

I B=J9FG=HM6 CCH< '5 H'8 5 H9 '&\$%#

The University Booth is organised during DATE and will be located in booth 1 of the exhibition area. All demonstrations will take place from Tuesday, March 28 to Thursday, March 30, 2017 during DATE. Universities and public research institutes have been invited to submit hardware or software demonstrators.

The University Booth programme is composed of 39 demonstrations from 13 different countries, presenting software and hardware solutions. The programme is organised in 11 sessions of 2 or 2.5 h duration and will cover the topics:

- Electronic Design Automation Prototypes
- Hardware Design and Test Prototypes
- Internet-of-Things Prototypes
- Wearable Electronics and Smart Medical Prototypes

The University Booth at DATE 2017 invites you to booth 1 to find out more about the latest trends in software and hardware from the international research community.

Several demonstrators will be shown more than once, giving visitors more flexibility to come to the booth and find out about the latest innovations.

We are sure that the demonstrators will give an attractive supplement to the DATE conference program and exhibition. We would like to thank all contributors to this programme.

More information is available online at <http://www.date-conference.com/exhibition/u-booth>. The University Booth programme is included in the conference booklet. The following demonstrators will be presented at the University Booth.

A FRAMEWORK FOR VARIATION-AWARE ANALOG CIRCUITS SIZING

Presenter:

Ons Lahiouel, Concordia University, CA

Authors:

Mohamed H. Zaki and Sofiene Tahar, Concordia University, CA

Abstract: *Today's analog design faces significant challenges due to circuit complexity and short time-to market windows. The proposed demonstration presents new techniques for enhancing variation-aware circuits sizing. The sizing problem is encoded using nonlinear constraints. A new algorithm using Satisfiability Modulo Theory (SMT) solving techniques exhaustively explores the design space and computes a continuous set of feasible sizing solutions. Two methods for the computation of parametric yield are implemented. The first method combines the advantages of sparse regression and SMT solving techniques for reliable and accelerated yield estimation. The second approach employs a statistical classifier to reduce the number of simulations. An optimization process using a two-step exploration strategy is also integrated to find the feasible design point with the highest yield. Experimental results show that our approach locates higher quality of design point within less run time.*

A TOOL FOR STATIC INSTRUCTION SET ARCHITECTURE ANALYSIS

Presenter:

Peer Adelt, Paderborn University / C-LAB, DE

Authors:

Bastian Koppelman¹, Wolfgang Mueller¹, Bernd Kleinjohann² and Christoph Scheytt¹

¹Heinz-Nixdorf Institute, DE; ²Paderborn University / C-LAB, DE

Abstract: *Mutation based testing, which is applied to assess testbenches by mutations of designs under test, is well established in the hardware and software domain. We demonstrate our tool, which analyses all 798 instructions of the TriCore™ microcontroller architecture for mutations of their binary representation. The tool provides an interactive graphical interface, which indicates general and detailed instruction set (ISA) specific statistics of 1-, 2- and 3-bitflip mutations in opcode, data and address sections of the instructions and their impact on the program execution. The ISA analysis tool is applied as a front-end for automatic mutation generations of TriCore™ binaries. In this context, we also show how the CPU emulator QEMU can be applied as a framework for mutation based analysis based on the static analysis of the TriCore™ ISA and the software binary.*

A VOLTAGE-SCALABLE FULLY DIGITAL ON-CHIP MEMORY FOR ULTRA-LOW-POWER IOT PROCESSORS

Presenter:

Jun Shiomi, Kyoto University, JP

Authors:

Tohru Ishihara and Hidetoshi Onodera, Kyoto University, JP

Abstract: A voltage-scalable RISC processor integrating standard-cell based memory (SCM) is demonstrated. Unlike conventional processors, the processor has Standard-Cell based Memories (SCMs) as an alternative to conventional SRAM macros, enabling it to operate at a 0.4 V single-supply voltage. The processor is implemented with the fully automated cell-based design, which leads to low design costs. By scaling the supply voltage and applying the back-gate biasing techniques, the power dissipation of the SCMs is less than 20 μ W, enabling the SCMs to operate with ambient energy source only. In this demonstration, the SCMs of the processor operates with a lemon battery as the ambient energy source.

ACCELERATORS: RECONFIGURABLE SELF-TIMED DATAFLOW ACCELERATOR & FAST NETWORK ANALYSIS IN SILICON

Presenter:

Alessandro de Gennaro, Newcastle University, GB

Authors:

Danil Sokolov and Andrey Mokhov, Newcastle University, GB

Abstract: Many real-life applications require dynamically reconfigurable pipelines to handle incoming data items differently depending on their values or current operating mode. A demo will show the benefits of an asynchronous accelerator for ordinal pattern encoding with reconfigurable pipeline depth. This was designed, simulated and verified using dataflow structure formalism in Workcraft toolset. The self-timed chip, fabricated in TSMC 90nm, shows high resilience to voltage variation and configurable accuracy of the results. Applications with underlying graph models foster the importance of a fast and flexible approach to graph analysis. To support medicine discovery biological systems are modelled by graphs, and drugs can disconnect some of the connections. A demo will show how graphs can be automatically converted into VHDL designs, which are synthesised into a FPGA for the analysis: thousand times faster than in software. Single stand will be used for both case studies.

AF3-MC: DEVELOPMENT OF MIXED CRITICALITY SYSTEMS USING MBSE

Presenter:

Thomas Boehm, fortiss, DE

Authors:

Johannes Eder and Sebastian Voss, fortiss, DE

Abstract: AutoFOCUS3 (<https://af3.fortiss.org/>) is an open-source model-based development tool, including a number of different analysis- and verification tools as well as design space exploration functionality, task scheduling dependent on a number of system requirements (timing, resource, energy, etc.), and code generators targeting C-code or VHDL. The presented demonstrator illustrates both a SW tool demonstrator and a corresponding HW demonstrator setup to show how a seamless model-based system approach could look like, w.r.t. to mixed-critical applications integrated on a (COTS) MC-platform. A floating ball can be controlled by a person by moving his hand over an US sensor, providing input to the control loop implemented in the high criticality part of the system. The low criticality part of the system which is running on the same CPU consists of the computation of the digits of PI and of the Fibonacci sequence, providing computationally intensive neighbors to the control loop.

BRAIN TO COMPUTER CONNECTIONS: A FAST TIME-DOMAIN APPROACH FOR BCI TRAINING

Presenter:

Giovanni Mezzina, Politecnico di Bari, Italy, IT

Authors:

Valerio Francesco Annese and Daniela De Venuto, Politecnico di Bari, IT

Abstract: We present a P300-based Brain Computer Interface (BCI) approach for the brain control of external devices through an innovative approach. The herein proposed HW/SW system acquires the signal from 6 EEG channels and synchronizes them with ad-hoc designed visual stimuli that evokes the P300 signal. The BCI signal processing comprises: (i) a Machine Learning stage, which is based on an algorithm (t-RIDE), which calibrates the system in ~ 190 s (ii) a smart approach for the time-domain features extraction greatly reduces the computational effort, speeding up the classification and finally (iii) the on-line classification, which is entrusted to a linear classifier. Noteworthy results obtained in experimental setup are: (i) P300 spatio-temporal characterization in 1.95s, (ii) classification accuracy of $80.5 \pm 4.1\%$ on single-trial. (iii) real time classification in 22ms (WC). As a PoC, supporting videos will show how the BCI outcomes can pilot a prototype car.

COSSIM: A NOVEL, COMPREHENSIBLE, ULTRA-FAST, SECURITY-AWARE CPS SIMULATOR

Presenter:

Nikolaos Tampouratzis, Technical University of Crete, GR

Authors:

Antonios Nikitakis and Andreas Brokalakis, Synelixis Solutions Ltd, GR

Abstract: *One of the main problems Cyber Physical Systems (CPS) and Highly Parallel Systems (HPS) designers face is the lack of simulation tools and models for system design and analysis. This is mainly because the majority of the existing simulation tools can handle efficiently only parts of a system (e.g. only the processing or only the network) while none of them supports the notion of security. Moreover, most of the existing simulators need extreme amounts of processing resources while faster approaches cannot provide the necessary precision and accuracy. COSSIM is an open-source framework that seamlessly simulates, in an integrated way, the networking and the processing parts of the CPS and Highly Parallel Heterogeneous Systems. In addition, COSSIM supports accurate power estimations while it is the first such tool supporting security as a feature of the design process. The complete COSSIM framework together with its sophisticated GUI will be presented.*

DEMONSTRATION OF HW/SW CO-PROCESSING WITH FPGA FOR FAST VISUAL NAVIGATION OF ROVERS

Presenter:

Konstantinos Maragos, National Technical University of Athens, GR

Authors:

George Lentaris and Dimitrios Soudris, National Technical University of Athens, GR

Abstract: *Autonomy, speed and accuracy constitute vital factors for the successful rover-exploration missions. However, the extremely low performance of the on-board space-grade CPUs in conjunction with the increased complexity of the sophisticated computer vision algorithms become a serious bottleneck for fast rover navigation. In this work, we present a HW/SW co-design solution based on FPGA to accelerate visual odometry algorithms tailored to the needs of future Mars exploration missions being scheduled by European Space Agency. For demonstration purposes, we use a Xilinx Kintex-7 FPGA to process images and perform feature detection, description, and matching. The FPGA communicates via ethernet port with the host CPU, which performs filtering and egomotion estimation with absolute orientation. We present the navigation path of a hypothetical moving rover which processes successively stereo images acquired by a hypothetical Martian surface while live-recording the CPU-FPGA co-processing.*

EMU: RAPID FPGA PROTOTYPING OF NETWORK SERVICES IN C#

Presenter:

Salvator Galea, University of Cambridge, GB

Authors:

Nik Sultana¹, Pietro Bressana², David Greaves¹, Robert Soulé², Andrew W Moore¹ and Noa Zilberman¹

¹University of Cambridge, GB; ²Università della Svizzera italiana, CH

Abstract: *General-purpose CPUs and OS abstractions impose overheads that make it challenging to implement network functions and services in software. On the other hand, programmable hardware such as FPGAs suffer from low-level programming models, which make the rapid development of network services cumbersome. We demonstrate Emu, a framework that makes use of an HLS tool (Kiwi) and enables the execution of high-level descriptions of network services, written in C#, on both x86 and Xilinx FPGA. Emu therefore opens up new opportunities for improved performance and power usage, and enables developers to more easily write network services and functions. We demonstrate C# implementations of network functions, such as Memcached and DNS Server, using Emu running on both x86 and NetFPGA-SUME platform and show that they are competitive to natively written hardware counterparts while providing a superior development and debug environment.*

FLEXPOR: FLEXIBLE PLATFORM FOR OBJECT RECOGNITION & TRACKING TO ENHANCE INDOOR LOCALIZATION AND MAPPING

Presenter:

Marko Rößler, Technische Universität Chemnitz, DE

Authors:

Christian Schott, Murali Padmanabha and Ulrich Heinkel, TU Chemnitz, DE

Abstract: *Object detection plays a crucial role in realizing intelligent indoor localization and mapping techniques. With the advantages of these techniques comes the complexity of computing hardware and the mobility. While the availability of open source computer vision algorithms and High-Level-Synthesis framework accelerates the development, the hybrid processing architecture of an All Programmable System on Chip (APSoC) enables efficient hardware-software partitioning. Using these tools, a generic platform was designed for evaluating the computer vision algorithms. Open source components such as Linux kernel and OpenCV libraries were integrated for evaluation of the algorithms on the software while Vivado HLS framework was used to synthesize the hardware counter parts. Algorithms such as Sobel filtering and Hough Line transformation were implemented and analyzed. The capabilities of this platform were used to realize a mobile object detection system for enhancing the localization techniques.*

GNOCS: AN ULTRA-FAST, HIGHLY EXTENSIBLE, CYCLE-ACCURATE GPU-BASED PARALLEL NETWORK-ON-CHIP SIMULATOR

Presenter:

Amir CHARIF, TIMA, FR

Authors:

Nacer-Eddine Zergainoh and Michael Nicolaidis, TIMA, FR

Abstract: *With the continuous decrease in feature sizes and the recent emergence of 3D stacking, chips comprising thousands of nodes are becoming increasingly relevant, and state-of-the-art NoC simulators are unable to simulate such a high number of nodes in reasonable times. In this demo, we showcase GNoCS, the first detailed, modular and scalable parallel NoC simulator running fully on GPU (Graphics Processing Unit). Based on a unique design specifically tailored for GPU parallelism, GNoCS is able to achieve unprecedented speedups with no loss of accuracy. To enable quick and easy validation of novel ideas, the programming model was designed with high extensibility in mind. Currently, GNoCS accurately models a VC-based microarchitecture. It supports 2D and 3D mesh topologies with full or partial vertical connections. A variety of routing algorithms and synthetic traffic patterns, as well as dependency-driven trace-based simulation (Netrace), are implemented and will be demonstrated*

GREENOPENHEVC: LOW POWER HEVC DECODER

Presenter:

Menard Daniel, INSA Rennes, FR

Authors:

Julien Heulot¹, Erwan Nogues¹, Maxime Pelcat² and Wassim Hamidouche¹

¹INSA Rennes, IETR, UBL, FR; ²Institut Pascal, Université Clermont-Ferrand, FR

Abstract: *Video on mobile devices is a must-have feature with the prominence of new services and applications using video like streaming or conferencing. The new video standard HEVC is an appealing technology for service providers. Besides, with the recent progress of SoC, software video decoders are now a reality. The challenge is to provide power efficient design to fit with the compelling demand for long battery. We present here a practical set-up demonstrating that the new HEVC standard can be implemented in software on an embedded GPP multicore platform. Different techniques have been integrated to optimize the energy: data-level and thread level parallelisms, video aware Dynamic Voltage and Frequency Scaling. To push back the limits, algorithm level approximate computing is carried-out on the in-loop filtering. The subjective tests have demonstrated that the quality degradation is almost imperceptible. A mean power of less than 1 Watt is reported for a HD 1080p/24fps video decoding.*

HEPSYCODE: A SYSTEM-LEVEL METHODOLOGY FOR HW/SW CO-DESIGN OF HETEROGENEOUS PARALLEL DEDICATED SYSTEMS

Presenter:

Luigi Pomante, University of L'Aquila, IT

Authors:

Giacomo Valente¹, Vittoriano Muttillio¹, Daniele Di Pompeo¹, Emilio Incerto² and Daniele Ciambrone¹

¹University of L'Aquila, IT; ²Gran Sasso Science Institute, IT

Abstract: *Heterogeneous parallel systems have been recently exploited for a wide range of application domains, for both the dedicated (e.g. embedded) and the general purpose products. Such systems can include different processor cores, memories, dedicated ICs and a set of connections between them. They are so complex that the design methodology plays a major role in determining the success of the products. So, this demo addresses the problem of the electronic system-level hw/sw co-design of heterogeneous parallel dedicated systems. In particular, it shows an enhanced CSP/SystemC-based design space exploration step (and related ESL-EDA prototype tools), in the context of an existing hw/sw co-design flow that, given the system specification and related F/NF requirements, is able to (semi)automatically propose to the designer: - a custom heterogeneous parallel architecture; - an HW/SW partitioning of the application; - a mapping of the partitioned entities onto the proposed architecture.*

ITMD: RUN-TIME MANAGEMENT OF CONCURRENT MULTITHREADED APPLICATIONS ON HETEROGENEOUS MULTI-CORES

Presenter:

Karunakar Reddy Basireddy, University of Southampton, GB

Authors:

Amit Singh, Bashir M. Al-Hashimi and Geoff V. Merrett, University of Southampton, GB

Abstract: *Heterogeneous multi-cores often need to deal with multiple applications having different performance requirements concurrently, which generate varying and mixed workloads. Runtime management is required for adapting to such performance requirements and workload variabilities, and to achieve energy efficiency. It is challenging to efficiently exploit different types of cores simultaneously and DVFS potential of cores. We present a run-time management approach that first selects thread-to-core mapping based on the performance requirements and resource availability. Then, it applies online adaptation by adjusting the voltage-frequency (V-f) levels to achieve energy optimization. We demonstrate the proposed run-time management approach on the Odroid-XU3, with various combinations of multi-threaded applications from PARSEC and SPLASH benchmarks. Results show an average improvement in energy efficiency up to 33% compared to existing approaches.*

LABSMILING: A FRAMEWORK, COMPOSED OF A REMOTELY ACCESSIBLE TESTBED AND RELATED SW TOOLS, FOR ANALYSIS AND DESIGN OF LOW DATA-RATE WIRELESS PERSONAL AREA NETWORKS BASED ON IEEE 802.15.4

Presenter:

Marco Santic, University of L'Aquila, IT

Authors:

Luigi Pomante, Walter Tiberti, Carlo Centofanti and Lorenzo Di Giuseppe, DEWS - Università di L'Aquila, IT

Abstract: *Low data-rate wireless personal area networks (LR-WPANS) are even more present in the fields of IoT, wearable devices and health monitoring. The development, deployment and test of such systems, based on IEEE 802.15.4 standard (and its derivations, e.g. 15.4e), require the exploitation of a testbed when the network is not trivial and grows in complexity. This demo shows the framework of LabSmiling: a testbed and related SW tools that connect a meaningful (but still scalable) number of physical devices (sensor nodes) located in a real environment. It offers the following services: program, reset, switch on/off single devices; connect to devices up/down links to inject or receive commands/messages/packets in/from the network; set devices as low level packet sniffers, allowing to test/debug protocol compliances or extensions. Advanced services are: possibility of design test scenarios for the evaluation of network metrics (throughput, latencies, etc.) and custom application verification.*

MARGOT: APPLICATION ADAPTATION THROUGH RUNTIME AUTOTUNING

Presenter:

Gianluca Palermo, Politecnico di Milano, IT

Authors:

Davide Gadioli, Emanuele Vitali and Cristina Silvano, Politecnico di Milano, IT

Abstract: *Several classes of applications expose parameters that influence their extra-functional properties, such as the quality of the result or the performance. This leads the application designer to tune these parameters to find the configuration that produces the desired outcome. Given that the application requirements and the resources assigned to each application might vary at runtime, finding a one-fit-all configuration is not a trivial task. For this reason, we implemented the mARGOT framework that enhances an application with an adaptation layer in order to continuously tune the parameters according to the evolving situation. More in detail, mARGOT is composed of a monitoring infrastructure, an application-level adaptation engine and an extra-functional configuration framework based on the separation of concerns paradigm between functional and extra-functional aspects. At the booth, we plan to demonstrate the effectiveness of the proposed infrastructure on three real-life applications.*

MATISSE: A TARGET-AWARE COMPILER TO TRANSLATE MATLAB INTO C AND OPENCL

Presenter:

Luís Reis, University of Porto, PT

Authors:

João Bispo and João Cardoso, University of Porto / INESC-TEC, PT

Abstract: *Many engineering, scientific and finance algorithms are prototyped and validated in array languages, such as MATLAB, before being converted to other languages such as C for use in production. As such, there has been substantial effort to develop compilers to perform this translation automatically. Alternative types of computation devices, such as GPGPUs and FPGAs, are becoming increasingly more popular, so it becomes critical to develop compilers that target these architectures. We have adapted MATISSE, our MATLAB-compatible compiler framework, to generate C and OpenCL code for these platforms. In this demonstration, we will show how our compiler works and what its capabilities are. We will also describe the main challenges of efficient code generation from MATLAB and how to overcome them.*

MTA: MANCHESTER THERMAL ANALYZER

Presenter:

Scott Ladenheim, University of Manchester, GB

Authors:

Yi-Chung Chen, Vasilis Pavlidis and Milan Mihajlović, University of Manchester, GB

Abstract: *The Manchester Thermal Analyzer (MTA) is a fast thermal analysis tool to compute temperature profiles of integrated circuits (ICs) in 3-D. The thermal simulations use the finite element method to discretize the heat equation in space coupled to an implicit time-integration method and are implemented with the open-source C++ library deal.II. The MTA supports higher-order elements, several time-integration methods, and fully adaptive spatiotemporal refinement. State-of-the-art preconditioned iterative methods solve the linear systems arising from the discretized equations as efficiently as possible. Using shared memory parallelization, the MTA solves systems on the order of tens of millions enabling modeling ICs at the cell-level. We present a thermal simulation of an Intel Xeon processor within a FCLGA package with heatsink to show the diverse structures of modern ICs the MTA simulates. The MTA also models other 3-D structures such as bonded tiers, TSVs, heatsinks, and heat spreaders.*

MULTI-CORE VERIFICATION: COMBINING MICROTESK AND SPIN FOR VERIFICATION OF MULTI-CORE MICROPROCESSORS

Presenter:

Mikhail Chupilko, ISPRAS, RU

Authors:

Alexander Kamkin, Mikhail Lebedev and Andrei Tatarnikov, ISPRAS, RU

Abstract: *The complexity of modern cache coherence protocols (CCP) in multi-core microprocessors prevents from complete verification of shared memory subsystems by means of random test-program generators (TPG). The following steps are suggested to target the problem. The first step is to separately specify CCP features and generate CCP-specific events to be used in TPG when generating a test program (TP). The protocol is specified in Promela, with Spin making a test template (TT). Spin also produces UVM (or C++TESK) testbench to make the execution of the resulting TPs to be controllable and deterministic. The second step is to let TPG produce the memory access instructions causing desired CCP-specific behavior. As a TPG we use MicroTESK. Its Ruby-based TTs abstractly describe future TPs. MicroTESK processes that TT making TP with CCP-specific events. The resulting TP is executed together with the testbench to exactly reproduce the situation Spin had found to be important for such a protocol.*

NETFI-2: AN AUTOMATIC METHOD FOR FAULT INJECTION ON HDL-BASED DESIGNS

Presenter:

Alexandre Coelho, Université Grenoble Alpe, FR

Authors:

Miguel Solinas, Juan Fraire, Nacer-Eddine Zergainoh, Pablo Ferreyra and Raoul Velazco, TIMA, FR

Abstract: *Fault injection tools, which include fault simulation and emulation, are a well-known technique to evaluate the susceptibility of integrated circuits to the effects of radiation. This work presents a methodology to emulate Single Event Upsets (SEUs) and Single Event Transients (SETs) in a Field Programmable Gate Array (FPGA). The method proposed combines the flexibility of FPGA with the controllability provided by the MicroBlaze, to emulate HDL circuit and control the fault injection campaign. This approach has been integrated into a fault-injection platform, named NETFI (NETlist Fault Injection), developed by our research group, and received the name of NETFI-2. To validate this methodology fault injection campaigns have been performed in Leon3 and Stochastic Bayesian Machine. Results on an Artix-7 FPGA show that NETFI-2 provides accurate measurements while improving the execution time of the experiment by more than 300% compared with analogous simulation-based campaigns.*

NETWORKED LABS-ON-CHIPS

Presenter:

Andreas Grimmer, Johannes Kepler University Linz, AT

Authors:

Werner Haselmayr, Andreas Springer and Robert Wille, Johannes Kepler University Linz, AT

Abstract: *Labs-on-Chip (LoC) allow for the miniaturization, integration, and automation of medical and bio-chemical procedures. In recent years, different technologies have been considered. However, all of them have their drawbacks, e.g. electrowetting-based LoCs suffer from the evaporation of liquids, the fast degradation of the surface coatings, and the inferior biocompatibility, while flow-based LoCs require a complex and costly multilayer fabrication process. Hence, an alternative has recently been proposed in terms of Networked Labs-on-Chips. We present and demonstrate the NLoC technology where so-called droplets flow inside channels of micrometer-size. Networking functionalities enable the designer to dynamically select the operations to be conducted. These networking functionalities exploit hydrodynamic forces acting on droplets. Moreover, NLoC devices can be produced at low cost (e.g. using 3D printers). By this, drawbacks of established LoC-technologies are addressed.*

NNN: NEURAL NETWORKS DESIGNING NEURAL NETWORKS

Presenter:

Brett Meyer, McGill University, US

Authors:

Warren Gross, Sean Smithson, Ossama Ahmed and Guang Yang, McGill University, US

Abstract: *Modern artificial neural networks currently achieve state-of-the-art results in various difficult problems, including image classification and speech recognition. However, both the performance and computational complexity of such models are heavily dependent on the design of characteristic hyper-parameters (e.g., numbers of hidden layers or nodes per layer) which are often manually optimized. With neural networks penetrating low-power mobile and embedded areas, the need now arises to optimize not only for performance, but also for implementation cost. In our work, we present a multi-objective design space exploration method leveraging machine learning based response surface modelling to reduce the number of solutions trained and evaluated. Experimental results are presented for several image recognition datasets, demonstrating the evolution of the approximated Pareto-optimal hyper-parameters and corresponding GPU code; all while exploring only a small fraction of the design space.*

NOXIM-XT: A BIT-ACCURATE POWER ESTIMATION SIMULATOR FOR NOCS

Presenter:

Pierre Bomel, Université de Bretagne Sud, FR

Authors:

André Rossi¹, Johann Laurent² and Erwan Moreac²

¹LERIA, Université d'Angers, Angers, France, FR; ²Lab-STICC, Université de Bretagne Sud, Lorient, FR

Abstract: We have developed an enhanced version of Noxim (Noxim-XT) to estimate the energy consumption of a NoC in a SoC. Noxim-XT is used in a two-step methodology. First, applications are mapped on a SoC and their traffics are extracted by simulation with MPSocBench. Second, Noxim-XT tests various hardware configurations of the NoC, and for each configuration, the application's traffic is re-injected and replayed, an accurate performance and power breakdown is provided, and the user can choose different data coding strategies. With the help of Noxim XT, each configuration is bit-accurately estimated in terms of energy consumption. After simulation, a spatial mapping of the energy consumption is provided and highlights the hot-spots. Moreover, the new coding strategies allows significant energy saving. Noxim XT simulations and a FPGA-based prototype of a new coding strategy will be demonstrated at the U-booth to illustrate these works.

OPENCTMOD: AN OPEN SOURCE COLLABORATIVE MATLAB TOOLBOX FOR THE DESIGN AND SIMULATION OF CONTINUOUS-TIME SIGMA DELTA MODULATORS

Presenter:

Dang-Kiên Germain Pham, LTCI, Télécom ParisTech, Université Paris-Saclay, 75013, Paris, France, FR

Author:

Chadi Jabbour, LTCI, Télécom ParisTech, Université Paris-Saclay, FR

Abstract: Simulating Continuous Time (CT) Sigma Delta Modulators (SDM) is commonly done using block level systems such as Simulink which is a highly time consuming task even at system level. Therefore, the existing design tools for SDM are either discrete time oriented (Schreier toolbox) or proprietary (Ulm toolbox). In this work, we propose a new Matlab/C toolbox for the design of CT SDM. Simulation is based on state space representation thereby allowing to support most of the existing SDM architectures. Moreover, the main non-idealities of the main blocks are modeled (opamp DC gain, finite GBW, DACs mismatch, ISI and quantizer offset). Besides, thanks to the modular and open source approach for this toolbox, every user can easily implement additional features and include it. During the forum, designs and simulations for various architectures of CT SDM will be performed to demonstrate the accuracy and efficiency of the proposed toolbox. The collaborative aspect will be also shown.

PER: METHOD AND TOOL FOR ANALYZING THE INTERPLAY BETWEEN PERFORMANCE, ENERGY AND SCALING IN MULTI- AND MANY-CORE PLATFORMS

Presenter:

Fei Xia, Newcastle University, GB

Authors:

Ashur Rafiev, Alexander Romanovsky and Alex Yakovlev, Newcastle University, GB

Abstract: Parallelization has been used to maintain a reasonable balance between energy consumption and performance in computing systems. However, the effectiveness of parallelization scaling is different for different hardware platforms. This is because the reliable operation region (ROR), a region defined in the voltage-throughput space for any hardware platform, is platform-dependent and its shape determines how effective parallelization scaling is in improving throughput and/or reducing power consumption. Although many of the interlinked issues are known, a unifying analysis method has just now been proposed to study the interplay between performance, energy, reliability and parallelization scaling. The method of bi-normalization of the ROR is designed to help achieve a meaningful cross-platform analysis of this interplay. The PER tool brings all these issues together and helps designers reason about hardware parallelization, DVFS and software parallelizability.

PULP: A ULTRA-LOW POWER PLATFORM FOR THE INTERNET-OF-THINGS

Presenter:

Francesco Conti, ETH Zurich, CH

Authors:

Stefan Mach¹, Florian Zaruba¹, Antonio Pullini¹, Daniele Palossi¹, Giovanni Rovere¹, Florian Glaser¹, Germain Haugou¹, Schekeb Fateh¹ and Luca Benini²

¹ETH Zurich, CH; ²ETH Zurich, CH and University of Bologna, IT

Abstract: The PULP (Parallel Ultra-Low Power) platform strives to provide high performance for IoT nodes and endpoints within a very small power envelope. The PULP platform is based on a tightly-coupled multi-core cluster and on a modular architecture, which can support complex configurations with autonomous I/O without SW intervention, HW-accelerated execution of hot computation kernels, fine-grain event-based computation - but can also be deployed in very simple configuration, such as the open source PULPino microcontroller. In this demonstration booth, we will showcase several prototypes using PULP chips in various configuration. Our prototypes perform demos such as real-time deep-learning based visual recognition from a low-power camera, and online biosignal acquisition and reconstruction on the same chip. Application scenarios for our technology include healthcare wearables, autonomous nano-UAVs, smart networked environmental sensors.

RIMEDIO: WHEELCHAIR MOUNTED ROBOTIC ARM DEMONSTRATOR FOR PEOPLE WITH MOTOR SKILLS IMPAIRMENTS

Presenter:

Alessandro Palla, University of Pisa, IT

Authors:

Gabriele Meoni and Luca Fanucci, University of Pisa, IT

Abstract: People with reduced mobility experience many issues in the interaction with the indoor and outdoor environment because of their disability. For those users even the simplest action might be a hard/impossible task to perform without the assistance of an external aid. We propose a simple and lightweight wheelchair mounted robotic arm with the focus on the human-machine interface that has to be simple and accessible for users with different kind of disabilities. The robotic arm is equipped with a 5 MP camera, force and proximity sensors and a 6 axis Inertial Measurement Unit on the end-effector that can be controlled using an app running on a tablet. When the user selects the object to reach (for instance a button) on the tablet screen, the arm autonomously carries out the task, using the camera image and the sensors measurements for autonomous navigation. The demonstrator consists in the robotic arm prototype, the Android tablet and a personal computer for arm setup and configuration.

RUNNING CONVOLUTIONAL LAYERS OF ALEXNET IN NEUROMORPHIC COMPUTING SYSTEM

Presenter:

Yongshin Kang, Incheon National University, KR

Authors:

Seban Kim, Taehwan Shin and Jaeyong Chung, Incheon National University, KR

Abstract: Neuromorphic hardware has drawn attention as an approach to deal with the issues of today's computing platforms based on Von Neumann architecture when running deep learning models, but large-scale deep neural networks such as AlexNet have not been demonstrated yet in any neuromorphic systems. Since 2014, we have been developing a non-Von Neumann computing system called INSight based on data flow architecture that aims at running large-scale deep neural networks in the neuromorphic fashion. We have now reached a major milestone and will demonstrate INSight running the convolutional layers of AlexNet. The proposed system is implemented with Xilinx Virtex 7 FPGA and performs the processing using 100K synapses mapped on LUTs without any array-type memories. It processes 1552 images per second and consumes 7.2W, resulting in the state-of-the-art energy efficiency.

SCCHARTS: SYNCHRONOUS STATECHARTS FOR SAFETY-CRITICAL APPLICATIONS

Presenter:

Reinhard von Hanxleden, Kiel University, DE

Authors:

Michael Mendler¹, Christian Motika², Christoph Daniel Schulze² and Steven Smyth²

¹Bamberg University, DE; ²Kiel University, DE

Abstract: We present a visual language, SCCharts, designed for specifying safety-critical reactive systems. SCCharts use a statechart notation and provide determinate concurrency based on a synchronous model of computation (MoC), without restrictions common to previous synchronous MoCs. Specifically, we lift earlier limitations on sequential accesses to shared variables, by leveraging the sequentially constructive MoC. For further details, see [von Hanxleden et al., PLDI'14] and <http://www.sccharts.com>. The SCCharts demonstrator is an Eclipse Richt Client and part of KIELER (<http://www.rtsys.informatik.uni-kiel.de/en/research/kieler>). The demonstration shows how to write an SCChart model using a textual notation, from which a visual model is generated on the fly using the Eclipse Layout Kernel (ELK). We also present a compilation chain that allows efficient synthesis of software and hardware.

SEFILE: A SECURE FILESYSTEM IN USERSPACE VIA SECUBE™

Presenter:

Giuseppe Airofarulla, CINI, IT

Authors:

Paolo Prinetto¹ and Antonio Varriale²

¹CINI & Politecnico di Torino, IT; ²Blu5 Labs Ltd., IT

Abstract: The SEcube™ Open Source platform is a combination of three main cores in a single-chip design. Low-power ARM Cortex-M4 processor, a flexible and fast Field-Programmable-Gate-Array (FPGA), and an EAL5+ certified Security Controller (SmartCard) are embedded in an extremely compact package. This makes it a unique Open Source security environment where each function can be optimized, executed, and verified on its proper hardware device. In this demo, we present a Windows wrapper for a Filesystem in Userspace (FUSE) with an HDD firewall resorting to the hardware built-in capabilities, and the software libraries, of the SEcube™.

SELINK: SECURING HTTP AND HTTPS-BASED COMMUNICATION VIA SECUBE™

Presenter:

Airofarulla Giuseppe, CINI & Politecnico di Torino, IT

Authors:

Paolo Prinetto¹ and Antonio Varriale²

¹Politecnico di Torino, IT; ²Blu5 Labs Ltd., IT

Abstract: *The SEcube™ Open Source platform is a combination of three main cores in a single-chip design. Low-power ARM Cortex-M4 processor, a flexible and fast Field-Programmable-Gate-Array (FPGA), and an EAL5+ certified Security Controller (SmartCard) are embedded in an extremely compact package. This makes it a unique Open Source security environment where each function can be optimized, executed, and verified on its proper hardware device. In this demo, we present a client-server HTTP and HTTPS-based application, for which the traffic is encrypted resorting to the hardware built-in capabilities, and the software libraries, of the SEcube™. By doing so, we show how communication can be secured from an attacker capable of inspecting, and tampering, the regular communication.*

STACKADROP: A MODULAR DIGITAL MICROFLUIDIC BIOCHIP RESEARCH PLATFORM

Presenter:

Oliver Keszöcze, University of Bremen, DE

Authors:

Maximilian Luenert and Rolf Drechsler, University of Bremen & DFKI GmbH, DE

Abstract: *Advances in microfluidic technologies have led to the emergence of Digital Microfluidic Biochips (DMFBs), which are capable of automating laboratory procedures. These DMFBs raised significant attention in industry and academia creating a demand for devices. Commercial products are available but come at a high price. So far, there are two open hardware DMFBs available: the DropBot from WheelerLabs and the OpenDrop from GaudiLabs. The aim of the StackADrop was to create a DMFB with many directly addressable cells while still being very compact. The StackADrop strives to provide means to experiment with different hardware setups. It's main feature are the exchangeable top plates, supporting 256 high-voltage pins. It features SPI, UART and I2C connectors for attaching sensors/actuators and can be connected to a computer using USB for interactive sessions using a control software. The modularity allows to easily test different cell shapes, such as squares, hexagons and triangles.*

TFA: TRANSPARENT CODE OFFLOADING ON FPGA

Presenter:

Roberto Rigamonti, HEIG-VD/HES-SO, CH

Authors:

Anthony Convers, Baptiste Delporte, Xavier Ruppen and Alberto Dassatti, HEIG-VD/HES-SO, CH

Abstract: *Genomics, molecular dynamics, and machine learning are just the most recent examples of fields where FPGAs could provide the means to achieve interesting breakthroughs. However, HDL programming requires considerable multi-disciplinary skills, experience, large budgets, time, and a bit of wizardry. Given that most implementations are short-lived, the investment simply does not pay off. In this demo we propose a multi-vendor LLVM-based automated framework that can transparently - without the user or developer being aware of it - offload computing-intensive code fragments to FPGAs. The system relies on a performance monitor to detect computing-intensive code sections and, if they are suitable for offloading, extracts the Data Flow Graph and uses it to program an overlay pre-programmed on the FPGA, which then interacts with the Just-In-Time compiler executing the program. The overall process requires hundreds of microseconds, and can be easily reverted should the outcome be unsatisfactory.*

TGV: TESTER GENERIC AND VERSATILE FOR RADIATION EFFECTS ON ADVANCED VLSI CIRCUITS

Presenter:

Miguel Solinas, TIMA, FR

Authors:

Alexandre Coelho Coelho, Juan Fraire, Nacer Eddine Zergainoh and Raoul Velazco, Univ. Grenoble Alpes, CNRS, Grenoble INP, FR

Abstract: *The purpose of this work is to describe a novel tester for radiation effects experiments, called TGV (Tester Generic and Versatile) based on a commercial development board ZEDBOARD. The main idea is to implement the whole DUT (Device Under Test) board architecture controlled by an FPGA, whose configuration is obtained from compiling the description of key features of the DUT in a high-level language such as C. This tester constitutes a powerful tool with generic capabilities for the functional validation and test under radiation of any digital circuit, with a particular focus on processor-like circuits. In this way, there is only a minor hardware development, limited to wiring the DUT pins to the ones of the tester connector. During the demonstration will be shown details of TGV platform, its use being illustrated by means of fault injection experiments which reproduces in a realistic way the random occurrence in time and location of SEUs in sensitive targets of the considered circuit.*

TIDES: NON-LINEAR WAVEFORMS FOR QUICK TRACE NAVIGATION

Presenter:

Jannis Stoppe, University of Bremen, DE

Author:

Rolf Drechsler, University of Bremen / DFKI, DE

Abstract: System trace analysis is mostly done using waveform viewers -- tools that relate signals and their assignments at certain times. While generic hardware design is subject to some innovative visualisation ideas and software visualisation has been a research topic for much longer, these classic tools have been part of the design process since the earlier days of hardware design -- and have not changed much over the decades. Instead, the currently available programs have evolved to look practically the same, all following a familiar pattern that has not changed since their initial appearance. We argue that there is still room for innovation beyond the very classic waveform display though. We implemented a proof-of-concept waveform viewer (codenamed Tides) that has several unique features that go beyond the standard set of features for waveform viewers.

TTOOL5G: MODEL-BASED DESIGN OF A 5G UPLINK DATA-LINK LAYER RECEIVER FROM UML/SYSML DIAGRAMS

Presenter:

Andrea Enrici, Nokia Bell Labs France, FR

Authors:

Julien Lallet¹, Imran Latif¹, Ludovic Aprville², Renaud Pacalet² and Adrien Canuel²

¹Nokia Bell Labs France, FR; ²Télécom ParisTech, FR

Abstract: Future 5G networks are expected to provide an increase of 10x in data rates. To meet these requirements, the equipment of baseband stations will be designed using mixed architectures, i.e., DSPs, FPGAs. However, efficiently programming these architectures is not trivial due to the drastic increase in complexity of their design space. To overcome this issue, we need to have unified tools capable of rapidly exploring, partitioning and prototyping the mixed architecture designs of 5G systems. At DATE 2017 University Booth, we demonstrate such a unified tool and show our latest achievements in the automatic code generation engine of TTool/DIPLODOCUS, a UML/SysML framework for the hardware/software co-design of data-flow systems, to support mixed architectures. Our demonstration will show the full design and evaluation of a 5G data-link layer receiver for both a DSP-based and an IP-based designs. We will validate the effectiveness of our solution by comparing automated vs manual designs.

WE DARE: WEARABLE ELECTRONICS DIRECTIONAL AUGMENTED REALITY

Presenter:

Davide Quaglia, University of Verona, IT

Authors:

Gianluca Benedetti¹ and Walter Vendraminetto²

¹Wagoo LLC, IT; ²EDALab srl, IT

Abstract: Current augmented reality (AR) eyewear solutions require large form factors, weight, cost and energy that reduce usability. In fact, connectivity, image processing, localization, and direction evaluation lead to high processing and power requirements. A multi-antenna system, patented by the industrial partner, enables a new generation of smart eyewear that elegantly requires less hardware, connectivity, and power to provide AR functionalities. They will allow users to directionally locate nearby radio emitting sources that highlight objects of interest (e.g., people or retail items) by using existing standards like Bluetooth Low Energy, Apple's iBeacon and Google's Eddystone. This booth will report the current level of research addressed by the Computer Science Department of University of Verona and the company Wagoo LLC. In the presented demo, different objects emit an "I am here" signal and a prototype of the smart glasses shows the information related to the observed object.

WORKCRAFT: TOOLSET FOR FORMAL SPECIFICATION, SYNTHESIS AND VERIFICATION OF CONCURRENT SYSTEMS

Presenter:

Danil Sokolov, Newcastle University, GB

Abstract: A large number of models that are employed in the field of concurrent systems' design, such as Petri nets, gate-level circuits, dataflow structures have an underlying static graph structure. Their semantics, however, is defined using additional entities, e.g. tokens or node/arc states, which collectively form the overall state of the system. We jointly refer to such formalisms as interpreted graph models. This demo will show the use of an open-source cross-platform Workcraft framework for capturing, simulation, synthesis, and verification of such models. The focus of our case study will be on synthesis from technology-independent formal specifications to verifiable circuit implementations.

XBARGEN: A TOOL FOR DESIGN SPACE EXPLORATION OF MEMRISTOR BASED CROSSBAR ARCHITECTURES.

Presenter:

Marcello Traiola, LIRMM, FR

Authors:

Mario Barbareschi¹ and Alberto Bosio²

¹University of Naples Federico II, IT; ²University of Montpellier - LIRMM laboratories, FR

Abstract: *The unceasing shrinking process of CMOS technology is leading to its physical limits, impacting several aspects, such as performances, power consumption and many others. Alternative solutions are under investigation in order to overcome CMOS limitations. Among them, the memristor is one of promising technologies. Several works have been proposed so far, describing how to synthesize boolean logic functions on memristors-based crossbar architecture. However, depending on the synthesis parameters, different architectures can be obtained. In this demo, we show a Design Space Exploration (DSE) that we use to select the best crossbar configuration on the basis of workload dependent and independent parameters, such as area, time and power consumption. The main advantage is that it does not require any simulation and thus it avoid any runtime overheads. The demo aims to show the tool prototype on a selected set of benchmarks which will be synthesized on a memristor-based crossbar circuit.*

See you at the University Booth!

University Booth Co-Chairs

Elena Ioana Vatajelu, TIMA Laboratory, FR and

Andreas Vörg, edacentrum, DE