Collaboration workshop

ADVANCED COMPUTING AND CYBER-PHYSICAL SYSTEMS 2016

ADEPT

Addressing Energy in Parallel Technologies
http://adept-project.eu
1 September 2013 – 31 August 2016

OVERALL OBJECTIVES

The objective of the Adept project is to better understand energy and power use in existing parallel software and hardware, and to use this knowledge to predict power use and energy efficiency. Adept brings together experts from both the high-performance and Embedded computing sectors, and utilizes their expertise to advance knowledge about the efficiency and power profiles of systems. Adept uses this knowledge to predict how parallel technologies will use power and energy – even if the system being predicted does not exist yet. This allows for the development of more economical, energy-efficient systems without the need for speculation.

MAIN RESULTS

The major results of the project to date are:

THE ADEPT POWER MEASUREMENT SYSTEM

One of the key outcomes of the project is our sophisticated Adept Power Measurement System (APMS). This fine-grained measurement infrastructure reads the current and voltage from the powerlines that feed the different components of a computer system, e.g. CPU, memory or disk.

THE ADEPT BENCHMARK SUITE

To complement the APMS, the Adept project has also developed a diverse suite of benchmarks that can be used to test and evaluate existing systems. The benchmarks are designed to be used for system characterisation and target specific operations and common computational patterns.

THE ADEPT PERFORMANCE AND POWER PREDICTION TOOL

Another important outcome of the project is our Performance and Power Prediction Tool. Using detailed statistical modelling that examines a software binary, we can predict how well a CPU and memory hierarchy system will perform and how power efficient it will be, even if we do not have access to that system or even if that system does not yet exist. The Adept tool will impact on software developers and system designers, by freeing them from making poorly informed decisions about how to implement changes to their systems.

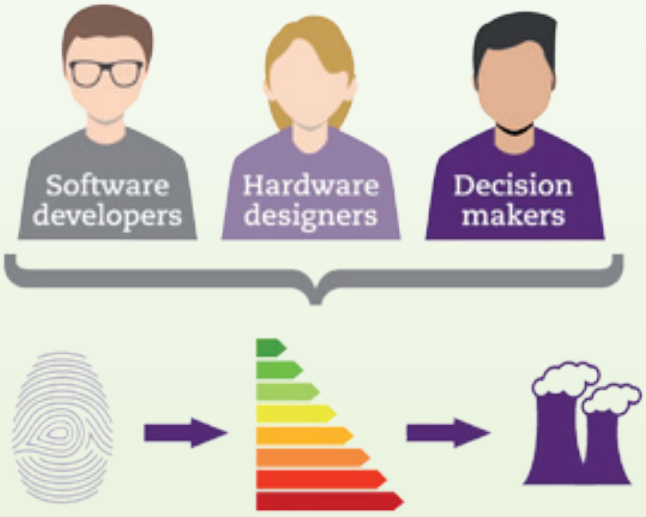

POTENTIAL IMPACT

Adept's advances, and the tools developed within the project, allow for the design of smarter, cheaper, and more efficient systems. The cost of ownership of these systems is high – and energy is a significant portion of the total cost. Thus, lowering consumption in these systems is of significant interest to the community, as the development of energy-efficient technologies could lower financial barriers. Additionally, giving owners and developers the freedom and flexibility to know how their equipment will perform prior to porting their workloads means giving them the ability to make better choices about what they implement, how, and when.

COORDINATOR + CONTACT INFORMATION

Michèle Weiland, EPCC, Email: m.weiland@epcc.ed.ac.uk,
Phone: 0131 651 3580

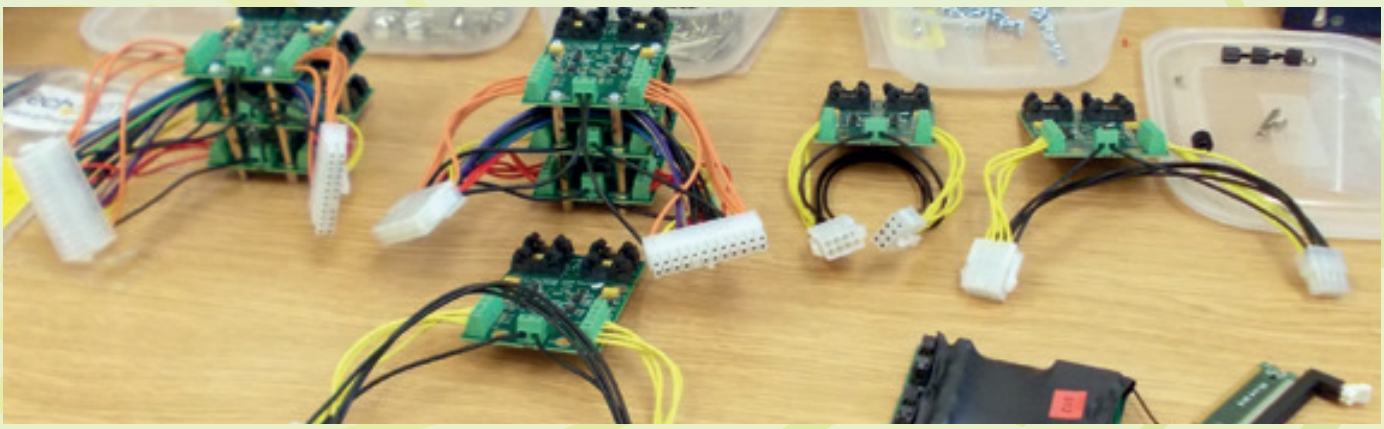

AMADEOS

Architecture for Multi-criticality Agile Dependable Evolutionary Open System-of-Systems
http://amadeos-project.eu/
1 October 2013 – 30 September 2016

OVERALL OBJECTIVES

The objective of AMADEOS is to bring time awareness and evolution into the design of System-of-Systems (SoS), to establish a *sound conceptual model, a generic architectural framework and a design methodology*, supported by some prototype tools, for the *modeling, development and evolution of time-sensitive SoSes with emergent behaviors*. Special emphasis is placed on evolution, emergence, dependability (e.g. safety, availability) and security. The viability of the framework is being validated on a case study of a CPS, a small smart grid application, where guaranteed responsiveness, evolution, dependability and security are essential requirements.

MAIN RESULTS

The main results AMADEOS project include:

- **A Conceptual Model** and an architectural framework for SoSes based on judiciously formulated basic concepts for this important new domain. The architecture is based on the assumption that a global time is available in every CS and will focus on CPS applications.

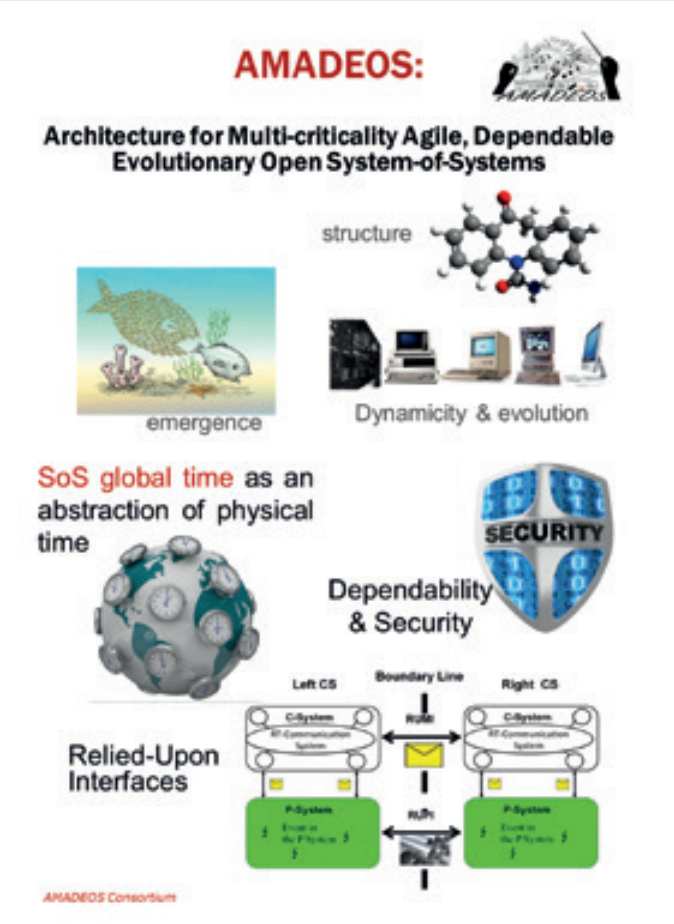

- **A proof-of-concept prototype of a Resilient Master Clock (RMC)**, to guarantee a resilient global time base across the SoS, thus validating the previous architectural assumption. RMC provides local clock corrections to compensate local clock deviations caused by physical environment variations, self-estimation of time uncertainty and fault tolerant synchronization solutions.
- Thorough investigations of the **emergence phenomena** in cyber-physical SoSes, and advances in their understanding to manage, control, predict their occurrence and avoid or mitigate the detrimental ones.
- Deep understanding of the key role of **stigmergic** (i.e., physical) **channels** between Constituent Systems, which represent primary means through which interactions (and emergent phenomena) are established and determine feedback loops, where constituent systems also affect the physical environment and vice-versa.
- A **methodology for the design of dynamic SoSs**, supported by a *SysMl profile* for SoSes and Eclipse-compliant editor extensions, that puts a focus on timeliness, evolution, dependability, security and emergent properties of the developing artifact.

POTENTIAL IMPACT

Reinforce of European technological leadership: With the AMADEOS framework industries can ensure a long term managed evolution of large scale SoSs. The approach exhibits long term vision while reducing and mastering complexity, providing practical means to improve current SoSs incorporating legacy systems.

Improved systems characteristics: By focusing on the temporal aspects, AMADEOS provides key concepts/methods for the improvement of the real time characteristics of cyber physical systems. The methodology further supports the long term evolution of new and legacy systems and improves their short term dynamicity and control of emergent properties, while providing guarantees for best adaptation to changing circumstances.

COORDINATOR + CONTACT INFORMATION

Andrea Bondavalli
Full Professor
Dip. di Matematica e Informatica "U. Dini"
Viale Morgagni 65
50134 – Firenze
Email: bondavalli@unifi.it

ARGO

WCET-Aware Parallelization of Model-Based Applications for Heterogeneous Parallel Systems

www.argo-project.eu

1 January 2016 – 31 December 2018

PROBLEM:

In aerospace, automation, and automotive technologies, smart electronic computer systems have to meet a number of security and real-time requirements. In case of critical incidents, for instance, the software's response time has to be very short. To meet the requirements, these safety-critical embedded electronic solutions are more and more based on high-performance, energy-efficient and heterogeneous multi-core processors. Programming corresponding applications for such heterogeneous computing systems is time- and cost-consuming. Especially for legacy code, engineers need to have a deep knowledge about the highly complex underlying hardware architecture in order to be able to manually adapt the application to an optimized and parallelized version running efficiently on the hardware.

OVERALL OBJECTIVES

To overcome this challenge, the ARGO project intends to provide a cross-layer programming approach to exploit the full potential of next generation heterogeneous parallel embedded systems respecting real-time constraints. The cross-layer programming combines the following technologies in a holistic approach:

- Cross-layer programming user interface
- WCET-aware automatic parallelization
- WCET-analysis for heterogeneous multi- and many-core architectures
- Model-based development and testing

The goal is to drastically increase the productivity and shorten the time-to-market, reduce porting effort to ease the use of legacy code, increase worst-case performance and reduce the gap between worst- and average-case performance. The major objective is to enable and prototype a holistic integrated toolchain, which enables real-time programmability of multi- and many-core architectures in a model-based design workflow.

MAIN RESULTS

Within the ARGO project, a new standardizable tool chain for programmers is being developed. The parallelization process is starting with the model-based development of applications based on Scilab/Xcos models. Therefore, users are able to rapidly develop, test and validate their applications and algorithms. Even without precise knowledge of the complex parallel processor hardware, the programmers can control the process of automatic parallelization in accordance with the real-time requirements. The user is offered guidance in the parallelization process that guarantees deterministic analysis and keeps the control on the parameters during the whole process. This results in a significant improvement of performance and a reduction of costs.

In the future, the ARGO toolchain can be used to manage the complexity of parallelization and adaptation to the heterogeneous multi-core target hardware in a highly automated manner with a small expenditure. During the project run-time, real-time critical applications in the areas of real-time flight dynamics simulation and real-time image processing are investigated and as realistic use cases. However, the approach can be transferred to any other application domain demanding for safety-critical and guaranteed real-time constraints.

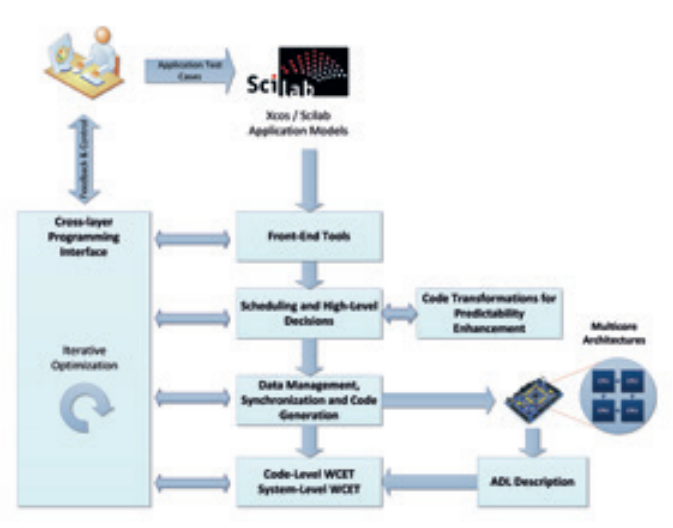

POTENTIAL IMPACT

The ARGO project encompasses system modelling, Software in the Loop (SIL), Hardware in the Loop (HIL), parallelization and Worst Case Execution Time (WCET) in a single framework with customization and visualization at each level of the embedded system design.

The WCET approach aims to guarantee the real-time response for critical systems in the European key sectors of aerospace, aircraft, transportation, automation (e.g. Industry 4.0) and automotive. Because of the critical impacts of the tasks performed by embedded systems, any development aiming at simplifying the code generation has a strong impact on productivity and competitiveness in these sectors.

The ARGO project brings together industrial players and chip manufacturers in the WCET approach while enabling the full usage of multi-core capabilities and latest parallelization techniques to systems engineering. Furthermore, the ARGO project enables great potential for ongoing projects, since the results can be used as a basis or addition in other projects related to multi- or many-core technology and programming.

COORDINATOR + CONTACT INFORMATION

Prof. Dr.-Ing. Dr. h. c. Jürgen Becker - becker@kit.edu
Karlsruhe Institute of Technology (KIT)
Institute for Information Processing Technology (ITIV)
Engesserstr. 5 - 76131 Karlsruhe - German

AXIOM

Agile, eXtensible, fast I/O Module for the cyber-physical era
Link to the web site of the project: http://www.axiom-project.eu/
1 February 2015 – 31 January 2018

OVERALL OBJECTIVES

- Goal: European-designed and -manufactured single board computer: The heart of future smart applications
- Flexible, energy efficient and multi-board
  - Flexibility: FPGA, fast-and-cheap interconnects based on existing connectors like SATA
  - Energy efficiency: low-power ARM, FPGA
  - Modularity: board-to-board fast interconnects
  - ...
- Easily Programmable FPGA
  - Programming model: Improved OmpSs
  - Runtime & OS: improved thread management
  - Compiler: BSC Mercurium, OS: Linux, Drivers: provided as open-source by partners
- Easy Interfacing with the Cyber-Physical Worlds
  - Platform: integrating also Arduino support for a plenty of pluggable board (so-called "shields")
  - Platform: building on the UDOO experience from SECO

MAIN RESULTS

As of the first year of the project we demonstrated:

- Dataflow-based execution model that spawns threads across single and multiple boards
- OpenMP based programming model to both accelerate applications on FPGAs and on multiple boards
- Inexpensive and fast board-to-board interconnect

POTENTIAL IMPACT

- Open-Source, Open-Hardware models
- Production of AXIOM based boards (more advanced that Raspberry, Zynqberry, and similar ones)
- Extending the Programming Model based on OpenMP

COORDINATOR + CONTACT INFORMATION

Roberto Giorgi – Giorgi@dii.unisi.it

CLERECO

Cross Layer Early Reliability Evaluation for the Computing Continuum
http://www.clereco.eu
1 October 2013 – 30 September 2016

OVERALL OBJECTIVES

The FP7 Project CLERECO investigates new design methods for early reliability evaluation of digital systems in the forthcoming computing continuum. Reliability is becoming an ever-increasing challenge for the ICT and must be guaranteed without penalizing or slowing down the characteristics of the final products. CLERECO research project recognizes the importance of accurately evaluating the reliability of systems early in the design cycle to be one of the most important and challenging tasks toward this goal. Being able to precisely evaluate the reliability of a system means being able to carefully plan for specific countermeasures rather than resorting to worst-case approaches.

MAIN RESULTS

The CLERECO design framework offers a comprehensive tool-suite that allows systems designers to precisely evaluate the reliability of their systems in the early design phases when key decisions must be taken. The CLERECO tool-suite offers tools, models and technologies that cover the four main design dimensions of a system: technology (RELTech Tools), Hardware Architecture (RelHW Tools), Software (RelSW Tools) and System integration (RelSyst Tools). Additional details are available on the project website.

There is a set of key innovations that CLERECO provides over existing competitors in academia and industry:

a) Assessment of **all system layers** (from technology to System Software).

b) Tools that work at **all design refinement stages**, from early conceptual and specification phases, to architectural design phases.

c) A design methodology oriented to the **computing continuum**, thus able to provide tools that are not restricted to a single application domain.

d) Introduction of **statistical models** for the evaluation of the system reliability that strongly reduce the need for extremely time consuming and costly fault injection campaigns.

e) Instruments for **design exploration and inspection** that enable to automatically identify weaknesses of the system and to provide automation for the optimization of the design.

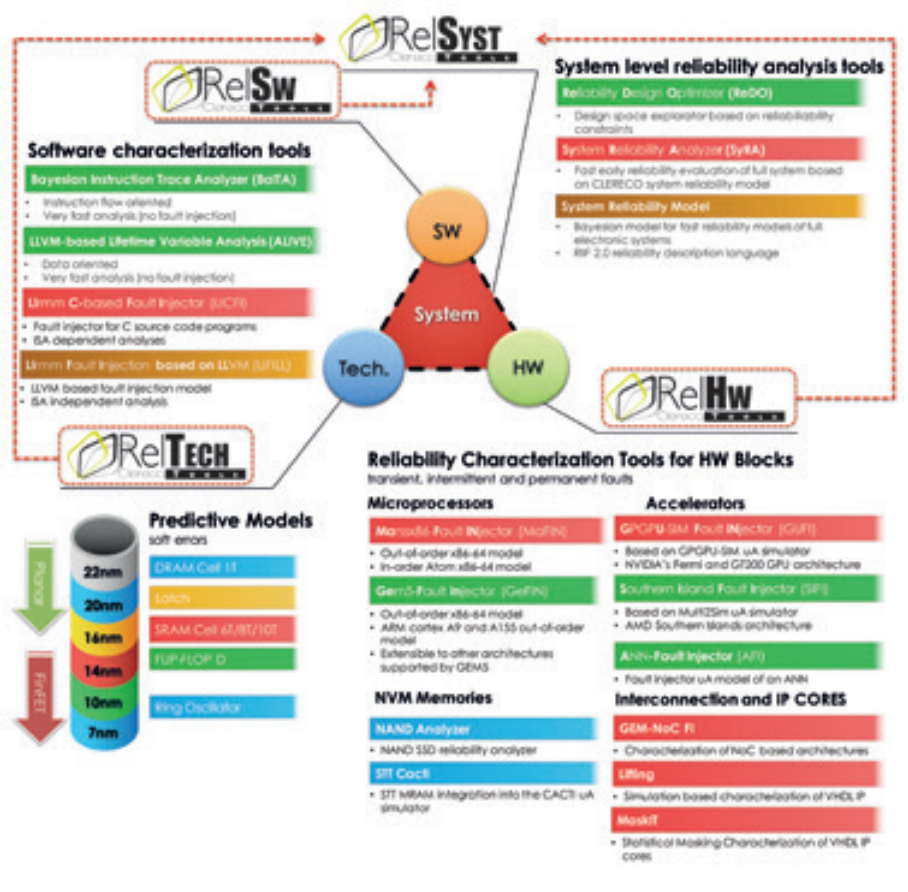

POTENTIAL IMPACT

The CLERECO framework takes a significant step forward in the introduction of reliability as a design dimension to trade-off with performance, power-consumption and costs starting from the beginning of the design process. This will impact reliability of future systems for the computing continuum reducing costs contrary to existing worst-case reliability design solutions. This will indirectly impact all applications where such systems play a major role for our society ranging from avionics, automobile, smartphones, mobile systems, Personal Computers WPCsY and future servers utilized in the settings of Data Centers, Grid Computing, Cloud Computing and other types of HPC systems.

COORDINATOR + CONTACT INFORMATION

Stefano Di Carlo
Politecnico di Torino –
Control and Computer Department
Email: stefano.dicarlo@polito.it
Phone: +39 011 090 7080

CONTREX

Design of embedded mixed-criticality CONTRol systems under consideration of EXtra-functional properties
Web: http://contrex.offis.de/
1 October 2013 – 30 September 2016

OVERALL OBJECTIVES

The increasing processing power of today's HW/SW platforms leads to the integration of more and more functions in a single device. Additional design challenges arise when these functions share computing resources and belong to different criticality levels. CONTREX complements current activities in the area of predictable computing platforms and segregation mechanisms with techniques to consider the extra-functional properties, i.e., timing constraints, power, and temperature. CONTREX enables energy efficient and cost aware design through analysis and optimization of these properties with regard to application demands at different criticality levels.

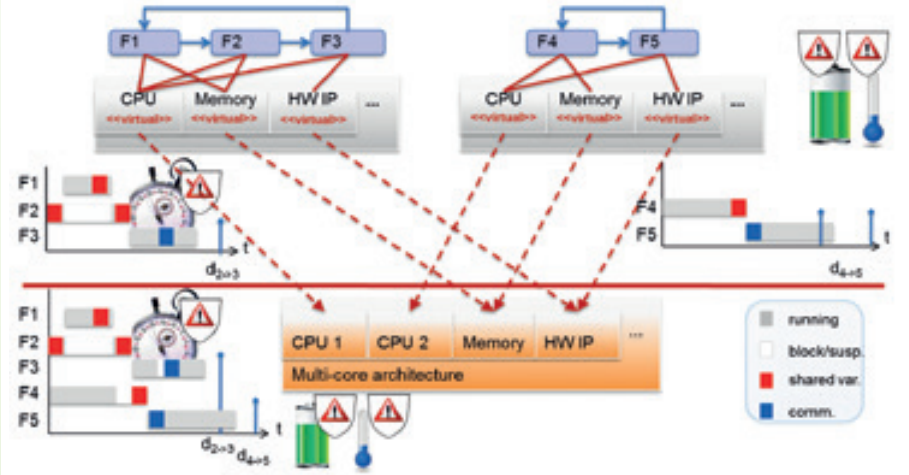

MAIN RESULTS

The UML/MARTE meta-model has been extended for the modelling of extra-functional properties, criticalities of components and the description of networks. Furthermore, code generators and analysis tools for functional, timing, power and temperature properties have been implemented. To support decisions and optimization for costs in the design process, virtual integration testing and a Joint Analytical Simulative Design Space Exploration framework can be used. A tracing framework facilitates virtual integration testing for power and temperature based on virtual platforms. It allows the extraction of activity traces from virtual platforms and their processing for simulation and analysis of functional and extra-functional properties. In addition, a methodology for the automatic generation of virtual platforms starting from components described with different languages and at different abstraction levels has been developed. The results are completed by resource management and service abstraction layers to consider extra-functional properties at runtime. The methods and tools have been successfully used with three industrial mixed-criticality systems from different domains: Avionics Computer of a remotely piloted aircraft, Ethernet Over Radio System, and Automotive Telematics System. The application of the developed methods and tools led to time and cost savings in the development, improved power consumption, size, and weight, as well as functional improvements made possible by better power/performance characteristics.

POTENTIAL IMPACT

Expected final results and potential impact:

- Increase of energy efficiency of next-generation embedded systems and reduction of their development cost per unit through the improvement of the design flow and the consideration of extra-functional properties in the verification phase.
- Reinforce European technological leadership and industrial competitiveness in the design, operations, and control of embedded systems with mixed criticalities and System-of-Systems.

COORDINATOR + CONTACT INFORMATION

Dr. Kim Grüttner
OFFIS - R&D Division Transportation
Escherweg 2 - 26121 Oldenburg - Germany
Phone/Fax.: +49 441 9722-228/-278
Email: contrex-mgt@offis.de

COSSIM

Novel, Comprehensible, Ultra-Fast, Security-Aware CPS Simulation
www.cossim.org
1 February 2015 – 31 January 2018

OVERALL OBJECTIVES

One of the main problems the CPS designers face is the lack of simulation tools and models for system design and analysis as well as the necessary processing power for executing the CPS simulations. The COSSIM framework address all those needs by providing an open-source system seamlessly simulating, in an integrated way, both the networking and the processing parts of the CPS while provide significantly more accurate results, especially in terms of power consumption, than existing solutions and report the, critical for many applications, security levels of the simulated CPS.

MAIN RESULTS

The COSSIM tool is an open-source framework that:

- Seamlessly simulates, in an integrated way, the networking and the processing parts of the CPS.
- Provides significantly more accurate results, especially in terms of power consumption, than existing solutions.
- Reports the, critical for many applications, security levels of the simulated CPS.

The novel COSSIM framework combines a state-of-the-art processing simulator (i.e. a "full-system simulator") with an established network simulator. These tools are integrated with high-level power estimators and the overall framework provides appropriate interfaces to security testing tools.

Moreover, COSSIM will be evaluated when simulating two high-end CPS applications: a novel building management infrastructure and an innovative Visual Search system.

POTENTIAL IMPACT

The COSSIM simulator will enable new business models for numerous service providers that will be able to utilize the unique features and the simulation speed provided by HPC-run COSSIM framework so as to deliver sustainable, high quality services based on novel CPS infra-structures to the citizen at home, in the road and everywhere. Moreover, the technology implemented in COSSIM will enable the development of applications in a number of areas that utilize CPS; for example new ambient assisted living, surveillance and security services will be available anywhere and anytime increasing the safety and the well being of the European citizens.

COORDINATOR + CONTACT INFORMATION

Yannis Papaefstathiou (Technical Manager)
Synelixis Solutions Ltd.
10 Farmakidou St.
Chalkida, GR34100,
GREECE
Email: ygp@synelixis.com

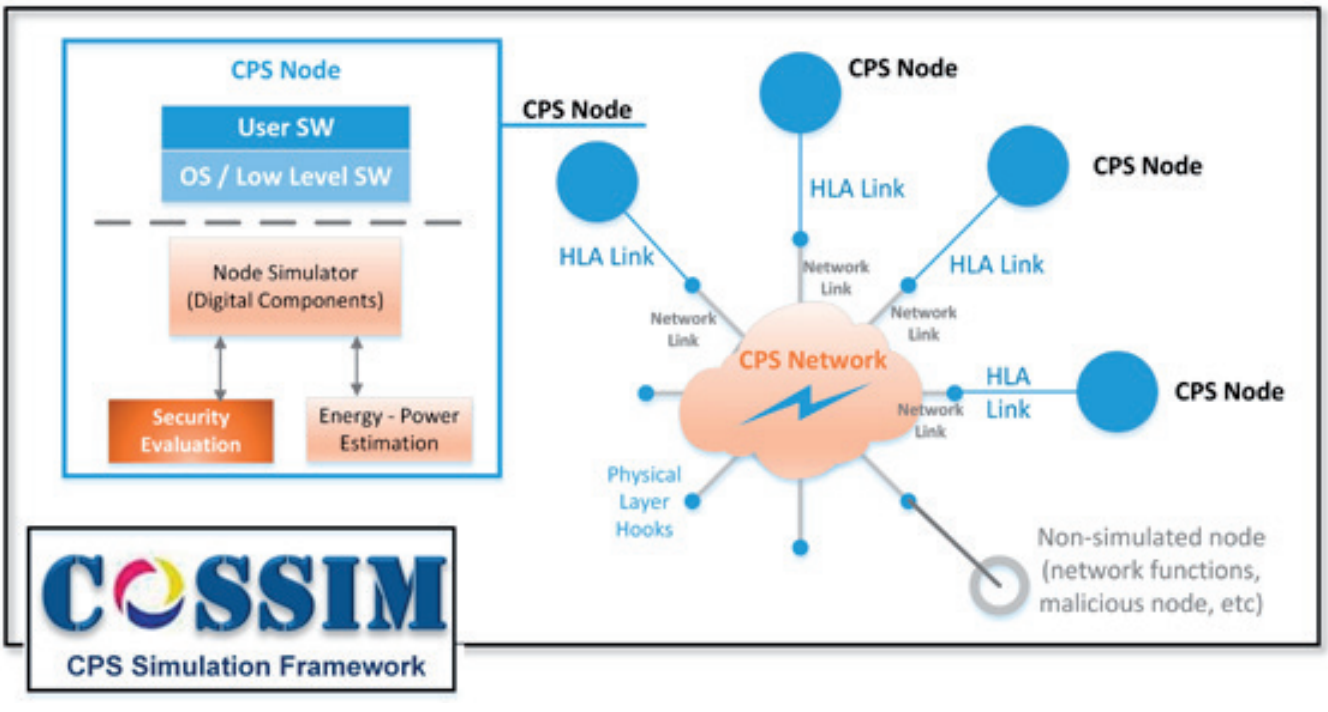

CP-SETIS

Towards Cyber-Physical Systems – Engineering Tools – Interoperability Standardisation
http://www.cp-setis.eu/
1 March 2015 – 28 February 2017

OVERALL OBJECTIVES

Cyber-Physical Systems (CPS) require multiple engineering competences across various disciplines. Because of the huge complexity at the development process, past and ongoing EU research projects have developed the basis for an International Open Standard for Development Tool Interoperability, the so called Interoperability Specification (IOS).

The main goal of CP-SETIS is to conceive and set up a sustainable organisational structure as a coordination platform joining all stakeholders, to coordinate all IOS-related activities, especially the formal standardisation and further extensions of the IOS.

CP-SETIS will ensure the support of all stakeholders for this structure, its implementation and their commitment to coordinate all IOS-related activities within this structure.

MAIN RESULTS

CP-SETIS has developed a model for a sustainable organisational structure called **ICF (IOS Coordination Forum)** as a cooperation platform in which all IOS stakeholders – CPS development organisations, Tool Provider, Research Organisations, Standardisation Bodies – can meet to synchronise and coordinate their IOS activities. Specifically, the ICF will

(a) collect and make available the current baseline of the IOS, together with information about the concrete specifications, maturity level, status of formal standardisation, current versions, etc., and update this information according to results from projects, standardisation activities, etc.

(b) facilitate and give organisational support for stakeholders to coordinate their activities to extend and further develop the IOS – for example by incubating new R&D projects,

(c) facilitate and give organisational support for stakeholders to synchronise their activities for formal standardisation of parts of the IOS, and

(d) support the building of an IOS community by collecting and proving all information related to IOS (from technical specifications and contacts to experts to workshop and event notifications) and organising workshops, coordination meetings, etc.

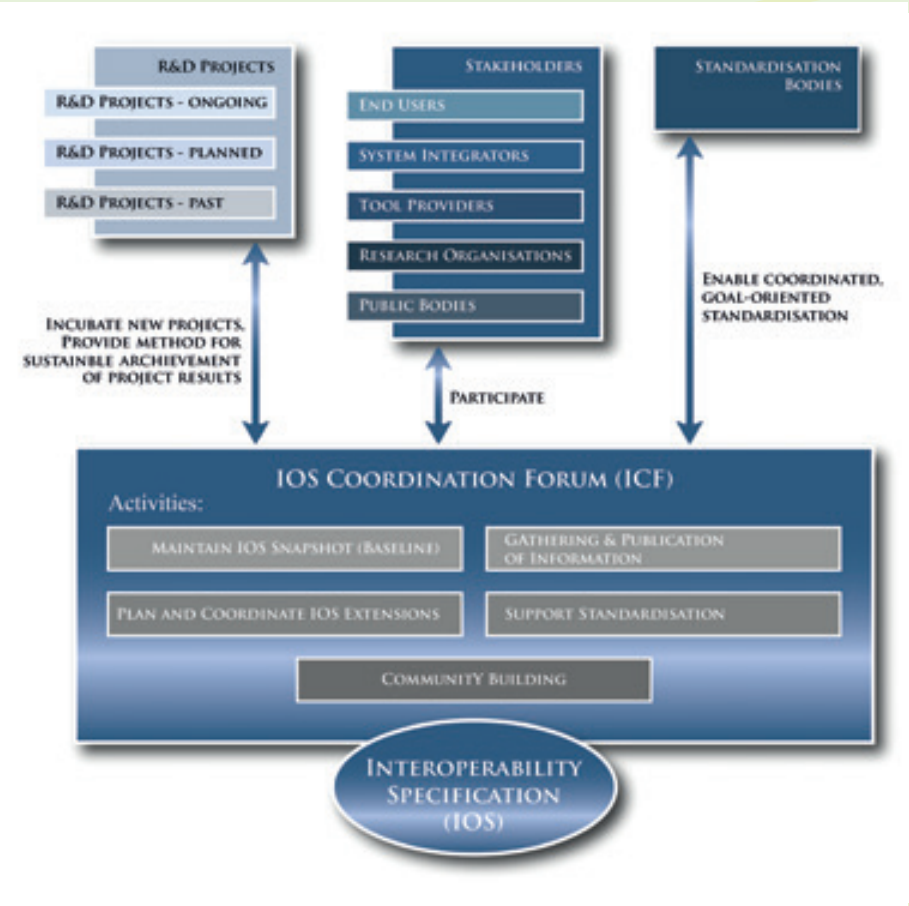

POTENTIAL IMPACT

CP-SETIS will push formal standardisation and industrial take-up of the IOS as an open standard for development tool interoperability, which will reduce the complexity and risk of installing and maintaining increasingly complex software infrastructures (Engineering Environments), avoid costly and inefficient in-house-developments and vendor lock-ins, and enable Tool Providers to focus energy and resources on higher-value functionality and customisation, thus providing time and cost savings.

The ICF set up within CP-SETIS will be an ideal forum for all stakeholders to coordinate IOS related activities, find allies and cooperation partners, and extend and shape those parts of the IOS that are relevant for them.

COORDINATOR:
SafeTRANS
Jürgen Niehaus
Escherweg 2
D-26121 Oldenburg
Phone: +49 441 / 9722 503
Email: juergen.niehaus@safetrans-de.org

CPSE LABS

CPS Engineering Labs - expediting and accelerating the realization of cyber-physical systems
www.cpse-labs.eu
1 Febraury 2015 – 31 January 2018

OVERALL OBJECTIVES

CPSE Labs' mission is to provide support for European companies to move into new markets and new application domains with innovative cyber-physical products & services, and to stimulate the uptake of advanced ICT technologies amongst Europe's SMEs. CPSE Labs has the ambition to

- Foster an open, pan-European network of design centres committed to transitioning science and technology for engineering trustworthy and dependable CPS into the marketplace.
- Identify, define, and execute focused and fast-track experiments with a specific innovation focus.
- Spread best CPS engineering practices and promote cross-regional and cross-sectoral learning among industry and academia.
- Establish a marketplace for CPS engineering assets.

MAIN RESULTS

CPSE Labs partners provide access to a broad range of state-of-the-art CPS design technologies and expertise in various CPS fields, including Internet of Things, industrial automation and control, autonomous vehicles, e-maritime applications, and model-based techniques for collaborative engineering and safety analysis and monitoring.

Since its start in 2015, CPSE Labs has initiated a number of fast-track industrial experiments that have been selected through competitive open calls. In these experiments, engineering and technology firms evaluate, use, and adapt the CPS design technologies and platforms supported by CPSE Labs, and build demonstrators to validate specific innovation objectives. Topics of our experiments include:

- Safe deployment of a mobile robot in an airfield runway for lighting maintenance
- Flexible development of an industrial control system for energy load management in a biogas plant using IEC 61499
- Decentralised architecture for traffic management systems

- Testing framework for efficient CPS tool chain integration
- Integration of co-simulation methods into a real-time platform for automotive CPS
- Augmenting legacy machine tools to cloud manufacturing environment
- Shore-based voyage planning using the Maritime Architecture Framework
- Improving Water Efficiency and Safety in Living Areas
- The first results from experiments are expected in late 2016.

POTENTIAL IMPACT

Through our portfolio of innovative experiments CPSE Labs stimulates the innovation capacity of the SMEs and other businesses working with us, and expects to facilitate lower cost and faster adoption of cyber-physical systems (CPS). The network of Design Centres established by CPSE Labs fosters stronger pan-European collaboration across value chains and technology levels, building an ecosystem around the CPSE Labs Design Centres committed to transitioning CPS technology into the marketplace. The Design Centres provide physical and virtual meeting points for all relevant stakeholders for CPS innovations, promoting best practices for sharing and learning.

COORDINATOR + CONTACT INFORMATION

Dr. Holger PFEIFER
fortiss
Guerickestraße 25
80805 München, Germany
Email: pfeifer@fortiss.org
Phone: +49 – 89 – 3603522-29

CPSOS

Towards a European Roadmap on Research and Innovation in Engineering and Management of Cyber-Physical Systems of Systems
http://www.cpsos.eu/
October 2013 – June 2016

OVERALL OBJECTIVES

CPSoS – Towards a European Roadmap on Research and Innovation in Engineering and Management of Cyber-physical Systems of Systems (CPSoS) – is a Support Action supported by the European Commission under the FP7 programme. Its main goal is to define a European Research and Innovation Agenda on CPSoS. CPSoS also provides a forum and an exchange platform for systems of systems related communities and ongoing projects. CPSoS has set up three working groups, two on application domains and one on methods and tools. Its approach is applications-driven and integrative, aiming at bringing together knowledge from different communities.

MAIN RESULTS

Cyber-physical systems of systems are large systems that consist of physical systems and computing, control and communication systems where the individual units have a certain degree of autonomy but are connected by physical couplings, and have overarching performance goals. Examples are railway systems, large industrial plants, or the electric power grid. The CPSoS project has performed an in-depth study of the state of the art in cyber-physical systems of systems, collecting input both from industry, including more than 15 SMEs, and from research institutions. It organized a number of workshops to discuss the challenges that arise in different application domains, to identify gaps in the existing methods and tools, and to communicate the proposed research agenda. The findings have been synthesized in the position paper "European research and innovation agenda on Cyber-physical Systems of Systems" (http://cpsos.eu/roadmap/). The document contains a number of medium-term research topics that should receive funding within the next 5 years in order to progress towards better design and operation of cyber-physical systems of systems. A volume of accompanying technical papers is in preparation. Results of CPSoS have been included in the recent edition of the ARTEMIS Strategic Research Agenda.

POTENTIAL IMPACT

The CPSoS project has proposed priorities and new avenues for the analysis, design and control of cyber-physical systems of systems. Implementing research and development on the priority topics that were defined by CPSoS will substantially reinforce Europe's scientific excellence and technological leadership in the area of designing and managing large cyber-physical systems that consist of a number of partly autonomously operating units and can show emerging behaviours.

COORDINATOR + CONTACT INFORMATION

Prof. Sebastian Engell
Head of the Process Dynamics and Operations Group (DYN)
Technische Universität Dortmund,
Emil-Figge-Str. 70, 44221 Dortmund, Germany
Phone: +49 (0) 231/755-5127
Email: s.engell@bci.tu-dortmund.de

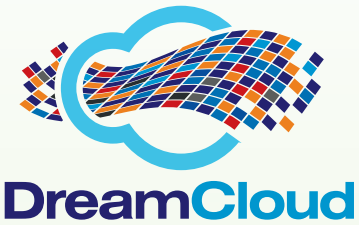

DREAMCLOUD

Dynamic Resource Allocation in Embedded and High-Performance Computing
www.dreamcloud-project.org
1 September 2013 – 31 August 2016

OVERALL OBJECTIVES

DreamCloud overall objective is to enable dynamic resource allocation in many-core embedded and high performance systems while providing appropriate guarantees on performance and energy efficiency. DreamCloud seeks to: 1) provide complex embedded systems with cloud-like capabilities available in today's high-performance computing, allowing them to dynamically tune resource usage without sacrificing critical constraints in performance and energy; and 2) enable HPC and cloud computing systems to balance workload and manage resources so they can offer more meaningful guarantees of performance and energy, focussing not only on improving average behaviour but also reducing variability and upper bounds of timing and energy metrics.

MAIN RESULTS

DreamCloud has developed novel load balancing mechanisms that can be applied during runtime in a wide range of many-core systems allowing for a fine-tuning of the trade-off between performance guarantees and system efficiency according to the specific needs of applications. Such mechanisms have been organised in distinct types of cloud-like system software infrastructure within DreamCloud that manage the workload on different kinds of systems. Embedded Clouds are used in systems with time-critical behaviour, allowing for restricted load balancing and privileging strict performance guarantees. Micro Clouds utilise novel extensions to operating systems and virtual machines, allowing for the dynamic migration of threads or full virtual machines between cores. High Performance Clouds are able balance highly dynamic workloads, aiming for full utilisation of the underlying platform while providing performance guarantees to selected applications.

DreamCloud brings together industrial partners from deeply embedded systems (e.g. automotive), consumer embedded systems (e.g. household media), and high performance computing (e.g. HPC platforms), as well as academic partners from embedded systems, real-time systems and HPC. This combination of expertise has enabled DreamCloud to develop innovative approaches to addressing complex resource allocation challenges by cross-fertilising expertise and experience from multiple industrial domains and academic communities.

POTENTIAL IMPACT

The key impact from DreamCloud technologies will be increased efficiency in terms of resource usage, which can bring about substantial benefits in terms of increased capacity for a many-core platform to reliably carry out computing tasks, as well as reductions in energy usage. This is a critical commercial issue across many different industries (e.g. transportation, entertainment, automation, financial services) and the DreamCloud approach to increased resource and energy efficiency centres around exploring the trade-off between the flexibility of the resource allocation and the strength of the guarantees that must be provided to the application end user.

COORDINATOR + CONTACT INFORMATION

The Open Group
www.opengrou.org
Name: Scott Hansen
Email: s.hansen@opengroup.org

DREAMS

Distributed REal-time Architecture for Mixed Criticality Systems
http://dreams-project.eu/
October 2013 – September 2017

OVERALL OBJECTIVES

The objective of DREAMS is to develop a cross-domain architecture and design tools for networked complex systems where application subsystems of different criticality, executing on networked multi-core chips, are supported. DREAMS will deliver architectural concepts, meta-models, virtualization technologies, model-driven development methods, tools, adaptation strategies and validation, verification and certification methods for the seamless integration of mixed-criticality to establish security, safety, real-time performance as well as data, energy and system integrity.

The project partners include major European companies (Alstom, STMicroelectronics, Thales, TÜV Rheinland, FENTISS, RealTime-at-Work, TTTech, Virtual Open Systems) and leading research organizations (fortiss, IKERLAN, ONERA, UPV, SINTEF, TEI, TUKL and University of Siegen).

MAIN RESULTS

DREAMS results in a flexible platform and associated design tools for embedded applications where subsystems of different criticality, executing on networked multi-core chips, can be integrated seamlessly. The objective is a unified view of the system through systematic abstraction of the underlying hierarchic network topology and related constraints. The platform will encompass both the chip and cluster-level and provide support for security, safety and real-time performance as well as data, energy and system integrity. DREAMS will offer pervasive support for design, modelling, verification, validation and certification up to the highest criticality levels in multiple domains (e.g., avionics, wind power, healthcare). The cross-domain architecture will enable synergies between application domains and exploitation of the economies of scale.

The major results include

- Architectural style und modelling methods based on waistline structure of platform services
- Virtualization technologies to achieve security, safety, real-time performance as well as da-ta, energy and system integrity in networked multi-core chips
- Adaptation strategies for mixed-criticality systems to deal with unpredictable environment situations, resource fluctuations and the occurrence of faults
- Development methodology and tools based on model-driven engineering
- Certification strategies for mixed-criticality product lines
- Feasibility of DREAMS architecture in real-world scenarios
- Community building for widespread adoption

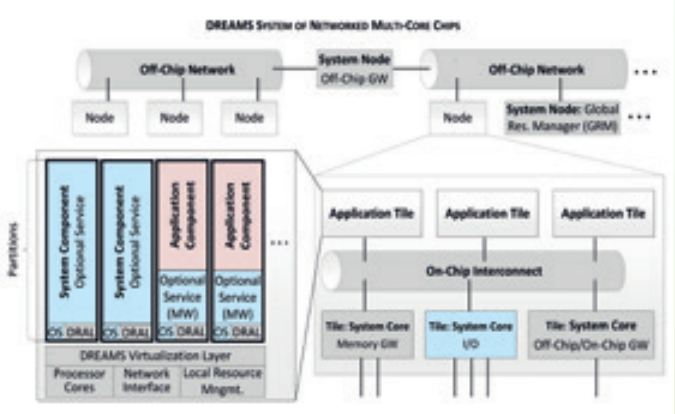

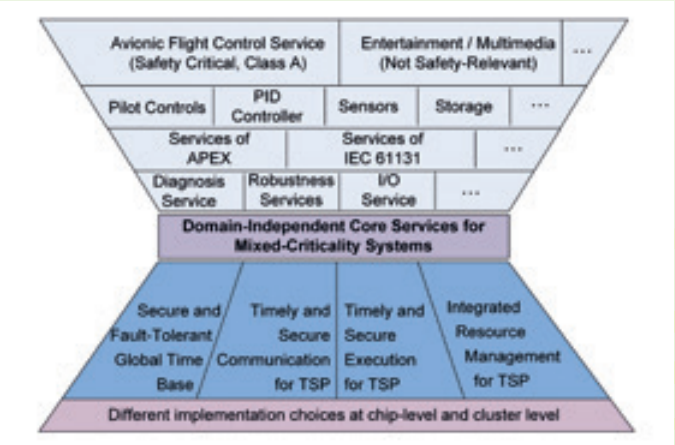

POTENTIAL IMPACT

The expected gains from DREAMS project are numerous. The most important will be a significant reduction of development lifecycle and certification efforts. The impact further includes a reduction of time-to-market, decreased development, deployment and maintenance cost, and the exploitation of the economies of scale through cross-domain components and tools. The architecture will be cross-domain in nature and support multiple domains (e.g., avionic, industrial, and healthcare). Furthermore, DREAMS will lead to improved system characteristics (e.g., reliability, safety, security, resource efficiency, adaptability).

COORDINATOR + CONTACT INFORMATION

Roman Obermaisser
Chair for Embedded Systems
Faculty of Natural Sciences and Engineering
University of Siegen
Hölderlinstraße 3
57076 Siegen
Germany

DREDBOX

disaggregated Recursive Datacenter in a Box
http://www.dredbox.eu
1 January 2016 – 31 December 2018

OVERALL OBJECTIVES

Cloud datacenters comprise individual server units with processing, memory, accelerators and storage resources. This arrangement incurs significant waste of resources and power, due to the inherent inability to closely match user IT requirements to the resources available within a single or set of servers.

dReDBox aspires to shift from today's server-as-the-unit model to a pooled-computing model, enabling arbitrary sizings of disaggregated IT resources, deploying them where and when required, perfectly matching cloud user demands. dReDBox plans to deliver a vertically integrated datacenter-in-a-box prototype to showcase the superiority of disaggregation in terms of scalability, efficiency, reliability and energy efficiency.

MAIN RESULTS

Crossing its fifth month (start date: 1 January 2016), dReDBox is progressing smoothly towards laying out a first version of its vertical architectural blueprint, including specification of server tray and components, intra-tray and rack high performance interconnects and protocols, operating system and hypervisor amendments/extensions, platform management/orchestration tools, and interfaces among the various layers.

In parallel, the project has completed a first round on analyzing the three target use-cases (NFV, video analytics, cloud analytics) and mapping their key performance indicators and value adds to dReDBox platform requirements. The project is also progressing in building targeted, custom emulation and simulation platforms for first principles design space exploration; results of the latter are planned to feed into the platform prototyping phase.

POTENTIAL IMPACT

dReDBox aspires to address the challenges stemming from the shifts happening in how and where computing is consumed today, principally by materializing the vision of disaggregation. This can have a tremendous impact in increasing IT utilization of the 3rd Platform, making cloud/edge computing more accessible through better economies of scale and tighter fitness to application and user needs, while guaranteeing sustainability in terms of energy efficiency roadmaps. On the technical innovation side, dReDBox aspires to spearhead the compartmentalization of IT resources at the component level and through that to break interdependencies and lock-in effects in the innovation cycles and usage lifetimes of individual IT components.

COORDINATOR + CONTACT INFORMATION

Dr Kostas Katrinis
Research Staff Member & Manager
IBM Research – Ireland
Email: katrinisk@ie.ibm.com
Phone: +353 (0)1 – 8269164
Address
IBM Dublin Technology Campus
Damastown Industrial Estate
Damastown, Mulhuddart
Dublin 15, Dublin
Ireland

DYMASOS

Dynamic Management of Physically Coupled Systems of Systems
http://www.dymasos.eu
October 2013 – September 2016

OVERALL OBJECTIVES

The main objective of DYMASOS is to develop new management methods and engineering tools for systems of systems (SoS), large interconnected systems with (partially) autonomous sub-units. These systems, which include electric power systems, transportation systems, and industrial production complexes, are essential for the future well-being of the citizens of Europe. DYMASOS explores three approaches for SoS management: population-control techniques, market-like mechanisms, and coalition games. In addition, engineering tools for validation and the information management of cyber-physical systems of systems are developed. All methods and tools are validated on realistic simulations of case studies in industry and in electric grids that are provided by leading European companies.

MAIN RESULTS

The main results of DYMASOS will be novel distributed management methods and tools that are suitable for complex cyber-physical systems of systems with partially autonomous subsystems. They will enable significant reductions of emissions, of the CO2 footprint, and of the resource consumption of industrial complexes and of the generation of electric power. The methods will be validated on realistic large-scale models that are provided by leading European companies in the fields of chemical production (BASF and INEOS, both among the largest chemicals producers in the world), and operation and engineering of electric power distribution systems (HEP ODS, Croatia, and Ayesa, Spain).

To facilitate the model-based validation and the industrial transfer of the novel management methods and tools, the DYMASOS engineering platform has been developed that enables a systematic simulation-based validation of distributed management architectures on SoS models via generic interfaces and plug-and-play capabilities.

The industrial relevance of the DYMASOS work was reinforced by a thorough analysis of markets, industrial needs, and industrial transfer challenges, as well as comprehensive dissemination and exploitation efforts, which are supported by the DYMASOS Industrial Advisory Board that consists of nine representatives of external companies.

POTENTIAL IMPACT

Overall improved management and control of interconnected systems of systems will have a positive impact on the everyday life of European citizens, leading to increased energy and resource efficiency of industrial production and power system stability with smaller carbon footprints and a more efficient use of renewables, fully in line with the "Smart Sustainable Growth" objective of the Horizon 2020 strategy of the European Commission. The results of the project will improve tomorrow's infrastructure: smart buildings, smart power grids, and smart industrial production sites, which will be characterized by adaptability, resource efficiency, and stable provision of services.

COORDINATOR + CONTACT INFORMATION

Prof. Sebastian Engell
Head of the Process Dynamics and Operations Group (DYN), Technische Universität Dortmund, Germany
Emil-Figge-Str. 70
44227 Dortmund
Phone: +49 (0) 231/755-5127
Email: s.engell@bci.tu-dortmund.de

EOT

Eyes of Things
www.eyesofthings.eu
1 January 2015 – 31 December 2017

OVERALL OBJECTIVES

Computer vision is rapidly moving beyond academic research and factory automation. The emerging possibilities are endless in terms of wearable applications, augmented reality, surveillance, ambient-assisted living, etc. However, vision is arguably the most demanding sensor in terms of power consumption and required processing power.

Our objective in this project is to build a power-size-cost-programmabilty optimized core vision platform that can work independently and also embedded into all types of artefacts. This will not only mean more hours of continuous operation, it will allow to create novel applications and services that go beyond what current vision systems can do.

MAIN RESULTS

Hardware for the device has evolved from a primitive bulky device made up of a set of circuit boards. Gradually, the fundamental components (camera, processor, WiFi...) are being squeezed into a pen-drive sized module board. Hardware development runs in parallel to the development of associated software, which shall run both in the device and in external controlling computers (smartphones, tablets or PCs). The overall objective of the software is to make the device usable in multiple scenarios, thus making it flexible (i.e. optimizing 'programmability').

At the time of writing, the platform is being finalized in both hardware and software. During the second half of the project, the platform will be demonstrated in 4 example applications:

- Peephole surveillance: When you go on vacation, you leave the device attached to your door's peephole and it will send you images/alarms
- Audio museum guide: Headset that automatically recognizes the painting that you are looking at and then provides audio information
- Facial expression recognition: The device embedded in a doll's head can recognize the girl's facial expression
- Wearable lifelogging camera: Lifelogging only interesting events/scenes in your life

POTENTIAL IMPACT

Computer vision is a discipline where we use computer software to analyse images for content. Traditionally focused on factory automation (inspect parts, measure distances, ...all in carefully controlled factory conditions) , this field is now booming with applications everywhere: videogames, automotive, drones, wearable headsets, automatic surveillance, intelligent toys and companions, etc. In EoT we want to develop a flexible open platform to contribute to this trend. Our tiny intelligent camera is targeted at original equipment manufacturers (OEMs), the idea being to provide the right tool to develop visually intelligent products and services with short time-to-market.

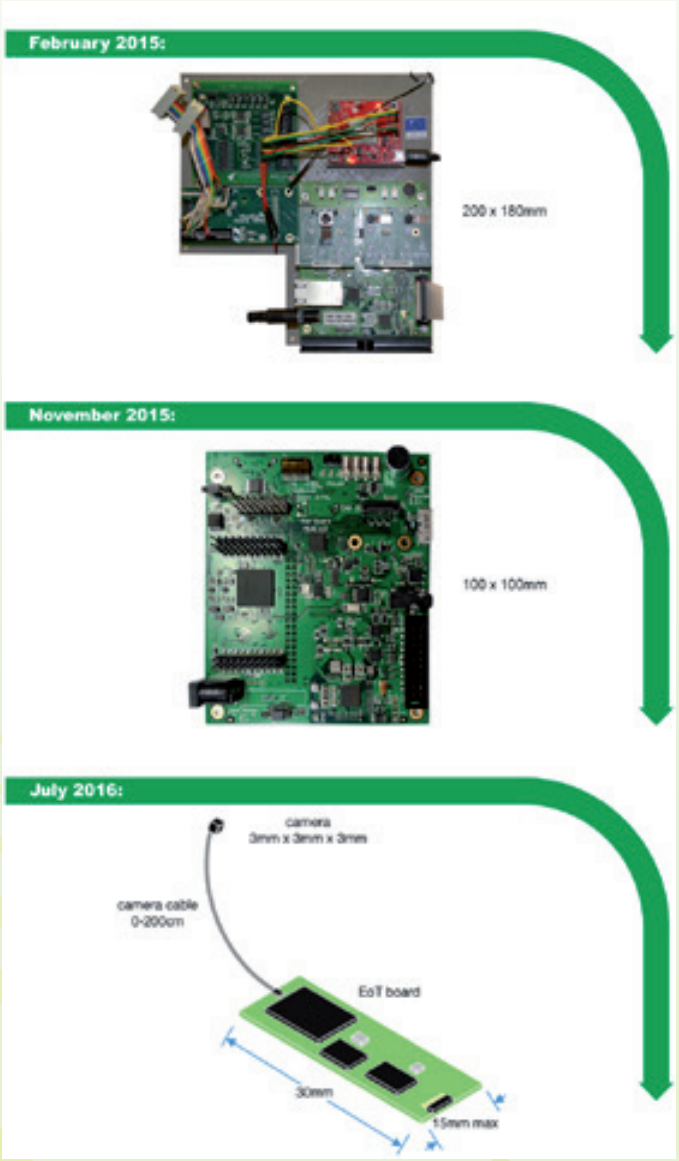

COORDINATOR + CONTACT INFORMATION

Dr. Oscar Deniz Suarez
Email: Oscar.Deniz@uclm.es
Phone: +34 926295300 ext 6286
VISILAB Grupo de Visión y Sistemas Inteligentes
UNIVERSIDAD DE CASTILLA-LA MANCHA
Av. Camilo José Cela, s/n 13071 · Ciudad Real · Spain

EUROCPS

European Network of competencies and platforms for Enabling Small and Medium Size Enterprises (SME) from any sector building Innovative Cyber Physical System (CPS) products to sustain demand for European manufacturing – www.eurocps.org
20 January 2015 – 31 December 2018

OVERALL OBJECTIVES

EuroCPS (www.eurocps.org) is targeting outcome 'b' "Innovation Actions" of the H2020 objective ICT-2014 "Smart Cyber-Physical Systems". The project aims to arm Europe with a network of design centers in order to initiate and boost synergies between SMEs, major CPS-platforms, and CPS-competency providers. The expected outcome is to capture the emerging CPS markets and create sustained demand for European manufacturing. To that end, EuroCPS design centers act as one-stop-shop, providing technical expertise, coaching and access to advanced industrial CPS platforms in order to get SMEs up to speed on the innovation ecosystem of CPS products by facilitating access to the leading edge technologies and their implementation. In the process, design centers tap existing regional ecosystems in several countries to bring the full value chain from hardware/software platforms to high value-added CPS products and services.

MAIN RESULTS

EuroCPS has started in February 2015. The first period of the project was meant to define all the materials necessary to manage the open call procedure, to monitor and report the granted industrial experiments, to enhance the networking and the communication through all the dedicated partners. Besides all these fundamental tasks, two open calls have been successfully launched and managed resulting in the selection of 24 IEs over 72 proposals. The 9 selected IEs on the first call have all started and some of them have already reached their first milestones. Regarding the 15 selected IEs on the second call, the Standard Agreements signature are in progress and the projects should start in a close future. Moreover more than 4 000 SMEs are aware of the EuroCPS initiative (800 targeted initially) and more than 100 SMEs have been interested in proposing an IE for the two first calls. It is worth to note that 20% of the submitted proposals come from European countries not represented by a EuroCPS partners. In overall 80% of the targets have been reached within the two first calls in terms of proposed and selected industrial experiments.

POTENTIAL IMPACT

The main impact will be measured when the first wave of Industrial Experiments will come to their achievements (end of 2016). Intermediate results show the benefit and ability of EuroCPS design centers to enable and accelerate the development of innovative CPS products by facilitating access to leading edge technologies and knowhow. Industrial experiments speed up SMEs to reach state of the art for their development and bring technology breakthroughs thanks to RTO and academic advanced solutions. Beside, business networking events and introductions facilitated by EuroCPS networking partners has already helped to raise business profile of certain SMEs leading to new business opportunities.

COORDINATOR + CONTACT INFORMATION

Olivier Thomas
MINATEC CAMPUS, 17 rue des Martyrs, F-38054 Grenoble Cedex
Olivier.thomas@cea.fr

EUROSERVER

Green Computing Node for European micro-servers
www.euroserver-project.eu
1 September 2013 – 31 January 2017

OVERALL OBJECTIVES

EUROSERVER prepares the European industry for a leadership position in building the end-to-end solution for next generation datacenters. To support the current data growth, these datacenters need to scale both in total processing capability and storage capacity across the entire cloud infrastructure. EUROSERVER proposes to design and build a drastically improved energy- and cost-efficient solution (micro-server) using a combination of components and various holistic-design and manufacturing techniques: 64-bit ARM cores, 2.5D nanotechnology integration, FD-SOI process technology together with new software techniques four efficient resource management including resource sharing and workload isolation. EUROSERVER will build two integrated full-system prototypes based on a common micro-server boards, validated under both real cloud applications and embedded market scenarios.

MAIN RESULTS

The major achievements reached by EUROSERVER so far are:

- A systematic requirement analysis conducted and used to refine the system specifications for Cloud/Enterprise, Embedded and Communications scenarios.
- Two prototype development, based on discrete components and programmable devices which demonstrated memory sharing as well as optimized Linux kernel and hypervisor extensions.

Several strategic choices have consolidated EUROSERVER architecture:

- An innovative structure based on interconnected compute "coherent island" adopted for an optimal balance between data locality and transfer efficiency.
- Compute SoC internal structure organized around several independent "chiplets" implementing the islands.
- Coupling between these chiplets realized by high speed serial links.
- Physically, integration of the system onto a cost effective organic interposer solution.

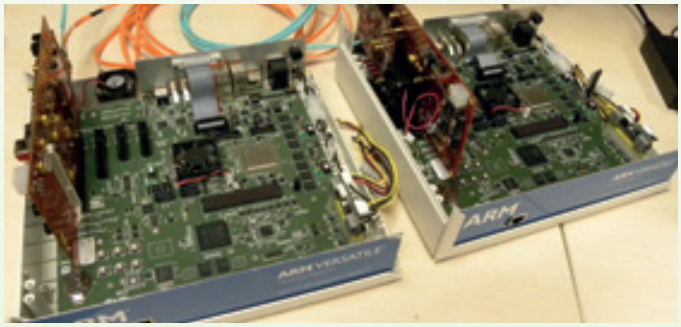

- Selection of Hybrid Memory Cubes technology for the central memory (to increase the density by stacking DDRAM modules in order to offer better bandwidth and memory storage density that are directly connected to the cores.)

These hardware innovations would not be exploitable without software support. Hardware agnostic improvements have been carried out for micro-servers in general, and applied on the prototypes. To differentiate EUROSERVER from micro-server and typical server designs the key software technologies being worked on are:

- Shared memory model (all cores can access RAM from remote units through a controlled interface).
- Highly efficient Hypervisor platforms (with a focus on reducing the virtualisation overhead).
- Specific features such as scalable M2M messaging service, dataflow based task parallel runtime environment to support RAN functionality.

POTENTIAL IMPACT

The emerging key differentiator for EUROSERVER is improved resource utilisation. Just as Cloud computing and virtualisation enables companies to converge workloads from many distributed and under-utilised hardware platforms into smaller numbers of servers, EUROSERVER proposes to more efficiently exploit micro-server and low power hardware in order to pave the way towards the next generation of more power efficient servers:

- Improvement by at least one order of magnitude of server energy/cost efficiency, performance x density through integration, total cost of ownership;
- Improvement of software application efficiency through IO rationalization (mutualisation and virtualization) and open source ecosystem reinforcement for embedded and cloud computing;
- Growth of European technology suppliers' competitiveness.

COORDINATOR + CONTACT INFORMATION

CEA / DRT / Leti
Isabelle.dor@cea.fr

EXCESS

Execution Models for Energy-Efficient Computing Systems
http://www.excess-project.eu/
1 September 2013 – 31 August 2016

OVERALL OBJECTIVES

EXCESS addresses the problem of energy efficiency on HPC and embedded platforms. EXCESS aims at providing radically new energy-aware execution models forming foundations for energy-efficient computing paradigms that will enable substantial improvements in energy efficiency for computing systems. EXCESS has embraced a holistic approach that involves both hardware and software aspects together. The scientific and technological concept is defined by novel execution models covering both common High Performance Computing infrastructures and Embedded Systems. EXCESS is developing energy, platform and component models that are applicable to both embedded processors and general purpose ones.

MAIN RESULTS

EXCESS main results are driven by the EXCESS-designed, component based programming framework for energy efficient computing that is shown below.

EXCESS main contribution is resource- and energy-aware programming models, a portable generic tool chain for generating energy-optimized code and adaptive libraries, which together can address energy efficiency issues for both classes (i.e. HPC and Embedded) of systems. On a high level EXCESS main results can be summarized as follows:

1. A software toolchain (including programming/components models, libraries/algorithms monitoring and runtime support) for energy-efficient computing.
2. Generic development methodology and prototype software tools that enable leveraging additional optimization opportunities for energy-efficient computing.

In more detail the main results of EXCESS include:

i) A language for energy-aware platform modelling.
ii) A component model that involves multi-variant components.
iii) An auto-tunable skeleton programming library.
iv) A generic portable light weight energy measurement abstraction API together with an energy monitoring framework for energy analytics at run-time.
v) Power models for embedded and HPC systems.
vi) Models that analytically predict the energy behavior of data structures implementations.
vii) A library of programming abstractions that include novel designs of data structures that exhibit significant energy saving compared to previously known ones.

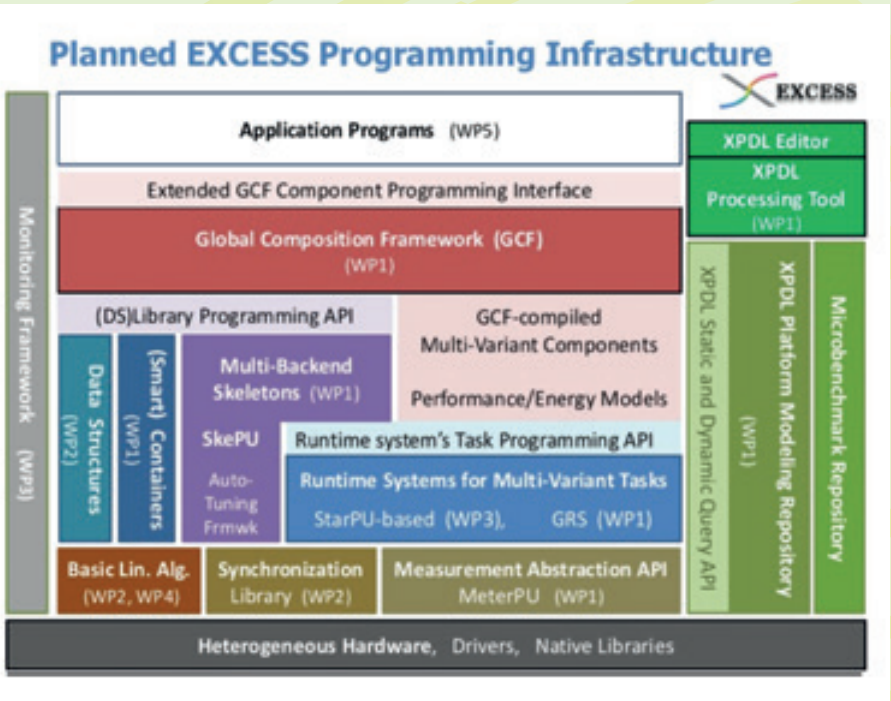

POTENTIAL IMPACT

***This may overstate some parts we haven't, in fact, touched much. ***

- The EXCESS programming framework with its re-targetable generic tool chain for generating energy-optimized code and its adaptive libraries, offers portability and best effort automated adaptation. This allows for rapid development of energy efficient software production.
- Energy modeling that supports high-level program optimizations such as selection of implementation variants, tuning of synchronization methods and data-structures, or context-adaptive strategies, allowing potentially new innovative software designs for energy quality.
- Many of the components of the EXCESS tool chain (including programming/components models, libraries/algorithms monitoring and runtime support) are generic and can be used and included in other frameworks that target energy software development.

COORDINATOR + CONTACT INFORMATION

Prof. Philippas Tsigas
Philippas.tsigas@chalmers.se
+46 31 772 5409
Chalmers University of Technology
S-412 96 Göteborg
Sweden

FIPS

Developing Hardware and Design Methodologies for Heterogeneous Low Power Field Programmable Servers
https://www.fips-project.eu
September 2013 – August 2016

OVERALL OBJECTIVES

The FiPS project aims to integrate FPGA-based computing nodes into a modular heterogeneous server platform and to develop a corresponding design methodology, which enables a user to choose the best hardware configuration for given software and to optimize that software for best efficiency and performance. With these technologies, the project plans to increase the energy efficiency of super-computing programs by 40% and at the same time to reduce the effort needed for application programming for the new hardware by a factor of 2.

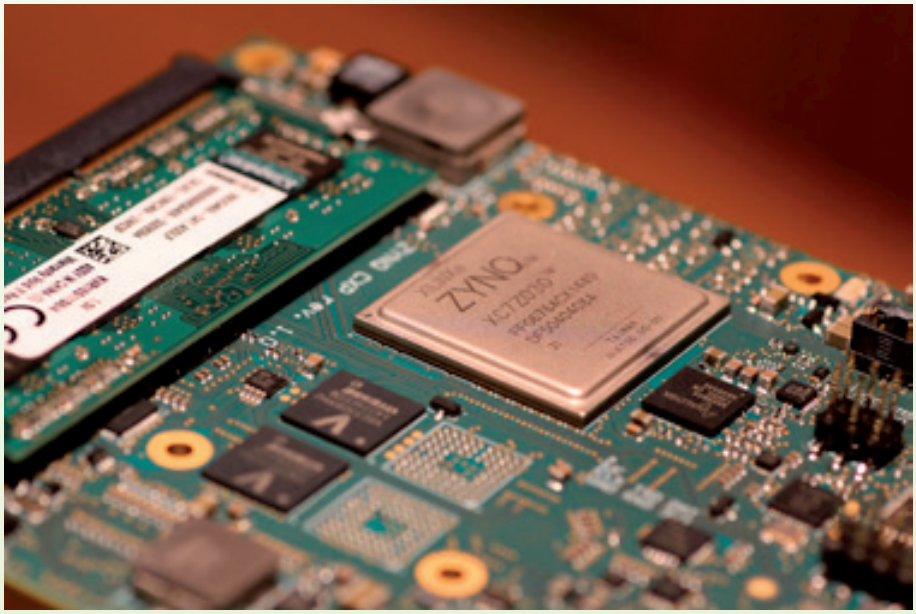

© Bielefeld University, Cognitronics and Sensor Systems group

MAIN RESULTS

During the runtime of the project, new FPGA-based compute nodes have been developed and successfully integrated into the RECS®|Box Compute Unit of the project beneficiary Christmann. This heterogeneous server platform now supports x86, GPGPU, ARM and FPGA nodes, which can be mixed to build hardware matching different kinds of requirements. The newly created design methodology consists of a holistic exploration, simulation, and prediction tool flow. It enables a developer to optimize super-computing programs at an early state without actually running the software on the hardware. As all communication and artifacts are simulated, the developer can get predictions of how changes made to the software affect energy and performance.

With these new technologies the project beneficiaries were able to optimize a top-of-the-class DNA sequence alignment program. Compared to the original GPGPU implementation the resulting FPGA-assisted implementation is almost 1.5 times as power efficient. Other applications (Artificial Neural Networks and CFD simulations) under evaluation already show promising benefits of the new technologies and will greatly profit from the combination of optimized hardware and software.

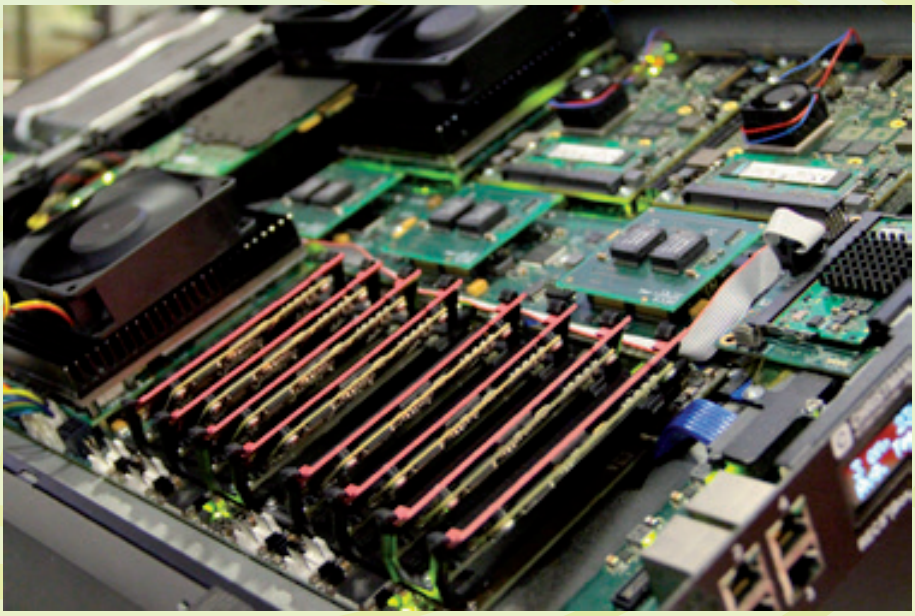

© Bielefeld University, Cognitronics and Sensor Systems group

POTENTIAL IMPACT

With the FiPS technologies the project beneficiaries foresee a great ecological impact by reducing energy demand (and thus carbon dioxide emissions), but also an economic impact by cutting major energy costs. Super-computing will become cheaper, more accessible and thus affordable for many other applications and even whole new groups of users.

COORDINATOR + CONTACT INFORMATION

Dr. Domenik Helms, OFFIS e.V.,
domenik.helms@offis.de

HERCULES

High-Performance Real-time Architectures for Low-Power Embedded Systems
http://hercules2020.eu/
1 January 2016 – 31 December 2018

OVERALL OBJECTIVES

In the tomorrow's society, automated systems will replace safety-critical human activities, with power consumption as a primary concern. Next-generation of commercial-of-the-shelf (COTS) heterogeneous multi-core platforms could capture this need for high-performance at reduced power budgets, but, unfortunately, a "super-computing platform" having impressive average performances with no guaranteed bounds on the response times is of little if no use to critical applications. Project HERCULES will develop the technology to allow achieving predictable performance on top of cutting-edge COTS. The framework will be applied to a pioneering autonomous driving system, and a visual recognition system for the avionic domain.

MAIN RESULTS

The HERCULES project aims at achieving multiple technological breakthroughs, which are summarized by the following strategic goals.

G1. Demonstrate and implement the first industrial-grade framework to provide real-time guarantees on top of cutting-edge commercial-off-the-shelf (COTS) platforms for the embedded domain.

G2. Obtain an order-of-magnitude improvement in the energy efficiency and cost of next generation real-time systems.

G3. Provide a simple programming interface to support the development of real-time software on top of COTS platforms.

These ambitious goals will be achieved according to the following verifiable objectives:

O1. Identify, characterize and select the most suitable next-generation computing platform for the target domains

O2. Provide a clean programming interface to ease software development

O3. Implement an integrated Real-Time capable Operating System (RTOS) and software stack that seamlessly manages the available hardware resources in an efficient way

O4. Develop smart real-time algorithms to manage resources (processors, memories) available on the selected board

O5. Implement lightweight execution support for time-critical software

O6. Port an autonomous driving system (ADAS) for self-driving cars to the HERCULES framework

O7. Enhance an application for object detection for avionic system

O8. Contribute to Open Source software

O9. Promote the adoption of HERCULES framework to key European industrial players

POTENTIAL IMPACT

The HERCULES project has been devised starting from clear industrial requirements, with particular relation to the new kinds of time-critical applications that will invade the embedded market in the near future. Its impact is related to the innovation potential brought by the framework with relation to the capability of executing these applications, with a higher performance/Watt compared to existing industrial solutions. The consortium has long standing relationships with main automotive OEMs, TIER1 and TIER2, while industrial partners from avionic domains of are also interested in adding innovative capabilities to their product lines to create added value for their customers.

COORDINATOR + CONTACT INFORMATION

Prof. Marko Bertogna, University of Modena and Reggio Emilia, Italy – marko.bertogna@unimore.it
Paolo Gai, Evidence srl, Italy – pj@evidence.eu.com
Paolo Burgio, University of Modena and Reggio Emilia, Italy – paolo.burgio@unimore.it

HIPEAC

High Performance and Embedded Architecture and Compilation
http://www.hipeac.net
1 January 2016 – 28 February 2018

OVERALL OBJECTIVES

HiPEAC is a support action that aims to structure and strengthen the European academic and industrial communities in computing systems: (i) by increasing innovation awareness and by encouraging researchers to engage in innovation activities; (ii) by professionally disseminating program achievements beyond the traditional scientific venues; (iii) by producing a vision document including recommendations on how to improve the innovation potential of H2020 projects in computing, and (iv) by growing the computing systems community beyond 2000 active members in Europe.

MAIN RESULTS

The project consists of four activities:

- **Community structuring** focuses on improving the conditions for innovation by creating stronger links between the academic and the industrial world. It stimulates (i) the mobility of talent between academia and industry, (ii) the sharing of information (presentations, demonstrations, ...) between projects and members, (iii) the sharing and transfer of knowledge and technology from academia to industry (IPR, data), and finally (iv) educating the community about the ins and outs of innovation.
- **Result dissemination** focuses on increasing the impact of the dissemination of research achievements with the help of a professional communication officer. This staff member works with the projects and the members to find the best way to disseminate their messages, and designs and run a roadshow on conferences and trade shows.
- **Vision building** is a work package with longer-term goals. It contains the impact analysis of finished FP7 projects on computing systems, the roadmapping activities that will lead to the HiPEAC Vision 2017 and a set of consultation workshops on a variety of topics.
- **Constituency building** focuses on growing and strengthening the computing systems community by organizing a major conference for disseminating results, networking events, mobility support and training.

POTENTIAL IMPACT

The impact of the HiPEAC support action is the creation of a confident and visible community in computing systems and to steer that community towards solving the challenges that are at the core of digitizing the European industry. Besides the technical challenges that need to be tackled, there is also the challenge of how to become a visible community beyond the computing community, the challenge of how to organize mobility of talent across Europe, and the challenge of how to turn research results into innovative products and services.

COORDINATOR + CONTACT INFORMATION

Prof. Koen De Bosschere, Ghent University,
Koen.DeBosschere@UGent.be

IMMORTAL

Integrated Modelling, Fault Management, Verification and Reliable Design Environment for Cyber-Physical Systems
http://www.h2020-immortal.eu
1 March 2015 – 28 February 2018

OVERALL OBJECTIVES

The objective of IMMORTAL is to develop an **integrated, cross-layer modelling based tool framework and a methodology for fault management, verification and reliable design of dependable Cyber-Physical Systems** (CPS). The framework will include *cross-layer modelling of CPSs and a holistic model for different error sources*. It will contain *reasoning engines for reliability analysis/certification* of CPS components and systems, as well as for *automated debug* of CPS models. IMMORTAL will also develop architectures and tool support for fault management providing *cost-efficient, ultra-low-latency detection, isolation and recovery for system faults* in dependable CPSs based on many-core architectures.

MAIN RESULTS

Currently the first year of the project is completed. The preliminary results include developing a cross-layer modelling framework and verification, debugging and testing methods for CPSs as well as developing models for CPS network configurations with the goal to support system self-organisation capabilities via dynamic reconfiguration. During the first project year, an IMMORTAL workshop at the MEDIAN Finale Event, November 2016 has been organised and the project was also presented at the HPCS Conference in Amsterdam, July 2015 and ICT1 - clustering and communication event – Vienna, April, 2016.

Up to now, more than 15 scientific publications have been published and 25 presentations held within the project.

POTENTIAL IMPACT

IMMORTAL seeks to attain a number of ambitious technological goals:

1) Minimising the verification effort in CPS by a factor of 2;
2) Speeding up fault detection, isolation and recovery in CPSs by a factor of 4;
3) Resumption of correct operation with up to 15% of CPS network resources failed;
4) Up to 40% reduction in the effort designers put in reliability related tasks;
5) Up to 10% savings in the area/power by optimising hardware protection logic overhead.

The expected overall impact of these quantitative objectives is twofold, addressing future CPSs' design effort as well as maintenance costs.

COORDINATOR + CONTACT INFORMATION

Tallinn University of Technology
Prof. Dr. Jaan Raik
Akadeemia 15a, Tallinn, Estonia
Phone: +372 620 2257
Email: Jaan.Raik@ati.ttu.ee

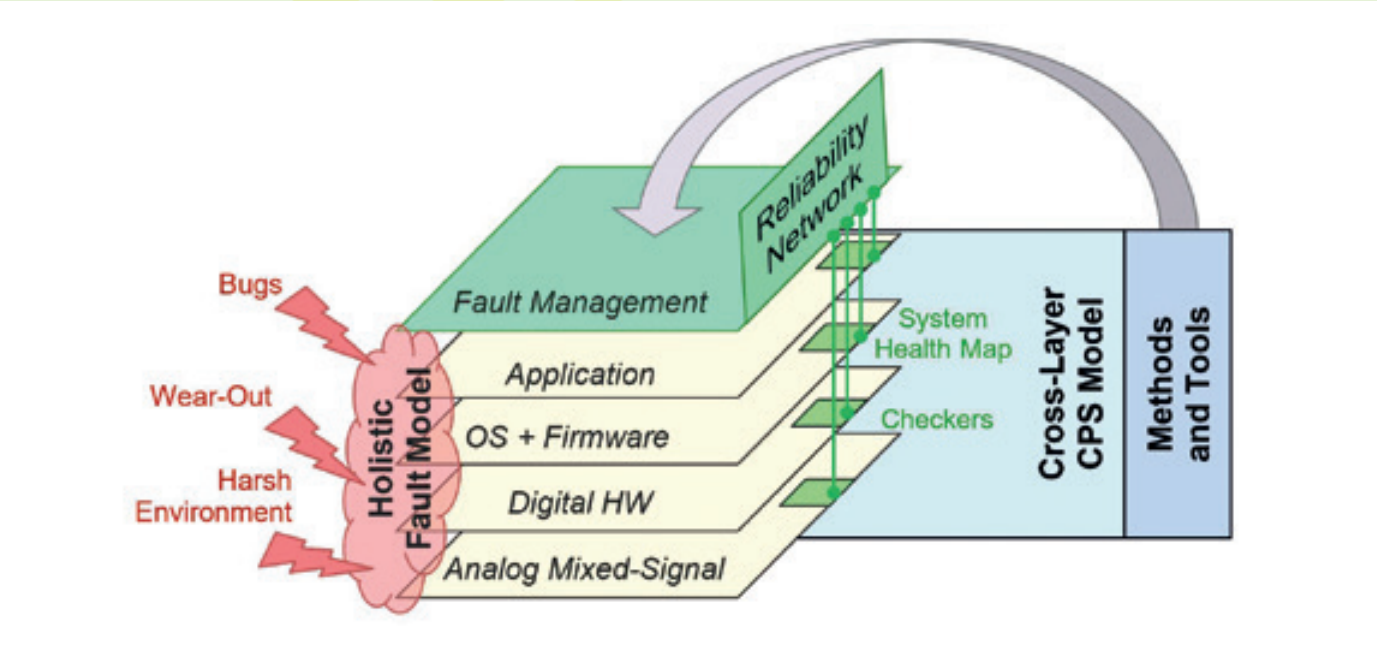

INTO-CPS

Integrated Tool Chain for Model-based Design of Cyber-Physical Systems
www.into-cps.au.dk
1 January 2015 – 31 December 2017

OVERALL OBJECTIVES

The aim of the INTO-CPS project is to create an integrated "tool chain" for comprehensive Model-Based Design of Cyber-Physical Systems. The tool chain will support the multidisciplinary, collaborative modelling of CPSs from requirements through design, down to realisation in hardware and software. This will enable traceability at all stages of the development. INTO-CPS will support the holistic modelling of CPSs, allowing system models to be built and analysed that would otherwise not be possible using stand-alone tools. We will integrate existing industry-strength tools with high Technology Readiness Levels in their application domains based centrally around Functional Mockup Interface compatible co-simulation.

MAIN RESULTS

- Foundations for SysML, the Discrete Event formalism VDM-RT, the Continuous-Time formalism OpenModelica and the Functional Mockup Interface (FMI). Initial semantics has been produced for all areas and all the deliverables from these are publically available.
- CPS methodology including guidelines, public pilot studies. The initial methodology guidelines and initial publically available pilot studies has been produced and published in public deliverables.
- Heterogeneous toolchain all the way from requirements over models in different formalisms to their corresponding realization including different analysis and traceability features. Documentation about the initial toolchain have been publically released and it is expected that external stakeholders will get access the tools summer 2016.
- Establishment of an Industrial Follower Group (IFG) with a large number of companies. Almost 50 members of the IFG have been included so far and it is expected that more will follow due to increased dissemination to the systems engineering community.
- Using the INTO-CPS on industrial case studies in agriculture, automotive, railways and building automation. A public deliverable including all four case studies have been released.

POTENTIAL IMPACT

- **Reduction in CPS development time and significant reduction in maintenance costs**: INTO-CPS will reduce CPS development time by 30% by integrating the right combination of CPS analysis tools to allow comprehensive and rapid early stage analysis of new CPSs, from component- to system-level.
- **Stronger pan-European collaboration across value chains**: The INTO-CPS technology can act as a catalyst for the formation of ecosystems within these domains, and the opportunity to demonstrate this will grow as the IFG grows in size.
- **Development in Europe of a competitive offer**: INTO-CPS will contribute a new tool-supported methodology that will maintain Europe's position at the forefront of multidisciplinary design for CPSs.
- **Uplifting Europe's innovation capacity and competitiveness**: INTO-CPS strengthens the innovation capacity of industry by supporting the rapid model-based assessment of CPS-based products.

COORDINATOR + CONTACT INFORMATION

Prof. Peter Gorm Larsen, pgl@eng.au.dk
Department of Engineering, Aarhus University, Denmark

LOCAL4GLOBAL

SYSTEM-OF-SYSTEMS THAT ACT LOCALLY FOR OPTIMIZING GLOBALLY
http://local4global-fp7.eu/
1 October 2013 – 1 October 2016

OVERALL OBJECTIVES

Develop, test and evaluate a groundbreaking, generic and fully-functional system for controlling Technical Systems-of-Systems (TSoS) where autonomous constituent systems interact, to optimize the TSoS performance at the global level. Within each constituent system is embedded a: (i) self-learning mechanism - providing "just-enough" estimate of the overall TSoS dynamics - (ii) situation awareness mechanism - responsible for extracting the necessary overall TSoS information - (iii) distributed optimizer - to calculate the constituent system's optimal actions.

MAIN RESULTS

The Local4Global (**L4G**) system is based on a previously developed, by some of the consortium partners, centralized Cognitive Adaptive Optimization approach (CAO). Three different versions of this approach (a. model-assisted optimization, automated periodic fine-tuning of existing control parameters, fully adaptive model independent approach) served as the basis for developing three distributed versions respectively within the Local4Global framework context.

The integrated control design system is tested and evaluated into two complex real-life TSoS Cases: (i) E.ON Energy Research Center at the RWTH Aachen University, Germany, (ii) Federal Road B13 in north Munich, Germany.

In both cases elaborate simulation models were available for lab tests, while the Local4Global system presented fault-tolerant control properties under emulated highly uncertain circumstances, satisfying convergence rate as well as substantial savings (around 15-20% in both cases) with respect to commonly used and widely accepted control strategies.

Moreover, the **first real-life application tests** showed very promising potential of the Local4Global system, **leading to improvements of 20% or even higher in terms of the overall system performance while neglecting the need for the deployment of an expensive infrastructure or the need for a tedious and time-consuming programming and verification effort**.

We also note that the system developed within L4G by the partner CERTH was among the 4 winners of the international innovation contest **Intelligent Energy Management Challenge** organized by the Swedish Energy Agency and Swedish Incubators & Science Park (with the participation of the municipalities of **Arvika, Göteborg, Eskilstuna, Herrljunga** and **Uppsala**).

POTENTIAL IMPACT

Local4Global has the potential to pave the way towards the development of "plug-and-play" control systems, similarly to the way plug-and-play systems for communication operate: no need for expensive infra-structure, no need for modelling, control calibration and periodic re-tuning responsible engineer(s). The consequences of such a system will be tremendous and not only limited to areas and systems where no sophisticated control is currently employed (due to the requirement for an elaborate infrastructure). It will also be of great significance to areas and systems where, current modern control and management systems "cannot always do the job".

COORDINATOR + CONTACT INFORMATION

Pf. Elias B. Kosmatopoulos
Email: kosmatop@iti.gr
Information Technologies Institute
Centre of Research & Technology - Hellas
1st km Thermi - Panorama, 57001, Thessaloniki, Greece
Phone.: +30 2310 464160 (ext. 132)

LPGPU2

Low-Power Parallel Computing on GPUs 2
www.lpgpu.org
1 January 2016 – 30 June 2018

OVERALL OBJECTIVES

Low-power GPUs can be found in domains ranging from wearables and IoT devices to mobile computing and automotive systems. This places an ever-increasing demand on their expected performance and power efficiency. The LPGPU2 project aims to develop a framework for quantifying the power efficiency of applications on mobile GPUs. This framework will help developers to improve the power efficiency of their applications and increase their productivity. The LPGPU2 project builds upon the successful LPGPU project and will bring the technology developed in this project closer to the market.

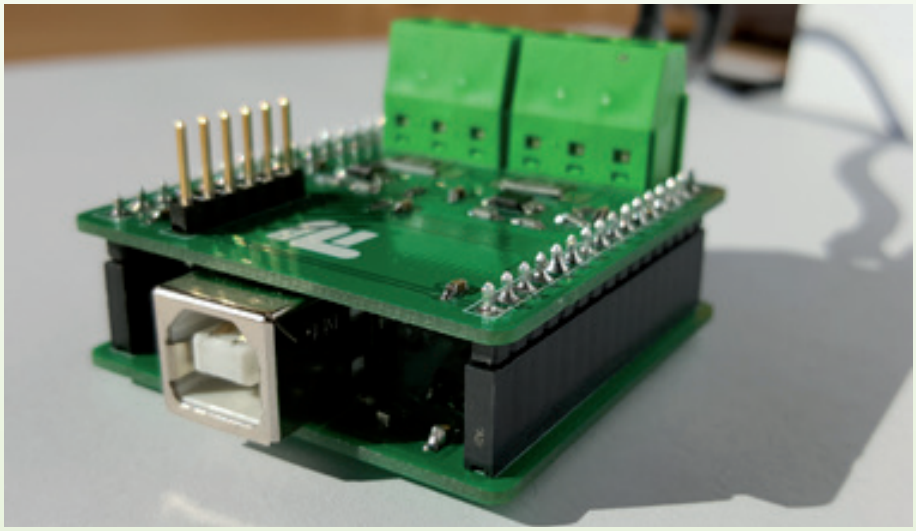

MAIN RESULTS

While the project started less than 6 month ago, significant progress has already been achieved. Work on important applications such as ISP camera processing and H.265 8K video decoding using GPUs has started and early results are already available. In addition hardware and firmware for a new power measurement testbed for mobile GPUs has been developed and is already working as expected with a higher resolution and sample rate than the previously used testbed. A new suit of micro-benchmarks as well as novel tools to analyze the results of these benchmarks are being developed. When these tools are finished they will allow assessing the performance and power consumption characteristics of mobile GPU applications in an automated manner. LPGPU2 will also develop a novel DVFS mechanism for better power management and sustained performance. Furthermore, the insights gained in the project will help us define new industry standards for resource and performance monitoring. This will allow the members of the LPGPU2 project as well as third parties to develop innovative tools that help software developers increase their productivity and improve the performance and power efficiency of their applications.

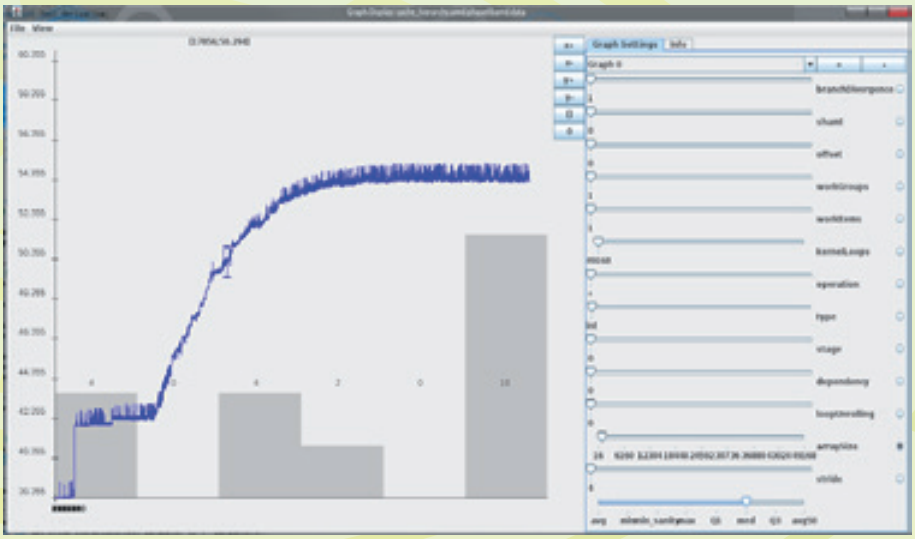

POTENTIAL IMPACT

The project will contribute to defining new industry standards for resource and performance monitoring to be widely adopted by the embedded hardware GPU community. The performance, power measurement and prediction framework built upon these standards will help decrease power consumption and increase the battery life of mobile devices. The productivity of developers will be improved with a better understanding of the power consumption of their applications and targeted hints on how to improve their energy efficiency. New and innovative applications will be made possible with and on mobile GPUs.

COORDINATOR + CONTACT INFORMATION

Prof. Dr. Ben Juurlink
Technische Universität Berlin
Phone: +49 30 314 73130 / 73131
Einsteinufer 17 / Sekr. EN12
Email: b.juurlink@tu-berlin.de
10587 Berlin, Germany
Web: www.aes.tu-berlin.de

M²DC

Modular Microserver Data Centre
http://m2dc.eu
1 January 2016 – 31 December 2018

OVERALL OBJECTIVES

The main goal of the M2DC project is to develop a new class of low-power TCO-optimised appliances with built-in efficiency and dependability enhancements. The appliances will be easy to integrate with a broad ecosystem of management software and fully software-defined to enable optimisation for a variety of future demanding applications in a cost-effective way. The M2DC flexible server architecture with heterogeneous hardware including ARM CPUs and FPGAs will enable customisation and smooth adaptation to various types of data centres, while advanced management strategies and system efficiency enhancements (SEE) will be used to achieve high levels of energy efficiency, performance, security and reliability.

MAIN RESULTS

The main M2DC result will include a set of turnkey appliances based on a microserver system enabling deployment of use-case driven, modular, high-density data centres. Appliances will be low cost, low power and energy efficient, dependable by design, versatile and scalable, easy to use and integrate with data centre ecosystems, and applicable to a variety of real-life applications.

M2DC will demonstrate turnkey appliances tailored to meet requirements from various application domains such as photo finishing system serving (more scalable photo finishing), IoT data processing (data analytics for vehicles' sensors), cloud computing (enhanced IaaS, PaaS solutions exploiting heterogeneity) or even HPC (efficient meteorological simulations).

POTENTIAL IMPACT

M2DC results should transform into significant benefits for data centres that are struggling with rising energy and generally OPEX costs.

M2DC will also have impact on the market of low-power hardware and microservers significantly strengthening Europe's position in the IT.

Improvements in server and data centre efficiency and reliability will also have impact on specific application markets. In particular, M2DC will enable competitive advantage with respect to performance, scalability, and cost-savings to photo finishing several million images per year, analysing large real-time sensor data from vehicles, and IaaS, PaaS cloud solutions.

COORDINATOR + CONTACT INFORMATION

Project Coordinator: Ariel Oleksiak - ariel@man.poznan.pl
Dissemination Manager: Mariano Cecowski - Mariano.cecowski@xlab.si

Twitter: @M2DC_Project
Facebook: M2DC project
LinkedIn: m2dc

NANOSTREAMS

Name: A Hardware and Software Stack for Real-Time Analytics on Fast Data Streams
Link to the web site of the project: http://www.nanostreams.eu/
01 September 2013 – 28 February 2017

OVERALL OBJECTIVES

NanoStreams co-designs a micro-server architecture and a software stack that address the unique challenges of hybrid transactional-analytical workloads, which are encountered by emerging applications in real-time big-data analytics. The project achieves its goals of energy and performance efficiency by bringing together embedded system design principles, application-specific compiler technology, and HPC software practices. Exemplar applications include: latency sensitive computation of option prices in the financial markets, improving patient outcomes in intensive care units by processing streams of physiological data from patient sensors and establishing linkages in patent portfolios using real-time graph analytics.

MAIN RESULTS

The project, now in the final year of its lifespan, is delivering integrated demonstrators in several application areas. These are in turn built on research output delivering new component technologies which have been co-designed to work efficiently together.

These are:

Nanowire: a scalable and low latency protocol running on top of Ethernet, a commodity technology, which allows for optimal communication among hosts and accelerators. This replaces the need for specialized data buses.

Nanocore: an analytics on chip micro server substrate which is programmable in a dialect of the C language. It a real-silicon prototype, based on the Xilinx Zynq platform and ARM-Linux and which delivers transactional throughput and improved system energy-efficiency and programmability. Nanocore accelerators exchange data with Intel and ARM host CPUs via Nanowire.

NanoRuntime: an interposer between an application and the POSIX threads library which is cognizant of thread management by different components within that application and seeks to optimize the threads in both the spatial and temporal domains on CPU hosts.

NanoLibDAG: enhancements to the C and C++ languages to allow sets of functions to be marked for compilation and execution in the form of a directed acyclic graph. The system includes the possibility of some nodes in the graph being located on Nanocores.

POTENTIAL IMPACT

The NanoStreams optimized stack with its energy efficient Nanocore accelerator offers the potential to drive down cost in data centres and at the same time improve application performance and programmer productivity. The medical use case offers clinicians new diagnostic real-time tools to improve patient outcomes while research into the programming language is influencing the next set of ISO standards for C++. Enhanced graph analytics is a cornerstone in the new field of cognitive computing and will unlock new knowledge across many sectors of the economy

Figure 1 The software ecosystem for development of an optimized analytics on chip Nanocore

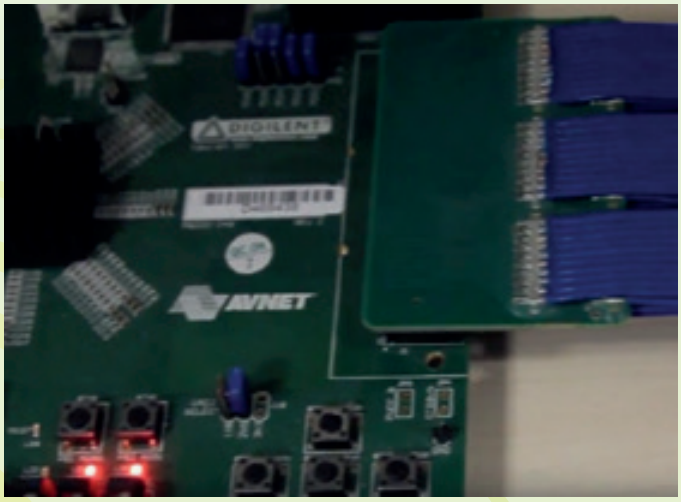

Figure 2 ZedBoard implementing Nanocore with Nanowire connectivity

COORDINATOR + CONTACT INFORMATION

Professor Dimitrios Nikolopoulos
The School of Electronics, Electrical Engineering and Computer Science
Queen's University Belfast, Northern Ireland BT7 1NN, United Kingdom
d.nikolopoulos@qub.ac.uk

OPERA

lOw Power heterogeneous architecture for nExt generation of smart infrastructure and platforms in industrial and societal Applications

www.operaproject.eu

1 December 2015 – 30 November 2018

OVERALL OBJECTIVES

OPERA project aims at supporting these ambitious challenges with technological innovation on three main aspects:

- design next generation Low Power (LP) and Ultra-Low Power (ULP) systems
- improve energy efficiency in computing by means of heterogeneous architectures
- provide smart and energy efficient solutions for the interaction between embedded smart systems and remote small form-factor data centers.

The objective is to implement an innovative computing infrastructure that covers different levels of the computing continuum by means of miniaturization and integration of existing cutting-edge technologies, such as Ultra-Low Power and Low Power architectures, next generation servers, 3D integrated circuits, optical interconnections.

MAIN RESULTS

The main results of the project will be represented by different platforms for the Ultra-Low Power computing applications accurately selected and defined in the project, and the deployment of a scalable small form-factor data center. The usage of these platforms foreseen in the defined scenarios will be applied in traditional field as well as in new areas of deployment.

The agile methodology adopted in the project will ensure the availability of preliminary results that will be evaluated and refined for the provided uses cases, as well as for applications in many other fields, not directly targeted by the sedimentation phase of the project.

The preliminary results in the project are already highlighting new possibilities, combining the reduction of power consumption of the components used, as well as a variety of methodologies for the power harvesting which are suitable for real-life scenarios and the optimization of communication through new technological concepts, merging the new wireless network of sensors with the existing infrastructure. Looking at the data center infrastructure, the project aims at showing not only the benefits of using FPGA accelerators to increase performance/Watts ratio, but also new ways for cloud applications to exploit such enormous computational power.

POTENTIAL IMPACT

The foreseen impact of the project is to drive the creation of a common working ground across key players in innovation across Europe, bridging for the first time extremely specialized skills such as Systems-on-Chip, highly performing FPGA, VHDL design, processing load optimization for multicore chips, multi-technology integration under low power constrains.

The project will leverage on a shared objective among all participants to shape the future deployment of Low Power technology without compromising the incremental processing and features enhancement.

OPERA would significantly contribute to reinforce the Europe position in leading Low Power Computing, as well as in Cloud Computing.

COORDINATOR + CONTACT INFORMATION

Giulio URLINI, STMicroelectronics,
Email: giulio.urlini@st.com

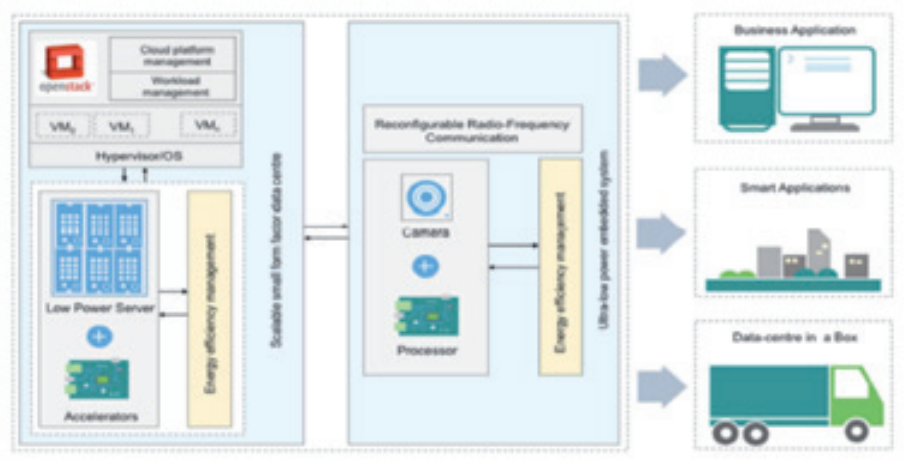

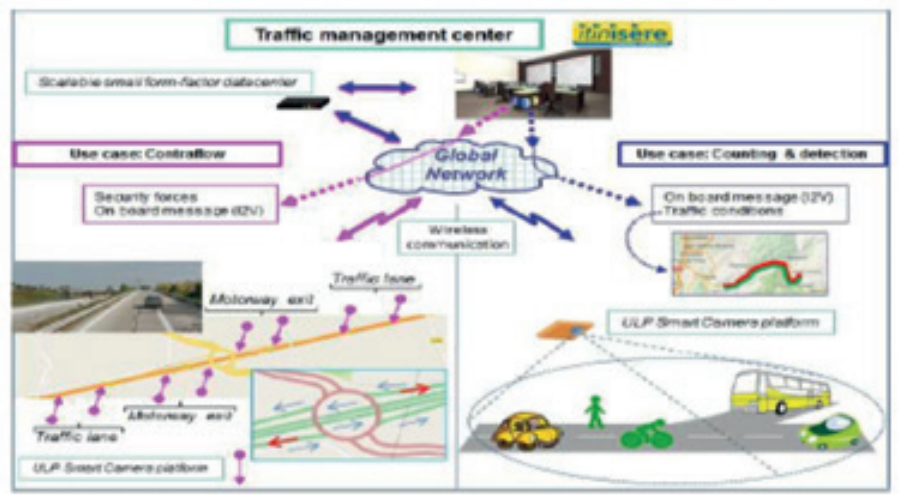

P-SOCRATES

Parallel Software Framework for Time-Critical many-core Systems
http://www.p-socrates.eu
October 2013 – September 2016

OVERALL OBJECTIVES

The aim of P-SOCRATES is to allow demanding applications with high-performance and real-time requirements to fully exploit the huge performance opportunities brought by the most advanced commercial off-the-shelf many-core embedded processors, whilst ensuring predictable performance and maintaining (or even reducing) development costs. P-SOCRATES focuses on combining the newest high-performance software techniques for exploiting task parallelism with the most advanced scheduling methodologies and timing and schedulability analysis techniques of real-time embedded systems. This is a fundamental step towards the convergence of high-performance computing (HPC), real-time and embedded domains, providing predictable performance to HPC systems and increasing performance of real-time embedded systems.

MAIN RESULTS

P-SOCRATES provides both a complete and coherent software system stack, able to bridge the gap between the design of high-performance real-time applications, and the many-core embedded processor, as well analysis methodology and tools to provide the required timing guarantees.

For this, the project provided scientific and technical advances in novel techniques for extraction of control and data flow task dependency information from annotated parallel programs, new schedulability analysis of parallel programs represented as task dependency graphs, lightweight mechanisms for run-time managing of fine-grained parallel computations, efficient managing of parallel code execution in advanced many-core embedded processors, and a new methodology for measurement-based timing analysis of parallel real-time applications.

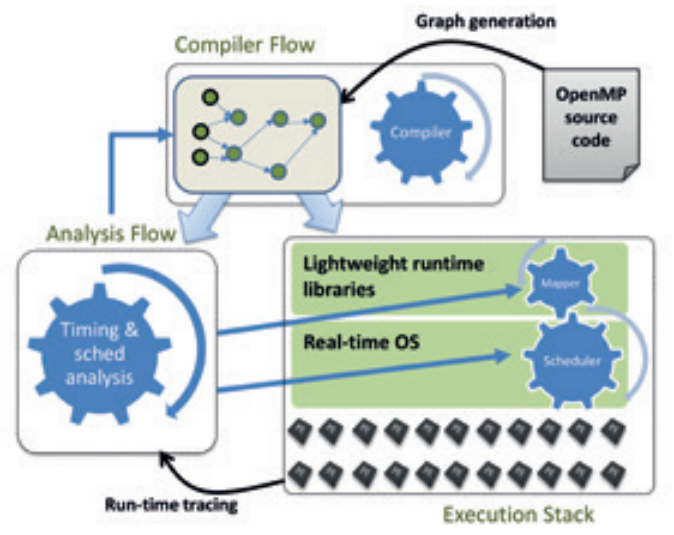

These advances have been integrated in the P-SOCRATES Software Development Kit (SDK), a set of tools to enable the development of high-performance real-time applications, in a vertical stack including compiler, runtime libraries and operating systems. This SDK also includes a feedback loop of timing and schedulability analysis tools to support the design process.

P-SOCRATES technology and tools are integrated into, and evaluated with, three real world industry stakeholder's end user applications to provide a thorough and complete evaluation of the proposed framework.

POTENTIAL IMPACT

Industrial companies will benefit from the project outcomes, allowing European technology suppliers to properly exploit the capabilities of next-generation hardware platforms in a predictable way. Impacts are foreseen in the development of enabling technologies for both the high-performance and embedded computing domains. From an applicative point of view, P-SOCRATES will represent a reference point for the implementation of workload-intensive applications with time-criticality requirements, enabling a more efficient smart society. The computing technology developed in the project will allow a deeper understanding of many-core off-the-shelf systems, enabling new kinds of applications to be developed on top of these platforms.

COORDINATOR + CONTACT INFORMATION

Coordinator: Prof. Luis Miguel Pinho
E-mail: LMP@isep.ipp.pt
Address:
CISTER Research Centre
Instituto Superior de Engenharia do Porto
Rua Dr. Antonio Bernardino Almeida, 431
4249-015 Porto
Portugal

PHANTOM

Cross-Layer and Multi-Objective Programming Approach for Next Generation Heterogeneous Parallel Computing Systems
www.phantom-project.org
1 December 2015 – 30 November 2018

OVERALL OBJECTIVES

The overall objective of the PHANTOM project is to deliver an integrated cross-layer (hardware and system software/programming environment), multi-objective and cross-application approach that will enable the rapid development of next generation heterogeneous, parallel and low-power computing systems, while hiding the complexity of computing hardware from the programmer, thus fostering greater programming productivity. The PHANTOM project will take a holistic approach by addressing the complete system view, both of hardware and software, as well as being applicable to development of applications for different market sectors.

MAIN RESULTS

The results from PHANTOM will be structured in three layers. First, parallel programming and productivity tools will be provided including application-driven APIs for programming and annotations, a parallelization toolset for maintaining intrinsically the code parallelization, and model based testing techniques for early parallel program verification. Second, multi-dimensional optimization will be addressed through an adaptive and multi-objective scheduler, deciding on where to execute each application component, which is supported by runtime monitoring/data analytics and security implementations. Finally, low power, heterogeneous hardware platforms are built together with system software for enabling their management as a service. In particular, the PHANTOM system provides a hardware-agnostic software platform that operates over reconfigurable multi-core and heterogeneous (GPU, FPGA, CPU) hardware platforms.

PHANTOM brings multi-disciplinary expertise through an ecosystem of academia, industries and a strong number of SMEs. The outcome will be validated in three use cases from the Automotive,

Telecoms and Surveillance industries in order to prove a cross-market approach, while demonstrating the benefits the PHANTOM technologies provide in improved productivity for developing systems for heterogeneous platforms.

POTENTIAL IMPACT

In technology driven industries "to out compute is to out compete" and Europe's key industrial sectors are highly dependent on technological progress in computing. Although computing has achieved unparalleled progress, new challenges including the evolution towards cyber physical systems, the proliferation of devices and the big data they produce, and requirements to reduce energy footprint have substantially increased programming complexity for systems development. PHANTOM will deliver an economically and energetically sustainable solution for next generation computing systems that hides the complexity of systems from the programmer and offers multi-dimensional optimisation thereby increasing the rate of innovation for Europe's technology driven industries.

COORDINATOR + CONTACT INFORMATION

The Open Group
www.opengrou.org
Name: Scott Hansen
E-mail: s.hansen@opengroup.org

POLCA

Programming Large Scale Heterogeneous Infrastructures
http://www.polca-project.eu
1 September 2103 – 31 August 2016

OVERALL OBJECTIVES

The POLCA project addresses the programmability concerns of embedded and high performance computing. POLCA proposes a hybrid programming model that decisively increases efficiency, performance and enables realisation of multi domain use cases.

The model allows efficient parallelisation and distribution of the application code across a highly heterogeneous infrastructure, through exploitation of fundamental mathematical axioms behind the execution logic.

To maintain controllability and ensure proper, reliable execution of the non-functional criteria, POLCA can generate stand-alone code that does not require support through virtualisation technologies, but addresses the specifics of the destination platform directly.

MAIN RESULTS

The POLCA project will enable more efficient and portable development of highly performant parallel applications in a variety of domains and infrastructures. Concretely, POLCA provides

(1) a mathematical data-flow oriented Programming Model that and allows transformation of the code, so as to optimise for different algorithms and infrastructures. The model is basing primarily on annotations to standard C/C++ code, thus ensuring maximum compatibility with existing source code and a minimal entry barrier.

(2) the Compilers and Analysis Tools will provide the necessary means to interpret POLCA's annotations and assess them against the exposed programming structure, in particular in terms of non-functional criteria, but also to extract dependency information, communication needs etc. The tools incorporate expertise from both domains in order to make common challenges addressable for the developer. POLCA acts on a source-to-source level, thus generating hardware-optimised C/C++ code that follows common programming models for that platform, such as CUDA or MaxJ. This allows full exploitation of low-level optimization through system-specific compilers and increases maintainability.

POTENTIAL IMPACT

POLCA offers a better abstraction from the complex platform specificities while enabling efficient and automated transformations that feed the back-end compilers. This allows for enhanced productivity, usability and performance exploitation on the target hardware.

POLCA is to be firstly exploited by partners Recore Systems and Maxeler as design entry tools for their customers: Coupled to the respective development frameworks, POLCA facilitates and accelerates application development on the respective hardware platforms.

Other specific components of the POLCA toolchain, such as Clash compiler and Poroto tool, give potential for exploitation via spin-off consultancy/training services (UT) or tech-transfer activities (CETIC).

COORDINATOR + CONTACT INFORMATION

Coordinator: Lutz Schubert, University of Ulm, Albert-Einstein-Allee 43, 89081 Ulm, Germany.
Email: lutz.schubert@uni-ulm.de

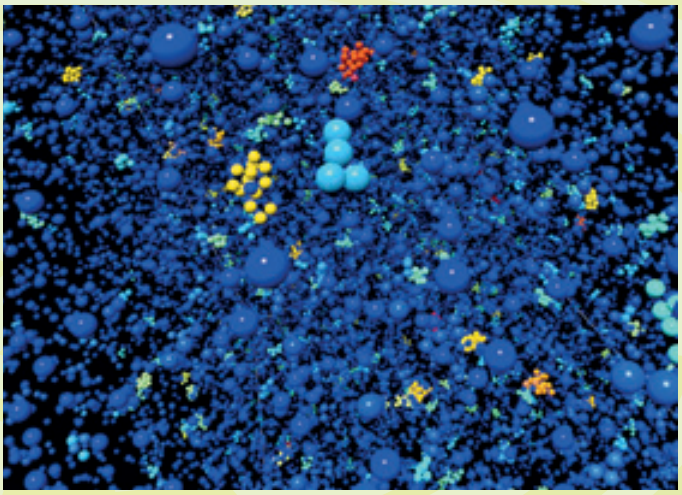

Picture 1: nucleation process

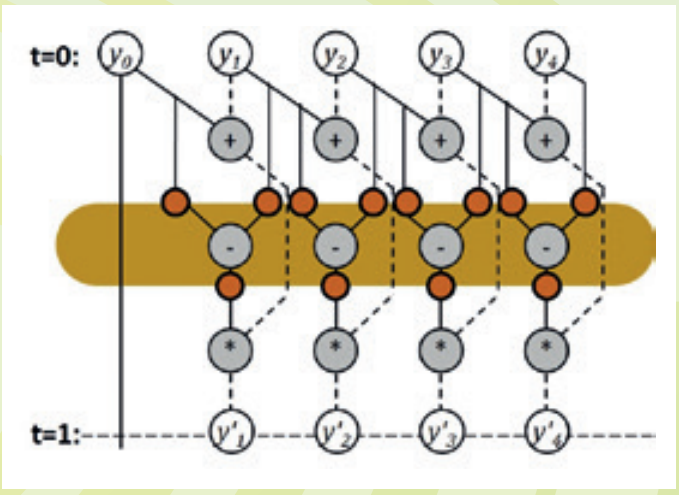

Picture 2: dependency graph breakdown

PROXIMA

Probabilistic real-time control of mixed-criticality multicore and manycore systems
http://proxima-project.eu/
1 October 201??? – 30 September 2016

OVERALL OBJECTIVES

PROXIMA aims at facilitating and reducing the overhead in the application of measurement-based timing analysis – the industrial most-used approach for timing. PROXIMA further helps to increasing the confidence obtained on timing bounds.

To reach its objectives PROXIMA pursues the development of measurement-based probabilistically time analyzable (MBPTA) techniques and tools for multicore platforms. PROXIMA selectively introduces randomization in the timing behavior of certain hardware/software resources as a way to facilitate the use probabilities to predict the overall timing behavior of the software and its likelihood of timing failure. PROXIMA considers COTS technology and address the disruption caused by non-PTA-conformant components.

MAIN RESULTS

The main results of the project include:

- A set of hardware designs principles, some of which have been implemented at the FPGA level, showing how to achieve with low implementation cost the desired timing behaviors required for MBPTA. These include the randomization of hard-to-analyse resources and the removal of jitter from other resources.
- Several software approaches have been designed for COTS architectures in which the desired timing behavior is achieved through a set of specialized randomization libraries that allocate programs code and data at random memory locations across runs.
- In conjunction with randomization injection, several probabilistic timing analysis techniques have been developed inside the project to provide tight and reliable probabilistic WCET estimates. The integration of those techniques into a commercial timing analysis framework has been achieved.
- Time composability principles have been pursued in the real-time operating systems (RTOS) - both academic and commercial - as a way to easily factor in the latency incur by the RTOS.
- Evaluations with industrial case studies have been performed to assess the benefits of the PROXIMA approach in avionics, space, railway and automotive case studies.

POTENTIAL IMPACT

The PROXIMA project is supporting the adoption of advanced hardware in critical software systems by providing the capability to analyze the timing behavior of multi-core and complex microprocessors. Targeting the aerospace, space and automotive industries, PROXIMA enables cost-effective verification of software timing analysis including worst case execution time for the next generation of processors. This will provide an increase in the computational power available and a reduction in the weight vehicle. PROXIMA provides industry ready software timing analysis using probabilistic methods for complex multi-core critical real-time embedded systems.

COORDINATOR + CONTACT INFORMATION

Francisco J. Cazorla
Barcelona Supercomputing Center – IIIA-CSIC
francisco.cazorla@bsc.es

REPARA

Reengineering and Enabling Performance and power of Applications
www.repara-project.eu
1 September 2013 – 30 August 2016

OVERALL OBJECTIVES

The REPARA project aims to help the transformation and deployment of new and legacy applications in parallel heterogeneous computing architectures while maintaining a balance between application performance, energy efficiency and source code maintainability.

Specific objectives are:

O1: Create a language representation allowing code transformations.
O2: Provide mechanisms for application partitioning for parallel heterogeneous architectures.
O3: Transform software components to specific programming models.
O4: Compile software components for execution on reconfigurable hardware.
O5: Create predictive performance and energy models.
O6: Provide a run-time system to integrate transformed components.
O7: Validate the framework with applications from embedded and high performance sectors.

MAIN RESULTS

- A platform description language that allows representing the details of a parallel heterogeneous architecture and a tool for extracting information from a real platform.
- A set of ISO C++11/14 code annotations to support source code transformations and a tool to support developers in the annotation process.
- A partitioning and mapping mechanism that allows the identification of the computational kernel and helps to decide the best mapping of different software components to the available devices.
- A refactoring tool integrated into an Eclipse based IDE (Cevelop) to support the transformations from annotated code to parallel code for multi-cores, GPUS and DSPs as well as to FPGAs.
- A run-time library (FastFlow) to support the coordination of software components running in different computing devices from a parallel heterogeneous architecture.
- A set of performance and energy monitoring components integrated in the REPARA run-time to get feedback as well as prediction mechanisms.

POTENTIAL IMPACT

The technology demonstrated in REPARA supports (semi-) automatic transformations of existing sequential C++ code and to parallel code targeting state-of-the-art parallel heterogeneous architectures. The approach, based in the use of fully standard C++ code and extendable to more programming models, has been tested with real use cases as a railway monitoring system and a protein docking software package with results comparable to hand written code but at a fraction of original development costs. Properly engineered and integrated with existing compiler tool chains the REPARA results may drastically reduce time-to-production of parallel applications with a notable impact on software costs.

COORDINATOR + CONTACT INFORMATION

José Daniel García.
University Carlos III of Madrid
Av. Universidad Carlos III, 22. 28911 Leganés Madrid
Email: josedaniel.garcia@uc3m.es

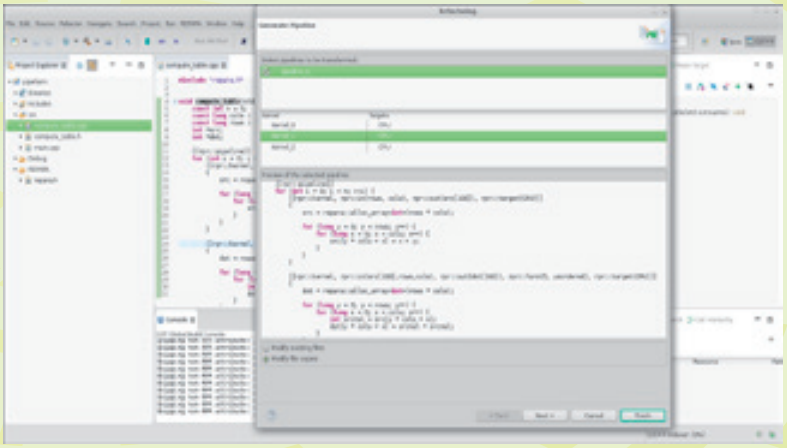

ROAD2CPS

Strategic action for future CPS through roadmaps, impact multiplication and constituency building
http://www.road2cps.eu/
1 February 2015 – 31 January 2017

OVERALL OBJECTIVES

- To develop technology, application and innovation strategy roadmaps, to perform case studies and to derive recommendations for future research and innovation strategies.
- To assess and multiply the impact of past and ongoing projects in CPSs and related fields, accompanied by raising awareness and disseminating programme achievements to support the timely uptake of novel approaches.
- To bound and build a Constituency aware of – and united by – their commonly faced CPS challenges and demands and to develop task forces for specific actions (CPS and society; CPS and business; CPS towards platforms; CPS connection).

MAIN RESULTS AND ACHIEVEMENTS

54 past and ongoing CPS related projects were analysed (results, gaps, impacts) and the results were integrated into Road2CPS deliverables and roadmap building activities.

CPS-roadmapping projects (e.g. CPSoS, CyPhERs, Road2SoS, T-AREA-SoS, Compass, Road4FAME, sCorPiuS, ProcessIT, ATOS vision, ARTEMIS-SRA) were involved in a consensus workshop to compare, discuss and prioritise emerging technologies and implementation barriers. The highest ranked research priorities included i) integration, interoperability, standards; ii) safety, reliability, resilience, fault tolerance; iii) modelling and simulation. Main barriers for CPS deployment next to lack of interoperability included i) skills, knowledge training, ii) policy, regulatory, security, and safety, iii) business models and financial burdens.

A workshop focusing on future platforms was held in Turin (October 2015), to match supply and demand and inform the EC/community on the latest developments and current needs. In April 2016, Road2CPS organised a clustering and communications event in Vienna, which provided fruitful grounds for connecting Horizon2020 ICT-1 and ARTEMIS projects.

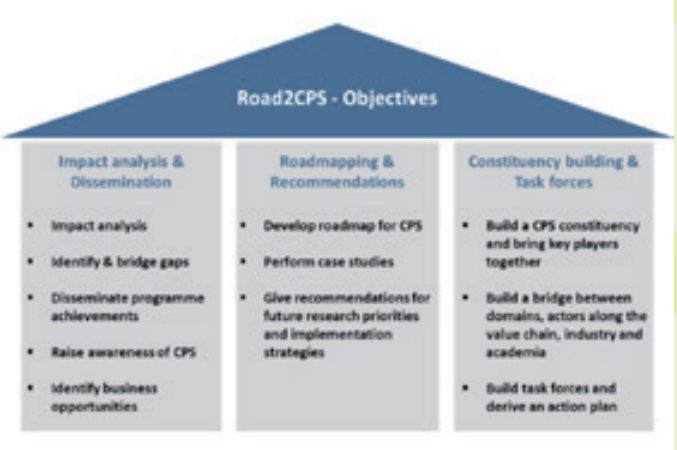

The results and insights gained from all activities were broadly disseminated, taking into account a variety of domains such as manufacturing, energy, transport, smart city, and health. In the course of the project case studies will be performed to show applicability of CPSs specifically to SMEs.

POTENTIAL IMPACT

Even though tremendous progress has been made in advancing CPS technology over the last couple of years, there is still a huge gap between theoretical concepts, technical developments and successful application, as well as considerable differences with regard to propagation and maturity of CPS between application domains. Road2CPS seeks to close these gaps by analysing the CPS landscape, bridging efforts and facilitating mutually beneficial collaborations between the related stakeholders. Moreover, Road2CPS roadmaps serve as orientation and catalyst for early adoption of CPS technologies. Road2CPS recommendations support the implementation of the EC's Strategy for "Digitsing European Industry" and give thematic input to the ICT-Work Programme.

COORDINATOR + CONTACT INFORMATION

Meike Reimann
reimann@steinbeis-europa.de

Dr. Meike Reimann has worked as a project manager at SEZ specializing in EU-ICT, EEB and FoF projects for 5 years. She has a decade of experience in the scientific and administrative coordination of EC-funded IPs, CPs, IAs and CSAs. Furthermore, she has an expertise in project management & coordination, roadmapping, workshop organisation, dissemination and exploitation activities. She works as a management partner in various projects e.g. CPSELab, as an exploitation partner in BRICKER, (concluded project INTUITEL), as a roadmapping and dissemination partner in TAMS4CPS and as the coordinator of CSAs e.g. Road2CPS (ICT-1).

SAFEPOWER

Safe and secure mixed-criticality systems with low-power requirements
http://safepower-project.eu/
1 January 2016 – 31 December 2018

OVERALL OBJECTIVES

The main objective of SAFEPOWER is to enable the development of cross-domain mixed-criticality systems with low power, safety and security requirements by means of the next key contributions:

- The definition and development of a cross-domain mixed-criticality and low power reference architecture upon multicore/heterogeneous processors
- The definition, implementation and demonstration of a set of mixed-criticality compliant low-power techniques and power management procedures that can be used in the development of mixed criticality real-time systems with safety and security requirements
- The development of platforms and tools to enable and facilitate the development of low power mixed criticality real-time systems

MAIN RESULTS

The main results of the project include:

- The definition and implementation of a reference architecture for low-power and mixed-criticality application
  - A catalogue of safe and low-power features and services
  - Safety standard (IEC-61508/EN-5012x) complaint architecture
- The development of tools for the architecture
  - A virtual platform technology including power/energy/temperature estimation
  - A custom SAFEPOWER PCB design: highly instrumented, communicated and low-power by design
  - An hypervisor that provides the interface to the architecture (with mixed-criticality and low-power services extension)

SAFEPOWER will demonstrate the benefits through two industrial use-cases in the railway and aerospace domain and a cross-domain public demonstrator.

POTENTIAL IMPACT

The SAFEPOWER innovative architecture on low power and safety-related hardware and software tools will enable EU industry to reinforce its technological and market leadership on the expanding autonomously powered systems increasingly including safety-critical certification requirements.

This includes increasing the programming productivity of such systems by the use of low-power architectural services and power-aware virtual platforms that enable the early development of the applications.

Additionally, as the scale of the development of autonomously energy powered system shrink, SMEs are more and more suitable to exploit the outcome of the SAFEPOWER project.

COORDINATOR + CONTACT INFORMATION

Dr. Mikel Azkarate-askasua
P°. J. Mª. Arizmendiarrieta, 2
20500 Arrasate-Mondragón
Spain
Tel: +34 943 712 400
Fax: +34 943 796 944
EMail: MAzkarateAskasua@ikerlan.es

SAFURE

SAFety and secURity by dEsign for interconnected mixed-critical cyber-physical systems

www.safure.eu

1 February 2015 – 31 January 2018

OVERALL OBJECTIVES

The project SAFURE aims at addressing the security of safety-critical cyber-physical systems by implementing a holistic approach to safety and security by construction. For this purpose, extensions of tools and system capabilities are developed to prevent, detect and protect against possible vulnerabilities and attacks. Efficient system configurations and reconfigurations, keep critical subsystems within their safety and security boundaries without inflicting performance impairments for best-effort applications. Thus, the SAFURE Framework will extend system capabilities to preserve the system integrity from time starvation, massive energy dissipation and data corruption, seamlessly integrating security requirements into safety systems from a new point of view.

MAIN RESULTS

SAFURE's mission is to design a cyber-physical system by implementing a methodology that ensures safety and security by construction. This methodology is enabled by a framework developed to extend system capabilities so as to control the concurrent effects of security threats on the system behaviour. With this in mind, the project aims at allowing European suppliers of safety-critical embedded products to develop more cost and energy-aware solutions.

Holistic approach to safety and security by construction targets to implement a holistic approach to safety and security by construction of embedded dependable systems, preventing and detecting potential attacks and increasing end-to-end system performance for security and safety-critical domains.

Empowering designers and developers is about empowering designers and developers with analysis methods, development tools and execution capabilities that jointly consider security and safety, communications and runtime system support requirements.

Opportunity to extend current standards aims at providing extensions to current safety-related standards that will set the ground for the development of SAFURE-compliant safe and secure mixed-critical embedded products.

POTENTIAL IMPACT

- Reduction of development time for CPS by 30% as compared to the state-of-the-art in 2013 and significant reduction in maintenance costs.
- Stronger pan-European collaboration across value chains and technology levels from the components and hardware to higher systems level creating open innovation eco-systems and stimulating consensus building on open tools, platforms and standards.
- Development of a next generation core ICT platforms spanning from operating systems and middle ware to application development and deployment tools with built-in security.
- Uplifting Europe's innovation capacity and competitiveness across all economic sectors with the wider adoption of networked embedded ICT, notably in SMEs.

COORDINATOR + CONTACT INFORMATION

Dr. Klaus-Michael Koch
Technikon Forschungs- und Planungsgesellschaft mbH
Burgplatz 3a
9500 Villach
Austria
Phone: +43 4242 233 55 – 71 / Fax: +43 4242 23355 – 77
Email: coordination@safure.eu
Web: www.safure.eu

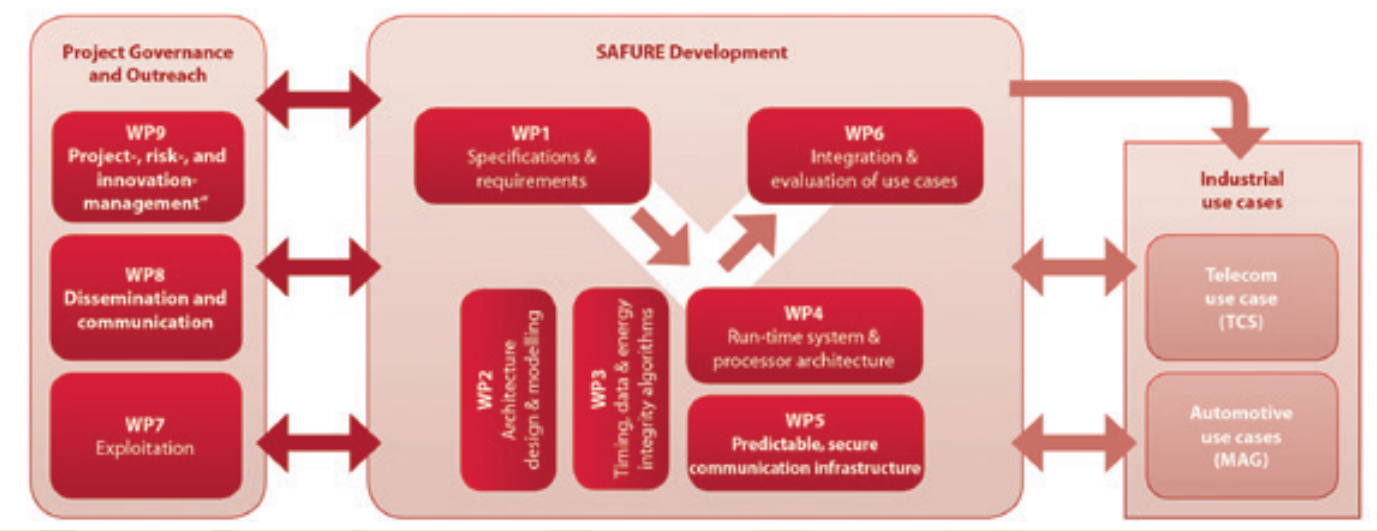

SAVE

Self-Adaptive Virtualisation-Aware High-Performance/Low-Energy Heterogeneous System Architectures
http://www.fp7-save.eu/
September 2013 – August 2016

OVERALL OBJECTIVES

SAVE aims at providing a set of software and hardware technologies for the exploitation of heterogeneous system architectures by means of dynamic, self-aware resource management for energy efficiency and power/performance optimization. More precisely, based on the current workload, the proposed technologies support the identification of the most appropriate kind of computing resource to be used (CPUs, GPUs or hardware accelerators) in order to meet the user's performance requirements while minimizing the energy footprint. To enable such transparent and simplified usage of heterogeneous resources, a virtualized environment is supported, thus decoupling applications from the heterogeneous architecture underneath.

MAIN RESULTS

SAVE has developed a set of hardware and software technologies for the exploitation of heterogeneous system architectures allowing for a dynamic, self-aware configuration of the execution of the workload with the aim to achieve the most interesting trade-off in terms of performance and power consumption. To decouple the underlying complex and heterogeneous architecture from the applications, the proposed solutions work also within a virtualized environment.

More precisely, the consortium has developed system components for i) enhanced runtime resource management, ii) generating code for accelerators (e.g., GPUs and DFEs) also providing autonomous migration of workload at runtime to such resources when deemed beneficial, iii) efficient communication among the heterogeneous resources, and iv) virtualization support for GPUs, DFEs and generic accelerators. These components work as standalone modules that can be integrated in a heterogeneous system architecture to better control/exploit some of the available features. Moreover, they have been successfully integrated for an "orchestrated" runtime management of the resources, aimed at optimizing user-defined goals (power minimization, performance, resiliency, ...). One of the key aspects of the proposed solutions is the ability of such technologies to autonomously adapt to changing application contexts at runtime.

POTENTIAL IMPACT

The developed technologies provide innovative solutions for the management of heterogeneous system architectures within the SAVE vision, but also constitute important building blocks. Patents have been filed for some of the developed technologies, that can be expected to be adopted in future-generation products, for an improved and simplified usage of the computing resources, for instance through an efficient communication, or via hypervisor extensions to access virtualized resources. Such new results and methods have also been disseminated in the scientific community and in the educational environment.

COORDINATOR + CONTACT INFORMATION
Politecnico di Milano - DEIB
P.zza L. da Vinci, 32 – 20133 Milano – Italy
cristiana.bolchini@polimi.it
+39-02-2399-3619

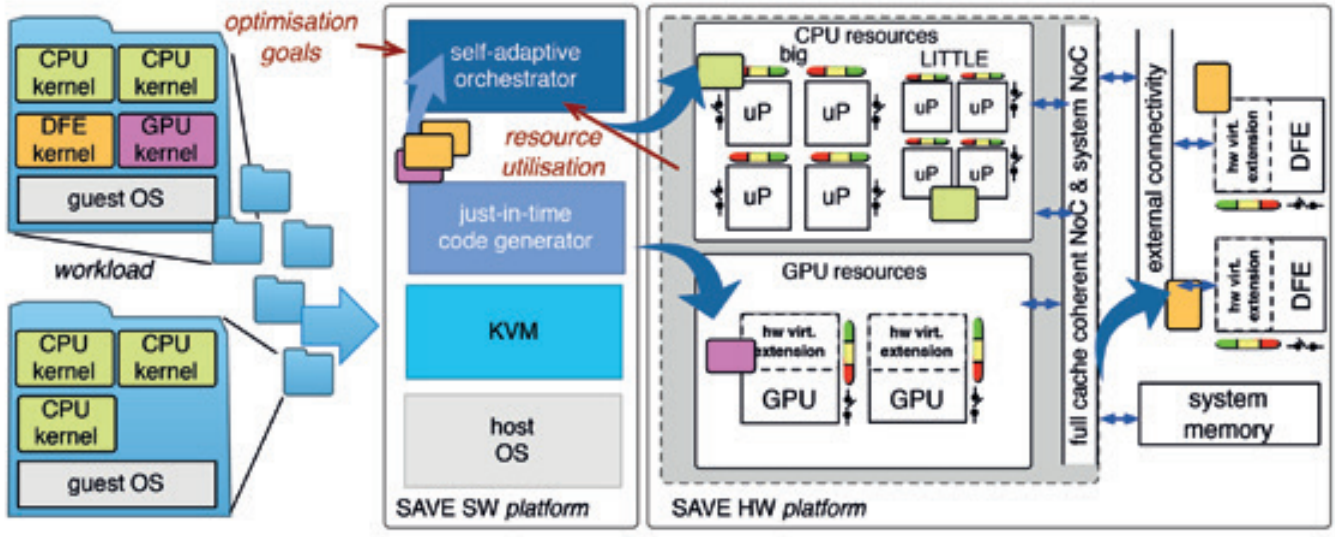

TAMS4CPS

Trans-Atlantic Modelling and Simulation for Cyber-Physical Systems
www.tams4cps.eu
February 2015 – January 2017

OVERALL OBJECTIVES

The TAMS4CPS project is a coordination and support action (CSA) funded by the European Commission under the grant agreement number 644821. The project runs from February 2015 to January 2017 and its overall aim is to lay the foundations for concrete EU-US collaboration in modelling and simulation for CPS.

To achieve this, the project partners (Loughborough University, Newcastle University and Steinbeis-Europe-Zentrum) work with five leading researchers in the field at top US universities (George Masson University, Georgia Institute of Technology, Purdue University, University of Texas at San Antonio and Stevens Institute of Technology) to create:

- A Strategic Research Agenda for Collaboration, endorse by researchers in EU and US,
- A set of Test Cases for use by model developers to perform collaborative evaluation
- A State of the Art web-based report to act as a baseline for collaborative research

MAIN RESULTS

During the first half of TAMS4CPS, eight dream projects have been formulated. These are descriptions of potential research activities that participants in the workshops have identified as plausible and useful for collaboration between EU and US. During the same period, nine potential test cases have been identified.

From the first workshops, it is clear that a major area of interest is access to sufficiently rich datasets to validate models and simulations is a high priority. Access may be through data being made available, but an important consideration is the setting up of testbeds that can be used for existing CPS and may form a building block towards new CPSs in the future.

- As an interim conclusion, TAMS4CPS would recommend establishing links between appropriate European and US partners to federate existing testbeds and, perhaps, establish new ones collaborative.
  In terms of funding collaborative research, it would seem to be the National Labs that align most suitably from a research perspective with H2020.
- As an interim conclusion, TAMS4CPS would recommend investigating the potential and opportunities of co-ordinated calls for research in M&S for CPS with a National Laboratory (such as Sandia National Lab).
- To be successful, calls based on EU-US collaboration must be highly co-ordinated so that both parties are funded (or not) and criteria for selecting projects are well-aligned.

POTENTIAL IMPACT

The TAMS4CPS, as a support action, is an enabler for the impacts. The impacts themselves will be the results of the research recommendation from TAMS4CPS being successfully implemented.

For example, it is becoming clear that a major obstacle, faced by researchers on both sides of the Atlantic, is a lack of test beds upon which to exercise CPS research advances. However, in the context of joint EU-US activity, the question arises: why should the European Commission invest in such facilities? This provides a useful example of the filter through which impact must be assessed. Although the impact achieved through research recommended by this programme may be significant, the question that will be addressed in this project concerns the impact of the European Commission investment in EU-US activity, rather than the more general question.

It is clear that a major concern for the development of EU-US projects is the funding mechanisms through which they would be enabled. TAMS4CPS began by focusing on the technical aspects of collaboration, but without appropriate mechanisms, collaboration will not be possible. To enhance the impact, consideration is being given to the mechanisms and US bodies through which collaboration may be enabled.

COORDINATOR + CONTACT INFORMATION

Michael J deC Henshaw
Professor of Systems Engineering
The Wolfson School
Sir David Davies Building
Loughborough University
Ashby Road
Loughborough, Leicestershire - LE11 3TU - United Kingdom
Phone: +44 (0)1509 635 269
e-mail: m.j.d.henshaw@lboro.ac.uk
Personal Webpage: http://www.lboro.ac.uk/departments/eese/staff/michael-henshaw/

TANGO

Transparent heterogeneous hardware Architecture deployment for eNergy Gain in Operation

Tango-project.eu

1 January 2016 – 31 December 2018

OVERALL OBJECTIVES

In the upcoming era of Mobile, IoT, Big Data and HPC, new and more demanding applications will soon show significant interest in exploiting the capabilities offered by emerging customized heterogeneous hardware such as FPGA, ASIP, MPSoC, heterogeneous CPU+GPU chips and heterogeneous multi-processor clusters.

The market shows evidence that exploiting parallelism is significantly increasing in relevance, as parallelization has become a dominant method of delivering higher performance and improved energy efficiency. Thus, TANGO's main goal is to understand the factors affecting power consumption in software development and operation for heterogeneous parallel environments.

Thus TANGO wants to take on the opportunity of designing more flexible software abstractions and improved system architectures to fully exploit the benefits of these heterogeneous platforms.

MAIN RESULTS

TANGO will address the total characterization of software and heterogeneous hardware with respect to the impact of software structure on power consumption, and performance among various other dimensions.

The novelties brought by TANGO are a reference architecture, and its implementation with mechanisms that allow control and optimization of the heterogeneous parallel infrastructures according to energy efficiency, performance, and other dimensions; and a programming model, a runtime, and IDE plug-ins with built-in support for various hardware architectures including heterogeneous clusters, heterogeneous chips and programmable logic devices. IT will also deliver a hardware and software energy model with metrics for heterogeneous architectures, which will lead the adaptive quality model for holistic system performance. The most important outcomes of the TANGO Toolbox will be released as Open Source.

TANGO also considers the foundation of a Research Alliance in which it will seek complementary efforts of other research projects, initiatives and IT community organizations to nurture a strong collaboration, integration and effective promotion of the results to advance in the future.

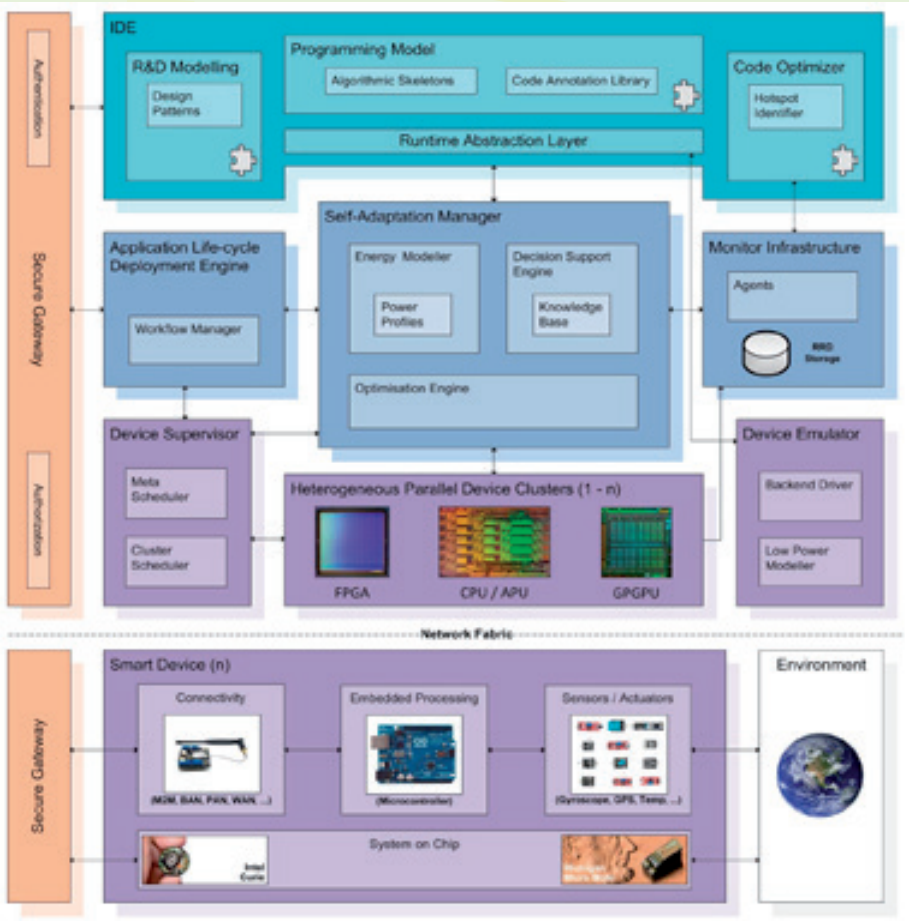

POTENTIAL IMPACT

TANGO will impact on IT industry and market; and the research community by advancing future application development processes to a new stage in which the development process for parallel architectures will be simplified; will be abstracted from underlying architectures and hardware; and will be empowered by tools and a programming model that consider optimized control and self-adaptation

Simply put, **TANGO can Simplify & Optimize Heterogeneity, thus simplifying the way developers approach the development of next-generation applications**. TANGO helps controlling and abstracting underlying heterogeneous hardware architectures, configurations and software systems including heterogeneous clusters, chips and programmable logic devices while providing tools to optimize various dimensions of software design and operations such as energy efficiency, performance, data movement and location, cost, time-criticality, security, dependability on target architectures.

COORDINATOR + CONTACT INFORMATION

Clara Pezuela, Atos (clara.pezuela@atos.net)

TAPPS

Trusted Apps for open CPSs
www.tapps-project.eu
1 January 2015 – 31 December 2017

OVERALL OBJECTIVES

The TAPPS project offers a new approach towards extensibility of cyber-physical systems (CPS) platforms, going beyond traditional solutions for safety, security and reliability in the CPS domain. The TAPPS approach is based on a separate, dedicated, real-time execution environment for distributed, safety-critical CPS applications offering multiple layers of security, and a holistic, open end-to-end tool chain for developing and deploying CPS apps. The multi-level trusted apps platform and tool chain is validated in applications from the automotive and healthcare domains using industrial, realistic use cases.

MAIN RESULTS

The architecture we propose within the TAPPS project addresses all layers from hardware over software to an app store ensuring security, safety and full real-time support for the applications. For ensuring trusted execution of CPS apps, we focus on four key features: the execution environments (EE) and apps platform, the inter-EE and inter-app communication, the system and network architecture, and the development / model-based tool-chain.

The main technical achievements includes

- spatial and temporal app isolation,
- an app development tool chain providing state-machine based modelling framework,
- verification trough model checking,
- access control to critical interfaces, e.g. restricted CAN bus access,
- multi-layered defense against malicious attacks using highly hardware-based security, as well as
- virtualization techniques and mechanisms for communication control at the middleware layer.

POTENTIAL IMPACT

The ambition is to impact in the automotive and healthcare domains by innovative solutions, which have the potential to rapidly enter the market for motorbike and smart health trolley products. These results shall have a showcase effect on other related domains, which can take up the TAPPS solution and exploit it in their target markets.

The innovative solutions by trusted apps can also improve the user experience and flexibility of such devices, as well as providing more resource-efficient, customized solutions. This has a general benefit to quality of life and resource efficient society. For instance using apps on the smart trolley, we can bring new treatments to the market much faster, compared to typical vertical solutions in the medical domain.

COORDINATOR + CONTACT INFORMATION

Coordinator: Christian Prehofer, christian.prehofer@fortiss.org
Contact information: Nora Koch, nora.koch@fortiss.org
Partners: Fortiss GmbH, ST Microelectronics, TTTech, Virtual Open Systems, Actility, Fondazione Centro San Raffaele, Technological Educational Institute of Crete, CRP GROUP /Energica Motor Company S.R.L

TETRACOM

Technology Transfer in Computing Systems
http://www.tetracom.eu
1 September 2013 – 31 August 2016

OVERALL OBJECTIVES

The mission of the TETRACOM Coordination Action is to boost European academia-to-industry technology transfer (TT) in all domains of Computing Systems. While many other European and national initiatives focus on training of entrepreneurs and support for start-up companies, the key differentiator of TETRACOM is a novel instrument called **Technology Transfer Project (TTP)**. TTPs help to lower the barrier for researchers to make the first steps towards commercialization of their research results. TTPs are designed to provide incentives for TT at small to medium scale via partial funding of dedicated, well-defined, and short term bilateral academia-industry collaborations that bring concrete R&D results into industrial use. This is implemented via open competitive calls for TTPs, whose coordination, prioritization, evaluation, and management are the major actions of TETRACOM. The project´s primary success metrics are the number and value of coordinated TTPs as well as the amount of newly introduced European TT actors.

MAIN RESULTS

TETRACOM is running, or has completed, 50 individual TTPs in different categories of ICT and computing systems, including e.g. communications and multimedia (12 TTPs), industrial automation (10), health (8), safety & security (5), automotive (5), and data analytics (10). The number of TTPs meets the initial expectation and clearly shows the existence of a European "technology transfer market" based on the TETRACOM model. The three open calls for TTPs received 107 proposals altogether, out of which 21 came from new EU member states. The average co-funding from TETRACOM is around 25k EUR per TTP, but with considerable variance. The average co-funding from the TTP industry partners side has been 33k EUR. Across all TTP proposals, the industry partners promised a total co-funding amount of more than 3.5M EUR, which indicates a significant "willingness-to-pay" for new computing technologies developed in academia. Approx. 67% of all company partners are SMEs. Moreover, TETRACOM has become a "brand name" in the European academic computing systems community. Hundreds of participants attended the tech transfer workshops and events organized by the project, which proves a significant community mobilization and interest in transfer opportunities and mechanisms.

POTENTIAL IMPACT

Within a survey, many TTP partners stated that the transfer would not have been possible without the TETRACOM incentive and support. During the TTP impact monitoring, various projects have indicated filing of new patents. Furthermore, most TTPs led to new and improved products and processes. This has been confirmed also by many company endorsements, e.g.

"The project has been very profitable for us. The optimization reduces the losses and it increases our competitiveness in the market. In gross numbers, we can save around 150,000€/year. In our opinion this kind of relationship and technology transfer should be encouraged."
David Rueda, Ingeniería & Mantenimiento, AGC Flat Glass Ibérica

Given a typical TTP budget of 50k EUR (25k from TETRACOM + 25k from the industry partner), the Return on Investment appears really significant. Moreover, TETRACOM also acts as an ICT job catalyst. In several cases, academic researchers performing individual TTPs have subsequently been hired by the industry partner as part-time or full-time staff member. Furthermore, at least four TTPs are known to have led to start-up company foundations or concrete plans.

COORDINATOR + CONTACT INFORMATION

Rainer Leupers
RWTH Aachen University
leupers@ice.rwth-aachen.de

TULIPP

Towards Ubiquitous Low-power Image Processing Platforms
www.tulipp.eu
1 February 2016 – 31 January 2018

OVERALL OBJECTIVES

The global objective of the Tulipp project is to provide vision-based system designers with a reference platform (hardware solution, operating system and programming tool chain) for high performance and energy-efficient image processing embedded systems. New standards derived from its reference platform to the industry will be proposed.

The four objectives of the Tulipp project are summarized below:

- Objective 1: Define a reference platform for low-power image processing applications
- Objective 2: Instantiate the reference platform through use cases applications
- Objective 3: Demonstrate and plan improvements of defined key performance indicators
- Objective 4: Start-up and manage an ecosystem of stakeholder to extend image processing norms

MAIN RESULTS

Tulipp proposed reference platform will achieve the definition of:

- Guidelines to select relevant combinations of computing and communication resources to be instantiated in the resulting platform while minimizing energy consumption and reducing development costs and time-to-market;
- A generic architecture ruling the organization of I/Os, computation, storage resources and management of idle phases with an objective of best utilization of resources;
- Resource management capabilities to coordinate and dynamically switch computing resources between a state of efficient activity and a state of non-consuming inactivity;
- Design tools to help the designer with the evaluation of different possible allocations of its application over the heterogeneous architecture and the selection of the most energy efficient solutions.

POTENTIAL IMPACT

Through the ecosystem built around Tulipp, any company producing Tulipp compliant components will benefit from the other Tulipp component to be able to deliver full platforms dedicated to power efficient image processing. The ecosystem will enable the production of affordable and powerful solutions.

During the project, the developed concepts are going to be proven against use cases from three different application domains:

- Medical: For the wellness of both patients and practitioners, we will divide the radiation dose required by X-Ray imaging by a factor of 4.
- Surveillance and Rescue UAVs: For sending rescue teams only where they are needed, accelerate the search of injured people e.g. after air crashes or nuclear accidents, we will bring intelligence to search drones.
- Advanced Driver Assistance: To lower the number of accident on the road, we will bring drivers with a safer experience.

COORDINATOR + CONTACT INFORMATION

Coordinator: Dr. Philippe Millet – Thales
Partners: Fraunhofer-IOSB, Hipperos, NTNU, RUB, Sundance, Synective Labs, Efficient Innovation
contact@tulipp.eu

Surveillance and Rescue UAVs

Medical X-Ray Imaging

Advanced Driver Assistance

U-TEST

Testing Cyber-Physical Systems under Uncertainty
www.u-test.eu
1 January 2015 – 31 December 2017

OVERALL OBJECTIVES

Cyber-Physical Systems (CPSs) are the next generation of highly connected embedded systems. These systems have applications in varied domains including industrial automation, healthcare, robotics, and maritime industry. Even in the presence of uncertainty, CPSs must be dependable, i.e., trustworthy, robust, efficient, and safe. Inappropriate handling of uncertainty in CPSs during their real operations may have devastating effects on their users and/or environment. The U-Test project aims at ensuring that CPSs are tested adequately under uncertainty using systematic and automated techniques such as model and search-based testing to guarantee their dependable operation in real environment.

MAIN RESULTS

At the current state-of-the-art and practice, Uncertainty in CPS, in general, is not explicitly studied. With this aim in mind, U-Taxonomy is designed relying on investigating the existing works on uncertainty from other fields, where uncertainty has been explicitly studied. In parallel, a set of uncertainty requirements were collected from the use case providers of U-Test.

We have carried out validation of the taxonomy and the requirements and studied with the following two aims in our mind, 1) To precisely define the requirements the requirements such that those can be manually transformed into test ready models, 2) Validating that the U-Taxonomy is sufficiently complete with respect to the two use cases. We have developed an initial version of Uncertainty Modeling Framework (UMF). At the core of framework is the implementation of U-Taxonomy as a UML profile, with which uncertainty can be modelled at the three levels of CPS using UML structural and behavioral models. The UMF also uses exiting standards including UML Profile for Modeling Real-Time and Embedded Systems (MARTE) and the UML Test Profile V.2. Finally, first version of the evaluation plan has been developed that will be used for assessing cost-effectiveness of test cases generated with the Uncertainty Testing Framework (UTF).

POTENTIAL IMPACT

26 potential sources of revenue have been identified and 13 of them have been prioritized with regards to U-Test's key predetermined results; Uncertainty Taxonomy (UTX), Uncertainty Modelling Framework (UMF) and Uncertainty Testing Framework (UTF). These potential sources of revenue, or value opportunities, take the form of products, services and collaborations. Also, key issues that can affect the potential market success of these value opportunities have been identified through the Technology, Market and Enablers-Barriers layers' analyses. These issues will be monitored along the project development.

COORDINATOR + CONTACT INFORMATION

Waqar Ahmed
PCoordinator
waqar.ahmed@oslomedtech.no
Oslo MedTech, Norway

UNCOVERCPS

Unifying Control and Verification of Cyber-Physical Systems
http://cps-vo.org/group/UnCoVerCPS
January 2015 – December 2018

OVERALL OBJECTIVES

The overall goal in UnCoVerCPS is to develop holistic model-based design methods of future cyber-physical systems with a special focus on researching essentially new methods to guarantee safety and reliability in (partially) unknown environments. This is realised by a cross-domain approach for synthesising and verifying controllers on-the-fly, i.e. during operation. In order to quickly react to situations that become critical, a tight integration between the control software and the verification software is realised.

MAIN RESULTS

In UnCoVerCPS, one of the most advanced algorithms for controlling and formally verifying cyber-physical systems are developed. We have already begun to strengthen this position by merging the capabilities of the state-of the-art tools for formal verification of cyber-physical systems: SpaceEx (http://spaceex.imag.fr/) and CORA (http://www6.in.tum.de/Main/SoftwareCORA). The tool SpaceEx also provides a modelling language for systems with mixed discrete and continuous systems, which has become a de-facto standard for exchange of cyber-physical systems in the academic community. The academic tools are integrated into the commercially available tool SCADE. We have further developed a tool for generating formal specifications from structured text, called formalSpec. We are also leading in the field of conformance checking of cyber-physical systems, i.e. we detect the maximum error between an implementation and the used models for developing the implementation.

POTENTIAL IMPACT

UnCoVerCPS will reduce development costs of smart cyber-physical systems used in safety and operation-critical applications. Examples are fully automated systems (e.g. autonomous cars), human-robot applications (e.g. manufacturing) and performance-focused systems (e.g. smart grids). Our deep integration and unification of control and verification techniques will help overcoming the formal verification barrier that exists for safety- and operation-critical cyber-physical systems. UnCoVerCPS develops cyber-physical systems that prove safety of their own actions during runtime, which is a key enabler for the successful deployment of systems like civil autonomous vehicles and systems with a tight interaction between humans and robots.

COORDINATOR + CONTACT INFORMATION

Matthias Althoff
Technische Universität München
Boltzmannstraße 3 - 85748 Garching - Germany

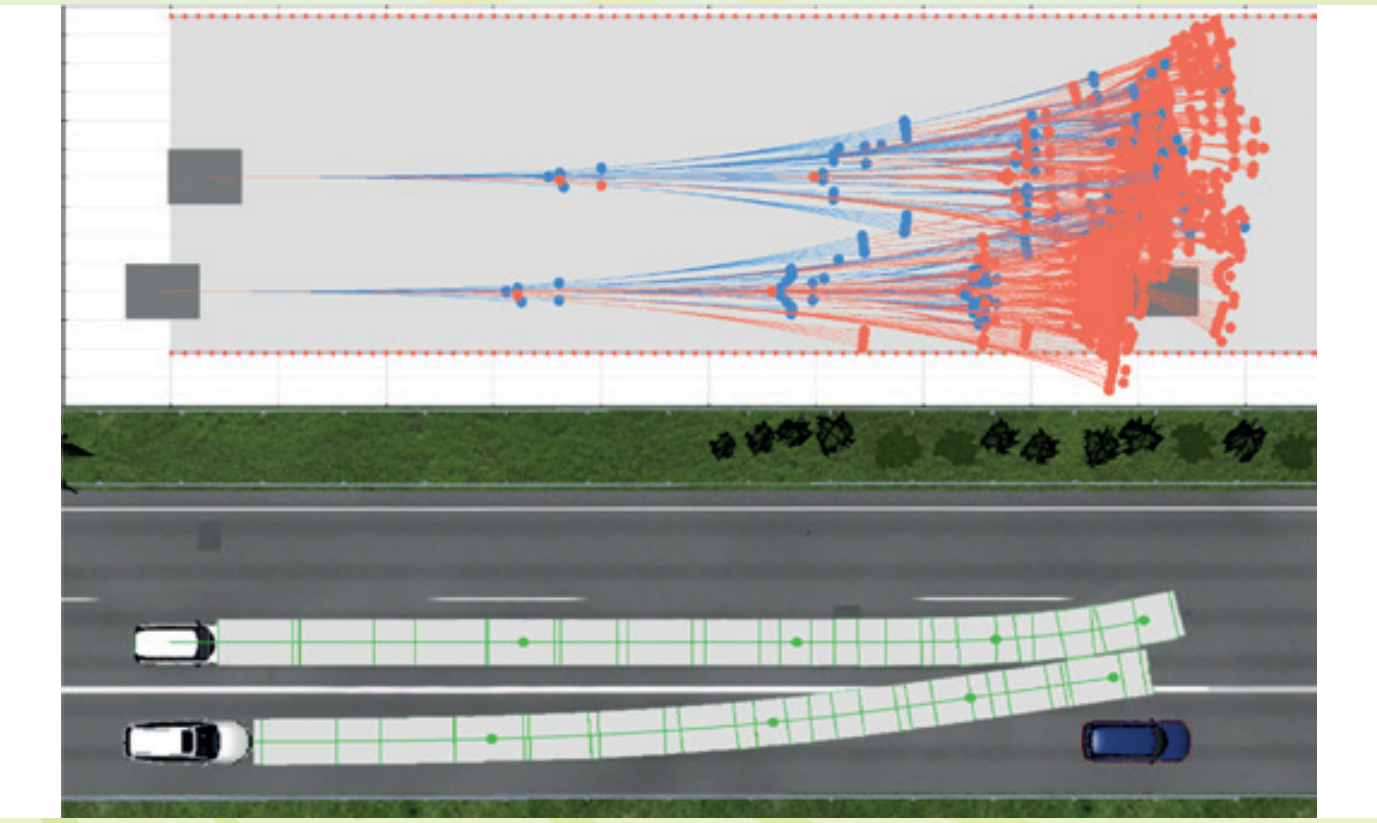

UNIFY-IOT

Supporting Internet of Things Activities on Innovation Ecosystems
www.unify-iot.eu
1 January 2016 – 31 December 2017

OVERALL OBJECTIVES

The UNIFY-IoT objectives are to stimulate the collaboration between IoT projects, between the potential IoT platforms and support these in sustaining the IoT ecosystems developed by focusing on complementary actions, e.g., fostering and stimulating acceptance of IoT technology as well as the means to understand and overcome obstacles for deployment and value creation. UNIFY-IoT is the "working partner" of the Alliance for Internet of Things Innovation (AIOTI) and the Internet of Things European Research Cluster (IERC) by coordinating and supporting the activities on innovation ecosystems, IoT standardisation, policy issues, research and innovation.

MAIN RESULTS

Value co-creation – bringing together the various stakeholders in the IoT ecosystem to work towards a mutually agreed outcome using IoT interoperable solutions and evaluate the value co-creation by analysing the results of the projects.

IoT Business Models – surveying and analysing existing business models related to IoT: from specific deployment in case of process optimisation in a company, to, at the opposite, providing a technological element to the open markets, and produce a taxonomy of business models.

Innovation Support – analysing existing IoT platform deployments and analyses at the innovation and other activities of those deployments. It assesses the relative success of the platform adoptions and identifies common innovation activities in the most successful platforms.

IoT Open platforms concepts – building upon on the open platforms activity chain started by the IERC an combine it with other initiative documenting project outcomes.

IoT Education platform – interacting with stakeholders to identify opportunities for interaction between IoT platforms and education institutions to ensure that future graduates are conversant with emerging IoT platforms and the opportunities they present.

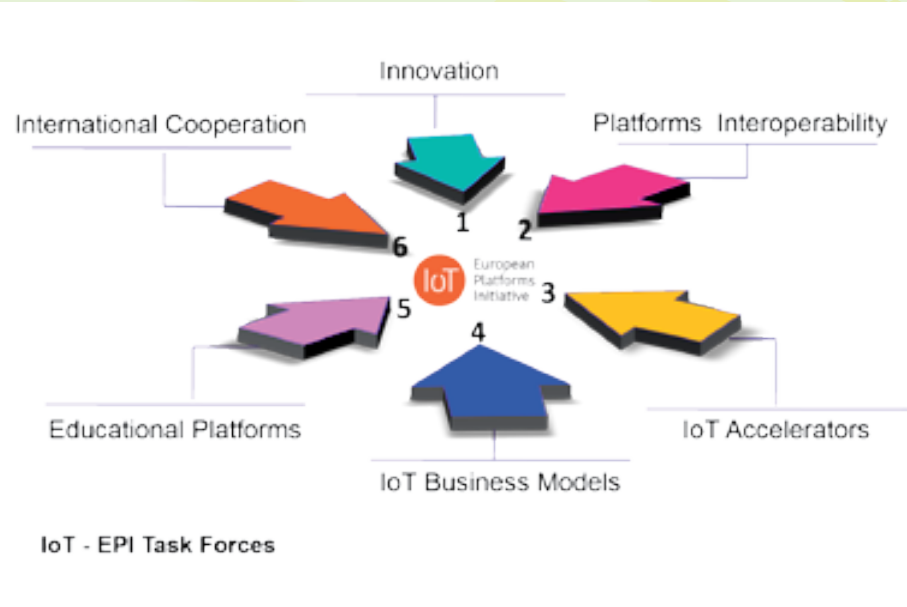

IoT - EPI Task Forces

UNIFY-IoT is leveraging the knowledge process supporting the emergence of an IoT Curricula and education platform in Europe.

Standardisation Support – sensing the global trends in term of interoperability and de-facto standards, and interacts with standardisation bodies including ETSI and CEN/CENELEC to systematise de facto standards emerging from the IoT-EPI projects. The project is working closely with the working group on standardisation of the AIOTI to ensure a coordinated approach to standardisation.

POTENTIAL IMPACT

- Stimulate the dialogue and collaboration between IoT projects and other stakeholders from the IoT ecosystem.
- Support the emergence of an integrated offer on IoT technologies and platforms at European level
- Support the development of architectures and methodologies useful to provide IoT turnkey solutions
- Define exploitation strategies to facilitate the uptake of successful IoT ecosystems targeting the major societal challenges for Europe.
- Foster technology transfer and pre-normative activities
- Provide a set of mechanisms for value co-creation around IoT platform ecosystems, involving end-users in an open innovation ecosystem.

COORDINATOR + CONTACT INFORMATION

Dr. Ovidiu Vermesan, Chief Scientist, SINTEF, Norway, Ovidiu.Vermesan@sintef.no

UNISERVER

A Universal Micro-Server Ecosystem by Exceeding the Energy and Performance Scaling Boundaries
http://www.uniserver2020.eu/
1 February 2016 – 31 January 2019

OVERALL OBJECTIVES

The principal aim of the UniServer project is the development of a universal system architecture and software ecosystem for servers targeting cloud data-centres as well as upcoming edge-computing markets. UniServer will realize its goal by greatly improving the energy efficiency, performance and dependability of the current state-of-the-art micro-servers, while maturing the corresponding system software. This will be achieved by exposing the intrinsic hardware-heterogeneity caused by process variations, to the system-software and enhancing it with new margin/fault-aware runtime and resource management policies. The UniServer technologies will be ported on the world's-first 64-bit ARM based Server-on-Chip and evaluated using smart emerging applications.

MAIN RESULTS

Having inaugurated on February 1st 2016, the project is in its early stages. However, given its aspiring objectives, the consortium has already started demystifying the existing micro-server architecture by revealing the pessimistic margins and by specifying the interfaces between the various layers on the system-stack.

Experimenting with commercial dynamic-memories and trading-off refresh-power and memory availability with reliability, it is being found that the refresh-rate currently adopted in DRAMs for maintaining data-integrity is by far very pessimistic. Remarkably, the initial results have indicated that refresh-rate can be relaxed even by 23x times for some of the tested workloads. Pessimistic voltage margins have also been revealed in competitive commercial multicore/multi-threaded processors, where initial experiments with SPEC2006 suite have shown that no single error occurs even after scaling-down the voltage by up-to 10%. In both cases, as the consortium predicted strong dependence of the magnitude of failures and the processed/stored workload has been observed, which will be further analyzed and exploited in UniServer's work-packages. To enable the propagation of the new extended margins to the software layers and the exchange of useful system-configuration information we have already extended libvirt, the de-facto communication-interface between cloud-management-frameworks (OpenStack) and the virtualization-layer (KVM).

POTENTIAL IMPACT

UniServer seeks to exemplify and materialize the evolution of the Internet from an infrastructure where data is aggregated to centralized data-centres to an infrastructure where data is handled in a distributed and localized manner closer to the data sources – at the Edge of the Cloud. It aspires to deliver a unique fully working micro-server prototype able to operate beyond the power/performance-scaling boundaries turning the opportunities in the emerging Big Data, Cognitive Computing and Internet of Things (IoT) markets into real, smarter products that can improve everyday life, while leading to a substantial financial and employment growth.

COORDINATOR + CONTACT INFORMATION

Georgios Karakonstantis
Institute of Electronics, Communications and Information Technology
Queen's University Belfast, United Kingdom
uniserver-info@qub.ac.uk

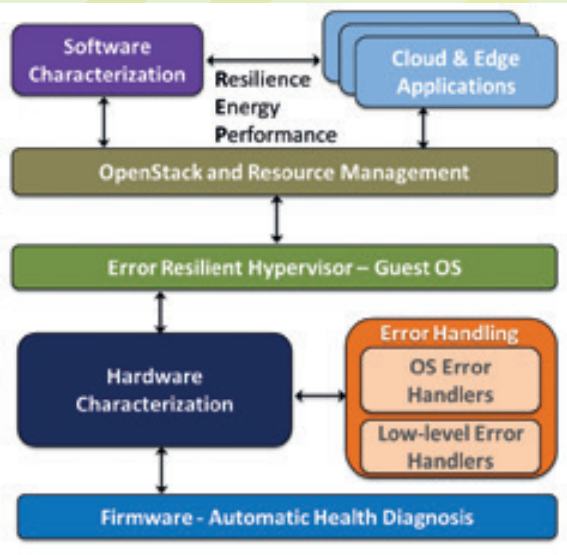

UniServer: Innovating across the system stack

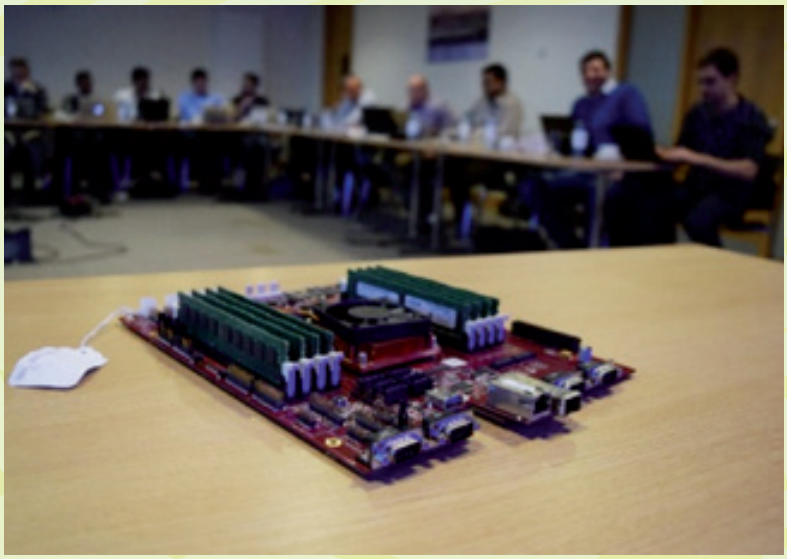

The chassis of UniServer technologies in the spotlight during the Kick-Off meeting (Feb. 2016)

VINEYARD

Versatile Integrated Accelerator-based Heterogeneous Data Centres
http://vineyard-h2020.eu
1 February 2016 – 31 January 2019

OVERALL OBJECTIVES

VINEYARD aims to:

- Build energy-efficient data centres based on novel programmable **hardware accelerators** (namely Dataflow engines and FPGA-coupled servers) that can speedup cloud computing and data analytic applications.
- Develop a **high-level programming framework** for allowing end-users to seamlessly utilize these accelerators in heterogeneous computing systems by employing typical data-centre programming frameworks (i.e. Spark).
- Foster the creation of a new eco-system in which hardware accelerators in the form of intellectual-property (IP) blocks hosted in a repository will be able to be instantiated seamlessly in heterogeneous data centres as pluggable modules that can be swapped in and out of the heterogeneous infrastructure in a similar way to software packages.

MAIN RESULTS

VINEYARD aims to increase significantly the **throughput** of several big data applications and to reduce substantially the **energy consumption** of the data center by utilizing seamlessly the hardware accelerators. The VINEYARD framework will be demonstrated in three real-world use case scenarios: *Financial applications, Data analytics and computational neuroscience.*

Specifically, early results (3 months after the start of the project) show that the hardware accelerators utilized in VINEYARD can provide significant higher performance and lower energy consumption in the domain of computational neuroscience. Specifically, the simulation of olivocerebellar brain circuit, a real-life neuroscientific application, compared the performance of the hardware accelerators based on dataflow engines (DFE) with the Intel Xeon Phi platform. The performance evaluation showed that the DFE can provide one to two orders of magnitude larger neural-network sizes than the Xeon Phi platform. Furthermore, when it comes to real time model execution, the Xeon Phi falls short of the 50 usec timing constraint for achieving real-time simulations, regardless of the network size used.

POTENTIAL IMPACT

VINEYARD's main goal is to design and develop an integrated framework that will allow the seamlessly utilization of customized hardware accelerators (dataflow engines) in data centres and parallel system architectures thus reducing significantly the energy consumption in the data centres. The VINEYARD programming framework will foster the development of an ecosystem that will empower open innovation based on hardware accelerators as data-centre IP plugins, thereby facilitating innovative enterprises (SMEs, and creative start-ups) to develop novel solutions using VINEYARDS's leading edge developments. The ecosystem will bring together existing communities from all relevant stakeholders including providers of hardware intellectual-property (IP) technologies, cloud computing application developers, data centre operators and more.

The VINEYARD framework will be demonstrated in three real-world applications:

Computational neuroscience, Financial applications and Data analytics.

COORDINATOR + CONTACT INFORMATION

Prof. Dimitrios Soudris, ICCS, GR, dsoudris@microlab.ntua.gr
TECHNICAL PROJECT MANAGER
Dr. Christoforos Kachris, ICCS, GR, kachris@microlab.ntua.gr


European
Commission