



Published on DATE 2019 (<https://past.date-conference.com>)

[Home](#) > [Printer-friendly PDF](#) > [Printer-friendly PDF](#)

FM01 Welcome Reception & PhD Forum, hosted by EDAA, ACM SIGDA, and IEEE CEDA

FM01

Welcome Reception & PhD Forum, hosted by EDAA, ACM SIGDA, and IEEE CEDA

Session Type: Fringe Meeting

Date: Monday, March 25, 2019

Time: 18:00-21:00

Location / Room: Lunch Area

Organiser

Robert Wille, Johannes Kepler University Linz, AT ([Contact Robert Wille](#) [1])

All registered conference delegates and exhibition visitors are kindly invited to join the **DATE 2019** Welcome Reception and subsequent PhD Forum, which will take place on Monday, March 25, 2019, from 1800 - 2100 at the DATE venue in the Lunch Area.

The PhD Forum of the DATE Conference is a poster session and a buffet style dinner hosted by the European Design Automation Association (EDAA), the ACM Special Interest Group on Design Automation (SIGDA), and the IEEE Council on Electronic Design Automation (CEDA). The purpose of the PhD Forum is to offer a forum for PhD students to discuss their thesis and research work with people of the design automation and system design community. It represents a good opportunity for students to get exposure on the job market and to receive valuable feedback on their work.

Agenda

Time	Label	Session
18:00	FM01	Reception & PhD Forum, hosted by EDAA, ACM SIGDA, and IEEE CEDA Chair: Robert Wille, Johannes Kepler University Linz, AT, Contact Robert Wille [1]
18:00	FM01-1	Adaptive Runtime Resource Management for Mobile CMPs through Self-awareness Bryan Donyanavard, University of California, Irvine, US Address and Affiliation: Department of Computer Science Donald Bren School of Information and Computer Sciences University of California, Irvine Irvine, CA 92697-3435, USA Advisor: Nikil Dutt Email: bdonyana@uci.edu Phone: +1-661-713-7479 Thesis Summary: We explore self-awareness in resource management for mobile systems by addressing autonomy and emergent behavior as they apply to mobile MPSoCs. In our initial investigation, we implement a software hierarchical resource manager as part of the operating system for multicore platforms that controls multiple distributed leaf controllers which implement management policies. The supervisory controller provides autonomy by monitoring the operating conditions and prioritizing a small set of simple goals dynamically. The supervisory controller manages emergent behavior by coordinating the leaf controllers and their objectives in unity toward the current goal. The final contribution of this thesis will explore self-awareness in a multicore system that integrates both hardware and software resource managers. We will address emergent behavior challenges that arise when resource managers at different layers of abstraction (HW/SW) exist, with both overlapping and independent policies. We will provide autonomy for distributed machine-learning-based hardware resource managers and create a holistic self-aware resource management infrastructure through software enhancement.
18:00	FM01-3	Optimization of Trustworthy Biomolecular Quantitative Analysis Using Cyber-Physical Microfluidic Platforms Mohamed Ibrahim, Duke University, US The thesis tackles important problems related to key stages of the biomolecular workflow. The results emerging from this book provide the first set of optimization and security methodologies for the realization of biomolecular protocols using microfluidic biochips.
18:00	FM01-5	Analysis and Optimization of Reliability Issues of VLSI Power Grid Networks Sukanta Dey, Indian Institute of Technology Guwahati, IN Designing reliable power grids for modern chips or System-on-Chip (SoC) is the main aim of the power grid designers for reliability and longevity of the chip. Generally, due to the ill-suited design techniques large power grid networks of a modern chip suffer from serious voltage drop noises which creates reliability issues and can malfunction the chip if voltage drop noise exceeds a particular threshold value. Locating and minimizing the voltage drop noise is very much required and it is a time-consuming process. Therefore, the first part of the report introduces a fast method for locating the voltage drop noise. In the later part, a framework was devised to minimize the voltage drop noise.
18:00	FM01-6	Computer-Aided Design for Quantum Computing Alwin Zulehner, Johannes Kepler University Linz, AT Currently, there is an ongoing "race" to build the first practically useful quantum computer between large companies like IBM, Intel, Rigetti, and Google. Although still limited by the number of available qubits and low fidelity, they provide the first step towards the dream of building a fault-tolerant quantum computer with the capability of running quantum algorithms for dedicated problems in domains such as quantum chemistry and physical simulation—or for factoring large numbers in polynomial time. The Computer-Aided Design (CAD) community needs to be ready for this revolutionizing new technology. While research on automatic design methods for quantum computers is currently underway, there is still far too little coordination between the CAD community and the quantum community. Consequently, many CAD approaches proposed in the past have either addressed the wrong problems, or failed to reach the end users. To overcome this issue, the thesis contributes solutions to actually relevant design problems that are required for making quantum computing accessible to end users. More precisely, the introduced CAD methods include the design of quantum circuits (given a high level description of the respective algorithm), their simulation, as well as technology mapping required to run the circuits on real hardware devices.

18:00 FM01- **New Views for Stochastic Computing: From Time-Encoding to Deterministic Processing**
7 M. Hassan Najafi, University of Louisiana at Lafayette, US

Stochastic computing (SC), a paradigm first introduced in the 1960s, has received considerable attention in recent years as a potential paradigm for emerging technologies and "post-CMOS" computing. Logical computation is performed on random bitstreams where the signal value is encoded by the probability of obtaining a one versus a zero. This unconventional representation of data offers some intriguing advantages over conventional weighted binary. Implementing complex functions with simple hardware (e.g., multiplication using a single AND gate), tolerating soft errors (i.e., bit flips), and progressive precision are the primary advantages of SC. The obvious disadvantage, however, is latency. A stochastic representation is exponentially longer than conventional binary radix. Long latencies translate into high energy consumption, often higher than that of their binary counterpart. Generating bit-streams is also costly. Factoring in the cost of the bit-stream generators, the overall hardware cost of an SC implementation is often comparable to a conventional binary implementation. This dissertation begins by proposing a highly unorthodox idea: performing computation with digital constructs on time-encoded analog signals. We introduce a new, energy-efficient, high-performance, and much less costly approach for SC using time-encoded pulse signals. Instead of encoding data in space, as random bit-streams, we encode values in time. We show how analog periodic pulse signals can be used in performing essential stochastic operations. We explore the design and implementation of arithmetic operations on time-encoded data and discuss the advantages, challenges, and potential applications. The approach is an excellent fit for low-power applications that include time-based sensors, for instance, image processing circuits in vision chips. Experimental results on image processing applications show up to 99% performance speedup, 98% saving in energy dissipation, and 40% area reduction compared to prior stochastic implementations. We further introduce a novel area- and power-efficient synthesis approach for implementing sorting network circuits based on unary bit-streams. The proposed method inherits the fault tolerance and low-cost design advantages of processing random stochastic bit-streams while producing completely accurate result. Synthesis results of complete sorting networks show more than 90% area and power savings compared to the costs of the conventional binary implementation. However, the latency increases. To mitigate the increased latency, we use our developed time-encoding method. Time-based encoding of data is exploited for fast and energy-efficient processing of data with the developed sorting circuits. The approach is validated by implementing a low-cost, high-performance, and energy-efficient implementation of an important application of sorting, median filtering. Poor progressive precision is the main challenge with the recently developed deterministic methods of SC. Relatively prime stream length, clock division, and rotation of bit-streams are the three deterministic methods of processing bit-streams that are initially proposed based on unary bit-streams. For applications that slight inaccuracy is acceptable, these unary stream-based approaches must run for a relatively long time to produce acceptable results. This long processing time makes the deterministic approaches energy-inefficient compared to the conventional random stream-based SC. We propose a high-quality down-sampling method which significantly improves the processing time and the energy consumption of the deterministic methods by pseudo-randomizing bit-streams. We also propose two novel deterministic methods of processing bit-streams by using low-discrepancy sequences. Significant improvement in the processing time and energy consumption is observed using the proposed methods. We further demonstrate that computation on stochastic bit-streams has another compelling advantage: circuits naturally and effectively tolerate very high clock skew. Exploiting this advantage, we investigate Polysynchronous Clocking, a design strategy for optimizing the clock distribution networks of SC systems. Clock domains are split at a very fine level, reducing power on an otherwise large global clock tree. Each domain is synchronized by an inexpensive local clock. Alternatively, the skew requirements for a global clock tree network can be relaxed. The proposed design approach allows for a higher working frequency and so lower latency. It also results in significant area and energy savings for a wide variety of applications. Next, we develop a low-cost SC-based hardware implementation of a large Restricted Boltzmann Machine (RBM) Classifier completely on a single FPGA. Conventional binary implementation of a fully parallel design of a large neural network is expensive, involves extra design overheads, and in most cases cannot be fit on a single FPGA. We also develop a new reconfigurable architecture and methodology for synthesizing any given target function stochastically using finite state machines. When the target function is relatively complex, such as the exponentiation, the hyperbolic tangent, or high-order polynomial functions, our developed sequential logic-based implementation is more efficient than the prior combinational architectures. Our synthesis method also has the ability to implement multi-input functions at a very low cost. Compared to prior combinational logic-based approaches, the proposed reconfigurable architecture can save hardware area and energy consumption by up to 30% and 40%, respectively, while achieving a higher processing speed. Finally, as the first study of its kind to the best of our knowledge, we rethink the memory system design for SC. We integrate analog memory with conventional stochastic systems to reduce the energy wasted in conversion units. We propose a seamless stochastic system, StochMem, which features analog memory to trade the energy and area overhead of data conversion for computation accuracy. Comparing to a baseline system which features conventional digital memory, StochMem can reduce the energy and area significantly at the cost of slight loss in computation accuracy.

18:00 FM01- **System-level Mapping and Synthesis of Data Flow-Oriented Applications on MPSoCs**
9 Tobias Schwarzer and Jürgen Teich, Friedrich-Alexander-Universität Erlangen-Nürnberg, DE

Design methodologies for modern Multi-Processor System-on-a-Chips (MPSoCs) need to capture the diverse and arising architectural innovations. This extended abstract presents two novel application mapping methodologies for data flow applications: First, (i) a system-level mapping and synthesis approach for multi-cores is presented that introduces a novel system compilation flow to combine system synthesis and the use of Quasi-Static Schedules (QSSs) within a Design Space Exploration (DSE). To reduce the exploration time of this approach, we additionally investigate DSE strategies that are able to dynamically trade off between (a) approximating heuristics and (b) accurate performance evaluation, i.e., compilation of the application and subsequent performance measurement on the target platform. Then, (ii) a hybrid application mapping (HAM) approach on many-cores is described that combines the strengths of design-time analysis and optimization with the flexibility of adaptive run-time resource management. The design-time analysis involves a novel meta-heuristic DSE that eliminates architectural symmetries by abstracting the problem of mapping tasks to concrete instances of processors to a clustering of tasks and their mapping to processor types which serve as generic mapping constraints. Finally, for dynamic resource management, we propose with (a) a problem-specific backtracking approach and (b) an approach that adopts a general-purpose SAT solver two exact techniques for solving the mapping constraints at run-time.

18:00 FM01- **Compositional Circuit Design with Asynchronous Concepts**
14 Jonathan Beaumont, Imperial College London, GB

Asynchronous Concepts is a domain-specific language used for capturing the behaviours of an asynchronous circuit. It features a library containing multiple levels of concepts, for signal-, gate- and protocol-level descriptions of circuits, and each concept is composable, allowing these behaviours to be automatically combined. A tool, Plato, automatically compiles Asynchronous Concept specifications, and translates these to Signal Transition Graphs, for use with existing EDA tools for verification and synthesis. A design flow for Asynchronous Concepts has also been developed, allowing a full design to be carried out through the use of Workcraft, into which Plato and Asynchronous Concepts has been integrated.

18:00 FM01- **Automatic Methods for the Design of Droplet Microfluidics**
15 Andreas Grimmer, Johannes Kepler University Linz, AT

Microfluidics deals with the manipulation of small amounts of fluids. For implementing a microfluidic chip, droplet microfluidics using closed microchannels provides a well-established and highly potential platform. However, when designing a microfluidic network implementing the required operations, a huge number of physical specifications need to be considered. This results in a complex task, where, thus far, the designer often has very few methods to derive a design. This thesis aims to change this state of the art. To this end, this thesis contributes (1) simulation and design methods which support the design process of droplet microfluidics in general and (2) design methods for a dedicated droplet routing mechanism, which eventually can be combined to a first integrated design process.

18:00 FM01- **Worst-Case Execution Time Guarantees for Runtime-Reconfigurable Architectures**
16 Marvin Damschen, Lars Bauer and Joerg Henkel, Karlsruhe Institute of Technology, DE

Real-time embedded systems need to be analyzable for timing guarantees. Despite significant scientific advances, however, timing analysis lags years behind current microarchitectures with out-of-order scheduling pipelines, several hardware threads and multiple (shared) cache layers. To satisfy the increasing performance demands, analyzable performance features are required. In this paper, a novel timing analysis approach is proposed that introduces processors with a runtime-reconfigurable instruction set as one way to escape the scarcity of timing-analyzable performance features. It is shown how runtime reconfiguration can be realized to adhere to timing constraints, and the problem of selecting reconfigurable custom instructions to optimize the worst-case execution time of an application is solved. An approach is presented that for the first time combines optimized static WCET guarantees and runtime optimization of the average-case execution (maintaining WCET guarantees) using runtime reconfiguration of hardware accelerators. Ultimately, runtime reconfiguration of accelerators is shown as a key feature to achieve predictable performance.

- 18:00 FM01- **Architecture and Programming Model Support For Reconfigurable Accelerators in Multi-Core Embedded Systems**
17 Satyajit Das, Université de Bretagne-Sud, FR

The research for the PhD thesis was conducted at the Université Bretagne Sud, France in collaboration with University of Bologna, Italy. The Supervisors of the thesis were Philippe Coussy and Luca Benini, and the co-supervisors were Kevin Martin and Davide Rossi. The submitted text is an extended abstract of my PhD thesis. It describes the novel CGRA design, implementation, integration in a computing system for ultra-low power IoT targets, and compilation for the CGRA. It also describes the performance, area and energy results compared to the state of the art CGRA architectures. The document contains the results showcasing the efficiency of the proposed compilation flow compared to the state of the art algorithms used in the classical compilation techniques for CGRAs. Since the thesis is about energy efficient programmable accelerators in a heterogeneous platform, the proposed CGRA is integrated as a programmable accelerator in an open source multi-core platform PULP. The submitted document contains performance and energy efficiency results while operating in a heterogeneous environment compared to the homogeneous solution. Furthermore, the document is self contained with all the necessary references. I also have submitted one paper published at the IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD) along with the abstract.

- 18:00 FM01- **Supervised Testing of Embedded Concurrent Software**
18 Jasmin Jahic, Fraunhofer IESE, DE

Finding bugs in multithreaded software is hard because of the non-deterministic interleaving of operations that threads perform on shared memory. Dynamic analysis of the execution trace is more precise than static analysis, but is still prone to false warnings due to instrumentation-based approaches for execution monitoring, ignoring implicit synchronization by scheduling, and not being able to properly cope with synchronization intentions. We describe an approach that is using non-intrusive execution monitoring, integrates scheduling synchronization with state-of-the-art Lockset algorithm for finding concurrency bugs, and considers synchronization intentions at the architectural level.

- 18:00 FM01- **Advanced 3D-Integrated Memory Subsystems for Embedded and High Performance Computing Systems**
19 Deepak Mathew, University of Kaiserslautern, DE

The efficiency in terms of compute power and energy of today's computing systems are more and more limited by the underlying memory architectures. This not only applies to embedded devices, such as smartphones and tablets, but also to servers. Various applications demand for higher bandwidth (e.g. image processing, artificial neural networks etc.), larger capacity (e.g. graph processing, deep learning etc.), and better energy efficiency in the main memory system. This puts Dynamic Random Access Memory (DRAM) in the focus to improve performance and energy efficiency of advanced computing systems. Therefore, the design space of existing main memory systems, which is composed of DRAMs must be revisited and optimized to meet these new computing demands. Furthermore, the main memory system can be augmented with new emerging Non-Volatile Memory (NVM) technologies such as Resistive Random Access Memory (RRAM), 3D-Xpoint, STT-MRAM etc. These NVMs offer larger capacities and better energy efficiency at the expense of higher access latencies in comparison to DRAM. These new memory technologies are still in research or in early stages of production and thus there exists no accurate architectural models. This is one of the major challenges to incorporate these new memory technologies into the existing main memory systems. Based on the above requirements, this thesis conducted research on the following two dimensions: •Optimize the existing DRAM main memory subsystem to further improve bandwidth and energy efficiency •Incorporate new emerging Storage Class Memory (SCM) technologies (e.g. RRAM) into the main memory subsystem to increase the total memory capacity

- 18:00 FM01- **Controlling Writes for Energy Efficient Non-Volatile Cache Management in Chip Multiprocessors**
20 Sukam Agarwal, Indian Institute of Technology Guwahati, IN

Large processing demands by many cores require larger on-chip caches. Conventional caches made up of SRAM fall short in fulfilling these demands in terms of power, performance, and scalability. The evolution of Non-Volatile Memories (NVM) draws the attention of computer architects to look beyond the conventional memory technologies in the memory hierarchy. The benefits offered by these NVMs are high density, good scalability, and low static power consumption. However, such NVM caches suffer from costly write operations and weak write endurance. We present three techniques to deal with these challenges of NVM caches. To overcome the costly write operations, our first technique considers the existence of private blocks and allocates dataless entries for such blocks in the non-volatile region of a hybrid cache. In order to deal with the weak write endurance, our other two policies reduce the inter and intra-set write variations present in the cache. To reduce the intra-set write variation, our policy logically partitions the cache into multiple windows and disperses the writes among these windows by using write restriction. On the other hand, our inter-set wear leveling technique exploits the concept of Dynamic Associativity Management to redirect the writes from one set to another. We implemented our proposed techniques on a full system simulator GEM-5 and the experimental evaluations show significant improvements over the existing techniques.

- 18:00 FM01- **Design Techniques for Energy-Quality Scalable Digital Systems**
21 Daniele Jahier Pagliari, Politecnico di Torino, IT

Energy-Quality (EQ) scalable digital systems systematically trade off the quality of computations with energy efficiency, by relaxing the precision, the accuracy, or the reliability of internal software and hardware components in exchange for energy reductions. This design paradigm is believed to offer one of the most promising solutions to the impelling need for low-energy computing. However, the current state-of-the-art in EQ scalable design suffers from important shortcomings. First, the great majority of techniques proposed in literature focus only on processing hardware and software components. Nonetheless, for many real devices, processing contributes only to a small portion of the total energy consumption, which is dominated by other components (e.g. I/O, memory or data transfers). Second, in order to fulfill its promises and become diffused in commercial devices, EQ scalable design needs to achieve industrial level maturity. This involves moving from purely academic research based on high-level models and theoretical assumptions to engineered flows compatible with existing industry standards. Third, the time-varying nature of error tolerance, both among different applications and within a single task, should become more central in the proposed design methods. This involves designing "dynamic" systems in which the precision or reliability of operations (and consequently their energy consumption) can be dynamically tuned at runtime, rather than "static" solutions, in which the output quality is fixed at design-time. This thesis introduces several new EQ scalable design techniques for digital systems that take the previous observations into account. Besides processing, the proposed methods apply the principles of EQ scalable design also to interconnects and peripherals, which are often relevant contributors to the total energy in sensor nodes and mobile systems respectively. Regardless of the target component, the presented techniques pay special attention to the accurate evaluation of benefits and overheads deriving from EQ scalability, using industrial-level models, and on the integration with existing standard tools and protocols. Moreover, all the works presented in this thesis allow the dynamic reconfiguration of output quality and energy consumption.

- 18:00 FM01- **MuTARE: A Multi-Target Adaptive Reconfigurable Architecture**
23 Marcelo Brandalero, UFRGS - Universidade Federal do Rio Grande do Sul, BR

Adaptability is a fundamental requirement in modern computing systems since a wide range of applications with distinct requirements must efficiently execute in a single hardware. This work proposes MuTARE, a Multi-Target Adaptive Reconfigurable architecture that combines different adaptability techniques in a single design to best optimize for different targets (performance or energy), enabling better Pareto-Optimal trade-offs. MuTARE employs a reconfigurable accelerator coupled to a set of heterogeneous cores, offering the ability to create customized data paths for critical application kernels, while also providing adaptability for those code sequences that cannot be mapped to the accelerator. MuTARE can either work transparently (i.e., with no changes to binaries already deployed) or manually through instruction set extensions. In manual mode, additional support for approximate computing is provided to improve performance and energy in error-tolerant applications.

- 18:00 FM01- **Bitstream-level Proof-Carrying Hardware**
25 Tobias Wiersema, Paderborn University, DE

Using a host of third party IP cores in reconfigurable designs is the de-facto standard today, to greatly increase productivity in reconfigurable hardware design. Trust in these modules is usually not warranted, however, and thus techniques that replace the need for trust with facts and proofs are sorely needed. In my project I have proposed and implemented one such technique, Proof-Carrying Hardware at the bitstream level, where it can unlock its full verification potential while maintaining its core advantage: To shift the considerable cost-of-trust from the receiving IP core customer to its vendor.

- 18:00 FM01- **Efficient Virtual Prototype Verification Techniques: Theory, Implementation and Application**
26 Vladimir Herdt, University of Bremen, DE

This thesis advances the state-of-the-art on Virtual Prototype (VP) verification and analysis techniques. We provide various techniques related to verification of functional as well as non-functional properties. A particular focus of this thesis is on formal verification of SystemC-based VPs.

18:00 FM01- **Hybrid-DBT: Hardware-Accelerated Dynamic Binary Translation targeting VLIW processors**
27 Simon Rokicki, Irlisa, FR

18:00 FM01- **Low-power Architectures for Automatic Speech Recognition**
28 Hamid Tabani, Barcelona Supercomputing Center, ES

Automatic Speech Recognition (ASR) is one of the most important applications in the area of cognitive computing. Fast and accurate ASR is emerging as a key application for mobile and wearable devices. These devices, such as smartphones, have incorporated speech recognition as one of the main interfaces for user interaction. This trend towards voice-based user interfaces is likely to continue in the next years which is changing the way of human-machine interaction. In the first place, we have done a thorough analysis to identify the performance bottlenecks and the sources of energy drain when running ASR application on mobile and desktop CPUs. This thesis introduces several techniques at the software level to improve the efficiency of ASR systems running on modern processors. In the second place, we propose a novel register renaming technique for Out-of-Order processors, which is able to reuse single-use registers to reduce the pressure on the register file. Unlike previous works, our scheme is able to precisely recover the state of the processor after the event of branch mispredictions, interrupts and exceptions. The proposed scheme is implemented by applying some changes to the renaming table, issue queue and register files while it does not need any changes in the compiler nor the ISA. In third place, we design and propose a hardware accelerator for GMM evaluation. Our accelerator consumes less energy and outperforms CPUs and GPUs by orders of magnitude. Our accelerator implements in hardware our scheme to predict the active senones in a batch of frames. We provide a comprehensive study of different lossy and lossless compression schemes and an analysis of GMM parameters. We propose a novel clustering scheme which provides significantly higher Compression/WER ratio in comparison with traditional schemes.

18:00 FM01- **Low Overhead & Energy Efficient Storage Path for Next Generation Computer Systems**
30 Athanasios Stratikopoulos, The University of Manchester, GB

The rise of social networks along with the growth of big data have driven storage systems to cope with large scale of data volume. In recent years, the emergence of the Non-Volatile Memory Express (NVMe) standard has enabled SSD drives to deliver high I/O rates by leveraging the fastest available interconnect (i.e. PCIe) to the processing chip. Additionally, the appearance of FPGAs in data centres is creating opportunities to accelerate not only application functionality but also OS operations. Although the majority of the servers in data centres have been connected to the FPGAs via the PCIe interconnect, there are now also available heterogeneous System on Chips (SoCs) with multi-cores and FPGAs integrated on the same die, resulting in low latency and energy efficient communication. This work analyses the source of performance overhead in existing state-of-the-art storage devices and proposes a novel low overhead and energy efficient storage path called FastPath, that operates transparently to the processing cores. The experimental results showed that FastPath can achieve up to 82% lower latency, up to 12x higher performance, and up to 10x more energy efficiency for standard microbenchmark on an Arm-FPGA Zynq 7000 SoC. Further experiments were conducted on a state-of-the-art SoC, such as the Zynq UltraScale+ MPSoC, using a real application, such as the Redis-in-memory database, which received requests by the Yahoo! Cloud Serving Benchmark (YCSB). The experimental evaluation showed that FastPath achieved up to 60% lower latency and 15% higher throughput than the baseline storage path in the Linux kernel.

18:00 FM01- **A Model driven Framework with Assertion Based Verification Support for Embedded Systems Design Automation**
31 Muhammad Waseem Anwar, National University of Sciences & Technology (NUST), PK

In this PhD thesis, a model driven framework with full Assertion Based Verification (ABV) support is introduced to perform both static as well as dynamic ABV. Firstly, a modeling methodology is proposed to model system design (structure + behavior) and verification constraints (assertions) of embedded systems. Particularly, UML and SysML based modeling approach is introduced to model system design i.e. Block Definition Diagram (BDD) is used to represent system structure and State Machine Diagram (SMD) is used to represent system behavior. Moreover, SVOCL (SystemVerilog in Object Constraint Language), an OCL temporal extension for SystemVerilog, is proposed to represent the verification constraints by means of SVA's. Furthermore, a Natural Language for Computational Tree Logic (NLCTL) is introduced to represent system assertions by means of CTL properties at higher abstraction level. Secondly, a complete transformation engine is implemented to generate synthesizable SystemVerilog RTL code, SystemVerilog Assertions code, Timed automata model and CTL properties from the source high level models. Particularly, transformation rules are developed to perform conceptual mapping between BDD / SMD constructs and SystemVerilog RTL / Timed Automata constructs. Furthermore, rules are also developed to convert SVOCL and NLCTL constraints into SVA's and CTL properties respectively. Finally, the implementation of transformation rules is carried out in JAVA language and Accelleo tool through Model-to-text transformation approach. As SystemVerilog RTL and assertions code is automatically generated from models through transformation engine, any UVM complaint simulator can be used to perform dynamic ABV. In this research, ABV is performed through QuestaSIM simulator. Similarly, UPPAAL tool is used to perform static ABV. The applicability of the proposed framework is demonstrated through eight benchmark case studies i.e. Traffic Lights Controller (TLC), Car Collision Avoidance System (CCAS), Arbitrator, Elevator, Unmanned Aerial Vehicle (UAV), ATM, Train Gate and Bridge Crossing system.

18:00 FM01- **Optimization and Analysis for Dependable Software on Unreliable Hardware Platforms**
32 Kuan-Hsun Chen, Technical University of Dortmund, DE

Due to the invention of semiconductor-based integrated circuits, embedded systems have become ubiquitous even in safety-critical domains. In these domains, the correctness of the system behaviors depends not only on the functional correctness but also upon the timeliness of the time instant at which the results are delivered. As chip technology keeps on shrinking towards higher densities and lower operating voltages, memory and logic components are now vulnerable to electromagnetic inference and radiation, leading to transient faults in the underlying hardware, which may jeopardize the correctness of software execution, so-called soft errors. Instead of solely addressing transient faults at the hardware level, embedded-software developers have started to deploy Software-Implemented Hardware Fault Tolerance (SIHFT) techniques. However, the main expenditure is significant amount of time from the overhead of using SIHFT techniques. To support safety-critical systems, real-time systems technology have been primarily used and widely studied. Even without considering any performance overhead incurred by SIHFT techniques, making a predictable real-time system is a challenge matter. While considering hardware faults and SIHFT techniques, classic stories in real-time systems might turn over a new leaf. In this dissertation, there are three main contributions providing analyses and optimizations for transient fault-tolerance of system software. The contributions presented in this dissertation have been published in peer-reviewed international conferences and journals, and have been used by researchers.

18:00 FM01- **Multiple NoC based Custom Implementation and Traffic Distribution to attain Energy Efficient CMPs**
33 Sonal Yadav, Vijay Laxmi and Manoj Singh Gaur, MNIT Jaipur, IN

Multi-NoCs are primarily implemented for application-specific processors e.g. SoC, MPSoC, and FPGAs. Contrary, general-purpose processors are less explored for multi-NoCs implementations. CMPs run a wide variety of applications with unpredictable low/high heterogeneous runtime traffic variations. For these processors, it is difficult to design a static power efficient customised multi- NoC that should dynamically adapt traffic distribution according to runtime variations of fine-grain messages from computation bound to communication and memory-bound applications. We have addressed a difficult challenge to design energy efficient multi-NoC for CMPs. To attain it, NoC power consumption should be proportional to the network demand without compromising communication delays. In our thesis, multiple NoCs itself hardware implementation is customised for static power along with fine-grain traffic distribution exploration for improving energy efficiency of CMP's runtime traffic. Our novel contributions are as follows: 1) Customised Multi-NoC Architecture to Attain Power Efficiency 2) Target Multi-NoC Architecture 3) Case Study: Message Distribution Problem of Multi-NoC in CMPs 4) Runtime Adaptive Fine Grained Message Distribution for improved Runtime Utilisation of Multi-NoCs

18:00 FM01- **True Random Number Generators for FPGAs**
37 Bohan Yang, ESAT/COSIC and iMinds, KU Leuven, BE

A True Random Number Generator (TRNG) circuit is designed to be sensitive to a particular physical phenomenon when it is in use, and to be resistant to process variations and other unwanted random physical phenomena. TRNGs are used in cryptography for generating session keys, nonces, and random challenges in various authentication protocols. The subject of my PhD thesis is the study of TRNG, which can be implemented on FPGA hardware. Our contributions to TRNG designs include two novel digital noise sources, a method to measure timing jitter, two design methodologies for online tests and the exploration of implementation tradeoffs of one post-processing algorithm.

18:00 FM01- **HW/SW Co-Design Methodology for Mixed-Criticality and Real-Time Embedded Systems**
38 Vittorio Muttolo, University of L'Aquila, IT

In the last years, the spread and importance of embedded systems are even more increasing, but it is still not yet possible to completely engineer their system-level design flow. The main design problems are to model functional (F) and non-functional (NF) requirements and to validate the system before implementation. Designers commonly use one or more system-level models (e.g. block diagrams, UML, SystemC, etc.) to have a complete problem view and to perform a check on HW/SW resources allocation by simulating the system behavior. In this scenario, SW tools able to support designers to reduce cost and overall complexity of systems development are even more of fundamental importance. Co-existence of functional and non-functional requirements is the most relevant challenge. Unfortunately, there are no general methodologies defined for this purpose and, often, the only option is to refer to experienced designer indications with respect to empirical criteria and qualitative assessments. In such a context, this Ph.D. work faces the problem of the HW/SW co-design of dedicated (possibly embedded and real-time) systems based on heterogeneous parallel architectures and presents a framework (with related methodology and prototypal tools), called Hepsycode, able to support the development of such kind of systems in different application domains considering mixed-criticality and real-time requirements.

18:00 FM01- **Improving Bundled-Data Handshake Circuits**
41 Norman Kluge, Hasso-Plattner-Institut, University of Potsdam, DE

Balsa provides an open-source design flow where asynchronous circuits are created from high-level specifications, but the syntaxdriven translation often results in performance overhead. The thesis presents an adopted design flow tackling core drawbacks of the original Balsa design flow: To improve the performance of the circuits, the fact that bundled-data circuits can be divided into data and control path is exploited. Hence, tailored optimisation techniques can be applied to both paths separately. For control path optimisation, STG-based resynthesis has been used (applying logic minimisation). To optimise the data path a standard combinatorial optimiser is used. However, this removes the matched delays needed for a properly working bundled-data circuit. Therefore, two algorithms to automatically insert proper matched delays are used. Circuit latency improvements of up to 44 % and energy consumption improvements of up to 60 % compared to the original Balsa implementation can be shown.

18:00 FM01- **IC Design of an Inductorless DC/DC Converter with Wide Input Voltage Range in Low-Cost CMOS**
43 Gabriele Ciarpì, University of Pisa, IT

This work shows the design and implementation of an inductorless DC/DC converter. It is able to convert a wide input voltage range (from 6 V to 60 V) in two regulated output voltages (5 V and 1.65 V) for low power applications. A 3D implementation of the converter is proposed for area and cost reduction. The achieved results, from experimental measurements, prove the suitability of this integrated DC/DC converter, using 3-D technology, for applications where low-power loads (e.g., sensors, memories, and processors) have to be supplied by high-voltage power supplies.

18:00 FM01- **Monolithic-3D Integration based Memory Design techniques towards Robust and in-memory computing**
44 Srivatsa Rangachar Srinivasa¹, John (Jack) Sampson², Meng-Fan (Marvin) Chang³ and Vijaykrishnan Narayanan⁴
¹penn state university, US; ²Penn State, US; ³National Tsing Hua University, TW; ⁴Penn State University, US

With the dominance of data-centric applications two important memory design challenges namely: 'speed and robustness' and 'energy efficiency' are of critical importance. Emerging 3D technology like Monolithic-3D integration (M3D-IC) has the potential to address both these issues in memories. In this work we exploit several key characteristics of M3D-IC to design novel SRAM based 3D memories offering reliability, multidimensional data read capability and in-memory compute support. We propose two flavors of memory designs for Boolean and search oriented in-memory computations. We also investigate architectural enhancements required to incorporate these memories into a processing environment with seamless computation offload.

18:00 FM01- **Cross-Layer Synthesis and Integration Methodology of Wavelength-Routed Optical Networks-on-Chip for 3D-Stacked Parallel Computing Systems**
47 Mahdi Tala, University of Ferrara, IT

This work has contributed to bridge the gap between developers of silicon photonic devices and system designers through a cross-layer synthesis and integration methodology of wavelength-routed optical networks-on-chip (WRONoC). In particular, the design space of WRONoC topologies has been populated with new solutions which outperform the state of the art topologies. The work has characterized quality metrics of design points by including a high-impact component that is typically overlooked, that is the bridge with the electronic section of the system. As a result, the work has identified the most energy-efficient configurations of the network as a whole, demonstrating feasibility of 1pJ/bit signaling, provided identified cost, signal integrity and fabrication challenges are overcome.

18:00 FM01- **Adaptive Knobs for Resource Efficient Computing**
48 Anil Kanduri, University of Turku, FI

Performance demands of emerging domains such as artificial intelligence, machine learning and vision, Internet-of-things etc., continue to grow. Given the increase in power densities, fixed power and energy budgets and thermal constraints, meeting performance requirements become challenging. This leaves an open problem on extracting the required performance within the power and energy limits, while also ensuring thermal safety. Architectural solutions including asymmetric and heterogeneous cores and custom acceleration improve performance-per-watt. Despite the efforts in architecture and run-time systems, satisfying applications' performance requirements under dynamic and unknown workload scenarios subject to varying system dynamics of power, temperature and energy requires intelligent run-time management. This dissertation proposes adaptive run-time strategies for resource efficient computing, considering unknown and dynamic workload scenarios, diverse application requirements and characteristics and variable effect of power actuation on performance. Our specific contributions are i) run-time mapping approach to improve power budgets for higher throughput, ii) thermal aware performance boosting for efficient utilization of power budget and higher performance, iii) approximation as a run-time knob exploiting accuracy-performance trade-offs for maximizing performance under power caps at minimal loss of accuracy and iv) co-ordinated approximation for heterogeneous systems through joint actuation of dynamic approximation and power knobs for performance guarantees with minimal power consumption.

18:00 FM01- **Device-Circuit Co-design Employing Phase Transitioning Materials for Low Power Digital Applications**
52 Ahmedullah Aziz, Purdue University, US

Phase transitioning materials (PTM) belong to the family of emerging technologies which have tremendous prospect but also pose unique challenges. My doctoral research reports several innovative ideas (with thorough analyses) for devices/circuits utilizing unique properties of these materials. I have worked on different levels of abstraction (materials to systems) to properly analyze the implications of the novel techniques. I have proposed solutions to several prevailing issues in concurrent digital electronics and sought to overcome the limitations of existing technologies. I have established design methodologies for my proposed devices and circuits to enable optimization and guide future exploration of useful PTMs.

18:00 FM01- **Advanced CAD Frameworks for Design IP Protection**
53 Satwik Patnaik, NEW YORK UNIVERSITY, US

Regular design IP may be duplicated without consent, resulting in financial loss for the IP owner. Besides, the tools and know-how for reverse engineering (RE) are becoming more widely available, thus rendering the scenario of malicious end users obtaining some IP a practical threat. Besides, adversaries in an untrustworthy fab can readily obtain the underlying IP from the design files given to them. We propose several advanced CAD frameworks, which allow concerned designers to properly protect their IP while exploring the inherent trade-offs for layout cost and commercial cost. Our techniques have been demonstrated to advance the state-of-the-art protection schemes. Besides covering regular 2D CMOS integration, we also explore emerging devices and 3D integration for their interesting prospects toward IP protection.

This proposal seeks to build systems capable of responding to the diverse needs in real time with orders of magnitude more energy efficient operation. We propose a novel hardware/software co-design of a hybrid Processing In-Memory platform which accelerates fundamental operations and diverse data analytic procedures using processing in-memory technology. In the hardware layer, the proposed platform has a hybrid structure comprising of two units: PIM-enabled processors and PIM-based accelerators. The PIM-enabled processors enhance traditional processors by supporting fundamental block-parallel operations inside processor cache structure and associated memory, e.g., addition, multiplication or bitwise computations. This capability will be implemented through the memory hierarchy in a similar way to the conventional architecture. To fully get the advantage of PIM for popular data processing procedures and machine learning algorithms, we design specialized accelerator blocks using in-memory processing technology. Our platform can process several applications including machine learning, graph, and query processing completely in-memory without using any processing cores.

Source URL: <https://past.date-conference.com/conference/fringe-meeting-fm01>

Links:

[1] https://past.date-conference.com/user/27317/contact_form