelligent transportation systems:
   - advanced road traffic control
   - energy consumption optimisation
   - urban parking and space management;
6. Interaction of human and autonomous vehicles:
   - information management
   - human factors
   - liability of autonomous vehicles;
7. Testing and validation of autonomous vehicles: relevant research subjects to the test track.

A new test track will also be established beside the research centre (close to the city of Zalaegerszeg, Hungary) through an investment of 130 million Euros by 2019. The planned proving ground (Figure 3) is specifically dedicated to testing autonomous vehicle functionalities in an urban environment (with general roadside objects, city traffic elements, building facades and traffic infrastructure). The 250 ha area will incorporate the following test features:

- standard vehicle dynamics testing and validation,
- fully integrated autonomous vehicle testing and validation,
- environment preparation (obstacles, traffic signs, traffic control, other vehicles, vulnerable road users),
- complex driving and traffic situations,
- Smart City features,
- from prototype testing through to series production testing and validation.

The use case vision is level 4 autonomous driving from suburban home to city office and valet parking (rural road, highway traffic, suburban area, city environment with continuous transition) [3].

Link: http://recar.bme.hu/eng

References:

[1] T. Tettamanti, I. Varga, Zs. Szalay: "Impacts of autonomous cars from a traffic engineering perspective", Periodica Polytechnica ser. Transp. Eng., 2016, Vol. 44. No.4. pp. 244-250, doi: 10.3311/PPtr.9464I

[2] A. Alessandrini, et al.: "Automated Vehicles and the Rethinking of Mobility and Cities", Transportation Research Procedia, 2015, Vol. 5, pp. 145-160, ISSN 2352-1465, http://dx.doi.org/10.1016/j.trpro.2015.01.002.

[3] Zs. Szalay: "Structure and Architecture Problems of Autonomous Road Vehicle Testing and Validation", 15th Mini Conference on Vehicle System Dynamics, Identification and Anomalies – VSDIA 2016, 2016.

Please contact:

Domokos Esztergár-Kiss
Budapest University of Technology and Economics, Hungary
+36-1-463-1029
esztergar@mail.bme.hu

Adding Autonomous Features to a Production Electric Car

by Péter Gáspár, Tamás Szirányi, Levente Hajder, Alexandros Soumelidis, Zoltán Fazekas and Csaba Benedek (MTA SZTAKI)

A project was launched at the Institute – relying on interdepartmental synergies – with the intention of joining the autonomous vehicles R&D arena. In the frame of the project, demonstrations of the added capabilities – appreciably making use of multi-modal sensor data – in respect of path-planning, speed control, curb detection, road and lane following, road structure detection, road sign and traffic light detection and obstacle detection are planned in controlled traffic environments.

The development of intelligent – and particularly self-driving – cars has been a hot topic for at least two decades now. It ignited heated competition initially among various university and R&D teams, later among vehicle engineering enterprises and more recently among high-profile car manufacturers and IT-giants. An autonomous car relies on many functions and comprises numerous subsystems that support its autonomous capabilities. These include route-planning, path-planning, speed control, curb detection, road and lane detection and following, road structure detection, auto-steering, road sign and traffic light detection and recognition, obstacle detection and adaptive cruise control, vehicle overtaking, auto-parking [1]. Apart from the mechanical state of the car and the geometrical and mechanical properties of the actual road, it is important to know the geographical context and the kind of environment the car moves in to achieve vehicular autonomy.

To gain first-hand experience in this hot R&D arena and to implement some of the aforementioned functions, a project for adding autonomous vehicle features – together with associated automatic control – to a production electric car was launched by the Computer and Automation Research Institute, Budapest, Hungary in 2016. Two of the institute's research laboratories are involved in the project: one which deals with the theory and the application of automatic control (among various other fields in the area of vehicle dynamics and control), and one which is engaged in machine perception. These closely related overlapping research interests and experiences

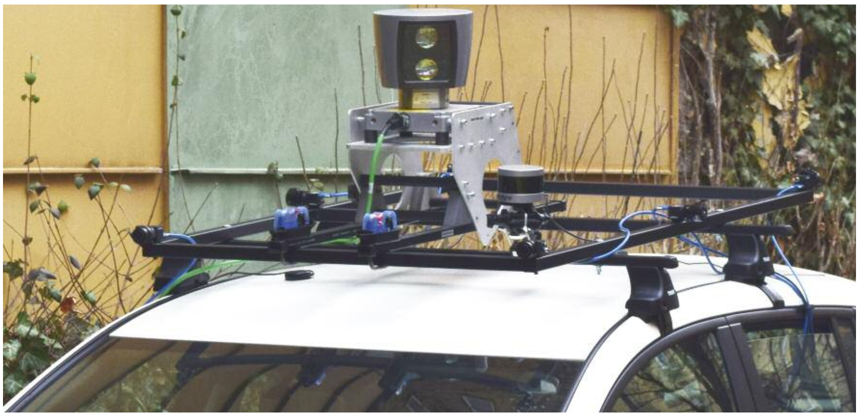

Figure 1: A host of image and range sensors are mounted onto the electric test car.

offer a great opportunity for R&D synergy.

A Nissan Leaf electric car was fitted up with a heavy-duty roof rack, which supports an array of sensors, namely six RGB cameras and two 3D scanning devices: a LIDAR with 16 and one with 64 beams. A further LIDAR with 16 beams is mounted onto the front of the car. These devices view and scan the spatial environment around the host car and explore the traffic and road conditions in the vicinity. The test car is also equipped with a high-precision navigation device – which includes an inertial measurement kit – and a professional vehicle-diagnostic device which can access data from the built-in vehicular sensors that are necessary for the autonomous control. Given that the car is fully electric and equipped with a host of computing and measurement facilities with considerable power consumption, provisions for recharging the batteries had to be made: a fast charge-point was installed in the institute's carport.

In order to obtain metric information from the image and the range sensors, they need to be calibrated to each other. However, the data coming from LIDARs and RGB cameras, respectively, are of different modalities: the former build up spatial point clouds, while the latter form colour image sequences. Recently, a fast and robust calibration method – applicable to either, or both of these modalities – was developed [2]. The method uses an easy-to-obtain fiducial (e.g., a cardboard box of uniform colour). This is an advantage over other methods as it allows the sensors to be transferred between vehicles when required. The sensor calibrations – required because of the different geometrical configurations that can be conveniently set up on different vehicles – can be carried out without much fuss. The idea is that the mentioned fiducial can be easily detected both in 2D and 3D spaces, thereby the calibration problem is transformed into one of 2D-3D registration. Two point clouds which have been registered using this approach are shown merged in Figure 2.

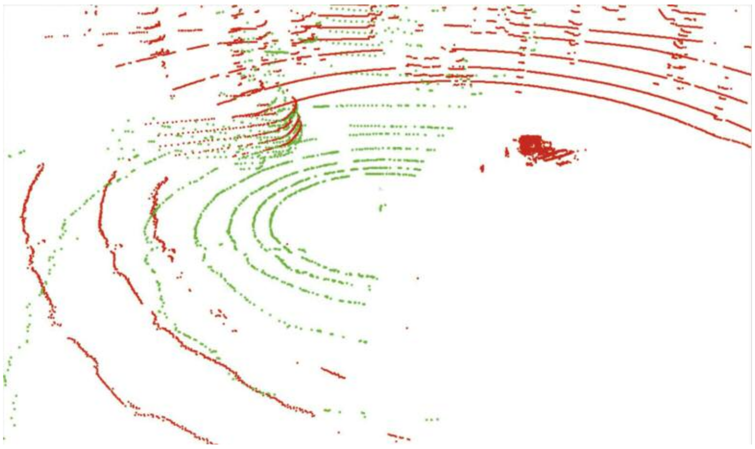

Figure 2: Two point clouds recorded simultaneously in street traffic with the two small LIDARs – mounted on the car shown in Figure 1 – are merged after registration. The green dots represent the point cloud taken with the front LIDAR (not visible in Figure 1), while the red dots are from the small (visible) LIDAR on the roof rack.

The data collection and recording of trajectory, video, point cloud and vehicular sensor data is performed in urban areas on public roads near Budapest. Presently, detection of cars, pedestrians, road structures and buildings are on the agenda [3]. The collected data is used in the design, simulation and testing of the algorithms and modules implementing the autonomous operation of the car. The test runs – testing the autonomous features and control developed for the project car – will initially take place in closed industrial areas and on test courses.

This kind of data can also be utilised in various mapping, surveying, measurement and monitoring applications in conjunction with roads (e.g., for assessing the size distribution of potholes, for measuring the effective width of snow/mud/water-covered roads), traffic (e.g., measuring the length and the composition of vehicle queues) and built environment (e.g., detecting changes in the road layout, detecting roadwork sites, broken guard-rails, skew traffic signs).

Link:
www.sztaki.hu/en/science/projects/lab-autonomous-vehicles

References:

[1] Z. Fazekas, P. Gáspár, Zs. Biró and R. Kovács (2014) Driver behaviour, truck motion and dangerous road locations – unfolding from emergency braking data. TRE, 65, 3-15.

[2] B. Gálai, B. Nagy and Cs. Benedek (2016) Cross-modal point cloud registration in the Hough space for mobile laser scanning data. IEEE ICPR, 3363-3368.

[3] D. Varga and T. Szirányi (2016) Detecting pedestrians in surveillance videos based on convolutional neural network and motion. EUSIPCO, 2161-2165.

Please contact:

Péter Gáspár, MTA SZTAKI, Hungary
+36 1 279 6171
peter.gaspar@sztaki.mta.hu

An Integrated Software Development Lifecycle for Intelligent Automotive Software: The W Model

by Fabio Falcini and Giuseppe Lami (ISTI-CNR)

Deep learning is becoming crucial to the development of automotive software for applications such as autonomous driving. The authors have devised a framework that supports a robust, disciplined development lifecycle for such software, and a comprehensive integration with traditional automotive software engineering.

Deep learning is a branch of machine learning based on a set of artificial neural network (ANN) algorithms that model high-level abstractions in input data by using a deep graph representation made of multiple processing layers [1].

The introduction of deep-learning technology on board cars is starting to have a significant effect on the automotive software engineering that needs to incorporate and harmonise the development of these technologies [2].

The software development process for on board automotive electronic control units (ECUs) is subject to proprietary OEM norms as well as several international standards. Among them, the most relevant and influential standards for deep learning are Automotive SPICE and ISO 26262. Needless to say, these standards are still far from addressing it with dedicated statements.

The Automotive SPICE standard – SPICE stands for “Software Process Improvement and Capability dEtermination” – provides a process framework that disciplines the software development activities. ISO 26262, titled ‘Road Vehicles – Functional Safety’, released in late 2011, targets safety-related development and its scope expectedly includes system, hardware and software engineering.

Both standards, as far as the software is concerned, rely conceptually on the traditional development lifecycle: the V-model. Also very relevant for deep learning is the ISO PAS ‘Safety of the Intended Functionality’ (ISO/TC 22 N356).

Because of its pervasive adoption and its holistic coverage of the automotive software development processes, Automotive SPICE standard can be recognised as the appropriate reference.

The software side of deep neural network (DNN) development is a highly iterative activity composed by a stream of steps in an end-to-end fashion, as shown in Figure 1.

Compared with traditional approaches, the deep learning development process needs the support of empirical design choices driven by heuristics. Development often starts from well-known learning algorithms, which have been proven effective in comparable problems or domains [3].

Automotive software engineering, while welcoming innovation and outstanding functional performances, remains strict in its request for a robust and predictable development cycle. For that reason, it is important to place deep learning in more controlled V-model perspective to address a lengthy list of challenges, such as requirements criteria for training, validation and test data sets, and much more.

The introduction of a more structured conception of the deep learning lifecycle is instrumental to reach a controlled development approach that cannot be addressed by the mere functional benchmarking obtained with validation activities. It is essential to pursue both fundamental directions: high performance at the functional level and a high-quality development process.

However, the central role that is played by data in this context has stimulated the authors to introduce of a new development lifecycle model: the W model. To support it, we employ the term ‘programming by example’ to highlight the

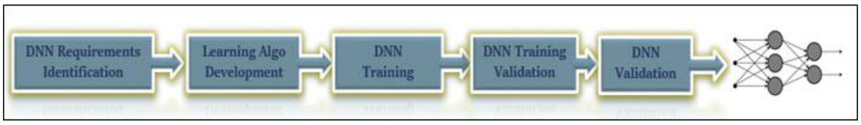

Figure 1: DNN development workflow.

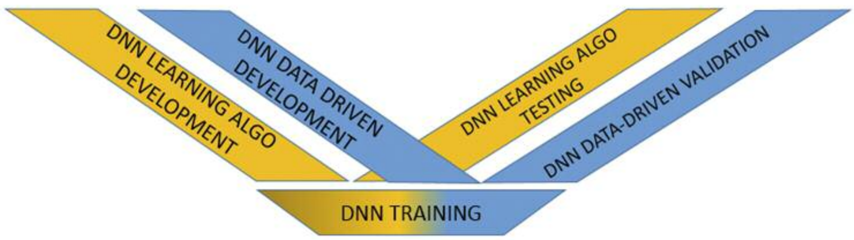
DNN LEARNING ALGO DEVELOPMENT
DNN DATA DRIVEN DEVELOPMENT
DNN LEARNING ALGO TESTING
DNN DATA-DRIVEN VALIDATION
DNN TRAINING

Figure 2: The W-model for deep learning.

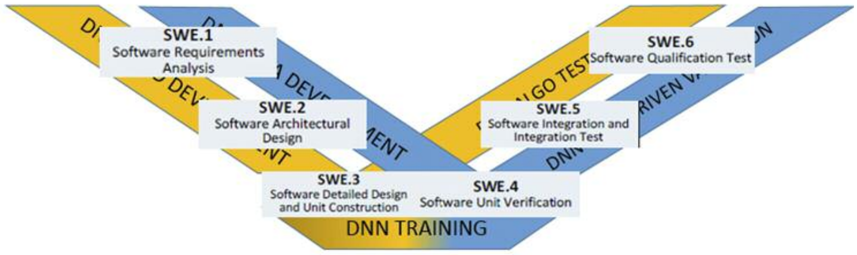
SWE.1 Software Requirements Analysis
SWE.2 Software Architectural Design
SWE.3 Software Detailed Design and Unit Construction
SWE.4 Software Unit Verification
SWE.5 Software Integration and Integration Test
SWE.6 Software Qualification Test
DNN TRAINING

Figure 3: The W-model in Automotive SPICE 3.0 perspective.

importance of data in developing systems based on deep-learning technology.

The deep learning W-model is a framework lifecycle that conceptually integrates a V model for data development in the standard V perspective (Figure 2).

This lifecycle model acknowledges that both software development and data development drive deep learning. The design and creation of training, validation and test datasets, together with their exploitation, are crucial development phases because the DNN's functional behaviour is the combined result of its architectural structure and its automatic adaptation through training. By definition, deep learning moves away from feature engineering. This aspect makes the W model an appropriate, useful representation of this sophisticated paradigm.

By placing the software part of the V-model of ASPICE 3.0 on top of the coined W-model, the following diagram hints at the integration of deep learning along the process requirements of the Automotive SPICE model (Figure 3). The same approach may apply for ISO 26262.

As deep learning ushers in radical changes to automotive software development (characterised by stringent requirements in terms of rigor, control and compliance with standards), the W model is a promising basis for the comprehensive integration of deep learning with traditional automotive software engineering.

References:

[1] S. Haykin: “Neural Networks and Learning Machines”, Prentice-Hall, 2009.

[2] F. Falcini, G. Lami: “Deep Learning in Automotive Software”, IEEE Software Vol. 34 No. 3, May-June 2017.

[3] J. Schmidhuber: “Deep learning in neural networks: An overview”, Neural Networks Vol 61, 85-117, Jan 2015.

Please contact:
Giuseppe Lami, Fabio Falcini
ISTI-CNR, Italy
giuseppe.lami@isti.cnr.it
fabio.falcini@isti.cnr.it

Reliable Vehicular-to-Vehicular Connectivity

by Lisa Kristiana, Corinna Schmitt, and Burkhard Stiller (University of Zurich)

Vehicular-to-vehicular (V2V) communications requires a reliable information interchange system. However, in a three-dimensional environment, inevitable obstructions (e.g., road topology and buildings) need to be considered in order to design data forwarding schemes. As one approach, Vehicle-to-Vehicle Urban Network (V2VUNet) introduces Vertical Relative Angle (VRA) as one of the significant factors in the case of a three-dimensional environment. VRA locates a participating vehicle's coordinates more precisely, when its position cannot be calculated based only on distance. Therefore, the increased location precision leads to several forwarding algorithms (e.g., prediction mobility and transmission area efficiency), which are expected to overcome frequent topology changes and re-routing.

The demand for reliable Vehicular-to-vehicular (V2V) communication increases on a daily basis [1]. V2V communication, as part of a Vehicular Ad hoc Network (VANET), offers advantages for safety (e.g., automatic braking system and traffic accident information) and non-safety applications (e.g., business and entertainment). V2V communication becomes interesting due to its flexibility to connect to other participating vehicles without any Road Side Unit (RSU) infrastructure. Additionally, V2V is a promising approach to extend the scalability issue due to high mobility behaviour and complex city environments.

Because of the high mobility behaviour of V2V, frequent disconnection of communication occurs due to overpasses, tunnels, and other obstructions leading to unstable routing issues for transmitted packets. Evaluating this unstable routing can be done by investigating short lifetime connections, low packet delivery ratio (PDR), and end-to-end (e2e) delays that are all highly influenced by the road topology. Therefore, the road topology impacts connection and re-connection establishment between vehicles. The road topology in a two-dimensional scenario, i.e., a road with intersections, and in a three-dimensional scenario, i.e., a road with overpasses, requires special calculations.

Vehicular-to-Vehicular Urban Network

Existing urban environments challenge the connectivity for V2V. Overcoming this challenge, the developed Vehicle-to-Vehicle Urban Network (V2VUNet) concept offers an optimisation for transmission between vehicles by selecting (1) the proper message route and (2) the best relay candidate in the network. Instead of using a flooding mechanism in route path establishment, the area of route request transmission is restricted by measuring the relative angles [2]. The Horizontal Relative Angle (HRA) is applied when the forwarding mechanism occurs in a two-dimensional environment and the Vertical Relative Angle (VRA) is applied in a three-dimensional environment to restrict the area for possible next relay candidates. This area restriction algorithm minimises the required time to build a complete route, which frequently changes, by broadcasting the route request to the closest vehicle (the relay candidate). Figure 1a shows the area restriction algorithms. The resulting number of relay candidates is now limited but still the best one has not been selected. Thus, the predictive forwarding al gorithm shown in Figure 1b is performed. It performs a more precise calculation with additional weight values θ_x and θ_z to select the best relay candidate. As an overall result the V2VUNet concept allows the lowest value HRA and VRA to be determined to find the best relay candidate by restricting the area of candidate

Algorithm 1

```
$S \leftarrow$ sender node
$R \leftarrow$ sender node
$I$ all neighboring nodes of $S$
$\theta_{x,max} \leftarrow$ maximum boundary of the horizontal angle
$\theta_{z,min} \leftarrow$ minimum boundary of the vertical angle
$\theta_{z,max} \leftarrow$ maximum boundary of the vertical angle
$\theta_x \leftarrow$ horizontal angle made by $n$ to $s$
$\theta_z \leftarrow$ vertical angle made by $n$ to $s$
$i_{filtered} \leftarrow$ only $i$ that is within $[-\theta_{x,max}, \theta_{x,max}]$ and $[-\theta_{z,min}, \theta_{z,max}]$
$d \leftarrow$ distance from $i_{filtered}$ to $R$, nexthop $\leftarrow$ $\arg_{min}$ $(d)$
```

(a) Area Restriction Algorithm

Algorithm 2

```
$s \leftarrow$ sender node, at position of $p_s$ and orientation of $v_s$
$i \leftarrow$ all neighboring nodes of $s$, at position of $p_i$ and orientation of $v_i$
$\theta_{solid} \leftarrow$ threshold of the solid angle for all $i$
$v_{si} = |p_s - p_i|$
$\theta_{si} = \text{atan}(||v_s \times v_{si}||, ||v_s . v_{si}||)$
$i_{filtered} \leftarrow i$ with $\theta_{si}$ within $[-\theta_{Solid}, \theta_{solid}]$
$d \leftarrow$ distance from $i_{filtered}$ to $R$
next hop $\leftarrow$ $\arg_{min}$ $(d)$
```

(b) Predictive Forwarding Algorithm

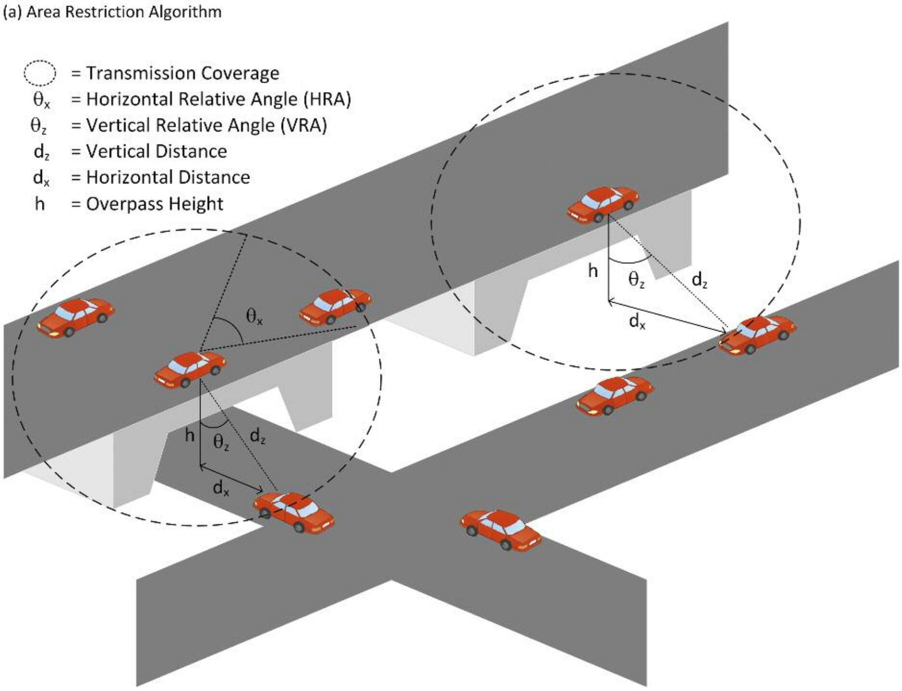

◌ = Transmission Coverage
θ_x = Horizontal Relative Angle (HRA)
θ_z = Vertical Relative Angle (VRA)
d_z = Vertical Distance
d_x = Horizontal Distance
h = Overpass Height
θ_x
h
θ_z
d_z
d_x
h
θ_z
d_z
d_x

Figure 1: V2VUNet concept [2].

searching and by maintaining the direction homogeneity assuming a scenario as shown in Figure 1c.

Conclusions

The V2VUNet concept supports two algorithms for link path stability and successful transmission, namely area restriction and predictive forwarding, including HRA or VRA as the weight values implemented in the forwarding decision. The goal of the proposed concept is to locate the best next relay candidate in the network to forward the message.

V2VUNet was implemented in two- and three-dimensional environments with obstructions. The HRA and VRA are introduced to distinguish the ‘real’ transmission distance between two vehicles on different road levels. The area restriction and predictive forwarding algorithms emphasize the height of the road topology and direction homogeneity. Current evaluation showed V2VUNet overcomes the frequent disconnection and three-dimensional road topology issues and improves the transmission ratio by around 25% compared to the general V2V transmission rate that is less than 50% [3].

Future work foresees investigations for data dissemination for variable data rates and higher data size. The variable data rate will be tested to evaluate the link lifetime when the vehicle’s direction is homogeneous, meaning the sender and receiver have the same driving direction. The higher data size transmission will be applied to obtain a non-safety application, since such application demands a massive data scale.

This work was supported in part by the Ministry of Research, Technology and Higher Education of the Republic of Indonesia.

References:

[1] H. Hartenstein, K. Laberteaux (Eds.): “VANET Vehicular Applications and Inter-Networking Technologies”, Vol. 1, John Wiley & Sons, ISBN: 978-0-470-74056-9.

[2] L. Kristiana, C. Schmitt, B. Stiller: “V2VUNet – A Filtering Out Concept For Packet Forwarding Decision in Three-dimensional Inter Vehicular Communication Scenarios”, IEEE 27th Annual International Symposium Personal, Indoor, and Mobile Radio Communications (PIMRC), pp 1-6, 2016, .

[3] L. Kristiana, C. Schmitt, B. Stiller: “Evaluation of Inter-vehicle Connectivity in Three-dimensional Cases”, IEEE Wireless Days, March 2017, pp 1-4.

Please contact:

Lisa Kristiana, Corinna Schmitt, Burkhard Stiller
Universität Zürich, Switzerland
+41 44 635 -7586/ -7585/ -6710
kristiana@ifi.uzh.ch, schmitt@ifi.uzh.ch, stiller@ifi.uzh.ch

Why is Autonomous Localisation Required in Lateral Cooperative Control?

by Tom van der Sande (Eindhoven University of Technology), Jeroen Ploeg (TNO) and Henk Nijmeijer (Eindhoven University of Technology)

What is the link between cooperative automated vehicles and autonomous vehicles? At the Eindhoven University of Technology, we recently started the i-CAVE (integrated cooperative automated vehicles, STW14893) project, focusing on the development of dual-mode operation of cooperative automated and autonomous vehicles, which will provide an answer to this question. Here we discuss the state of the art of cooperative automated and autonomous vehicles using the guidelines as provided by the Society of Automotive Engineers (SAE) and show that contemporary systems are far from being fully autonomous.

Although we have been driving cars for more than a century, human error remains the primary cause of accidents. One way of solving this problem is to automate the task of driving. This can be done on five levels, as specified in SAE-J3016. At level 1, the vehicle automates the longitudinal or the lateral motion, whereas the driver is always required to pay attention. Level 2 has similar characteristics, except that both longitudinal and lateral motion are automated. With level 3, the automation system will monitor the environment and warn the driver in case of danger, but the driver is used as fall-back. Beyond level 3, the system performs fall-back functionality.

The question is how to maintain the driver at a sufficient awareness level to perform the required monitoring and fall-back task at lower automation levels: level 3, in particular, might give a false sense of safety. Consequently, level 3 automation should only be done within a closed perimeter, designed exclusively for automated driving and careful driver instruction.

An example of a commercially available system functioning at level 1 is adaptive cruise control (ACC). We show that this system has only limited performance when encountering busy traffic [1] because large errors of the desired distance might occur. By incorporating wireless communication, referred to as cooperative adaptive cruise control (CACC), this issue can be solved. Vehicles equipped with CACC (shown in figure) communicate desired acceleration, in contrast to ACC, which only uses the relative position and velocity. In an ideal situation, with no communication delay and identical vehicles, the feedback system of the ACC-controller is now obsolete. As such, CACC is a typical example of cooperative automation, characterised by relying on information obtained via wireless communication.

The assumption that in a perfect world no additional information is required vanishes when considering lateral cooperative manoeuvring. To visualise this, consider a vehicle that is moving on a plane, giving it a longitudinal, lateral and yaw degree-of-freedom. For it to move in a lateral direction, it should have a longitudinal velocity and yaw-rate, indicating the coupled and non-linear dynamics and the resulting control problem. As long as the vehicle is moving, we show that input-output linearization by state-feedback decouples these dynamics. Consequently, a vehicle-following controller is developed by using PD-control [2]. This gives rise to a particular problem: corner cutting. In case of lateral vehicle automation, it may cause the follower vehicle to leave its lane. Thus, the vehicle can no longer rely solely on the information received from the preceding vehicle, but it needs to have self-awareness about its position and surroundings. This bridges the gap between cooperative and autonomous vehicles [3].

The analysis above shows that for combined longitudinal and lateral CACC, the vehicle needs to be aware of its own location with respect to infrastructure as the default controller promotes corner cutting. To this end, on board sensors such as radar, scanning lasers or cameras are employed.

In conclusion it is worth asking whether we can design a successful automated vehicle without communication.

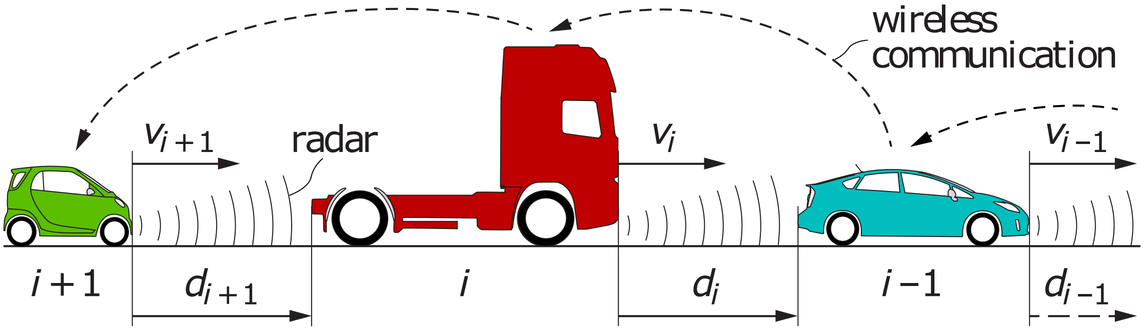

Figure 1: A homogeneous vehicle platoon. Our research shows that contemporary systems are still far from being fully autonomous, as they will benefit from information that can only be obtained through communication. Picture: J.PLoeg [1].

Obviously full autonomy must be a fallback mechanism of automation in combination with communication, however to be successful as an individual vehicle, do we need to know more about our environment than our on board sensors tell us? Classical control theory teaches us there is a large benefit in using information in the control system that can only be obtained through communication.

Link:
www.stw.nl/nl/content/p14-18-i-cave-integrated-cooperative-automated-vehicle

References:

[1] J. Ploeg, N. van de Wouw, H. Nijmeijer: "Lp string stability of cascaded systems: Application to vehicle platooning", IEEE Transactions on Control Systems Technology.

[2] A. Bayuwindra, Ø. Aakre, J. Ploeg, H. Nijmeijer: "Combined lateral and longitudinal CACC for a unicycle-type platoon", 2016 IEEE Intelligent Vehicles Symposium (IV).

[3] T. van der Sande, H. Nijmeijer: "From cooperative to autonomous vehicles": Springer LNCS, 2017. DOI: 10.1007/978-3-319-55372-6

Please contact:
Tom van der Sande
TU/e, The Netherlands
+31402475730
t.p.j.v.d.sande@tue.nl

Safe Human-Inspired Adaptive Cruise Control for Autonomous Vehicles

by Alessio Iovine, Elena De Santis, Maria Domenica Di Benedetto (University of L'Aquila) and Rafael Wisniewski (Aalborg University)

The safety issues associated with autonomous vehicles necessitate more robust and reliable solutions for Adaptive Cruise Control (ACC). A human-inspired hybrid model with the design goal of replacing and imitating the behaviour of a human driver is being developed, ensuring an appropriate safety level while respecting comfort.

The sixty year old research field of traffic control has recently been addressing the ultra-modern subject of autonomous vehicle control. The main goal of traffic control is to optimise traffic management with respect to a number of variables, including congestion, emissions, travel time reduction and safety. To this end, Adaptive Cruise Control (ACC) systems were envisioned. Once introduced, these tools need to conform to normal traffic dynamics, consequently they will need to resemble a human driver's behaviour.

Properties of hybrid systems, which encompass both continuous and discrete time dynamics evolution, allow for implementation of the complex modelling of the human way of driving. Our approach considers a number of control laws, one for each different situation, and embeds them in a unique model obtaining a mesoscopic hybrid model, i.e., a microscopic hybrid model that takes into account macroscopic parameters [1]. Indeed, environmental information is used to better represent human characteristics, defining a model closer to reality. The macroscopic information can either be provided by a centralised traffic supervisor, or it can be gathered, elaborated and transmitted by the vehicles themselves, when they are connected and can exchange information.

The resulting ACC system processes information about other vehicles and takes decisions about braking or throttle actions on the basis of a human-inspired model, with the purpose of safely controlling single vehicle dynamics in relation to a car in front. The framework considered in [1] is deterministic, and

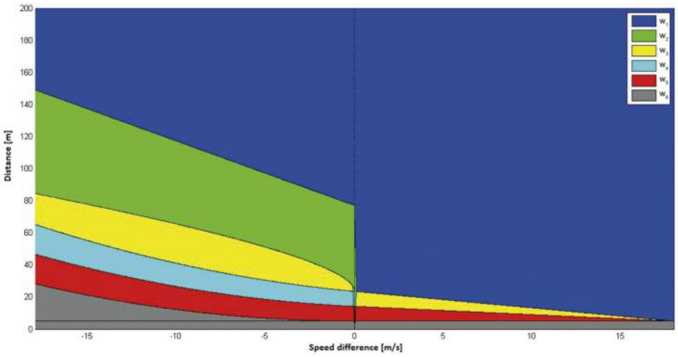

Figure 1: The interaction zone split in different regions.

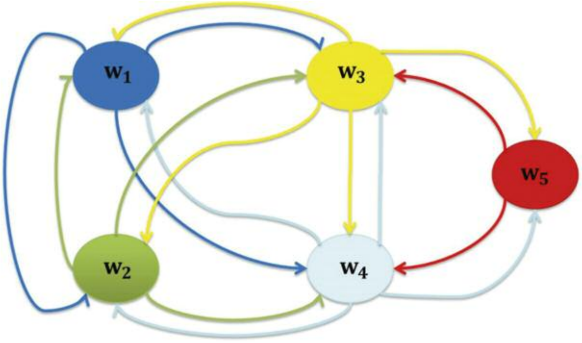

Figure 2: The possible discrete transitions resulting in the considered hybrid system.

the next step is to consider stochastic processes, giving a definition of safety with respect to a probability measure, and using new techniques producing barrier functions for a risk acceptance probability [2]. The objective is to define the minimum distance between two vehicles which allows collision avoidance with a given desired probability. The barrier functions calculated in [1] will then change, while the hybrid system structure imitating the human driver is preserved.

Our research is also investigating bio-inspired robots, with the aim of imitating and reproducing the natural behaviour of a living being. This allows us to include a human driver in the control loop, reproducing the adaptability to various conditions and comfort, while maintaining the full controllability of the system and avoiding human weaknesses, such as mistakes or distractions.

The first purpose of the model is to provide information to the human driver about which risk level applies to a given situation, which will enable the driver to select the appropriate balance between safety and performance. The autonomous vehicle will react according to the chosen setup.

The model has a range of potential applications. First, it makes it possible to fix a probability safety minimum range that should be respected by autonomous vehicles; this would lead to improvements both from a single vehicle point of view and for traffic flow (in fact, the presence of a single dangerous vehicle on the road has a negative impact on the entire flow). Furthermore, the considered model could reduce legal disputes between owners of autonomous vehicles and the car manufacturing industry: vehicles could have a completely safe setting as default with the option for the owner to modify it. In the case of an accident, the responsibility is in the hands of the driver who chose the risk level. It must be noted that in a competitive world where autonomous vehicle sales may depend in part on their potential to engage in truly dangerous situations, it is essential to set a minimum legal safety range to avoid a race among vehicle manufacturers.

Link:
http://safecop.deib.polimi.it/

References:

[1] A. Iovine, et al.: “Safe human-inspired mesoscopic hybrid automaton for autonomous vehicles”, Nonlinear Analysis: Hybrid Systems, http://dx.doi.org/10.1016/j.nahs.2016.08.008

[2] C. Sloth, R. Wisniewski: “Safety Analysis of Stochastic Dynamical Systems”, IFAC-PapersOnLine, Volume 48, Issue 27, 2015, Pages 62-67, ISSN 2405-8963, http://dx.doi.org/10.1016/j.ifacol.2015.11.153

Please contact:

Alessio Iovine, Elena De Santis, Maria Domenica Di Benedetto
Center of Excellence DEWS, University of L’Aquila, Italy
alessio.iovine@univaq.it,
elena.desantis@univaq.it,
mariadomenica.dibenedetto@univaq.it

Rafael Wisniewski
Aalborg University, Denmark
raf@es.aau.dk

Hello Human, can you read my mind?

by Jonas Andersson, Azra Habibovic, Maria Klingegård, Cristofer Englund and Victor Malmsten-Lundgren (RISE Viktoria)

For safety reasons, autonomous vehicles should communicate their intent rather than explicitly invite people to act. At RISE Viktoria in Sweden, we believe this simple design principle will impact how autonomous vehicles are experienced in the future.

Autonomous vehicles are now a reality, and several cities around the globe have seen their development in different shapes and contexts. A question is how these vehicles will interact with other road users in daily traffic, especially with vulnerable road users such as pedestrians and cyclists? [1]. According to the World Health Organisation, approximately 325,000 pedestrians and cyclists were killed in traffic worldwide in 2015, making safety of this group of road users a priority. Vehicle automation is expected to reduce such fatalities in the future.

As vehicles start driving on their own and drivers become riders who can do other things in the vehicle, new types of interactions are likely to emerge that are unrelated to the driving task. Thus, driver behaviour is no longer necessarily representing the actions of the vehicle by, for example, eye contact and gestures. A series of studies conducted by RISE Viktoria, RISE Interactive, Autoliv, Volvo Cars, Volvo Group and Scania show that pedestrians experience discomfort and confusion and feel unsafe when they encounter vehicles not having an active driver behind the steering wheel [2]. To compensate for the missing signals in driver behaviour, it was also found that pedestrians may require supporting information about the status of autonomous vehicles, including current operation mode (manual vs. autonomous) and imminent behaviour (e.g., ‘about to start’, ‘about to stop’ and ‘resting’).

Based on these findings, we have designed an example interface concept named AVIP [L1] to illustrate a potential universal, visual communication interface for autonomous vehicles. The AVIP interface consists of an outward-facing LED light strip affixed to the top of the vehicle wind shield that uses distinct patterns of light to communicate the vehicle’s status and intent to pedestrians [3]. In simple terms, it could be described as a type of “smart” frontal brake light that informs surrounding road users about what the autonomous vehicle is about to do, without explicitly telling the road users what they should

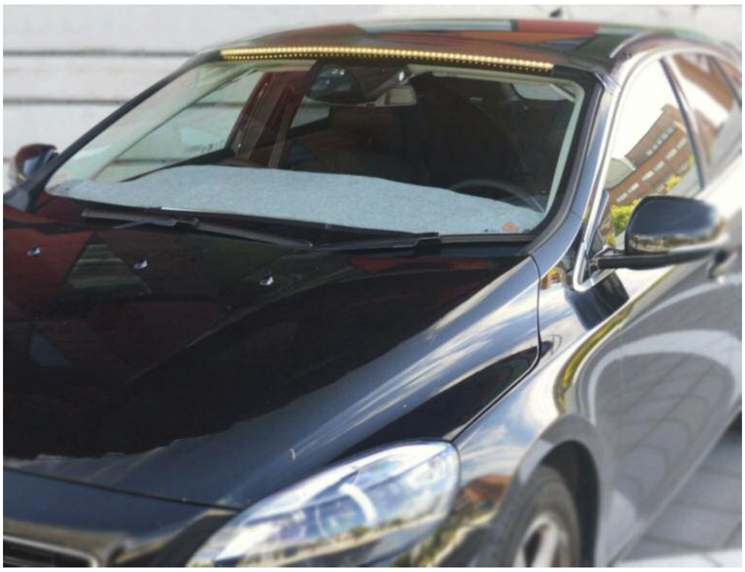

Figure 1: The AVIP concept. The changing LED light communicates the intention of an autonomous vehicle.

do and when. This is different to other interface concepts suggested by other players in the area, where the vehicle explicitly invites you to go ahead, for example with a text message, a picture or a projected zebra crossing.

We believe that inviting other road users to act may create false expectations of the surrounding traffic that the autonomous vehicles cannot necessarily account for. Instead, a design principle based on communicating intent rather than explicitly inviting people to act could avoid ambiguities due to a mismatch between the vehicle's invitation and the surrounding traffic. Regardless of what the communication interface looks like, we suggest that this design principle should be widely adopted, if not standardised, as autonomous vehicle technology matures.

Link:
[L1] www.viktoria.se/projects/avip-automated-vehicle-interaction-principles

References:

[1] M. Nilsson, T. Ziemke, S. Thill: "Action and intention recognition in human interaction with autonomous vehicles", in "Experiencing Autonomous Vehicles: Crossing the Boundaries between a Drive and a Ride" workshop in conjunction with CHI2015 (2015).

[2] V. M. Lundgren, et al.: "Will There Be New Communication Needs When Introducing Automated Vehicles to the Urban Context?", Advances in Human Aspects of Transportation, Springer, 485-497, 2017.

Please contact:

Jonas Andersson
RISE Viktoria, Research Institutes of Sweden
+46 (0)73 925 80 89
jonas.andersson@ri.se

Cross-Domain Fertilisation in the Evolution towards Autonomous Vehicles

by Christophe Ponsard, Philippe Massonet and Gautier Dallons (CETIC)

There is a strong move towards assisting the driver, and even relieving the driver from driving duties, in various modes of transportation, including autonomous car, train and aircraft operation. Successive levels of automation are progressively shifting responsibility from the driver to the vehicle. We present a brief comparative analysis of automation levels in the automotive, railway and aeronautic domains with a focus on how to keep the human in the loop both at design and run-time.

Mobility is a key dimension of our daily lives both for work and personal duties. Ensuring safe, rapid, predictable and affordable transportation is quite challenging because our transportation networks are under high load, especially at peak hours. The level of performance that can be achieved is often limited by the human driver; the inherent constraints being reaction time (resulting in triggering jams), reliability (90% of car accidents are a result of human error) and also economic cost. Airlines, railway and car manufacturers have consequently worked to assist the driver by reducing the driver's work and are even considering the ultimate goal of going driverless.

Different levels of automation have been defined in each transportation domain. Table 1 gives a comparative overview of those levels with related automation features and the respective split of responsibility between the system and human driver/pilot. The automotive SAE classification [url1] was used as reference framework to align railway grades of operation (GoA) [url2] and aircraft levels of automation [url3]. The lowest level is fully manual driving, although even this level generally includes some safety protection systems. Vehicles in the next level up provide assistance in specific or time-bounded tasks. While cars are currently only reaching partial automation, railways and airplanes have already achieved much higher levels, e.g., fully autonomous metro lines.

The automotive industry is now considering the critical transition to a level where the car is able to drive itself so the driver has no need to steer, brake, accelerate or watch the road. However, as the system cannot cope with all situations, the driver must be ready to

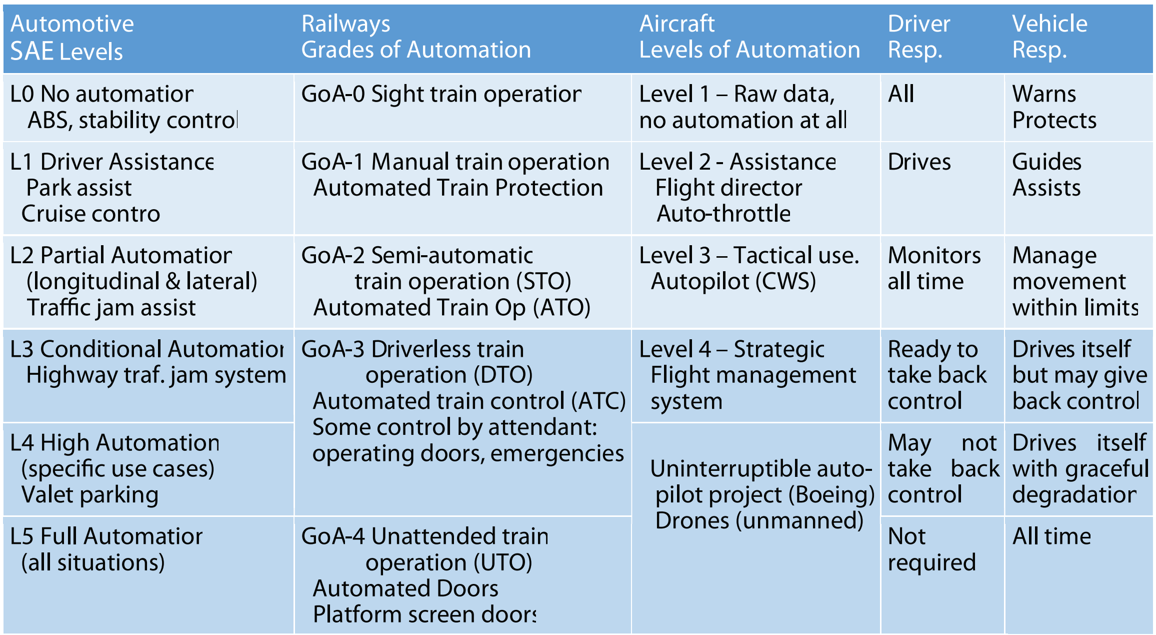

Automotive SAE Levels	Railways Grades of Automation	Aircraft Levels of Automation	Driver Resp.	Vehicle Resp.
L0 No automation ABS, stability control	GoA-0 Sight train operation	Level 1 – Raw data, no automation at all	All	Warns Protects
L1 Driver Assistance Park assist Cruise contro	GoA-1 Manual train operation Automated Train Protection	Level 2 - Assistance Flight director Auto-throttle	Drives	Guides Assists
L2 Partial Automation (longitudinal & lateral) Traffic jam assist	GoA-2 Semi-automatic train operation (STO) Automated Train Op (ATO)	Level 3 – Tactical use. Autopilot (CWS)	Monitors all time	Manage movement within limits
L3 Conditional Automatior Highway traf. jam system	GoA-3 Driverless train operation (DTO) Automated train control (ATC) Some control by attendant: operating doors, emergencies	Level 4 – Strategic Flight management system	Ready to take back control	Drives itself but may give back control
L4 High Automation (specific use cases) Valet parking		Uninterruptible auto-pilot project (Boeing) Drones (unmanned)	May not take back control	Drives itself with graceful degradatior
L5 Full Automatior (all situations)	GoA-4 Unattended train operation (UTO) Automated Doors Platform screen door:		Not required	All time

Table 1: Comparison of automation levels in automotive, railways and aeronautics.

resume operation when instructed. This is quite challenging because several obstacles have to be carefully addressed [1]. Our work on this topic relies on requirements engineering techniques [2], based both on characteristics from the human driver and on the transposition of known problems in other domains. Some interesting issues to address are the following:

- Situation awareness: the system needs to make sure the driver has a correct mental picture of the environment when handing back control. This is more difficult to achieve when driver focus has switched to non-driving duties.

- Human reaction abilities: raising the alarm at short notice is dangerous because the driver will not be able to fully analyse the situation or even go into panic. The system needs to anticipate known events (e.g., motorway exit in five kilometres), pre-crash scenarios (e.g., suspicious car behaviour ahead) and also report about degradation of its performance (e.g., due to weather or road conditions). This also enables the driver to increase their situation awareness as required.

- Warning annoyance: driver trust in the system can rapidly decrease if the system sounds alarms at inappropriate times or generates too many warnings.

- Task inversion: the driver might switch to a mode where his focus is only on monitoring alarms and not paying attention to the real world situation.

Based on such an analysis, specific strategies can be designed and experimented. The whole system should not be seen as static but rather an evolving cooperation between the driver and the vehicle with learning occurring on both sides. In this regard, machine learning techniques can play an important role for making sure the driver and the system are operating optimally together [3]. With the advent of connected car, this driving experience (especially problematic scenarios) will easily be reported, analysed and enhanced. Our future work will focus more specifically on the railway domain in the context of INOGRAMS project objectives related to automated train operation, especially for high efficiency [4].

Links:

[L1] www.sae.org/misc/pdfs/automated_driving.pdf
[L2] tiny.cc/uitp-goa
[L3] tiny.cc/nbaa-flight-automation
[L4] www.cetic.be/INOGRAMS-2104

References:

[1] S. M. Casner, et al.: “The Challenges of Partially Automated Driving”, Communications of the ACM, 59(5):70-77, 2016

[2] A. van Lamsweerde: “Requirements Engineering: From System Goals to UML Models to Software Specifications”, Wiley, 2009.

[3] P. Koopman, M. Wagner: “Challenges in Autonomous Vehicle Testing and Validation”, SAE Int. J. Trans. Safety 4(1):15-24, 2016.

Please contact:

Christophe Ponsard
CETIC, Belgium
+32 472 56 90 99
christophe.ponsard@cetic.be

Will the Driver Seat Ever Be Empty?

by Thierry Fraichard (Inria)

Self-driving vehicles are here and they already cause accidents. Now, should road safety be considered in a trial and error perspective or should it be addressed in a formal way? The latter option is at the heart of our research.

Self-driving technologies have improved to the point that major industrial players are now investing in, developing and testing self-driving vehicles in various countries. It is touted that fleets of self-driving vehicles will be operational within five years. Besides economic reasons, a major incentive behind self-driving vehicles is safety. Road traffic deaths across the world reached 1.25 million in 2015 and since the driver is responsible for most crashes, it seems natural to strive to design self-driving vehicles. Obviously, an autonomous vehicle is safe if it avoids collision. Now, collisions happen for different reasons: hardware or software failures, perceptual errors that result in an incorrect understanding of the situation (in May 2016, a self-driving vehicle crashed killing its driver because its camera failed to recognise a white truck against a bright sky), and reasoning errors, i.e., a wrong decision is made (in February 2016, a self-driving vehicle was for the first time responsible for a crash because of a wrong decision). Our research focuses on reasoning errors. At a fundamental level, a collision will happen as soon as a vehicle reaches an inevitable collision state (ICS), i.e., a state for which, no matter what the vehicle does, a collision will eventually occur [1].

The insights on collision avoidance resulting from our investigation of the ICS concept are expressed succinctly by the following quote: ‘One has a limited time only to make a motion decision, one has to globally reason about the future evolution of the environment and do so over an appropriate time horizon.’

As abstract and general as this quote may seem, it implicitly contains motion safety laws whose violation is likely to cause a collision [2]. ICS were initially investigated with the aim of designing motion strategies for which collision avoidance could be formally guaranteed. Assuming the availability of an accurate model of the future up to the appropriate time horizon, we were able to design motion strategies with guaranteed absolute motion safety i.e., no collision ever takes place.

In the real world, things are not so rosy since accurate information about the future evolution of the environment is not available. The choice then is between conservative or probabilistic models of the future. In conservative models, each obstacle is assigned its reachable set, i.e., the set of positions it can potentially occupy in the future. Conservative models solve the problem of the discrepancy between the predicted future and the actual future. In theory, they could allow for guaranteed motion safety. In practice however, the growth of the reachable sets as time passes by is such that every position becomes potentially occupied by an obstacle and it is then impossible to find a safe solution. In probabilistic models, the position of an obstacle at any given time is represented by an occupancy probability density function. Such models are well suited to represent the uncertainty that prevails in the real world. However, as sound as the probabilistic framework is, it cannot provide motion safety guarantees that can be established formally, minimising the collision risk is the only thing that can be done then. Today, most, if not all, self-driving vehicles rely upon probabilistic modelling and reasoning to drive themselves. The current paradigm is to test self-driving vehicles in vivo and, should a problem occur, to patch the self-driving system accordingly. In an effort to improve the situation and to provide provable motion safety guarantees, we would like to advocate an alternative approach that can be summarised by the following motto: “Better guarantee less than guarantee nothing.”

The idea is to settle for levels of motion safety that are weaker than absolute motion safety but that can be guaranteed. One example of such a weaker level of motion safety guarantees that, if a collision must take place, the self-driving vehicle will be at rest. This motion safety level has been dubbed

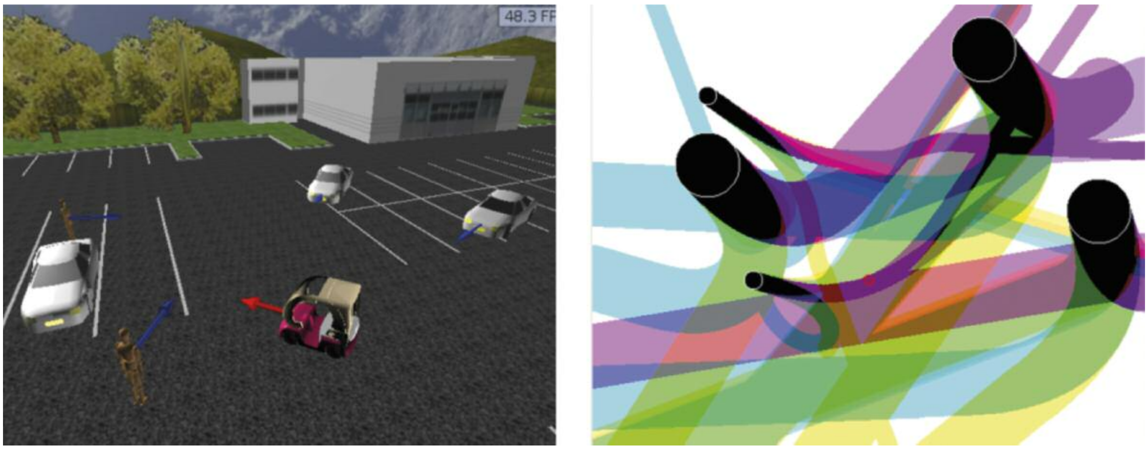

Figure 1: A self-driving vehicle among fixed and moving obstacles (left); 2D slice of the 5D state space of the self-driving vehicle, the black areas are the corresponding inevitable collision states that must be avoided (right).

passive motion safety and it has been used in several autonomous vehicles. Passive motion safety is interesting for two reasons: (i) it allows provision of at least one form of motion safety guarantee in challenging scenarios (limited field-of-view for the robot, complete lack of knowledge about the future behaviour of the moving obstacles [3]), and (ii) if every moving obstacle in the environment enforces it then no collision will take place at all.

We are currently exploring more sophisticated levels of motion safety. We are also studying the relationship between the perceptual capabilities of the self-driving vehicles and the levels of motion safety that can be achieved. The long-term goal of our research is to investigate if and how current self-driving technologies can be improved to the point that the human driver can safely be removed from the driving loop altogether, paving the way to truly self-driving vehicles whose motion safety can be formally characterised and guaranteed.

References:

[1] T. Fraichard, H. Asama: “Inevitable collision states. a step towards safer robots?”, Advanced Robotics, vol. 18, no. 10, 2004.

[2] T. Fraichard: “Will the driver seat ever be empty?” INRIA, Research Report RR-8493, Mar. 2014. [Online]. http://hal.inria.fr/hal-00965176

[3] S. Bouraine, T. Fraichard, H. Salhi: “Provably safe navigation for mobile robots with limited field-of-views in dynamic environments," Autonomous Robots, vol. 32, no. 3, Apr. 2012.

Please contact:

Thierry Fraichard
Inria Grenoble Rhône-Alpes, France
thierry.fraichard@inria.fr

Intelligent Environmental Perception for Navigation and Operation of Autonomous Mobile Machines

by Robert Rößler, Thomas Kadiofsky, Wolfgang Pointner, Martin Humenberger and Christian Zinner (AIT)

Research at AIT on autonomous land vehicles is focusing on transport systems and mobile machines which operate in unstructured and heavily cluttered environments such as off-road areas. In such environments, conventional technologies used in the field of advanced driver assistance systems (ADAS) and highly automated cars show severe limitations in their applicability. Therefore, more general approaches for environmental perception have to be found to provide adequate information for planning and decision making. As a special challenge, our focus is on vehicles that actively change their environment, e.g., by cutting high plants or manipulating piles of pellet materials. To operate such machines autonomously, novel approaches for autonomous motion planning are required.

In our latest research activity, a tractor (type Steyr 6230 CVT) was provided by the Austrian Armed Forces for the purpose of automating agricultural tasks on special areas. In a first step, this mobile machine was prepared in close cooperation with the manufacturer CNH Austria to serve as a mobile research platform. An electrical single-point interface provides access to the various functions of the tractor via actuators and the CAN bus system, which enables a set of additional computers and dedicated software to take over control. The vehicle is equipped with a multi-modal sensor system combining stereo vision, laser scanning, RTK-GPS and IMU to address difficult and varying outdoor conditions. A dedicated AIT stereo vision system mounted behind the windshield at the front and another on the rear of the vehicle provide dense 3D data of the tractor’s environment (Figure 1).

An important step to further increase efficiency in farming is to develop completely driverless machines which would, for example, enable one operator to monitor and control multiple machines at once. Thus, AIT investigates important computer vision based technologies and methods that improve the abilities of future autonomous vehicles and machines for agricultural applications. Furthermore, research is driven by a specific use-case defined by the Austrian Armed Forces: cultivation (e.g., mulching, mowing) of firebreaks on military training areas endangered by explosive ordnances. The technology developed by AIT allows the operator to safely monitor and control the semiautonomous machine from a safe distance with a high level of situational awareness.

The sensed 3D data of the stereo camera systems [1] and the laser scanner is continuously fused into an ‘elevation map’ – a 2.5D representation of the terrain. As the vehicle moves, the reconstruction of the vehicle’s surroundings becomes more and more complete. An essential task is estimating the ego-motion and the pose of the vehicle. The trajectories of the stereo visual odometries, the orientation of the inertial measurement unit and the position of the RTK-GPS are fused with an extended Kalman filter. The filter models the kinematics of the vehicle and computes a precise estimate of the pose [2]. Furthermore, the geometry of the elevation map is analysed in order to obtain a traversability map, which denotes drivable and non-drivable areas as well as obstacles around the vehicle. The map is organised using a tile based approach, which allows storing, loading and updating large scale maps. In that way the georeferenced mapped areas can also be accessed via a geographic information system.

The real-time traversability map is used by the path planning module to calculate a collision free trajectory along a predefined mission path. Dynamically appearing obstacles are avoided by calculating a local bypass route using a sampling-based motion planning algo-

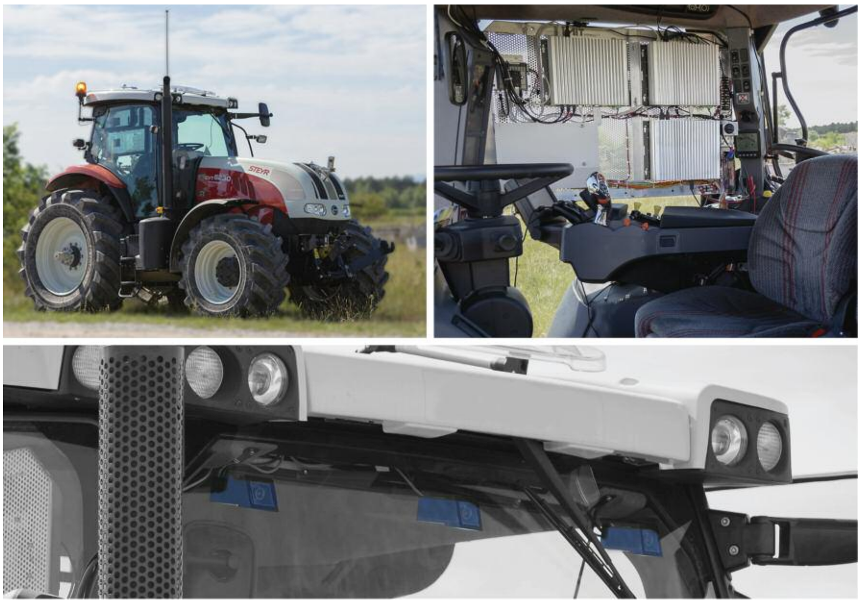

Figure 1: Top left: research platform Steyr 6230 CVT, Top right: additional computers running dedicated software (3x Intel i7 Ivy Bridge), Bottom: windshield mounted AIT stereo camera system (light blue).

rithm [3]. If rerouting is not successful, the vehicle stops safely. A pilot module generates drive commands based upon the relative position of the tractor to the planned path. These commands are executed by the interface module controlling the tractor actuators.

The information collected by the vehicle including camera streams, the vehicle's trajectory, the computed map data, etc. is transmitted to a base station, where the operator can access it through a graphical user interface. Choosing between different visualisation modes allows for better evaluation of complex situations. The possibilities to overlay satellite images with live map data and 3D reconstruction of the vehicle's environment increase the situational awareness of the operator and ensure efficient and safe handling (Figure 2).

Vehicle tasks are planned within the graphical user interface on the basis of satellite images and sensor data. For interaction and control of the system, common user-friendly input devices like a mouse and a game controller are used. The vehicle supports different modes of operation, which represent different levels of autonomy:

- manual teleoperation
- a semi-autonomous mode called 'Click&Drive'
- following a planned path fully autonomous.

In 'Click&Drive' mode the next waypoint can be set directly or special manoeuvres such as Y-turns can be triggered by the operator. The vehicle executes these commands using its autonomous capabilities of path planning and collision avoidance.

With the developed system, AIT is able to address a real safety problem (cultivating areas endangered by explosive ordnances) which currently lacks appropriate solutions within the available off-the-shelf technologies, thus showing the potential of autonomous systems that will be operating in the near future. Simultaneously, research on human machine interaction and environmental perception and segmentation for machines that are changing their environment during operation is pressed ahead.

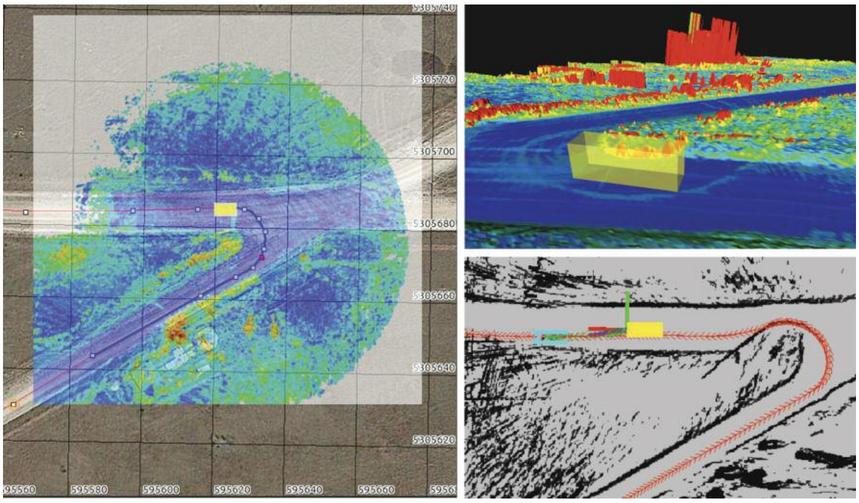

Figure 2: Left: satellite image (Google Maps) overlaid with real-time traversability map; planned path and waypoints (red), vehicle trajectory (blue), Top right: 3D reconstruction of vehicle surroundings coloured by traversability from blue (good) to red (bad/obstacle), Bottom right: real-time path planning; planned path (red arrows), next planned vehicle position (cyan).

Link:
http://www.ait.ac.at/themen/3d-vision/autonomous-land-vehicles/

References:

[1] M. Humenberger et al.: "A fast stereo matching algorithm suitable for embedded real-time systems", J.CVIU, 114 (2010), 11.

[2] S. Thrun et al., "Probabilistic Robotics". MIT Press, Cambridge 2005, ISBN 9780262201629.

[3] S. Karaman, E. Frazzoli: "Sampling-based Algorithms for Optimal Motion Planning", IJRR, 30 (2011), 7.

Please contact:
Robert Rößler, Thomas Kadiofsky, Wolfgang Pointner, Martin Humenberger, Christian Zinner
AIT Austrian Institute of Technology GmbH, Austria
E-mail: {robert.roessler, thomas.kadiofsky, wolfgang.pointner, martin.humenberger, christian.zinner}@ait.ac.at

Avoiding Gridlocks in the Centralised Dispatching of a Fleet of Autonomous Vehicles

by Franco Mazzanti, Alessio Ferrari and Giorgio O. Spagnolo (ISTI-CNR)

Autonomous vehicles will become pervasive in the near future. The TRACE-IT project is investigating how gridlocks may be avoided.

One of the pillars of current industry-related research in Europe is the development of intelligent transport systems, which can ease the safe movement of people and goods by means of smart computer platforms that require limited human support. In near future, fleets of autonomous vehicles will need to be supervised and managed. To this end, at ISTI-CNR, we have been developing an advanced automatic train supervision (ATS) system that is able to dispatch trains in a railway layout, while avoiding gridlocks [1]. A gridlock is a localised deadlock in which vehicles mutually block each other in a specific region, while other vehicles are still free to move. While safety is ensured by interlocking systems, efficient dispatching is currently performed by humans, and the TRACE-IT project lays a possible basis for automating this task.

Model-checking technology underpins our approach, which is composed of two main phases: a configuration phase and an actual dispatching phase. In the first phase, model checking is used to identify a set of constraints that the ATS should consider when dispatching the trains to avoid gridlocks. To identify these constraints, the model-checking algorithm initially takes as input a formal model that includes the railway layout, and the planned missions of the train. Then, it explores all the possible sequences of movements of the trains, and identifies potential gridlocks. For each gridlock found, the algorithm produces the critical regions of the layout in which the gridlock might occur, and the maximum number of trains allowed in each region that guarantees the absence of gridlock. These are the constraints that are used by the ATS during the actual dispatching phase. In this phase, before allowing a train to enter one of the constrained critical regions, the ATS checks that the maximum number of trains allowed in that region is not exceeded. When novel missions or vehicles enter into play, a reconfiguration cycle is performed through model checking, following the first phase. Hence, new regions and numerical constraints are produced for consideration by the ATS.

In TRACE-IT, the approach was experimented on a realistic layout for metro systems provided by ECM S.p.A., and reported in Figure 1. The figure depicts the physical layout in black solid lines, and the missions of eight trains in coloured lines. Within TRACE-IT, configuration and reconfiguration were performed offline by means of the general-purpose model checker UMC, belonging to the KandISTI family [2]. We have also experimented the implementation of a dynamic reconfiguration process, by developing a special-purpose model checker – integrated within the ATS – which focusses on the verification of these particular properties for this particular class of models.

This gridlock avoidance approach can be adapted to any kind of vehicle moving along pre-defined paths – either physical or virtual – that should satisfy safety distance requirements. The fleets of autonomous vehicles that can be managed with this approach range from cars to trucks, and from drones to swarms of robots. Of course, when the layout of movements is particularly large, it may become necessary to decompose it into smaller fragments to avoid exponential state-space explosion during the analysis [3].

This work was carried out within the PAR FAS 2007-2013 TRACE-IT Project, funded by the Tuscan Region, and involving the University of Florence, and railway signalling manufacturer ECM S.p.A.

Link:
http://fmt.isti.cnr.it/umc

References:

[1] F. Mazzanti, G.O. Spagnolo, S. Della Longa, A. Ferrari: "Deadlock avoidance in train scheduling: A model checking approach". FMICS 2014, LNCS Vol. 8718, Springer, 2014, 109-123.

[2] M.H. ter Beek, S. Gnesi, F. Mazzanti: "From EU projects to a family of model checkers". Software, Services, and Systems, LNCS Vol. 8950, Springer, 2015, pp. 312-328.

[3] A. Ferrari, G. Magnani, D. Grasso, A. Fantechi: "Model checking interlocking control tables". FORMS/FORMAT 2010. Springer, 2011, pp. 107-115.

Please contact:

Franco Mazzanti, ISTI-CNR, Italy
franco.mazzanti@isti.cnr.it

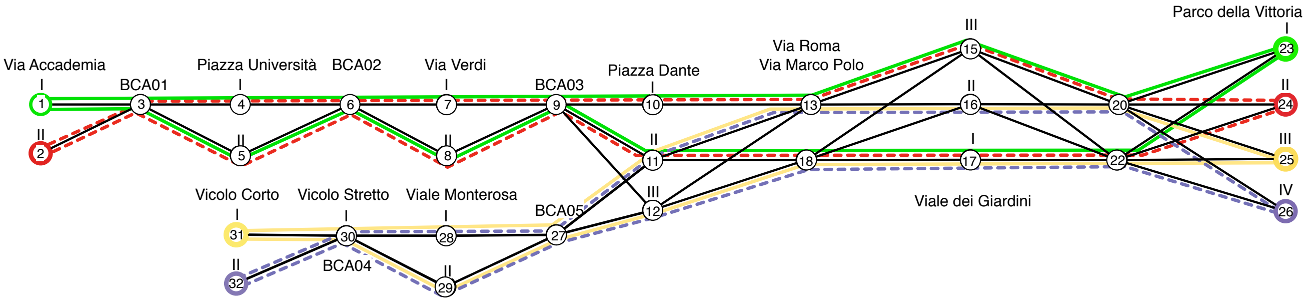

Figure 1: Metro layout for TRACE-IT project.

Diagnosis of Deviations in Distributed Systems of Autonomous Agents

by Farhad Arbab (CWI)

Components that comprise the agents, such as drones, in a distributed cyber-physical system may individually take seemingly reasonable actions, based on their available information Such actions may collectively lead to an eventual unacceptable system behaviour. How can we identify culpable components and agents? At CWI we develop formal models for compositional construction and analysis of such systems, along with tools and techniques to narrow the set of possible components that contribute to undesirable behaviour.

Events and activities in the real world of physical objects comprise too large a set of possibilities and contingencies to account for or control in any static plan or algorithm, meaning that cyber-physical systems simply need to be robust enough to deal with those contingencies as they come up. Thus, a software agent must tolerate less-than ideal conditions and cope with hindrances, while it stretches its resources and capabilities to still achieve its goals by sub-optimal measures. When something goes awry – as it invariably does – the ability to diagnose what went wrong and why is crucial for a system to take corrective or compensatory actions, and/or try to avoid similar undesirable behaviour in the future.

Consider a team of crop surveillance drones, deployed to survey a field and relay locations of possible signs of disease, drought, flooding, etc. The team may include flying drones as well as land lobbers. Maintaining an acceptable altitude, awareness of battery levels, avoiding birds of prey, circumnavigating unforeseen obstacles such as soil shifts and slides or weed growth on paths, are among concerns that such autonomous agents must handle gracefully. Engineering best practices call for isolating such separate concerns into different modules by design, thus allowing for their reuse, as well as their independent verification. Of course, the behaviour of each agent that is composed of such modules must satisfy certain crucial properties, such as ‘when energy level is low, the drone still manages to reach its charging station’. Moreover, we need to establish that a certain team of such agents exhibits collective behaviour that satisfies some other set of crucial properties, such as ‘every patch of the field gets surveyed within a bounded time window’.

Compositional techniques comprise the hallmark of all disciplines of engineering: they enable construction of complex systems by composition of simpler components or sub-systems, such that designers can obtain and analyse the properties of interest of a resulting system through composition and analysis of the corresponding properties of its constituent parts. To compositionally model cyber-physical systems, we employ an automata-based paradigm called Soft Component Automata (SCA) [3]. An SCA is a state-transition system where transitions are labelled with actions and preferences. The intuition to the preferences is that higher-preference transitions typically contribute more towards the goal; if a component is in a state where it wants the system to move north, a transition with action “north” has a higher preference than a transition with action “south”. Preferences thus provide a natural fall-back mechanism for the agent: in ideal circumstances, the agent would perform only actions with the highest preferences, but in some circumstances, the agent may be permitted to choose a

Figure 1: surveillance drone (source: Shutterstock).

transition of lower preference. The behaviour of an SCA is given by all transitions whose preference exceeds a threshold. By adjusting the threshold, the designer can include transitions of lower preference, thus obtaining a more flexible system by virtue of the newly available actions — even though those actions contribute less directly towards the goal.

Soft component automata constitute a generalisation of constraint automata [2], one of the many semantics of Reo [1], a language for verifiable coordination of interactions among components. Consequently, our framework possesses the interesting feature that, like Reo, it blurs the line between computation and coordination — both are formalised by the same type of automata, which allows us to reason about these concepts in a uniform fashion.

When the collective behaviour of such a system violates some of its verification requirements, it is useful to know not only the particular scenario that led or leads to such violation (i.e., a counterexample), but also the likely culprits. For instance, if a flying surveillance drone fails to maintain its target altitude, a counter-example can reveal that the drone attempted to reach the far side of the field, ran out of energy and was required to perform an emergency landing. Many existing LTL-based verification tools are able to provide such counterexamples. Taking this idea of diagnostics one step further in the context of a compositional design, we would like to be able to identify the component (or components) responsible for the first action in the counterexample that violated the requirement. Such an indication may assist the designer of a system in finding the component responsible for a failure (which, in our example, may turn out to be the route planning component), or, at the very least, rule out components that are not directly responsible (such as the wildlife evasion component).

Link:
http://projects.cwi.nl/dedea

References:

[1] F. Arbab: “Puff, The Magic Protocol”, Formal Modeling: Actors, Open Systems, Biological Systems 2011, SRI International, Springer LNCS, vol. 7000, pp. 169-206.

[2] C. Baier, M. Sirjani, F. Arbab, J. Rutten: “Modeling component connectors in Reo by constraint automata”, Science of Computer Programming, 61:75–113, 2006, http://dx.doi.org/10.1016/j.scico.2005.10.008.

[3] T. Kappé, F. Arbab, C. Talcott: “A Compositional Framework for Preference-Aware Agents”, CWI Technical Report FM-1603, 2016, https://repository.cwi.nl/noauth/search/fullrecord.php?publnr=24625.

Please contact:

Farhad Arbab, CWI, The Netherlands
farhad@cwi.nl

Robust Fault and Icing Diagnosis in Small Unmanned Aircrafts

by Damiano Rotondo and Tor Arne Johansen (NTNU)

Small unmanned aircraft working in harsh weather conditions such as those encountered in the Arctic will suffer from the consequences of icing. In order to assure high efficiency and autonomous operation, a robust diagnosis scheme that detects both faults and icing is required. Research at NTNU is focusing on developing these schemes and integrating them with fault tolerant control techniques.

The phenomenon of ice accretion on aircraft's surfaces is a critical and well recognised problem that affects aviation safety. The consequences of this phenomenon, commonly referred to as icing, include a significant decrease of the lift and manoeuvrability of the vehicle and a simultaneous increase in drag, weight and power consumption, which obstruct conducting autonomous operations in environments with harsh weather conditions, e.g., those typically encountered in the Arctic region.

Ice protection systems (IPS) have been proposed as a means of mitigating the effects of icing. Large aircraft are typically equipped with efficient anti-icing and de-icing devices. However, due to space and weight limitations, these devices are unsuitable for small aircraft and alternative solutions must be sought. At NTNU, very promising results have been obtained using carbon nanotechnology to develop thin layers of coating material to be painted on the surfaces of aircraft. These layers can be heated to melt the ice using onboard electricity. However, in order to assure high efficiency and limit energy consumption, icing diagnosis schemes are needed for optimal activation of the icing protection scheme. At the same time, the diagnosis scheme must be able to robustly distinguish icing from actuator faults, e.g., loss of efficiency in the thrust or in the elevator, and tolerance against these types of faults must also be guaranteed.

A robust fault/icing diagnosis and tolerance scheme must take into account changes in the aircraft's dynamics due to the variation of the operating point, the presence of uncertainty in the available mathematical model and other undesired effects, such as exogenous disturbances, e.g., the wind, and measurement noise in the sensors.

The changes in the aircraft's dynamics due to the variation of the operating point can be addressed using a linear parameter varying (LPV) methodology. The LPV paradigm provides an elegant way to apply linear-like tech-

Figure 1: Flight test at Svalbard, Norway. Photo by Kjell Sture Johansen, NORUT.

niques to nonlinear systems, with theoretical guarantees of stability and performance. Unlike linearization techniques, LPV methods do not involve any approximation, since they rely on an exact transformation of the original nonlinear system into a quasi-linear one, by embedding all the original nonlinearities within some varying parameters. As a consequence, it is possible to develop LPV fault/icing diagnosis schemes, which are consistent with a wide range of operating conditions. To this end, the research at NTNU is focusing on two approaches. The first involves using unknown input observers (UIOs) [1], i.e., observers that allow the state of a given system to be estimated independently of some unknown inputs. UIOs can be made insensitive to certain inputs if some structural conditions on the system are fulfilled, which is a property that can be exploited in order to achieve a successful fault/icing diagnosis. The second approach involves multiple model adaptive estimators (MMAE) [2], which comprise a collection of local observers, each of which provides the state estimation which would correspond to a predefined value of some unknown parameter that can be related to the presence/absence of faults or icing. Under some conditions, the diagnosis can be performed by analysing which observer exhibits the smallest output estimation error energy.

On the other hand, the presence of uncertainty, disturbances and noise is being addressed at NTNU using the theory of interval observers. In contrast with classical observers, which try to provide an exact estimation of the state, interval observers are an appealing alternative that aims at providing the set of admissible values for the state at each instant of time. In other words, assuming that the different sources of uncertainties are constrained in known bounded sets, an interval observer computes the lower and upper bounds for the state, which are compatible with the uncertainty. By merging the aforementioned UIO/MMAE with the theory of interval observers, robust fault diagnosis approaches that assure the absence of false alarms can be obtained. At the same time, interval observers can be integrated with fault tolerant control (FTC) techniques in order to obtain a theoretical framework for achieving robust fault/icing tolerance in small unmanned aircraft.

The theoretical development is supported by numerical simulations with high-fidelity simulators, and field experiments are planned for the future in order to test and validate the efficiency of the proposed techniques.

References:

[1] D. Rotondo et al., "Detection of icing an actuators faults in the longitudinal dynamics of small UAVs using an LPV proportional integral unknown input observers", Proc. of the 3rd Conf. on Control and Fault-Tolerant Systems (SysTol), Barcelona, Spain, 2016, pp. 690-697.

[2] D. Rotondo, V. Hassani, A. Cristofaro, "A multiple model architecture for the adaptive estimation of discrete-time systems", accepted for publication in Proc. of the IEEE American Control Conference (ACC), 2017.

Please contact:

Damiano Rotondo, Tor Arne Johansen
Department of Engineering Cybernetics, NTNU, Norway
damiano.rotondo@yahoo.it,
tor.arne.johansen@itk.ntnu.no

European Research and Innovation

Bringing Business Processes to the Cloud

by Kyriakos Kritikos and Dimitris Plexousakis (ICS-FORTH)

Many organisations are considering moving their business processes to the cloud to save money by also taking advantage of its infinite and affordable resources. However, this migration requires aligning the business and cloud levels as well as adaptively managing the business processes in the cloud. CloudSocket is a project dealing with these issues, which facilitates the transformation of business processes into business process as a service (BPaaS) products offered by organisations playing the role of a broker. A respective platform has been produced which enables the management of the whole lifecycle of a BPaaS.

In order to survive in a competitive and dynamic business world, organisations need to reduce costs and efficiently manage and scale their core and supporting business processes. As such, the cloud seems to be the right medium for deploying and adaptively provisioning these business processes. However, moving a business process in the cloud can be quite challenging. First, the business and IT levels must be aligned in such a way that the gap between business requirements and technical cloud-based capabilities can be bridged. Second, there is a need to appropriately manage an executable business process in the cloud such that it can be dynamically provisioned to guarantee the service level agreed with the customer.

CloudSocket is a H2020 project [L1] which aims to address these needs. It ambitions to support the role of the broker which has the ability to either offer BPaaS products in the market or consult potential customers on how they can migrate their business processes in the cloud. This project offers a BPaaS management platform [L2] that can deal with the whole lifecycle of a BPaaS, thus constituting the main instrument for realising the business of a broker. Such a platform comprises different environments which focus on different activities of the BPaaS lifecycle.

The BPaaS Design Environment enables the broker to design business processes and map them to abstract workflows. This mapping is facilitated by two semantic approaches which also cater for the translation between business requirements and technical capabilities: (a) one [1] focusing on annotating business process activities and workflow fragments with ontological specifications to facilitate their alignment; (b) another focusing on a questionnaire-based approach which guides the user to provide the right answers in the pursuit of discovering cloud services that can realise the functionality of a business process.

The abstract workflow developed by the previous environment becomes concrete and deployable via the use of the BPaaS Allocation Environment. This environment is able to concretise the abstract workflow by selecting SaaS services that can realise workflow tasks as well as IaaS services to support the provisioning of the BPaaS workflow software

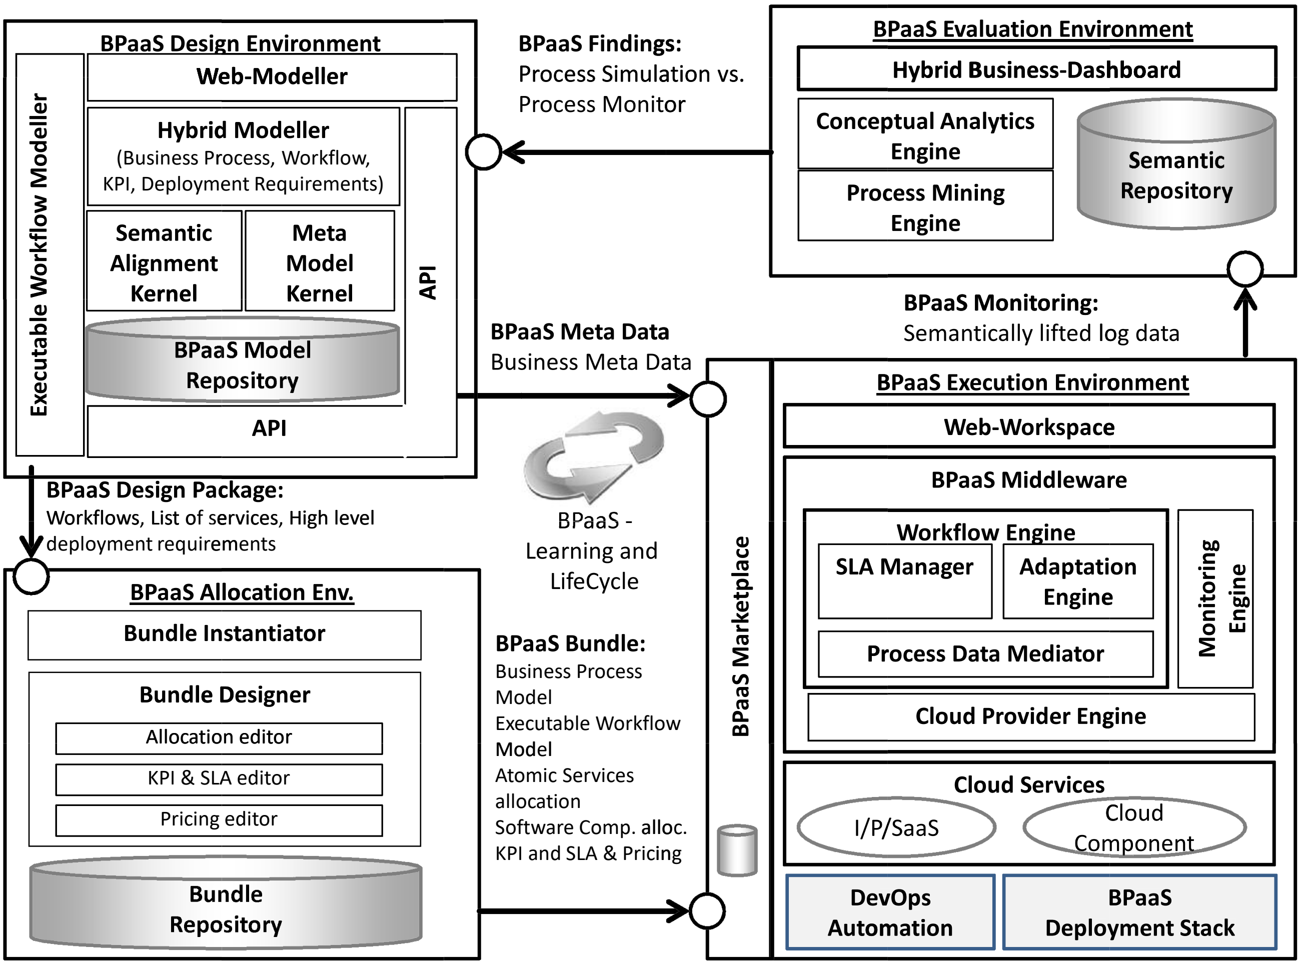

Figure 1: CloudSocket BPaaS management platform architecture.

components. This concretisation is facilitated by two complementary research approaches: (a) one [2] which relies on semantics and the concurrent selection of both IaaS and SaaS services to optimise the broker non-functional objectives; (b) another [3] which uses a DMN [L3]-based approach mapping decisions to quite technical description artefacts in the CAMEL [L4] language.

Deployable workflows are then offered in the form of a BPaaS bundle to potential customers via a Marketplace. Once such bundles are purchased by customers, they are deployed in the cloud and adaptively managed via the integration of three main components: (a) a Workflow Engine dealing with the workflow execution and management; (b) a cloud orchestration engine (Cloudiator [L5]) dealing with the cross-cloud workflow deployment; (c) a cross-layer service monitoring and adaptation framework.

Finally, the lifecycle loop is closed via the BPaaS Evaluation Environment which attempts to offer various types of analysis over the BPaaS performance to produce added-value knowledge that can be used to optimise the BPaaS. This environment relies on a semantic approach which integrates and semantically lifts information coming from the three components in the BPaaS Execution Environment into a Semantic Knowledge Base (SKB) and adopts a transformation approach to map high-level descriptions of KPIs to SPARQL [L6] queries to be assessed over the SKB so as to support KPI evaluation and drill-down.

Links:
[L1] www.cloudsocket.eu
[L2] www.cloudsocket.eu/download
[L3] www.omg.org/spec/DMN/
[L4] www.camel-dsl.org
[L5] github.com/cloudiator/
[L6] www.w3.org/TR/rdf-sparql-query/

References:

[1] K. Hinkelmann, et al.: “A Modelling Environment for Business Process as a Service”, in CAiSE, Springer, 2016.
[2] K. Kritikos, D. Plexousakis: “Multi-Cloud Application Design through Cloud Service Composition”, In Cloud, 686-693, 2015.
[3] F. Griesinger, et al.: “A DMN-based Approach for Dynamic Deployment Modelling of Cloud Applications”, in International Workshop on Cloud Adoption and Migration (CloudWays) at ESOCC, 2016.

Please contact:

Kyriakos Kritikos and Dimitris Plexousakis
ICS-FORTH, Greece
{kritikos,dp}@ics.forth.gr

Shaping Future – A Method to Help Orient Science and Technology Development toward Public Preferences

by Martina Schraudner and Marie Heidingsfelder (Fraunhofer IAO)

Our research project Shaping Future presents a methodology for setting need-oriented research agendas in the field of human-machine interaction.

Innovative technologies must reflect their socio-technical context and conditions to be transferred into technical innovations; emerging from a techno-scientific world, they must be embedded into social practices to be successful. Agenda-setting processes have conventionally been determined by specialists, largely because technology foresight faces a double-bind problem known as the 'Collingridge dilemma'. The full functionality and impact of a new technology cannot be easily predicted until this technology is sufficiently developed and widely used, after which time it is difficult to make substantial changes [1]. In addition, current language-based methods predispose one to rely on pre-existing terminology and schemas and thereby restrict one's openness and creativity while trying to picture desirable future developments.

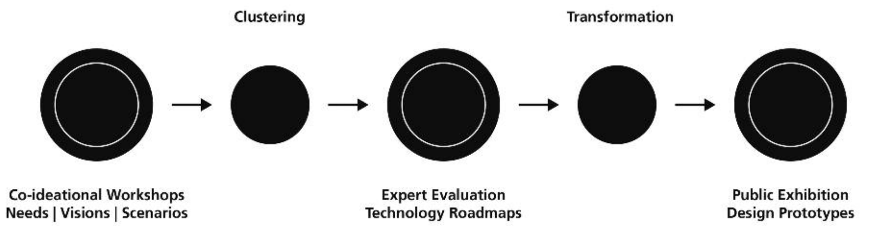

Figure 1: The Shaping Future process.

In the project 'Shaping Future' we developed an original and interdisciplinary methodology to integrate laypersons into this otherwise mainly expert-driven process [2]. This process comprises three major stages connected by two transfer stages, clustering and transformation (see Figure 1). Each of these stages requires its own methodological approaches. (i) In a series of co-ideational workshops, we needed to enable laypersons to articulate their preferences toward technological advances and the directions they might take. (ii) For the professional utilisation of these preferences, we needed to develop a method to enable researchers from different backgrounds to understand each other, to interact on equal terms and to find interdisciplinary solutions based on the layperson's needs. (iii) Finally, for the dissemination of results, we publicly displayed a number of design fiction prototypes, developed by professional designers, in order to initiate a widespread discourse about desirable technological advances. Thus, the developed methodology can open the processes of technology development and agenda-setting to the general public.

The current results of Shaping Future suggest that methods of participatory design and the social sciences can be adopted and combined for participatory agenda-setting processes far beyond the field of human-machine interaction [3]:

- Participatory prototyping provides a particularly effective technique for enabling laypersons to realise and explore their preferences toward prospective, as yet unknown, technological advances and the methods by which they are produced. Overall, workshop participants considered this technique to be the key to the entire process. Building upon approaches from both participatory design and the social sciences, we developed an original, methodologically sound procedure for data analysis. With the help of this procedure, we extrapolated laypersons' preferences from the aggregate data, comprising tangible objects, written descriptions, audios, and videos, and clustered these preferences. This refinement rendered the lay input not only understandable and appealing to participating specialists but also generally suitable for the purposes of technology development.
- By transforming this input, participating specialists were, in turn, able to outline promising research trajectories. Going beyond the limits of language, our design-based method enabled them to transcend everyday practicality, professional terminology and 'silo-knowledge'[4].
- By publicly displaying design fiction prototypes, that give a physical shape to the layperson's needs and the expert's evaluation, we intend to reach a larger audience and ultimately engage a larger number of people in technology foresight.

The methodology allows us to go beyond preference articulation and to open up a process of collaborative sense-making in the context of future technologies.

The Shaping Future project is funded by the German Federal Ministry for Education and Research (BMBF), project ID 16I1639.

References:

[1] A. F. Blackwell, et al.: "Radical innovation: crossing knowledge boundaries with interdisciplinary teams", No. UCAM-CL-TR-760, University of Cambridge, Computer Laboratory, 2009.
[2] M. Heidingsfelder, et al.: "Shaping future – Adapting design know-how to reorient innovation towards public preferences", Technological forecasting and social change, 101, 291-298, 2015.
[3] M. L. Heidingsfelder, F. Schütz, S. Kaiser: "Expanding participation participatory design in technology agenda-setting", in Proc. of the 14th Participatory Design Conference: Short Papers, Interactive Exhibitions, Workshops-Volume 2, 25-28, ACM, 2016.

Please contact:

Martina Schraudner, Fraunhofer IAO- Fraunhofer Center for Responsible Research and Innovation, Berlin Germany
Martina.schraudner@iao.fraunhofer.de

User-Centric Analysis of the CAPTCHA Response Time: A New Perspective in Artificial Intelligence

by Darko Brodić, Sanja Petrovska, Radmila Janković (Univ. of Belgrade, Serbia), Alessia Amelio (Univ. of Calabria, Italy) and Ivo Draganov (Technical University of Sofia, Bulgaria)

The Completely Automated Public Turing Test to tell Computers and Humans Apart (CAPTCHA) test is designed to differentiate between human users and bots. The time spent to find a solution to the CAPTCHA test is correlated with particular demographic characteristics of Internet users. This knowledge can be used to predict which CAPTCHA tests are best suited to different users.

CAPTCHA is an acronym of 'Completely Automated Public Turing Test to tell Computers and Humans Apart'. It is a test-puzzle verified by computers that has particular characteristics that make it only solvable by humans. Essentially, the CAPTCHA is used to differentiate computer users from bots when they access a particular website. A bot (abbreviation of robot) is a software program that automatically runs tasks on the Internet, which represent a security risk. If a user correctly solves the CAPTCHA task, then the program will recognise it as a human. Hence, the aim of the CAPTCHA test is to identify attacks by bots.

Different CAPTCHA types have been proposed in the literature. The most widespread types are: (i) text-based, characterised by text and/or numbers, and (ii) image-based, related to image recognition tasks. In text-based CAPTCHA types, the user is required to recognise text and/or numbers and report them inside a text field. Figure 1 shows an example of text-based CAPTCHA types.

In image-based CAPTCHA types, the user is required to identify a certain image among a collection of images. It can include the identification of image types in a specific context, i.e. a house with numbers, an animated character, a picture of the CAPTCHA or an old woman, and the identification of images related to psychological states in the context, i.e. a worried face or a surprised face. Figure 2 depicts an example of image-based CAPTCHA types.

There has been relatively little research into the usability of the CAPTCHA test: how easy it is for a user to solve the CAPTCHA. The main limitations of the proposed approaches to analysis are: (i) the number of users involved in the analysis, which is quite small (i.e., the statistical significance of the tested population), (ii) the reduced user population features, and (iii) the low number of considered CAPTCHA types [1].

We extend the way of analysing CAPTCHA usability based on the time spent for a user to find a solution to the CAPTCHA. We recruited approximately 200 volunteer users for this study. Each user was interviewed to evaluate their time spent to solve the CAPTCHA test. Each user provided anonymous personal data including: age, gender, education level and Internet experience (in number of years). Users included students, engineers, teachers, and employees, with ages ranging from 18 to 52 years, with different levels of experience in Internet use. Each user was required to solve the text and image-based CAPTCHA types and their response times were recorded.

This data was analysed to determine whether a correlation exists between each demographic variable and the user's response time to different CAPTCHAs. It is pursued by formulating statistical hypotheses describing such a correlation. Analysis should accept or reject the formulated hypotheses. Accordingly, the correlation coefficient and statistical tests

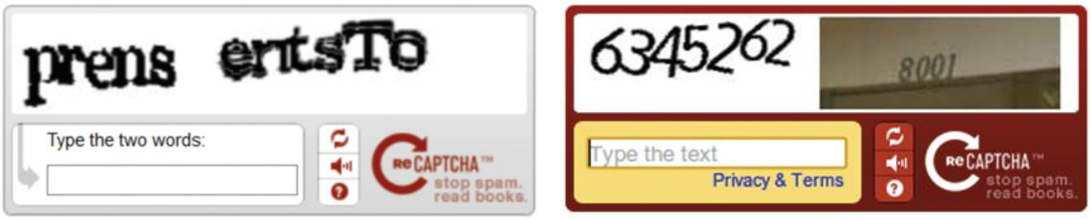

Figure 1: Text-based CAPTCHA samples: only text (left), and only numbers (right).

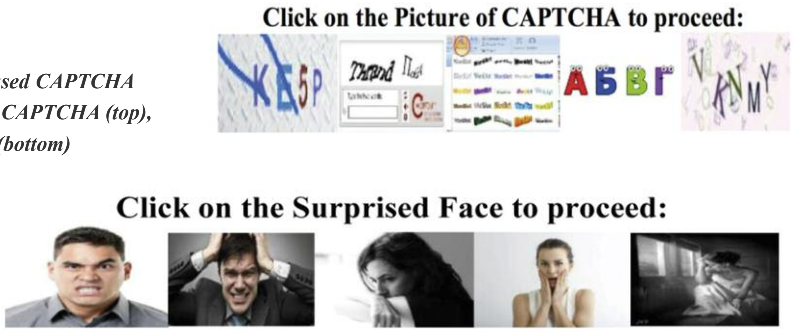
Click on the Picture of CAPTCHA to proceed:
KE5P
АБВГ
Click on the Surprised Face to proceed:

Figure 2: Image-based CAPTCHA samples: picture of CAPTCHA (top), and surprised face (bottom)

are employed for analysis of this correlation [2]. This allows us to determine whether education level or gender, for example, has an impact on CAPTCHA response time, and the strength of the impact.

Statistical analysis was extended further by investigating the correlation between a set of demographic features and the response time to different CAPTCHAs. It was performed by extracting association rules from the population data [3]. They are rules expressing the dependence of the response time (e.g., high response time) from the co-occurrence of some demographic feature values (e.g., age below 35 years and female gender), and the strength of such a dependence.

The analysis demonstrates that number of years of Internet use, young age and a higher education level help users to quickly solve the CAPTCHA test. Gender also has a small impact on CAPTCHA solving time. Older users require a higher education level to solve the CAPTCHA with surprised and worried faces. Finally, CAPTCHA with only numbers is solved more quickly than CAPTCHA with only text, while image-based CAPTCHA has the lowest resolution time, in particular the animated character images.

This study offers invaluable insights to help predict which CAPTCHAs are best suited to different Internet users. It involves the University of Belgrade, Serbia, the University of Calabria, Italy, and the Technical University of Sofia, Bulgaria.

References:

[1] L. Ying-Lien, H. Chih-Hsiang: “Usability study of text-based CAPTCHAs”, Displays 32(2): 81-86, 2011.

[2] D. Brodić, S. Petrovska, M. Jevtić, Z. N. Milivojević: “The influence of the CAPTCHA types to its solving times”, MIPRO: 1274-1277, 2016.

[3] D. Brodić, A. Amelio, I. R. Draganov: “Response Time Analysis of Text-Based CAPTCHA by Association Rules”, AIMSA: 78-88, 2016.

Please contact:

Darko Brodić
Technical Faculty in Bor, University of Belgrade
E-mail: dbrodic@tfbor.bg.ac.rs

How to Secure Internet of Things Devices in an Energy Efficient Way

by Zeeshan Ali Khan and Peter Herrmann (NTNU)

In our daily lives, we are surrounded by a plethora of connected devices. But can we be sure that they are all secure and will not start malicious attacks against us? Many devices are small and have only limited processing and battery power such that traditional security mechanisms cannot be used. In this article, we propose a trust-based intrusion detection system that is resource-friendly and offers energy-efficient media access.

Innovative technology plays an increasingly important role in our daily lives. We deal not only with independent devices, but with devices that communicate with each other and form an Internet of Things (IoT). According to most predictions, several billions of ‘things’ will be connected with the internet in the coming years [L1]. Many of these connected devices will be very small and cheap such that they can be placed wherever they are required. The interconnectedness of the IoT devices and networks, however, poses a significant risk since the systems will be subject to malicious attacks. An example is denial of service attacks precluding the devices from communicating with other stations. Therefore, security issues must be considered for the engineering and deployment of IoT networks. Further, the applied security mechanisms have to address the limitations of many devices in terms of memory, power, and bandwidth. Otherwise, the devices can be quickly depleted.

In the scope of an ERCIM fellowship, we propose the use of Intrusion Detection Systems (IDS) to minimise the chances of Denial of Service (DoS) attacks. Our solution concentrates on attacks that actively try to omit crucial communication between nodes. In ad-hoc networks, this can occur via insider attacks in which the intruder manages to compromise the routing capability of devices. To reduce the computing effort for the small devices, we use trust management [1] which allows the emulation of trust relations in other entities by relatively lightweight but powerful mechanisms. In a computer, one can describe the trust of an entity in another one by a trust value. Using special metrics, the trust values can be adapted based on the good and bad experiences one has with a trustee over time. This reflects the development of trust relations in quite a natural way. Moreover, operators exist that allow us to consider recommendations by third parties and to aggregate the trust values of several entities with a trustee to a more general reputation.

In [2], we introduced a trust-based IDS for the popular Routing Protocol for Low power and Lossy networks (RPL [L2]). An example of RPL based IoT network for Healthcare applications is explained in Figure 1. In this approach, network nodes monitor the behaviour of adjacent nodes, e.g., by listening to whether they forward messages as demanded by the protocol. When a node forwards a message correctly

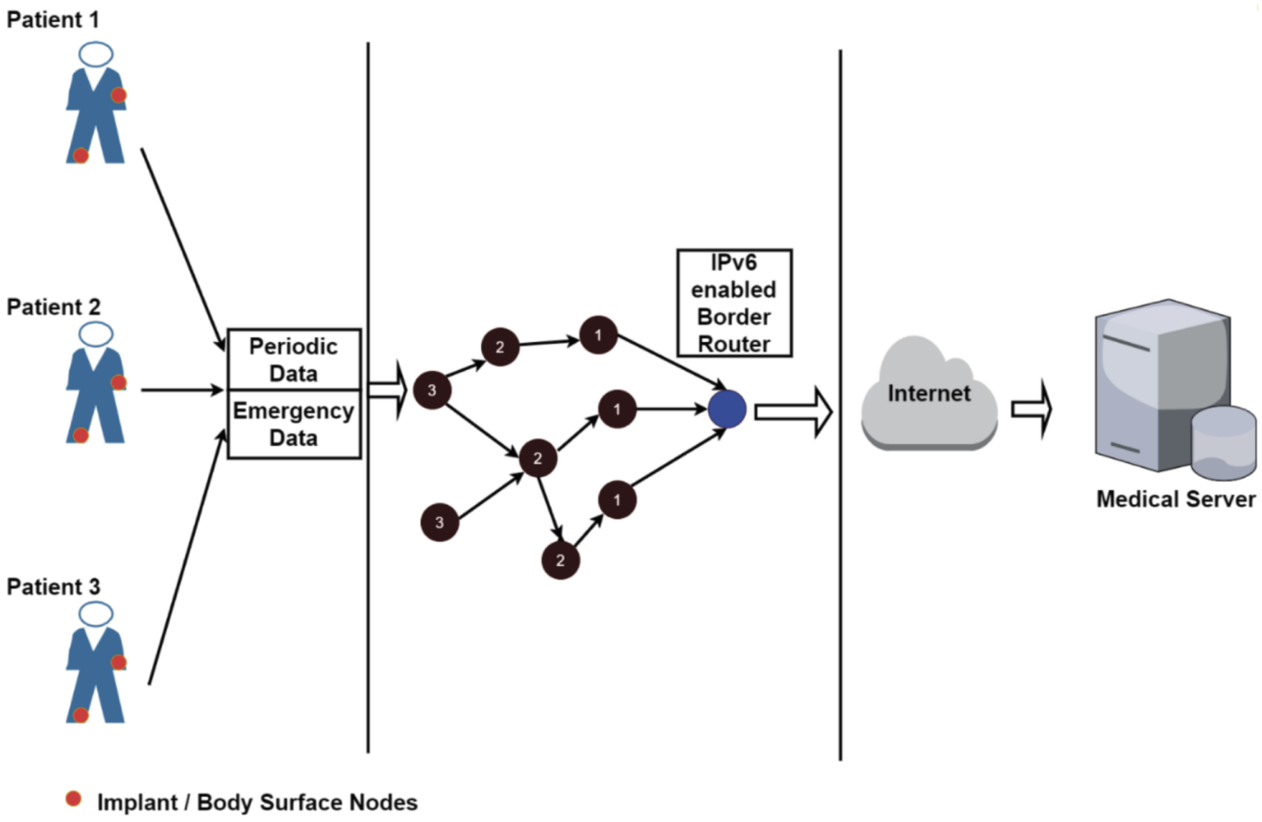

Figure 1: An Example of RPL based IoT network, for Healthcare Application.

within a certain period of time, this counts as a positive experience and otherwise as a negative one. Based on that, the personal trust values are amended and, by exchanging them with other nodes, a more general reputation of a node is built. If the trust-building process reveals that a certain node has a bad reputation among its neighbours, it is assumed to be infected and can be quarantined.

To get practical evidence of our approach, we are in the process to create a test-bed consisting of Z1 devices running the operating system Contiki. First results show that the computations needed to build trust values are not computing intensive. On the other side, however, the trust based mechanisms seem to drain the often scarce battery supplies of a node. The main reason for that is that the monitoring of the neighbouring nodes is achieved by active channel listening, i.e., by constantly listening to the wireless network channel which is an energy intensive activity.

To overcome this problem, we have extended our work and propose to use our approach to the media access protocol IEEE 802.15.4 [L3] that offers mechanisms to reduce active channel listening. A first outcome of this extension is that a beacon-based mode with guaranteed time slots, offered by this protocol, can reduce the active channel listening dramatically. Our simulations and analysis showed that the reduction will be between 50 % and 75 %. Like with the RPL protocol, the extension shall be further tested with our Z1-based test-bed.

Finally, we are in the process of developing a study reviewing the state of the art of IDS existing for the IoT networks. Here, we will also look on existing systems for Mobile Ad-hoc Networks (MANET), Wireless Sensor Networks (WSN) and Cyber Physical Systems (CPS) that have properties close to IoT and might also be usable in this area.

Links:

[L1] www.woodsidecap.com/wp-content/uploads/2015/03/WCP-IOT-M_and_A-REPORT-2015-3.pdf
[L2] tools.ietf.org/html/rfc6550
[L3] www.ieee802.org/15/pub/TG4.html

References:

[1] A. Jøsang: “A Logic for Uncertain Probabilities”, International Journal of Uncertainty, Fuzziness and Knowledge-based Systems, vol. 9, pp. 279-311, 2001.
[2] Z. A. Khan, P. Herrmann: “A Trust Based Distributed Intrusion Detection Mechanism for Internet of Things”, in 31st IEEE International Conference on Advanced Information Networking and Applications (AINA), Taipei, Taiwan, 2017.
[3] G. Swamynathan, K. C. Almeroth, B. Y. Zhao: “The design of a reliable reputation system”, Electronic Commerce Research, vol. 10, no. 3, pp. 239-270, 2010.

Please contact:

Zeeshan Ali Khan, Peter Herrmann
NTNU Trondheim, Norway.
zakhan@ntnu.no, herrmann@ntnu.no

Security Testing for Mobile Applications

by Peter Kieseberg, Peter Frühwirt and Sebastian Schrittwieser (SBA Research)

In recent years, standard end-to-end encryption protocols have become increasingly popular for protecting the security of network communication of smartphone applications, as well as user privacy. On the whole, this has been a good thing, but this reliance has also resulted in several issues related to testing.

Software products have been becoming increasingly closely connected, particularly since the rise of mobile environments. Even programs doing menial tasks now require the user to go online, either to use cloud-based resources, or simply because the software provider does not charge users directly, but generates revenue by collecting and selling user preferences and advertising space. While the latter is often to be considered unethical, a study indicates that even those who complain are actually quite willing to give up privacy at a rather small price (see also for the 'privacy paradox' [1]).

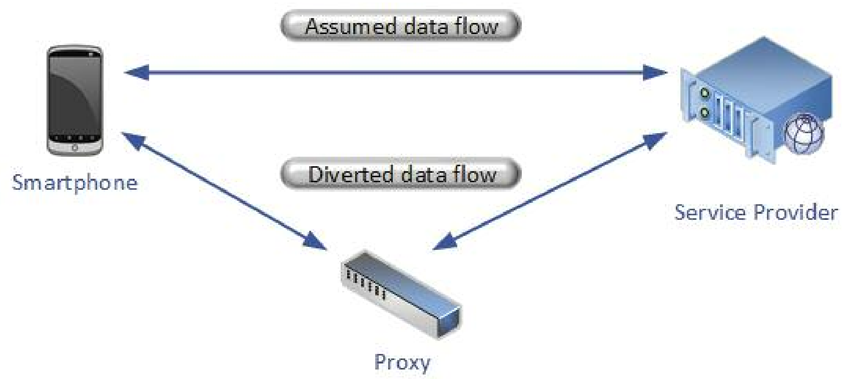

Figure 1: Testing approach.

This increase in connectivity is accompanied by a larger attack surface open to potential malicious actors. Furthermore, the time to market of software products is getting shorter as the market is getting increasingly competitive, resulting in a virtual omnipresence of security issues in current productive software. Even more problematic is the fact that most of the issues currently impairing software security are not based on new scientific findings or large-scale (criminal) organisations conducting intense research for new vulnerabilities to put to criminal use, but rather exploit known weaknesses introduced either by ignorance, or by putting trust into well-known security measures that were never designed for the problem at hand.

While the first problem refers to the problem of diminishing resources during the software development process and can be fixed using training, appropriate management and, most importantly, by setting aside resources, the latter is far more problematic and often touches at the very core of the design process of the software: Developers intrinsically assume certain aspects of their software and its use without further analysing the real attack surface. One very prominent example is the use of transport layer protection in mobile environments, with Transport Layer Security (TLS)[L1] being the de-facto standard for transport encryption. During a case study we analysed a set of mobile apps with respect to the usage of this protocol. We focussed on mobile environments since they possess some characteristics that are different from typical web-based applications. Condensing the results, we typically encountered two major issues with the use of TLS:

1. TLS was designed to provide secure end-to-end communication, always assuming the two endpoints to be trusted entities, i.e., removing them from the attacker model. Likewise, many software designers seem to assume that they possess full control over the client side of the software, which is clearly a problematic view in the case of mobile apps, where the client needs to be identified as a potential attacker. Many reasons for client side attacks can be found, e.g., in the realm of mobile messenger apps, users could try to impersonate other people in order to access or spoof messages. Particularly in apps that require the user to pay for the services provided, the motivation for attacking from the client side is rather straightforward.

2. Developers seem to rely far too much on the powers of TLS, even assuming capabilities, a protocol for transport layer protection can never hold up to. All TLS does is allow encrypted communication. ITLS cannot, for example, fix a protocol for authentication that is fundamentally flawed from a logical perspective. Nevertheless, we got the impression that in many apps TLS is seen as the silver bullet to fix any security-related issues.

These issues should typically be uncovered during the test phase, but our study on many popular apps, including mobile messengers, mobile games, social networks and even online ticketing shops, revealed an astonishing number of insecure protocols and a general misconception of the attacker surface. Our testing approach for uncovering insecure protocols (see [2] and [3]) relied on the fact that the smartphone the app is running on is under full control of us as adversaries, especially since any attacker having physical access to the phone can potentially manipulate hard- and software, including the operating system (or even run the whole app on a emulator without any actual smartphone). Since we controlled the smartphone we could make the phone route all of the (encrypted) traffic through a proxy, which was also controlled by us (see Figure 1). While TLS protected the communication from the phone to the proxy and from the proxy to the server of the service provider of the app, all the information is visible on the proxy, thus making every aspect of the protocols visible for analysis. For this approach, only standard tools are required, not only making this a good strategy for attacking apps, but also a viable and cost-effective measure for testing.

In conclusion, it is clear that the topic of software testing needs far more focus, especially considering the new dangers of fully integrated and interconnected environments as envisioned by ubiquitous computing and the 'internet of things'. Based on the experiments carried out in order to grasp the major problems and assess them, we will put our efforts into defining a testing approach suitable for interconnected environments, starting with the design phase of the software. Furthermore, we plan to support this approach with tools that will speed up the testing process for well-known and important protocols. Nevertheless, one of the major issues still prevalent with software testing, the issue of resources, cannot, in our opinion, be fixed on a technical level, but requires far more awareness on the management level.

Link:
[L1] https://tools.ietf.org/html/rfc5246

References:

[1] P. A. Norberg, D. R. Horne, D. A. Horne. "The privacy paradox: Personal information disclosure intentions versus behaviors", in Journal of Consumer Affairs 41, no. 1, 100-126, 2007.

[2] P. Kieseberg,et al.: "Security tests for mobile applications – Why using TLS SSL is not enough", in 2015 IEEE Eighth International Conference on Software Testing, Verification and Validation Workshops (ICSTW), 2015.

[3] R. Mueller, et al.: "Security and privacy of smartphone messaging applications", International Journal of Pervasive Computing and Communications, vol. 11, 2015.

Please contact:

Peter Kieseberg, SBA Research, Vienna, Austria
pkieseberg@sba-research.org

u'smile – Secure Mobile Environments

by Georg Merzdovnik, Damjan Buhov, Artemios G. Voyiatzis and Edgar Weippl (SBA Research)

Protecting user security and privacy on the internet takes more than transport layer security (TLS) and strong cryptographic algorithms: utilizing TLS notary services and certificate pinning for improved defences against prevalent third-party tracking.

Smartphones and mobile devices are becoming an integral part of life in the connected world. They allow easy access to information and content generation as well as consumption. A significant driving factor for their wide acceptance is the enormous number of available applications (commonly referred to as 'apps'). For example, by the beginning of 2017, the Google Play Store offered more than 2.6 million apps for smartphones, ranging from games and weather apps to office suites and banking apps. While these applications offer huge potential, they can also, either intentionally or inadvertently, pose a security and privacy risk to their users.

At the Josef Ressel Centre for User-friendly Secure Mobile Environments (u'smile) we are analysing security issues in current and future mobile applications. We are working on: (i) the design, development, and evaluation of concepts, methods, protocols, and prototype implementations for addressing security risks, and (ii) communication and coordination with industry partners and standardisation organisations with the goal of establishing globally accepted standards for secure, interoperable mobile services.

One area that we are working on is the security and privacy of network communications of mobile apps and their cloud-based services. All published applications go through inspection for malware behaviour before being published in the Google Play Store. However, user tracking and other forms of privacy leaks are still feasible, especially through advertisements included in apps and web applications to provide a revenue stream for the developers of free or low-priced software. A large number of tracker-blocking applications and apps are available for privacy-conscious users. We performed a large-scale study of the effectiveness of available tracker-blocking tools [1], which showed that despite their sophistication, tools cannot defend against all privacy leaks. Furthermore, a significant portion of applications transmit information over HTTP without establishing a secure connection over the TLS protocol (i.e., use HTTPS). This bad practice allows traffic interception and content inspection in-transit, increasing the threat of user tracking.

Digital certificates are used by TLS and allow verification of the identity of the server an app is connecting to as well as setting up a connection for exchanging information securely. Numerous certificate authorities come pre-installed in modern mobile operating systems, all of which are equally trusted by an app to offer a valid certificate for any internet server. TLS notary services collect and distribute information from TLS certificates presented at various points on the

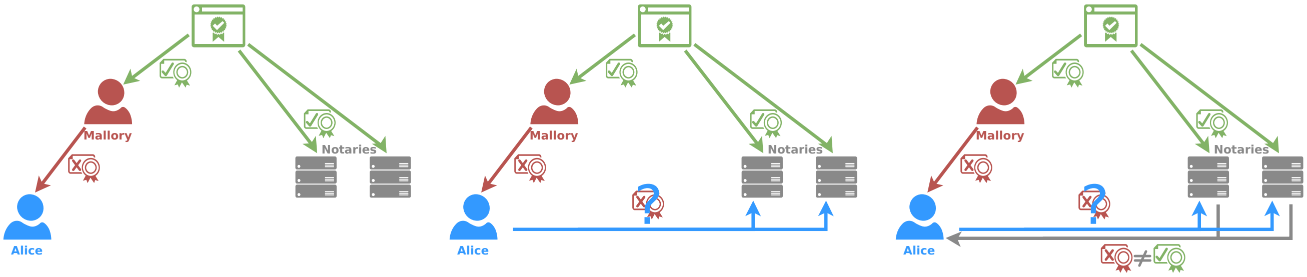

Figure 1: Example of TLS notary service operation – the connection to the server is only trusted if the two certificates match.

internet. Upon connecting over HTTPS to a website, an app can consult a TLS notary service and check if the presented certificate is the same as the one presented to other users. If not, it might be a good indicator that the connection with the server has been intercepted and a fake-but-valid certificate presented to the user for nefarious purposes (i.e., a man-in-the-middle attack). We performed a longitudinal study of openly-accessible TLS notary services spanning one year [2]. Our findings indicate that while the offered service is indeed valuable for internet hygiene, there is currently little incentive for service operators to cover the incurred costs and thus, their availability is dropping.

‘Certificate pinning’ is an approach used by highly-sensitive apps (e.g., banking apps) to ensure that only selected, app-specific certificates are used to validate server identity. However, developers often fail to implement certificate pinning correctly. In some cases, this even breaks the normal TLS certificate validation procedure, rendering the app more insecure compared to not implementing certificate pinning at all. We designed, implemented, and demonstrated a system-level approach that transparently forces certificate pinning for installed apps and thus, reduces their exposure to related attacks [3]. Our approach does not require developers to update their app. As a next step, we extended our design towards a notary-assisted approach. This allows proper certificate pinning even when the application is connecting with a web resource for the first time by relying on the collective knowledge already contributed to the TLS notary services. Also, the new design requires minimal if any intervention by users; users are effectively freed from any uninformed system-level security decisions, which are otherwise a common source of security problems.

Secure mobile environments are key building blocks for the Internet of Things (IoT) world. The need for security is becoming increasingly apparent as self-managing mobile devices and apps are deployed and operate transparently to the users. Researchers at u’smile have been working on this topic since October 2012. The Josef Ressel Center u’smile is funded by: the Christian Doppler Gesellschaft (CDG); A1 Telekom Austria AG; Drei-Banken-EDV GmbH; LG Nexera Business Solutions AG; NXP Semiconductors Austria GmbH; and Österreichische Staatsdruckerei GmbH.

Links:
[L1] usmile.at/
[L2] www.sba-research.org/

References:

[1] G. Merzdovnik et al.: “Block Me If You Can: A Large-Scale Study of Tracker-Blocking Tools”, in 2nd IEEE European Symposium on Security and Privacy (IEEE Euro S&P), 2017.
[2] G. Merzdovnik et al.: “Whom you gonna trust? A longitudinal study on TLS notary services”, IFIP DBSec 2016.
[3] D. Buhov, et al.: “Pin It! Improving Android network security at runtime”, IFIP Networking 2016.

Please contact:

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

ARIES: Reliable European Identity Ecosystem

by Nicolás Notario, Alberto Crespo (Atos), Antonio Skarmeta, Jorge Bernal and José Luis Cánovas (Universidad de Murcia)

ARIES (ReliAble euRopean Identity EcoSystem) will provide secure privacy-preserving identity management, effectively reducing the risk of identity fraud and crime.

Personal data and individual identities are becoming increasingly vulnerable in the virtual world, which facilitates interaction between international stakeholders and the globalisation of crime. Public trust in online security is waning owing to the current lack of adequate solutions, including applied technologies and processes for trusted enrolment, identification and authentication processes. For example, most common authentication means are usernames and passwords, a solution that has been demonstrated to be vulnerable. Furthermore, there is a lack of a common EU wide concept of identity theft and room for improvement in the area of reporting mechanisms (especially across borders) [1] which costs companies, countries and citizens billions of Euros in fraud and theft. As acknowledged in the European Agenda on Security, 'cybercrime is an ever-growing threat to citizens' fundamental rights and to the economy, as well as to the development of a successful Digital Single Market' [2]

In this context, ARIES H2020 European research project [L1, L2] will enable secure, reliable and privacy-preserving identity management and derivation techniques, both to allow a secure user interaction with services and to prevent / reduce risks of identity theft and fraud crimes. ARIES ecosystem appears in the context of delivering a comprehensive framework and holistic approach of innovative technologies, improved processes and security features capable of enhancing the European eID ecosystem and achieving a tangible reduction in levels of identity theft, fraud and associated crimes.

ARIES ecosystem (see Figure 1) will empower its users with a mechanism (identity virtualisation process) allowing them to generate virtual identities, simultaneously linked to the citizens' biometrics and to existing digital or physical identities possessing a high level of identity assurance such as an eID or ePassport. These virtual identities can be stored and managed through a secure wallet usually installed in the citizens' smartphones. The usage of a convenient secure wallet to manage and present virtual identities will positively impact the usage of highly assured identities as it will avoid the usability issues and technological fragmentation (multiple standards) of physical identities and related technologies across Europe.

The ARIES ecosystem allows and encourages users to combine different virtual identities that have a high contextual value (e.g., national eID and bank information for an online transaction) into derived identities that minimises the data disclosed (e.g., there may be no need to disclose the sex of the person) and that can be reused in further interactions. Cryptographic proofs for the identity enrolment, virtualisation and derivation processes can be stored, with the consent of users, in a secure vault that provides guarantees of integrity, confidentiality, auditability and compliance only accessible to law enforcement authorities when a cybersecurity incident occurs and when their assistance is required to help recover from an identity theft or loss [3].

The ecosystem will be demonstrated in two orthogonal real world use cases. In the first, ARIES will allow users to

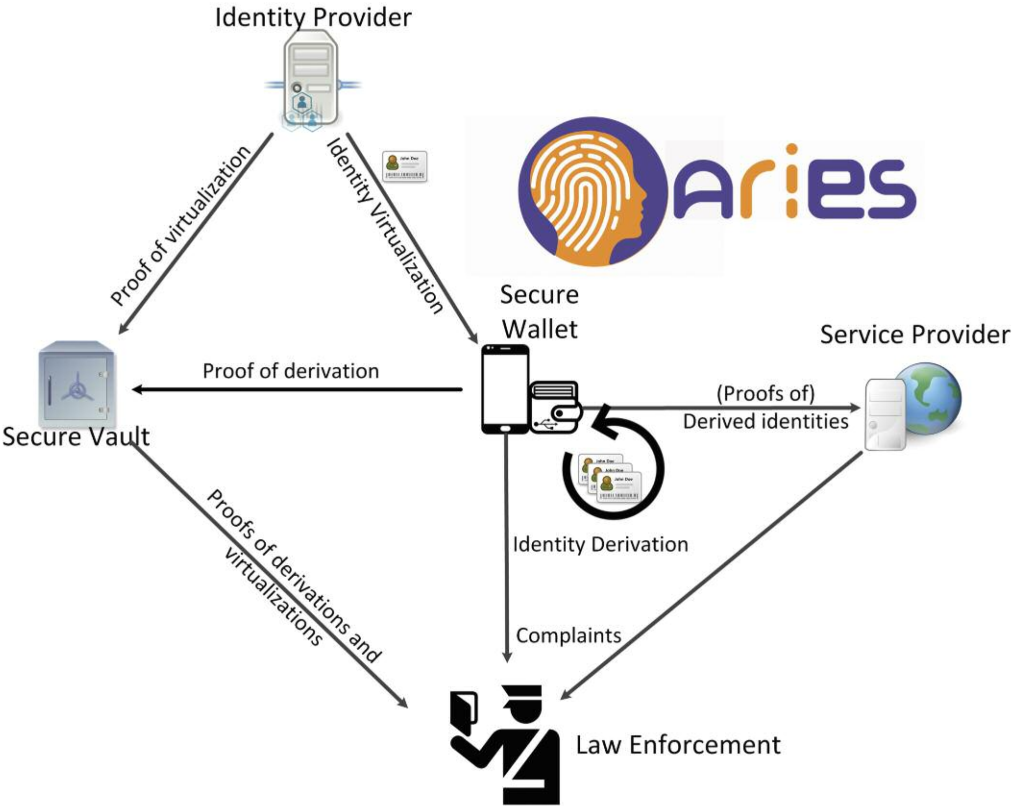

Figure 1: ARIES ecosystem.

securely connect to an online commerce site, enabling a mutually trusted relationship without the need of disclosing non-essential personal data. The second use case will focus on the enrolment process and physical access control to a secure-sensitive environment such as an airport. The airport scenario will effectively show how different derived identities, with different levels of privacy, can be used in different situations (e.g., airport access control, boarding control and duty free shops) combining attributes with different levels of assurance coming from different identity and attribute providers (e.g., national eID or passport and electronic boarding passes).

In terms of social transformation factors, ARIES will contribute to lower several barriers, including end-user acceptance, by providing a secure and privacy by design enabled solution. ARIES will endow users with the ability to anchor trust on a secure and high level assurance infrastructure that will be used to derive additional virtual identities supporting different levels of privacy-preserving and anonymization capabilities but relying on a law enforcement mechanism to obtain effective support in the event of identity-related crimes. This will make users feel more secure in these eID ecosystems, which ultimately will encourage the use of electronic identities and increase the trust in, and adoption of, ICT and online services across the EU by both citizens and businesses.

ARIES is a project that specifically aims to prevent and reduce the risk of identity theft and fraud crimes. This is achieved by the means of cryptographic links between derived, virtual and biometric identities and by the cryptographic proofs accessible at the secure wallet by law enforcement agencies that can leverage them when investigating identity crimes.

The ARIES project is a Research and Innovation Action funded by the European Commission’s Horizon 2020 programme and the consortium carrying it out consists of a well-balanced mixture from six European countries consisting of industry partners, SMEs, public law enforcement bodies and also one retailer.

Links:
[L1] aries-project.eu/
[L2] twitter.com/AriesH2020

References:

[1] N. Robinson et al.: “Comparative Study on Legislative and Non Legislative Measures to Combat Identity Theft and Identity Related Crime”
[2] European Commission: “The European Agenda on Security”, COM (2015) 185 final.
[3] I. Naumann, G. Hogben: “Privacy features of European eID card specifications” Network Security, Vol. 2008.

Please contact:

Nicolás Notario, Atos Research & Innovation, Spain
nicolas.notario@atos.net

Antonio Skarmeta, Jorge Bernal
Universidad de Murcia, Spain
skarmeta@um.es, jorgebernal@um.es

The KandISTI/UMC Online Open-Access Verification Framework

by Franco Mazzanti, Alessio Ferrari and Giorgio O. Spagnolo (ISTI-CNR)

ISTI-CNR provides an online open-access environment for the experimentation of design, analysis and verification of UML-based system models. Great as a didactic environment, it can successfully compete in terms of friendliness and usability with the most mainstream verification frameworks.

KandISTI [1] [L1] is an open-access, online, modelling and verification framework for software-intensive systems developed at ISTI by the FMT Laboratory. It is composed of a set of experimental analysis/verification tools (CMC, FMC, UMC, VMC) of which UMC is the most advanced component. To date, UMC has been applied to a range of case studies in the railway, automotive and telecommunications fields [2,3].

In UMC, a system is defined as a set of communicating state machines. Each state machine is described by a UML Statechart. The dynamic behaviour of a UML system can be: interactively explored; visualised as an evolutions graph; summarised by a minimal set of traces; model-checked using a parametric, branching-time state- and event-based, parametric, temporal logic.

The development of UMC started in 2001 and since then has been continuously improved with the support of several EU and regional projects (AGILE, SENSORIA, TRACE-IT). It has now reached version 4.4, and a maturity level that makes it usable not only for small prototypes but also for real-world systems.

The main feature of the framework is the high degree of usability of all its functionalities; also for this reason it has often been used with satisfaction as a didactic environment for teaching and experimenting formal verification principles, and for supporting PhD and master’s degree projects.

During the interactive exploration of the system behaviour, it is possible to observe all the internal details of the reached system configurations. For each system component we can observe: the values of its local variables; the status of the event queue; the set of currently active states and fireable transitions; and the set of possible next states reachable by a run-to-completion step performed by the component.

If we request the visualisation of the graph that models the possible system evolutions, we can click over a node of the graph to display all the internal details of the corresponding system state.

After the verification of a logic formula, it is possible to request a detailed explanation of how the evaluation result has been reached. Given that the supported logic is a branching-time logic, the counterexample does not have the shape of a simple execution trace. The explanation of the validity of a formula is indeed presented in an interactive way, and at any step of the explanation it is possible to observe all the internal details of the involved system configurations.

These characteristics of the UMC environment make it particularly suitable for the analysis and verification of designs in the early stages of system development, when the basic structures and ideas are being initially drawn, and are still likely to contain errors that the tool can discover, and allow the user to make sense of them.

The decision to make the overall framework publicly accessible through the web is driven by the desire to make UMC and the other tools accessible without overhead from any kind of platform (Unix, Linux, Windows, macOS) to all interested parties, while maintaining centralised control over its continuous improvement.

On the other hand, the web encapsulation allows a transparent integration of the locally developed tools (which are command-line oriented) with features provided from other frameworks (like minimisations with ltsmin and visualisation with graphviz), and allows the dynamic interactions with the user to be exploited in a natural and user-friendly way. For example, when the user clicks over a node of a visualised graph, a command is dispatched to a model exploration generation tool, which generates a new fragment of the state-space. This is saved as a '.dot' file, converted into '.svg' by another tool, embedded into an '.html' document and visualised again as a graph to the user. We are not aware of any other free verification environment that provides a comparable support for the dynamic analysis of UML system designs.

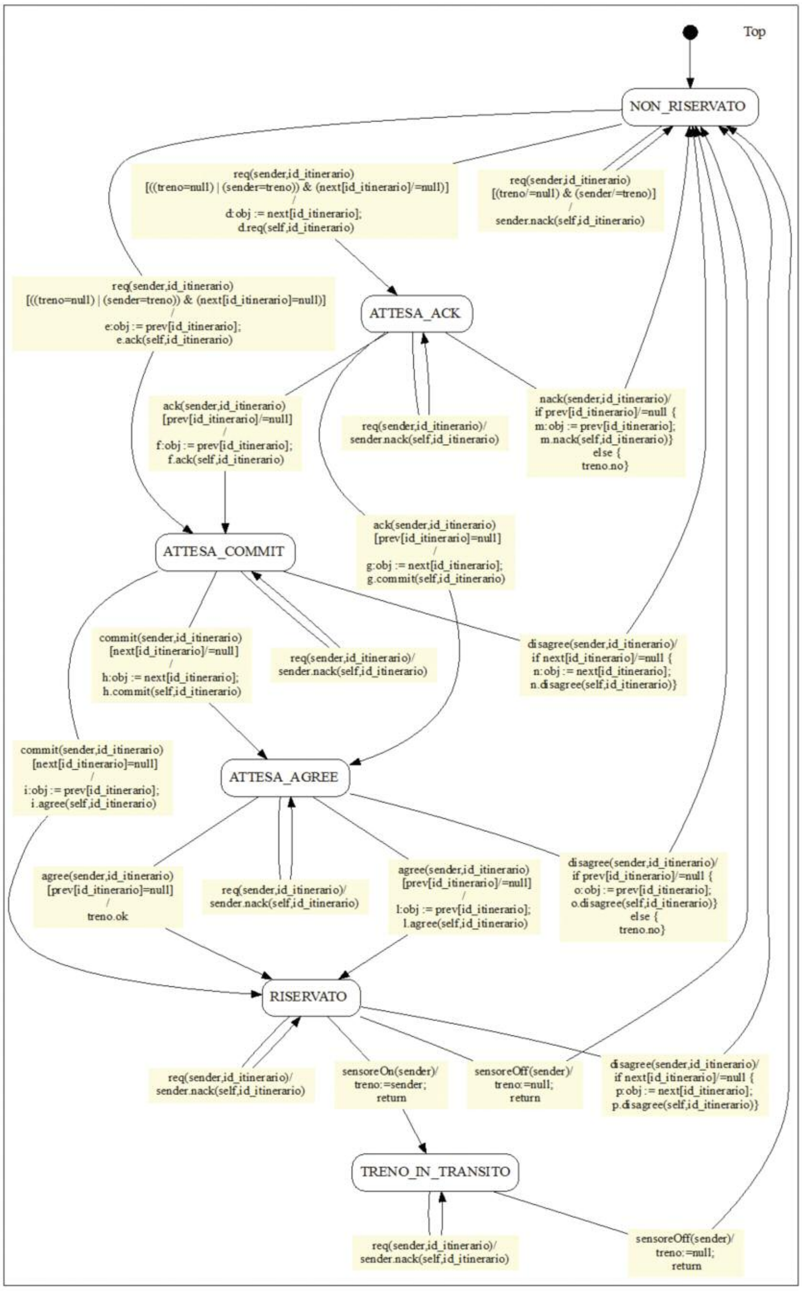
Top
NON_RISERVATO
req(sender,id_itinerario)
[((treno=null) | (sender=treno)) & (next[id_itinerario]/=null)]
/
d:obj := next[id_itinerario];
d.req(self,id_itinerario)
req(sender,id_itinerario)
[(treno/=null) & (sender/=treno)]
/
sender.nack(self,id_itinerario)
req(sender,id_itinerario)
[((treno=null) | (sender=treno)) & (next[id_itinerario]=null)]
/
e:obj := prev[id_itinerario];
e.ack(self,id_itinerario)
ATTESA_ACK
ack(sender,id_itinerario)
[prev[id_itinerario]/=null]
/
f:obj := prev[id_itinerario];
f.ack(self,id_itinerario)
req(sender,id_itinerario)/
sender.nack(self,id_itinerario)
nack(sender,id_itinerario)/
if prev[id_itinerario]/=null {
m:obj := prev[id_itinerario];
m.nack(self,id_itinerario)}
else {
treno.no}
ATTESA_COMMIT
ack(sender,id_itinerario)
[prev[id_itinerario]=null]
/
g:obj := next[id_itinerario];
g.commit(self,id_itinerario)
commit(sender,id_itinerario)
[next[id_itinerario]/=null]
/
h:obj := next[id_itinerario];
h.commit(self,id_itinerario)
req(sender,id_itinerario)/
sender.nack(self,id_itinerario)
disagree(sender,id_itinerario)/
if next[id_itinerario]/=null {
n:obj := next[id_itinerario];
n.disagree(self,id_itinerario)}
commit(sender,id_itinerario)
[next[id_itinerario]=null]
/
i:obj := prev[id_itinerario];
i.agree(self,id_itinerario)
ATTESA_AGREE
agree(sender,id_itinerario)
[prev[id_itinerario]=null]
/
treno.ok
req(sender,id_itinerario)/
sender.nack(self,id_itinerario)
agree(sender,id_itinerario)
[prev[id_itinerario]/=null]
/
l:obj := prev[id_itinerario];
l.agree(self,id_itinerario)
disagree(sender,id_itinerario)/
if prev[id_itinerario]/=null {
o:obj := prev[id_itinerario];
o.disagree(self,id_itinerario)}
else {
treno.no}
RISERVATO
req(sender,id_itinerario)/
sender.nack(self,id_itinerario)
sensoreOn(sender)/
treno:=sender;
return
sensoreOff(sender)/
treno:=null;
return
disagree(sender,id_itinerario)/
if next[id_itinerario]/=null {
p:obj := next[id_itinerario];
p.disagree(self,id_itinerario)}
TRENO_IN_TRANSITO
req(sender,id_itinerario)/
sender.nack(self,id_itinerario)
sensoreOff(sender)/
treno:=null;
return

Figure 1: A component of an interlocking model.

The KandISTI project is a long-term ISTI internal project. We have plans for improving the framework in several directions, among which a greater integration with other verification frameworks (like SPIN, LTSmin, CADP, NuSMV, mCRL2, DiVinE), a better exploitation of parallel/multicore architectures and the support of further specification/design languages.

Link:
[L1] http://fmt.isti.cnr.it/kandisti

References:

[1] M.H. ter Beek, S. Gnesi, F. Mazzanti: "From EU projects to a family of model checkers", Software, Services, and Systems, LNCS Vol. 8950, Springer, 2015, 312-328.

[2] F. Mazzanti, G.O. Spagnolo, S. Della Longa, A. Ferrari: "Deadlock avoidance in train scheduling: A model checking approach", in Proc. of the 19th International Workshop on Formal Methods for Industrial Critical Systems (FMICS'14), LNCS Vol. 8718, Springer, 2014, 109-123.

[3] M.H. ter Beek, S. Gnesi, F. Mazzanti, C. Moiso: "Formal Modelling and Verification of an Asynchronous Extension of SOAP", in Proc. of the 4th IEEE European Conference on Web Services (ECOWS'06), IEEE, 2006, 287-296.

Please contact:

Franco Mazzanti
ISTI-CNR, Italy
franco.mazzanti@isti.cnr.it

How Ecosystems of e-Infrastructures can support Blue Growth

The BlueBRIDGE H2020 project convened a one-day workshop on 11 January 2017 in Brussels entitled “Understanding how ecosystems of e-infrastructures can support Blue Growth”. The aim of the event was to showcase how the Blue Growth sector is benefitting from services that the BlueBRIDGE project has to offer. The latter primarily regard the combination of existing e-infrastructure services with an easy to use interface. The workshop also provided the opportunity to further discuss opportunities related to data management with e-Infrastructures.

Panel session at the BlueBRIDGE workshop.

The rapid developments of on-line data services and exponential growth of data quantity require collaborative and computing intensive solutions that can only be available through e-Infrastructures. A characteristic of e-Infrastructures is the re-use of existing solutions. Re-use of existing services is an effective way to save time and money. This is the approach adopted by the BlueBRIDGE project and the thread of the presentations of the event. BlueBRIDGE delivers specialised data services for aquaculture farm management, the ecosystem approach to fisheries, spatial data analysis, and marine spatial planning. These services are operated through Virtual Research Environments built on top of a hybrid-data infrastructure. This hybrid-data infrastructure is capable of dynamically federating computing and storage resources coming from third party providers, datasets belonging to heterogeneous data sources, analytical tools developed by different organisations, and to offer all of these services in a unique collaborative environment.

Several users have already adopted BlueBRIDGE services to solve their data management issues.

The workshop [L1] brought together stakeholders operating in the Blue Growth sector, all with data management challenges needing integrated solutions. Blue Growth is the long term strategy of the European Commission Directorate-General for Maritime Affairs and Fisheries to support sustainable growth in the marine and maritime sectors. The stakeholders ranged from international fisheries organisations, to organisations providing scientific advice and recommendations to policy makers and governments on marine protected areas and fish stocks, to private enterprises providing support to aquafarms in evaluating their economic performance, or in detecting areas to locate cages, to academics and scientists supporting international fisheries commissions in analysing marine species or performing environmental, societal and nutritional research. Representatives from the European Commission were also present. The full list of participants is available at [L2].

The Workshop discussion focused on three main themes:

1. How data-driven science is shaping Blue Growth practices: Practical Insights from BlueBRIDGE Users
2. How European e-infrastructures can support Blue Growth related data science and innovation
3. How the Blue Growth sector can benefit from the European Open Science Cloud

Agenda and workshop presentations are available at [L1].

The main findings of the report have been published a report, which is structured as follows:

- Section 1 presents the BlueBRIDGE offer and approach.
- Section 2 describes how users from the Blue Growth sector are benefitting from the adoption of the BlueBRIDGE services.
- Section 3 presents the main findings of the final panel discussion focusing on the role that data-driven science and innovation plays in the Blue Growth sector, on how European e-infrastructures concretely facilitate this role and on how the Blue Growth sector will benefit from the European Open Science Cloud.

The report is available for download at [L3].

The BlueBridge project is coordinated by ISTI-CNR, ERCIM EEIG provides administrative and financial support.

[L1] www.bluebridge-vres.eu/agenda-bluebridge-workshop-11-january-2017-brussels-belgium
[L2] bluebridge-vres.eu/participants
[L3] www.bluebridge-vres.eu/sites/default/files/january-workshopreport.pdf

Please contact:
info@bluebridge-vres.eu

Call for Submissions

2017 Best Practices in Education Award

The Informatics Europe Best Practices in Education Award is devoted to initiatives making Informatics education available to all. To continue promoting teaching quality in Informatics, Informatics Europe presents every year the Best Practices in Education Award, which recognizes outstanding European educational initiatives that improve the quality of Informatics teaching and the attractiveness of the discipline, and that can be applied beyond their institutions of origin. On its 6th edition, the Award is devoted to curriculum and professional development initiatives for making Informatics education available to all.

It will honor original contributions that emphasize successful initiatives for the teaching of informatics fundamentals in one of the following areas:

- Reaching out to non-traditional audiences, e.g., in continuing professional development or to senior citizens.
- Educating the general public, e.g., with respect to data security and privacy.
- Including Informatics education in other curricula, e.g., in general teacher training.

The Award, which carries a prize of 5.000 EUR, will be presented at the 13th European Computer Science Summit, in Lisbon, 23-25 October 2017, where the winner or winners will be invited to give a talk on their achievements.

www.informatics-europe.org/awards/education-award/call-for-submissions.html

Call for Nominations

Minerva Informatics Equality Award

The 2017 Informatics Europe Minerva Award, sponsored by Google, is devoted to initiatives supporting the transition of female PhD and postdoctoral researchers into faculty positions.

This is the second edition of the Minerva Informatics Equality Award. The Award seeks to celebrate successful initiatives that have had a measurable impact on the careers of women within the institution. Such initiatives can serve as exemplars of best practices within the community, with the potential to be widely adopted by other institutions. Nominations will need to demonstrate the impact that has been achieved.

For 2017 examples of impact could include an improved career development and better agreements on career planning for female PhD students and postdocs as recorded in objective surveys of staff experience, and increasing numbers of female faculty. The Award carries a prize of € 5,000

The Informatics Europe Minerva Informatics Equality Award recognizes best practices in Departments or Faculties of European universities or research labs that have been demonstrated to have a positive impact for women. On a three-year cycle the award will focus each year on a different stage of the career pipeline:

- Developing the careers of female faculty, including retention and promotion;
- Supporting the transition of female PhD and postdoctoral researchers into faculty positions;
- Encouraging female students to enroll in Computer Science/Informatics programmes and retaining them.

Deadlines:

- Full nominations: June 1, 2017
- Notification of winner(s): August 1, 2017

http://www.informatics-europe.org/awards/minerva-informatics-equality-award/cal-for-submissions-2017.html

A Digital Humanities Award for the EAGLE Project

The EU funded EAGLE Project – whose data aggregation and image processing infrastructure was implemented by the NeMIS Lab of ISTI-CNR – won the Digital Humanities Awards 2016 for the “Best DH Tool or Suite of Tools” category.

Digital Humanities Awards are a set of annual awards where the public is able to nominate resources for the recognition of talent and expertise in the digital humanities community. These awards are intended as an awareness raising activity, to help put interesting DH resources in the spotlight and engage DH users (and general public) in the work of the community.

EAGLE (Europeana network of Ancient Greek and Latin Epigraphy, a Best Practice Network partially funded by the EU) aggregates epigraphic material provided by some 15 different epigraphic archives for ingestion to Europeana (www.europeana.eu/). The aggregated material is made available to the scholarly community and the general public for research and cultural dissemination. EAGLE has defined a common data model for epigraphic information, into which data models from different archives can be optimally mapped. The data infrastructure is based on the D-NET software toolkit (www.d-net.research-infrastructures.eu) developed by the InfraScience Research Group of the NeMIS Lab. A novel search feature offered by EAGLE and developed by the Multimedia Information Retrieval Research Group of the NeMIS Lab is the possibility of visually searching for epigraphic information. A picture of an inscription can be used as a query to search for similar inscriptions or to obtain information on it. Visual inscription search leverages jointly on deep learning techniques and index structures for large-scale similarity searching.

www.eagle-network.eu
dhawards.org/dhawards2016/results/

ERCIM
European Research Consortium
for Informatics and Mathematics

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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