12th International Symposium on

Reconfigurable Communication-centric Systems-on-Chip (ReCoSoC 2017)

July 12-14, 2017, Madrid, Spain

Program Book

Welcome to ReCoSoC 2017!

ReCoSoC has established itself as a consolidated international event, acting as a reference for researchers in the areas of reconfigurable and communication-centric systems-on-chip. Its informal and dynamic philosophy encourages technical and scientific discussions between researchers at very different levels of expertise: from senior academics to young researchers are welcomed to attend ReCoSoC.

ReCoSoC 2017 is the Twelfth International Symposium on Reconfigurable Communication-centric Systems-on-Chip to take place. The previous ones were held in Estonia (2016), France (2005, 2006, 2007, 2011, 2014), Spain (2008), Germany (2010, 2013, 2015) and UK (2012).

ReCoSoC 2017 will take place in Madrid, the capital city of Spain. Madrid is located in the center of the Iberian Peninsula and has a population of almost 3.2 million. It is the third-largest city in the European Union, after London and Berlin. This is a modern, economic, safe, and cultural city with more than 500 years, which are still visible throughout the city. Madrid has a good touristic infrastructure with many interesting places to visit and fascinating things to see, such as the Royal Palace, the Royal Theatre, the Retiro Park, the National Library, and the Golden Triangle of Art: the Prado Museum, the Reina Sofia Museum, and the Thyssen-Bornemisza Museum.

ReCoSoC 2017 is organized by:
Centro de Electrónica Industrial
E.T.S. Ingenieros Industriales
Universidad Politécnica de Madrid (UPM)
José Gutiérrez Abascal, 2
28006, Madrid, Spain
http://www.cei.upm.es/

ReCoSoC 2017 is sponsored by:

Committees

General Chair:

Eduardo de la Torre
Universidad Politécnica de Madrid, Spain

Program Chairs:

Gilles Sassatelli
LIRMM, Montpellier, France

Thilo Pionteck
Universität Magdeburg, Germany

Andrés Otero
Universidad Politécnica de Madrid, Spain

Steering Committee:

Eduardo de la Torre (Universidad Politécnica de Madrid, Spain)
Alberto Garcia Ortiz (U. Bremen, Germany)
Manfred Glesner (TU Darmstadt, Germany)
Thomas Hollstein (Frankfurt Univ. of Applied Sciences, Germany)
Gert Jervan (Tallinn TU, Estonia)
Gilles Sassatelli (LIRMM-CNRS, France)
Leandro Soares Indrusiak (University of York, UK)
Lionel Torres (LIRMM-CNRS, France)

Program Committee:

Pascal Benoit (LIRMM, University of Montpellier 2)
François Berry (Université Clermont-Auvergne)
Lilian Bossuet (TELECOM Saint-Etienne, University Jean Monnet)
Florent Bruguier (LIRMM, CNRS)
Masoud Daneshtalab (KTH Royal Institute of Technology)
Jean-Luc Dekeyser (Inria)
Eduardo de La Torre (Universidad Politécnica de Madrid)
Christian De Schryver (University of Kaiserslautern)
Marco Domenico Santambrogio (Politecnico di Milano)
Peeter Ellervee (Tallinn University of Technology)
Abdoulaye Gamatie (LIRMM, CNRS)
Alberto Garcia-Ortiz (U. Bremen / TEM.IDS)
Diana Goehringer (Ruhr-University Bochum)
Guy Gogniat (Université de Bretagne Sud – UEB)
Daniel Gregorek (University of Bremen)
Jan Haase (Helmut-Schmidt-Universität der Bundeswehr)
Jim Harkin (University of Ulster)
Dominic Hillenbrand (KIT)
Michael Huebner (Ruhr-University Bochum)
Leandro Indrusiak (University of York)
Gert Jervan (Tallinn University of Technology)
Eduardo Juarez (Universidad Politecnica de Madrid)
Athanasios Kakarountas (University of Thessaly)
Chrysovalantis Kavousianos (University of Ioannina)
Loic Lagadec (ENSTA-Bretagne)
Sébastien Le Beux (Lyon Institute of Nanotechnology)
Vagelis Mariatos (TEI of Mesolonghi)
Fernando Moraes (PUCRS, Brasil)
Andrés Otero (Universidad Politécnica de Madrid)
Kyprianos Papadimitriou (Technical University of Crete, MHL)
Maxime Pelcat (IETR/INSA)
Sebastien Pillement (University of Nantes - IETR)
Thilo Pionteck (Universität Magdeburg)
Juha Plosila (University of Turku)

Thomas B. Preußer (Marie-Curie Fellow, Xilinx, Ireland)
Ruben Salvador (Universidad Politécnica de Madrid)
Gilles Sassatelli (LIRMM, Univ. Montpellier II / CNRS)
Tiberiu Seceleanu (ABB Corporate Research)
Johanna Sepulveda (Technical University of Munich)
Kostas Siozios (Aristotle University of Thessaloniki)
Dimitrios Soudris (NTUA)
Chao Wang (University of Science and Technology of China)
Peter Zipf (Uni Kassel)

Publication Chair:

Jorge Portilla (Universidad Politécnica de Madrid, Spain)

Local Organization Chair:

José M. Lanza-Gutiérrez (Universidad Politécnica de Madrid, Spain)

Finance Chair:

Teresa Riesgo (Universidad Politécnica de Madrid, Spain)

Technical Program

Wednesday, July 12th 2017

8:45 Registration

***9:15 - 9:30* Opening**

***9:30 – 10:30* Keynote talk** – *Extending Operating System Services over FPGAs.* Marco Platzner. Paderborn University.

10:30 – 11:00 Coffee break

11:00 – 12: 30 Session 1: Reconfigurable On-Chip Communication Architectures

> **Exploring the Performance of Partially Reconfigurable Point-to-Point Interconnects.** *El Mehdi Abdali, François Berry, Maxime Pelcat, Jean-Philippe Diguet and Francesca Palumbo.*
>
> **Fault-resilient NoC router with transparent resource allocation**. *Tsotne Putkaradze, Siavoosh Payandeh Azad, Behrad Niazmand, Jaan Raik and Gert Jervan.*
>
> **Adaptive and Reconfigurable Bubble Routing Technique for 2D Torus Interconnection Networks**. *Poona Bahrebar and Dirk Stroobandt.*

12:30 – 14:00 Lunch

***14:00 - 15:00* Keynote talk –** *"The rebirth of hardware: Open-source hardware initiatives for IoT- and MINT-based developments"*. Manfred Glesner. Technische Universität Darmstadt.

15:00 – 16:00 Session 2: Applications

Current Mode Detection in Hard Real-time Automotive Applications Dedicated to Many-Core Platforms. *Piotr Dziurzanski.*

Programmable SoC platform for Deep Packet Inspection using enhanced Boyer-Moore algorithm. *Adrian Dominguez, Pedro P. Carballo and Antonio Núñez.*

16:00 – 16:30 Coffee break

16:30 – 18:00 Session 3: Multi-core and Many-core Systems

ElasticSimMATE: a Fast and Accurate gem5 Trace-Driven Simulator for Multicore Systems. *Alejandro Nocua, Florent Bruguier, Abdoulaye Gamatié and Gilles Sassatelli.*

High-Level Test Generation for Processing Elements in Many-Core Systems. *Adeboye Stephen Oyeniran, Raimund Ubar, Siavoosh Payandeh Azad and Jaan Raik.*

Characterization and Optimization of Behavioral Hardware Accelerators in Heterogeneous MPSoC. *Yidi Liu, Monica Villaverde San Jose, Felix Moreno and Benjamin Carrion Schafer.*

19:00 – 21:00 **Welcome Cocktail**

Thursday, July 13th 2017

9:30 – 10:30 **Keynote talk** - *Space, the final frontier, the best testbed: Applications, Challenges and Needs for the application of reconfigurable SoCs to Space.* Ángel Álvaro. Thales Alenia Space

10:30 – 10:50 Coffee break

10:50 - 11:00 *"ReCoSoC in one year from now: insight on the 2018 edition"*

11:00 – 12:30 **Special Session 1: High Level Design Methodologies for Reconfigurable Computing and Adaptive Systems: tool flows and applications** (organized by Ruben Salvador, Jocelyn Sérot and Eduardo Juarez)

System-Level Design for Communication-Centric Task Farm Applications. *Daniela Genius and Ludovic Apvrille.*

Analysis of heterogeneous multi-core multi-HW accelerator based systems designed using PREESM and SDSoC. *Leonardo Suriano, Alfonso Rodriguez, Karol Desnos, Maxime Pelcat and Eduardo de La Torre.*

High-Level Design using Intel FPGA OpenCL: a Hyperspectral Imaging Spatial-Spectral Classifier. *Rubén Domingo, Rubén Salvador, Daniel Madroñal, Raquel Lazcano, Eduardo Juárez, Cesar Sanz, Himar Fabelo, Samuel Ortega and Gustavo M. Callicó.*

12:30 – 14:00 Lunch

14:00 -15:00 **Keynote talk** – *Networks-on-Chip for Real-Time and Mixed-Criticality Applications.* Leandro Indrusiak. University of York.

15:00 – 16:30 **Session 4: Reconfigurable and Self-Aware Systems-on-Chip**

Design and Scalability Analysis of Bandwidth-Compressed Stream Computing with Multiple FPGAs. *Antoniette Mondigo, Tomohiro Ueno, Daichi Tanaka, Kentaro Sano and Satoru Yamamoto.*

On-Demand Instantiation of Co-Processors on Dynamically Reconfigurable FPGAs. *Marcel Essig and Kurt Ackermann.*

Computational Self-Awareness as Design Approach for Visual Sensor Nodes. *Zakarya Guettatfi, Philipp Hübner, Marco Platzner and Bernhard Rinner.*

16:30 – 17:00 Coffee break

17:00 – 18:00 **Session 5: Emerging Technologies**

Design Method for Asymmetric 3D Interconnect Architectures with High-Level Models. *Jan Moritz Joseph, Lennart Bamberg, Sven Wrieden, Dominik Ermel, Alberto Garcia-Oritz and Thilo Pionteck.*

Energy Aware and Reliable STT-RAM based Cache Design for 3D Embedded Chip-Multiprocessors. *Fatemeh Arezoomand, Arghavan Asad, Mahdi Fazeli, Mahmood Fathy and Farah Mohammadi.*

19:00 Social Event and Gala Dinner

Friday, July 14th 2017

9:00 – 10:00 **Keynote talk –** *Safety, security and availability – How will this fly for autonomous cars?* Albrecht Mayer. Infineon Technologies

10:00 – 11:00 Special Session 2: Fault and Security Management in NoCs and MPSoCs (organized by Martha Johana Sepulveda and Thomas Hollstein)

Towards Trace-driven Cache Attacks by Exploiting Bus Communication on Systems-on-Chips. *Johanna Sepulveda, Mathieu Gross, Andreas Zankl and Georg Sigl.*

Fault Recovery and Adaptation in Time-Triggered Networks-on-Chips for Mixed-Criticality Systems. *Hamidreza Ahmadian, Farzad Nekouei and Roman Obermaisser.*

11:00 – 11:30 Coffee break

11:30 – 13:00 **Session 6: More Secure SoCs and NoCs**

Federated system-to-service authentication and authorization combining PUFs and tokens. *Marta Beltran, Miguel Calvo and Sergio González.*

Side-Channel Attack Resilience through Route Randomisation in Secure Real-Time Networks-on-Chip. *Leandro Indrusiak, James Harbin and Martha Johanna Sepulveda.*

SecBoot - Lightweight secure boot mechanism for Linux-based embedded systems on FPGA. *Peter Rouget, Benoît Badrignans, Pascal Benoit and Lionel Torres.*

13:00 – 13:15 **Closing Session**

13:15 – 14:45 Lunch

Paper Abstracts

Wednesday, July 12th 2017

Session 1: Reconfigurable On-Chip Communication Architectures

1. **Exploring the Performance of Partially Reconfigurable Point-to-Point Interconnects.** *El Mehdi Abdali, François Berry, Maxime Pelcat, Jean-Philippe Diguet and Francesca Palumbo.*

An ever larger share of FPGAs are supporting Dynamic and Partial Reconfiguration (DPR). A reconfigurable point-to-point interconnect (ρ-P2P) is a communication mechanism based on DPR that swaps between different precomputed configurations stored in partial bitstreams. ρ-Point-to-Point (P2P) is intended as a lightweight interconnect that suits the reconfigurable systems where a limited number of configurations are desirable. This paper assesses the pros and cons of ρ-P2P in terms of resource and performance depending on the number of input/output signals, their width and the number of supported configurations.

Experimental results, conducted on an Intel Cyclone V FPGA, compare ρ-P2P to an equivalently functional non-DPR solution called μ-P2P and to a full crossbar. They show that ρ-P2P is indeed lightweight but introduces performance limitations on operating frequency, memory footprint and reconfiguration time. However, ρ-P2P is in general the least resource intensive of the tested interconnects, except in the trivial case of low numbers of signals and configurations. In particular, an 18 × 18 full crossbar interconnect requires 75% more resources than an equivalent ρ- P2P. Interestingly, this resource difference between ρ-P2P and a full crossbar grows linearly with the interconnect size.

2. **Fault-resilient NoC router with transparent resource allocation**. *Tsotne Putkaradze, Siavoosh Payandeh Azad, Behrad Niazmand, Jaan Raik and Gert Jervan.*

The current trend of aggressive technology scaling results in a decrease in system's reliability. This motivates investigation of fault-resilient architectures which provide graceful degradation of system's functionality. In this paper, three novel fault-resilient Network-on-Chip (NoC) router architectures are proposed. These architectures, exploit the regularity of the router and reallocate available existing and spare units to maintain functionality of certain turns. The resource reallocation is performed transparently from system's resource manager and is based on predefined priorities. A new metric for architecture reliability comparison based on reliability block diagrams is introduced. In contrast to Silicone Protection Factor (SPF) metric, the proposed metric also takes into account the areas of different units. Area overhead and reliability of proposed architectures are compared with Triple Modular Redundancy (TMR) and Unit-Duplication mechanisms. All proposed architectures showed remarkable reliability improvement compared to original, TMR and Unit Duplication architectures; while at the same time, their area overhead is less than or equal to unit-duplication mechanisms.

3. **Adaptive and Reconfigurable Bubble Routing Technique for 2D Torus Interconnection Networks**. *Poona Bahrebar and Dirk Stroobandt.*

Networks with torus interconnection topology are widely used due to the symmetry in traffic distribution. In order to ensure deadlock-freedom and provide adaptive routing in torus, at least two Virtual Channels (VCs) per physical channel are required to break the cyclic channel dependencies. However, VCs increase the arbitration latency and consume large power/area overheads which is undesirable, particularly for on- chip networks with limited power/area budgets. In this paper, we propose a novel technique for routing in wormhole-switched 2D torus networks. The proposed method relies on the Abacus Turn Model (AbTM) and Worm-Bubble Flow Control (WBFC) to support adaptive and deadlock-free routing without using VCs. Furthermore, the network blocking is reduced by providing on-demand routing adaptiveness through reconfiguration. The experimental results demonstrate the efficiency of the proposed scheme in terms of performance and hardware overhead.

Session 2: Applications

1. **Current Mode Detection in Hard Real-time Automotive Applications Dedicated to Many-Core Platforms.** *Piotr Dziurzanski.*

This paper proposes a technique for determining the current mode in an electronic control unit (ECU) during run-time. We use a decision tree classifier which observes the latest execution times of processes (runnables). When a mode change is detected, the migration of runnables is performed to decrease the number of active cores leading to considerable energy savings while still not violating any of timing constraints. The proposed approach consists of both off-line and on-line steps, whereas more computational intensive steps are performed statically. In the presented automotive use case, the current mode is detected with 100% accuracy while observing execution time of a particular single runnable. The migration time of systems with dynamic mode detection based on the runnable execution time with various periods is also provided.

2. **Programmable SoC platform for Deep Packet Inspection using enhanced Boyer-Moore algorithm.** *Adrian Dominguez, Pedro P. Carballo and Antonio Núñez.*

This paper describes the work done to design a SoC platform for real-time on-line pattern search in TCP packets for Deep Packet Inspection (DPI) applications. The platform is based on a Xilinx Zynq programmable SoC and includes an accelerator that implements a pattern search engine that extends the original Boyer-Moore algorithm with timing and logical rules, that produces a very complex set of rules. Also, the platform implements different modes of operation, including SIMD and MISD parallelism, which can be configured on-line. The platform is scalable depending of the analysis requirement up to 8 Gbps. High-Level synthesis and platform based design methodologies have been used to reduce the time to market of the completed system.

Session 3: Multi-core and Many-core Systems

1. **ElasticSimMATE: a Fast and Accurate gem5 Trace-Driven Simulator for Multicore Systems.** *Alejandro Nocua, Florent Bruguier, Abdoulaye Gamatié and Gilles Sassatelli.*

Multicore system analysis requires efficient solutions for architectural parameter and scalability exploration. Long simulation time is the main drawback of current simulation approaches. In order to reduce the simulation time while keeping the accuracy levels, trace-driven simulation approaches

have been developed. However, existing approaches do not allow multicore exploration or do not capture the behavior of multi- threaded programs. Based on the gem5 simulator, we developed a novel synchronization mechanism for multicore analysis based on the trace collection of synchronization events, instruction and dependencies. It allows efficient architectural parameter and scalability exploration with acceptable simulation speed and accuracy.

2. **High-Level Test Generation for Processing Elements in Many-Core Systems.** *Adeboye Stephen Oyeniran, Raimund Ubar, Siavoosh Payandeh Azad and Jaan Raik.*

The advent of many-core system-on-chips (SoC) will involve new scalable hardware/software mechanisms that can efficiently utilize the abundance of interconnected processing elements found in these SoCs. These trends will have a great impact on the strategies for testing the systems and improving their reliability by exploiting system's re-configurability to achieve graceful degradation of system's performance. We propose a strategy of Software-Based Self-Test (SBST) to be used for testing of processing elements in many-core systems with the goal to increase fault coverage and structuring the test routines in a way which makes test-data delivery in many-core systems more efficient. A new high-level fault model is introduced, which covers a broad class of gate-level Stuck-at- Faults (SAF), conditional SAF, and bridging faults of any multiplicity in processor control paths. Two algorithms for high- level simulation-based test generation for the control path and a bit-wise pseudo-exhaustive test approach for data path are proposed. No implementation details are needed for test data generation. A novel method for proving the redundancy of high- level functional faults is presented, which allows for precise evaluation of fault coverage.

3. **Characterization and Optimization of Behavioral Hardware Accelerators in Heterogeneous MPSoC.** *Yidi Liu, Monica Villaverde San Jose, Felix Moreno and Benjamin Carrion Schafer.*

This work presents a method to characterize and optimize hardware accelerators (HWaccs) given as behavioral IPs (BIPs) mapped as loosely coupled HWaccs in heterogenous MPSoCs. The proposed HWacc exploration flow is composed of two main stages. The first stage characterizes each BIPs individually by performing a High-Level Synthesis (HLS) Design Space Exploration (DSE) on each of the BIPs to obtain a trade-off curve of Pareto-

optimal designs. It then continues by exploring the system-level design space using these Pareto-optimal designs and finding configurations with unique area vs. performance trade-offs. Our proposed system-level explorer makes use of cycle-accurate simulation models to explore the search space fast and accurately. Experimental results show that our proposed method works well for MPSoCs of different sizes ranging from systems with 1 to 4 masters and with 3 to 7 HWaccs.

Thursday, July 13th 2017

Special Session 1: High Level Design Methodologies for Reconfigurable Computing and Adaptive Systems: tool flows and applications

1. **System-Level Design for Communication-Centric Task Farm Applications.** *Daniela Genius and Ludovic Apvrille.*

Massively parallel applications such as telecommunication and video streaming have the par- ticularity that a large proportion of the time is spent on accessing communication channels between the tasks, due to contention on the on-chip interconnect. Moreover, the analysis of a given task deployment is often fastidious. Thus, we propose to extend an existing easy-to-use System-level Design methodology to task farm applications. The contribution first concerns adding relevant SysML modeling elements to take into account application code, hardware platforms and deployment constraints. Secondly, new modeling elements – including access techniques to communication channels – must be given a semantics in order to transform models into a well-defined SystemC virtual prototyping MPSoC platform. A telecommunication application serves as an example.

2. **Analysis of a heterogeneous multi-core multi-HW accelerator based systems designed using PREESM and SDSoC.** *Leonardo Suriano, Alfonso Rodriguez, Karol Desnos, Maxime Pelcat and Eduardo de La Torre.*

Nowadays, new heterogeneous system technologies are flooding the market: through the past years, it is possible to observe the move from single CPUs to

multi-core devices featuring CPUs, GPUs and large FPGAs, such as Xilinx Zynq- 7000 or Zynq UltraScale+ MPSoC architectures. In this context, providing developers with transparent deployment capabilities to efficiently execute different applications on such complex devices is important. In this paper, a design flow that combines, on one side, PREESM, a dataflow-based prototyping framework and, on the other side, Xilinx SDSoC, an HLS-based framework to automatically generate and manage hardware accelerators, is presented. This integration leverages the automatic, static task scheduling obtained from PREESM with asynchronous invocations that trigger the parallel execution of multiple hardware accelerators from some of their associated sequential software threads. An image processing application is used as a proof of concept, showing the interoperability possibilities of both tools, the level of design automation achieved and, for the resulting computing architecture, the good performance scalability according to the number of accelerators and sw threads.

3. **High-Level Design using Intel FPGA OpenCL: a Hyperspectral Imaging Spatial-Spectral Classifier.** *Rubén Domingo, Rubén Salvador, Daniel Madroñal, Raquel Lazcano, Eduardo Juárez, Cesar Sanz, Himar Fabelo, Samuel Ortega and Gustavo M. Callicó.*

Current computational demands require increasing designer's efficiency and system performance per watt. A broadly accepted solution for efficient accelerators implementation is reconfigurable computing. However, typical HDL methodologies require very specific skills and a considerable amount of de- signer's time. Despite the new approaches to high-level synthesis like OpenCL, given the large heterogeneity in today's devices (manycore, CPUs, GPUs, FPGAs), there is no one-fits-all solution, so to maximize performance, platform-driven optimization is needed. This paper reviews some latest works using Intel FPGA SDK for OpenCL and the strategies for optimization, evaluating the framework for the design of a hyperspectral image spatial-spectral classifier accelerator. Results are reported for a Cyclone V SoC using Intel FPGA OpenCL Offline Compiler 16.0 out- of-the-box. From a common baseline C implementation running on the embedded ARM® Cortex®-A9, OpenCL-based synthesis is evaluated applying different generic and vendor specific optimizations. Results show how reasonable speedups are obtained in a device with scarce computing and embedded memory resources. It seems a great step has been given to effectively raise the abstraction level, but still, a considerable amount of HW design skills is needed.

Session 4: Reconfigurable and Self-Aware Systems-on-Chip

1. **Design and Scalability Analysis of Bandwidth-Compressed Stream Computing with Multiple FPGAs.** *Antoniette Mondigo, Tomohiro Ueno, Daichi Tanaka, Kentaro Sano and Satoru Yamamoto.*

Stream computing in Field Programmable Gate Arrays (FPGAs) is seen as a promising solution in delivering the necessary performance and energy efficiency requirements of compute-intensive applications like numerical simulations. The inherent structure and customizability of FPGAs naturally make them the better alternative in achieving a highly-scalable computing design solution. This paper presents a scalable custom computing approach through temporal parallelism by increasing the depth of a computing pipeline in a 1D ring of cascaded FPGAs with high-speed, low-latency communication links. Spatial parallelism is also explored by replicating the computing core inside the FPGAs to further increase throughput. Due to communication bandwidth limitations, a hardware-based lossless bandwidth compression scheme was utilized in order to alleviate this bottleneck and transfer more data streams. A performance model is presented for the scalability analysis and performance estimation of this approach. For evaluation and verification, an actual numerical simulation was implemented on an Intel Arria 10 FPGA with spatially paralleled computing cores. Initial results show that the measured performance ratings are close to the predicted values using the performance model. Similarly, it was also demonstrated that the 1D ring topology of multiple FPGAs with bandwidth-compressed links can scale the performance when a sufficiently large data set is computed, even with a deeper pipeline and insufficient inter-FPGA bandwidth.

2. **On-Demand Instantiation of Co-Processors on Dynamically Reconfigurable FPGAs.** *Marcel Essig and Kurt Ackermann.*

State of the art FPGAs comprise various architectural features providing the performance and flexibility required to comply with growing real-time demands of today's industrial applications. Nevertheless, the requirements on engineering expertise in order to exploit these platform features significantly increased during the past few years, consequently raising product costs and the time-to-market as well. Especially the feature of dynamic partial

reconfiguration, enabling time- division multiplexing of resources within the reconfigurable fabric, is barely adopted by industry yet. This paper introduces a lightweight co-processing framework, taking advantage of an embedded processor closely coupled with the programmable logic inside the FPGA. The basic idea of this concept is to implement the sequential control flow of applications in software, while reconfigurable hardware accelerators may be utilized on- demand, in order to increase the performance on computation-intensive tasks. A hardware abstraction layer hides complex architectural processes and provides software engineers with a set of routines, enabling run-time requests and the interfacing of co-processors from within the code. Implementation details and sequences of operations are given and discussed.

3. **Computational Self-Awareness as Design Approach for Visual Sensor Nodes**. *Zakarya Guettatfi, Philipp Hübner, Marco Platzner and Bernhard Rinner.*

Visual sensor networks (VSNs) represent distributed embedded systems with tight constraints on sensing, processing, memory, communications and power consumption. VSNs are expected to scale up in the number of nodes, be required to offer more complex functionality, a higher degree of flexibility and increased autonomy. The engineering of such VSNs capable of (self-) adapting on the application and platform levels poses a formidable challenge.

In this paper, we introduce a novel design approach for visual sensor nodes which is founded on computational self-awareness. Computational self-awareness maintains knowledge about the system's state and environment with models and then uses this knowledge to reason about and adapt behaviours. We discuss the concept of computational self-awareness and present our novel design approach that is centred on a reference architecture for individual VSN nodes, but can be naturally extended to networks. We present the VSN node implementation with its platform architecture and resource adaptivity and report on preliminary implementation results of a Zynq-based VSN node prototype.

Session 5: Emerging Technologies

1. **Design Method for Asymmetric 3D Interconnect Architectures with High-Level Models.** *Jan Moritz Joseph, Lennart Bamberg, Sven Wrieden, Dominik Ermel, Alberto Garcia-Oritz and Thilo Pionteck.*

New 3D production methods enable heterogeneous integration of dies manufactured in different technology nodes. Asymmetric 3D interconnect architectures (A-3D-IAs) are the communication infrastructure targeting these heterogeneous 3D system on chips (3D SoCs), for which design methodologies and design tools are still missing. Here, a design method is proposed following an incremental approach enabled by high level models. Therefore, we present the first simulator and design framework covering the diverse requirements of A-3D-IAs. This includes an abstract model to estimate the application specific energy consumption of 2D metal wires and 3D through silicon vias (TSVs) in an A-3D-IA. It is validated by circuit simulations in combination with an electromagnetic field solver which is used for the extraction of the TSV array equivalent circuit. The model lays on a high abstraction level for fast simulations. Nonetheless, for real data stream scenarios it still shows a small maximum error of less than 8%. Additionally, a mathematical description is presented which enables a fast evaluation of low power coding schemes for A-3D-IA on a high level of abstraction.

2. **Energy Aware and Reliable STT-RAM based Cache Design for 3D Embedded Chip-Multiprocessors.** *Fatemeh Arezoomand, Arghavan Asad, Mahdi Fazeli, Mahmood Fathy and Farah Mohammadi.*

In Nano-scale technologies, static power consumption due to leakage current has become a serious issue in the design of SRAM based on-chip cache memories. To address this issue, non-volatile memory technologies such as STT-RAM (Spin Transfer Torque-RAM) have been proposed as a replacement for SRAM cells due to their near zero static power consumption and high memory density. Nonetheless, STT-RAMs suffer from some failures such as read disturb and limited endurance as well as high switching energy. One effective way to decrease the STT-RAMs' switching energy is to reduce their retention time, however, reducing the retention time has a negative impact on the reliability of STT-RAM cells. In this paper, we propose a hybrid cache layer for an embedded 3D- Chip Multiprocessor which employs two types of STT-RAM memory banks with retention time of 1s and 10ms to provide a beneficial

tradeoff between reliability, energy consumption, and performance. To this end, we also propose an optimization model to find the optimal configurations for these two kinds of memory banks. Simulation results using the Gem5 simulator through comparisons with fully SRAM and fully STT-RAM based cache show that the proposed hybrid cache consumes significantly less power while offering higher throughput (instructions per cycle) compared to a fully STT-RAM based cache.

Friday, July 14th 2017

Special Session 2: Fault and Security Management in NoCs and MPSoCs

1. **Towards Trace-driven Cache Attacks by Exploiting Bus Communication on Systems-on-Chips.** *Johanna Sepulveda, Mathieu Gross, Andreas Zankl and Georg Sigl.*

The growing complexity of Systems-on-Chips (SoCs) increases the risk of malware infections and trojans introduced at design time. A critical threat to system security are the so- called side-channel attacks based on the leakage due to the cache behavior during the execution of cryptographic algorithms. Recent publications have analyzed cache attacks on mobile devices and network-on-chip platforms. In this work, we investigate cache attacks on bus-like tile-based Multi-Processors Systems- on-Chips (MPSoCs). This work present two contributions. First, we demonstrate for the first time, a trace-driven cache attack on AES-128 based on the exploitation of the bus communication. Second, we integrate two countermeasures (Shuffling and Mini- table) and evaluate their impact on the attack and in the performance of the system. The results show that trace-driven attack are possible in SoC environments. Moreover, we show that the protection techniques are feasible and are able to mitigate the attack.

2. **Fault Recovery and Adaptation in Time-Triggered Networks-on-Chips for Mixed-Criticality Systems.** *Hamidreza Ahmadian, Farzad Nekouei and Roman Obermaisser.*

Adaptivity in terms of fault recovery and energy efficiency alongside with mixed-criticality support are demanded in today's embedded systems. Safety-critical system are desired to switch between precomputed resource allocations at runtime based on the monitored information from the platform. In addition, those systems are desired to adjust their internal behavior with regard to a change in the environment, while operating at a desired safety level. At the same time, resource requests in such systems can be highly dynamic and data dependent. Aiming at meeting a superset of all worst case demands leads to unaffordable overheads in terms of resource utilization. Hence, efficient resource management mechanisms are required to provide fault recovery and to make the system adaptive to the changes in the environmental or the resource requests, while keeping the system at a safe state. This paper introduces a solution for supporting resource management in networks-on-chips that fulfills the requirements of adaptive mixed-criticality systems and proposes an architecture that establishes fault recovery by switching between precomputed resource allocations based on the statistical and diagnostic information.

Session 6: More Secure SoCs and NoCs

1. **Federated system-to-service authentication and authorization combining PUFs and tokens.**
   Marta Beltran, Miguel Calvo and Sergio González.

Different application domains are challenging the still immature access control mechanisms currently used to authenticate and to authorize system-on-chip architectures to services deployed locally or in the cloud. These domains include Internet of Things, Smart Places or Industry 4.0 where different kinds of devices and objects, often poorly physically protected, low-cost and energy-constrained, interact with different kinds of services through lightweight communication protocols. These protocols usually guarantee basic data confidentiality and integrity, securing communication channels using cryptography, but there are still important challenges related to authentication and authorization. This work pro- poses a new system-to-service authentication and authorization mechanism based on the combination of a Physical Unclonable Function (PUF) and two tokens (one devoted to authentication and the other devoted to authorization), capable of working over HTTP or COAP relying on federated schemes and adapted to the specific requirements of this kind of environments. The new mechanism is validated and its efficiency and security are evaluated using a real healthcare case study.

2. **Side-Channel Attack Resilience through Route Randomisation in Secure Real-Time Networks-on-Chip.** *Leandro Indrusiak, James Harbin and Martha Johanna Sepulveda.*

Security can be seen as an optimisation objective in NoC resource management, and as such poses trade-offs against other objectives such as real-time schedulability. In this paper, we show how to increase NoC resilience against a concrete type of security attack, named side-channel attack, which exploit the correlation between specific non-functional properties (such as packet latencies and routes, in the case of NoCs) to infer the functional behaviour of secure applications. For instance, the transmission of a packet over a given link of the NoC may hint on a cache miss, which can be used by an attacker to guess specific parts of a secret cryptographic key, effectively weakening it.

We therefore propose packet route randomisation as a mech- anism to increase NoC resilience against side-channel attacks, focusing specifically on the potential impact of such an approach upon hard real-time systems, where schedulability is a vital de- sign requirement. Using an evolutionary optimisation approach, we show how to effectively apply route randomisation in such a way that it can increase NoC security while controlling its impact on hard real-time performance guarantees. Extensive experimental evidence based on analytical and simulation models supports our findings.

3. **SecBoot - Lightweight secure boot mechanism for Linux-based embedded systems on FPGA.** *Peter Rouget, Benoît Badrignans, Pascal Benoit and Lionel Torres.*

In recent years, the need in security for embedded devices and data centers has increased sharply. The possible consequences of attacks on these equipments make them privileged targets. In these fields, FPGA are increasingly used because of their flexibility and constantly decreasing power consumption and cost: they can embed several hard/soft processors running Linux enhancing system integration. This paper discusses the security issues related to operating system boot security on FPGAs. We show how the software early boot stages can be protected using FPGA built-in security mechanisms and user logic. We consider that external memories can be tampered by software attacks or board level attacks. By using open source elements and standard tools, we present and implement a lightweight solution. We show that the dynamic reconfiguration has nearly no impact on usable resources of the FPGA matrix at the end of the boot process.

Invited talks

Marco Platzner. Paderborn University

"Extending Operating System Services over FPGAs"

Abstract: *The idea of operating systems for FPGAs is around for some 20 years, and evolved from early concepts of hardware paging to more expressive software-centric abstractions that support partial reconfiguration and hardware/software migration. In this talk I will discuss the motivation and approaches for operating system integration of reconfigurable hardware, and then introduce to ReconOS, our open-source operating system for reconfigurable computers that offers a unified multi-threaded programming model for hardware and software threads. By semantically integrating hardware accelerators into a standard operating system environment, ReconOS supports a structured application development process, rapid design space exploration, and the creation of autonomous hybrid multi-core systems.*

Bio: Marco Platzner is Professor for Computer Engineering at Paderborn University. Previously, he held research positions at the Computer Engineering and Networks Lab at ETH Zurich, Switzerland, the Computer Systems Lab at Stanford University, USA, the GMD - Research Center for Information Technology (now Fraunhofer IAIS) in Sankt Augustin, Germany, and the Graz University of Technology, Austria. His research interests include reconfigurable computing, hardware-software codesign, and parallel architectures. Marco Platzner holds diploma and PhD degrees in Telematics (Graz University of Technology, 1991 and 1996), and a Habilitation degree for the area hardware-software codesign (ETH Zurich, 2002). Marco Platzner is member of the board of the Paderborn Center for Parallel Computing and the board of the Paderborn Institute of Advanced Studies in Computer Science and Engineering.

Manfred Glesner

"The rebirth of hardware: Open-source hardware initiatives for IoT- and MINT-based developments"

Abstract: *Open source software is a well-known topic. Open source hardware is a fairly new one: as we approach the limits of available electronic systems in Silicon, many specific designs in the past moved to generic standard chips. Best example are electronic components like ARDUINO and Raspberry-Pi. In this domain we see an increased offer of high performance boards from companies like INTEL. The offers are extended to sensor and actor systems. There is now an active Hacker-Community which is meeting in many places of the world and offering "hacks" in form of finished designs that can be openly exchanged. The whole development under the term DIY= Do it yourself has led to new offerings and companies (e.g. Adafruit et al) also for printed circuit board design in form of multi-project-realizations (OSH-park initiative). The talk is carefully reviewing the developments in this new area and gives an overview of typical killer applications in IT (= software defined radio), wearables, medical applications and Industry 4.0.*

Bio: Prof. em. Dr. Dr.h.c.mult. Manfred Glesner, IEEE Fellow (2000), is head of the microelectronic system research group at Technische Universitaet Darmstadt, Fachbereich Elektrotechnik und Informationstechnik. His research activities are in the areas of embedded systems design, high-level synthesis and physical design, especially for intelligent signal processing in mechatronics and smart systems in general. He was in 1982 the first in Germany to design with students a multiproject chip as part of the regular education. In recent time his focus was on the promotion of MOOC-based education in advanced technologies. He has organised many national and international conferences and is a member of the programme committees of several conferences and workshops. Together with

Gilles Sassatelli he is cofounder of the ReCoSoC-Conference Series starting in 2005 in Montpellier. He already successfully supervised 64 PhDs during his career. Prof. Glesner was founder and is one of the leading steering committee member of the European Workshop on Microelectronic Education. In 2007, he received the French Order "Chevalier dans l'Ordre des Palmes Académiques" for his collaborative work with multiple French research institutes. Recently, he received international awards from Estonia (Medal Terra Cross Mariana, 2014) and Mongolia (Kublai-Kahn-Medal, 2009) in recognition of his academic achievements with these countries. Prof. Glesner holds four Honorary Doctorships from the Technical University Tallinn (Estonia, 1996), Bukarest Polytechnic University (Romania, 1997), Mongolian Technical University Ulaanbaatar, (Mongolia, 2009) and the University of Liepaya (Latvia, 2014). Since 2013 Prof. Glesner is a member of the Advisory Board (=Hochschulrat) of the University of Applied Sciences in Dortmund. Further he is also member of the VDE-Rhein-Main advisory board which is actively promoting electrical engineering for the society in Germany. He is a consultant to many national, European and international organizations (IEEE; IFIP et al) and companies. In the last years he supported DAAD (Germany) to build up the GMIT = German Mongolian Institute for Resources and Technology in Nalaikh (Ulaanbaatar, Mongolia) where he was also a DAAD-guest professor in 2016. In recent times he devoted new activities to the promotion of Open-Source-Hardware activities, also for STEM/MINT-education.

Angel Álvaro Sánchez. R&D Manager. Thales Alenia Space España

"Space, the final frontier, the best testbed: Applications, Challenges and Needs for the application of reconfigurable SoCs to Space"

Abstract: *The space business is in turmoil, the rise of satellite constellations linked to the fierce evolution of the telecommunication market has changed the rules of a traditionally stable game. In the scientific domain, successes as Rosetta and Cassini-Huygens are making scientists grow bolder and ask for more ambitious missions. This means that satellites and planetary probes are evolving at an accelerated pace towards increased data processing needs, higher flexibility, higher system complexity, reduced cost and higher autonomy. This scenario is a natural application field for ReCoSoCs, that can provide the data processing capability with hardware and firmware flexibility at a reduced cost. But this will only be possible if we are able to make them compliant to the stringent constraints derived from the harsh space environment and strict reliability and dependability requirements, imposed by missions that fly there where no one can go to press a reset button or solder a last minute wire.*

This lecture will review the current and future needs of the space business in terms of data processing, assessing the applicability of the ReCoSoC devices for scientific, telecommunication or earth observation missions. The particularities of the space environment and the ways SoC reconfiguration and self-awareness can be used to overcome the challenges of radiation and mission reliability shall be presented, as well as the current trends on SoC architectures for space. Finally, the role of the space missions as technological demonstrators and testbeds will be discussed.

Bio: Head of R&D and Systems Project Manager: Telecommunications Engineer (1997) by UPM in Madrid, has more than 20 years of experience in the space sector. Starting at Alcatel Space as a designer of

digital equipment where he designed several FPGA based units for the Rosetta and Mars Express Missions. He has been responsible for the digital engineering group and director of operations. In 2009 returns to the technical activity and occupies a position within the R&D Management Area at Thales Alenia Space Spain. In his present position he coordinates the company R&D strategy including both internal R&D as well as the several H2020 projects where Thales Alenia Space Spain is involved.

Leandro Indrusiak

"Networks-on-Chip for Real-Time and Mixed-Criticality Applications"

Abstract: *Networks-on-Chip (NoC) provide packet-switching infrastructure for multiple types of system-wide communications, such as message passing between tasks running on different cores, data transfers between external memories and local scratchpads, or paging and coherency mechanisms for multi-level caches. In all cases, the performance of the NoC affects system timeliness and thus must be taken into account when analysing application-level real-time guarantees. This is specially true when application tasks and communication packets of different levels of criticality share the NoC infrastructure.*

This talk will present pros and cons of NoC architectures with priority-preemptive virtual channels, and will show their potential for handling real-time and mixed-criticality traffic. Then, it will review the state-of-the-art in real-time analysis for those NoC architectures, enabling safe upper bounds to end-to-end latency (tasks and packet flows) in single and mixed-criticality applications. Finally, it gives an insight on how real-time analysis can be used as a fitness function in the design space exploration of NoC-based embedded systems, aiming to meet performance guarantees and optimise energy dissipation.

Bio: Leandro Soares Indrusiak graduated in Electrical Engineering from the Federal University of Santa Maria (UFSM, Brazil) and obtained a MSc in Computer Science from the Federal University of Rio Grande do Sul (UFRGS, Brazil) in 1995 and 1998, respectively. He held a tenured assistant professorship at the Informatics department of the Catholic University of Rio Grande do Sul (Brazil) from 1998 to 2000. From 2001 to 2008 he worked as a researcher at the Technische Universitaet Darmstadt (Germany) where he worked towards a PhD and then lead a research team on the area of System-on-Chip design. His binational doctoral degree was jointly awarded by UFRGS and TU Darmstadt in 2003.

Since 2008, he is a permanent faculty member of University of York's Computer Science department (Lecturer 2008, Senior Lecturer 2013, Reader 2016), and a member of the Real-Time Systems (RTS) research group. His current research interests include on-chip multiprocessor systems, distributed embedded systems, resource allocation, cloud computing, and real-time networks, having published more than 120 peer-reviewed papers in the main international conferences and journals covering those topics (seven of them received best paper awards). He has graduated seven doctoral students, currently supervises three doctoral students and three post-doc research associates. He is a principal investigator of EU-funded SAFIRE project, and a co-investigator in a number of other funded projects. He serves as the department's Internationalisation coordinator, and has held visiting faculty positions in five different countries. He is a member of the EPSRC College, a member of the HiPEAC European Network of Excellence, and a senior member of the IEEE.

Albrecht Mayer

"Safety, security and availability – How will this fly for autonomous cars?"

Abstract: *Cost effective, safe and secure systems with high availability are the key building blocks for more and more automated driving. One architectural challenge is to structure such systems in a way that the complexity and design effort is manageable on all levels. Multi-core microcontrollers currently play a key role due to their dependability, safety and hard real-time behavior. This talk will look at the challenges, tradeoffs and some of the solutions for such system architectures.*

Bio: Albrecht Mayer is Senior Principal for Emulation Systems and Tooling at Infineon. In the past years he has been working on a multi-core microcontroller family architecture for a broad range of automotive applications, fulfilling highest requirements for safety, security, real-time behavior and availability under harsh conditions. He has many publications and holds more than 20 patents. Dr. Mayer received a Ph.D. degree in electrical engineering from the Technical University of Munich.

Social Program

Welcome cocktail
July 12th, 2017

The ReCoSoC 2017 Conference reception will be held on Wednesday July 12th, at Remigio (Paseo de la Castellana, 72, 28006 Madrid), a cool restaurant and bar where is possible to enjoy in one of the most exclusive areas of Madrid, located in Castellana Street. It is at two minutes walking distance from ReCoSoC Conference place.

Social event
July 13th, 2017

The ReCoSoC 2017 social event will take place on Thursday, July 13. We will start by visiting the Sorolla Museum (Paseo del General Martínez Campos, 37, 28010 Madrid).

The Sorolla Museum aims to promote maximum awareness and enjoyment of Joaquín Sorolla's legacy by as many people and types of public as possible in the conviction that this legacy, and particularly the works of Joaquín Sorolla, can provide a positive and enriching emotional and aesthetic experience, stimulate their senses, make them

aware of the role of the sense of sight as an instrument of knowledge and encourage them to develop their own creative abilities.

Next, the Conference dinner will take place al "La Máquina de Chamberí" (Calle de Ponzano, 39-41, 28003 Madrid), a fantastic restaurant in the heart of one of the most attractive areas for tapas and nightlife in Madrid.

Access to the Wireless Internet

We provide WiFi for you, being guest of the ReCoSoC. No additional software needs to be added to your notebook, just follow a few simple steps:

1. Connect to WIFI with SSID “InvitadosUPM”.
2. Open your favorite web browser and type this data:
   Usuario (user): reco.soc@invitadosupm
   Contraseña (password): recosoc17
3. Enjoy ;P

Notes:

ReCoSoC at a glance

Wednesday, July 12th		Thursday, July 13th		Friday, July 14th	
08:45 - 09:15	Registration	**09:30 - 10:30**	Keynote talk - Ángel Álvaro	**09:00 - 10:00**	Keynote talk - Albrecht Mayer
09:15 - 09:30	Opening	**10:30 - 10:50**	Coffee Break	**10:00 - 11:00**	Special Session 2
09:30 - 10:30	Keynote talk - Marco Platzner	**10:50 -11:00**	ReCoSoC in one year from now: insight on the 2018 edition	**11:00 - 11:30**	Coffee Break
10:30 - 11:00	Coffee Break	**11:00 - 12:30**	Special Session 1	**11:30 - 13:00**	Session 6
11:00 - 12:30	Session 1	**12:30 - 14:00**	Lunch	**13:00 - 13:15**	Closing Session
12:30 - 14:00	Lunch	**14:00 - 15:00**	Keynote talk - Leandro Indrusiak	**13:15 - 14:45**	Lunch
14:00 - 15:00	Keynote tallk - Manfred Glesner	**15:00 - 16:30**	Session 4		
15:00 - 16:00	Session 2	**16:30 - 17:00**	Coffee Break		
16:00 - 16:30	Coffee Break	**17:00 - 18:00**	Session 5		
16:30 - 18:00	Session 3	**19:00**	Social Event and Gala Dinner		
19:00 - 21:00	Welcome Cocktail				