

# Guest Editorial: *IEEE Transactions on Computers* and *IEEE Transactions on Emerging Topics in Computing* Joint Special Section on Innovation in Reconfigurable Computing Fabrics from Devices to Architectures

Ronald F. DeMara, *Senior Member, IEEE*, Marco Platzner, *Senior Member, IEEE*, and Marco Ottavi, *Senior Member, IEEE*

As facilitated by increases in integration density from technology scaling and the potential of beyond-CMOS device technologies, emerging innovations in reconfigurable computing fabrics span from the circuit-level through the application level. This Special Section focuses on the advancement of the associated performance and reliability objectives via technology and functional heterogeneity, as well as advanced resilience properties in which reconfigurable fabrics are embarking. Emerging device characteristics such as non-volatility and new static versus dynamic energy consumption profiles, as well as novel interconnect mechanisms, and storage blocks within reconfigurable fabrics, in turn innovate architectural advances that enable new applications. The manuscripts appearing in this Special Section focus on one or more of these innovations during the design-cycle phases including modeling, simulation, design, fabrication, and applications.

The main motivations for embracing reconfigurable fabrics such as Field Programmable Gate Arrays (FPGAs) still solidly remain along two categories of attributes: (1) *flexibility and accessibility*: fabric flexibility allows realization of logic elements at medium and fine granularities while incurring low non-recurring engineering effort and rapid deployment to market, and (2) *resiliency and runtime adaptability*: reconfigurable fabrics have been demonstrated to provide a viable solution for process-voltage-temperature variation induced concerns and can be utilized effectively for fault recovery. However, these advantages have traditionally been achieved at some cost of increased fabric area and power consumption, as well as a decrease in throughput compared to Application-Specific Integrated Circuits (ASICs). Innovations using emerging devices within reconfigurable fabrics which are highlighted in this Special Section have sought to bridge the gaps needed to provide more of these benefits at reduced energy and area costs.

- R. DeMara is with the Department of Electrical and Computer Engineering, University of Central Florida, 4328 Scorpilus Street, Orlando, FL, USA, 32816-2450.  
E-mail: [ronald.demara@ucf.edu](mailto:ronald.demara@ucf.edu)
- M. Platzner is with the Department of Computer Science at Paderborn University, Warburgerstrasse 100, 33098 Paderborn, GERMANY.  
E-mail: [platzner@upb.de](mailto:platzner@upb.de)
- M. Ottavi is with the Department of Electronic Engineering, University of Rome Tor Vergata, Via del Politecnico 1, 00133 Rome ITALY.  
E-mail: [ottavi@ing.uniroma2.it](mailto:ottavi@ing.uniroma2.it)

For information on obtaining reprints of this article, please send e-mail to: [reprints@ieee.org](mailto:reprints@ieee.org), and reference the Digital Object Identifier below.

Currently, Static Random Access Memory (SRAM) cells form the basis for most commercial FPGAs. SRAM cells are employed within programmable switching elements to control the interconnection between logic building blocks. Moreover, they are utilized in Look-Up Tables (LUTs) to store the logic function configuration data, which constitute the primary components in reconfigurable fabrics. In particular, in conventional fabrics, each LUT is comprised by a memory with  $2^m$  cells in which the truth table of an  $m$ -input Boolean function is stored. The re-programmability of the SRAM cells, and the fact that they can be readily implemented by highly-scaled CMOS technology, have traditionally resulted in SRAM-based FPGAs being the most popular reconfigurable fabrics on the market. Current challenges addressing SRAM cell are (1) *high static power*: due to the existence of intrinsic leakage current having a proportion that increases significantly with technology scaling, (2) *volatility*: SRAM is volatile, therefore all functions must be reprogrammed upon each power-up cycle, and (3) *low logic density*: SRAM cells are composed of six transistors which ultimately limits the logic density, as well as impacting the energy profile.

This issue highlights innovations in reconfigurable computing fabrics continuing their transition towards embracing the benefits of increased heterogeneity of SRAM and beyond-SRAM technologies, along several cooperating dimensions. These include new device technologies, such as Resistive Random Access Memory (RRAM), Spin-Transfer Torque Magnetic Random Access Memory (STT-MRAM), and inkjet configurable devices, as well as architectures and methods for fault-tolerance, and interconnect advancements including Via-based reconfigurability and networks-on-chip, while embracing applications spanning from physically unclonable functions to audio processing on a neuromorphic fabric.

This Special Section attracted roughly 30 submissions along with the involvement of nearly 120 reviewers who were selected for their expertise on the precise topics of each submission. Thus, it represents a collective effort from the research community and industry participants on an international scale. Among the submissions, four articles focused on topic areas including beyond-CMOS and alternative device technologies are being published in *IEEE Transactions on Emerging Topics in Computing*. Meanwhile, seven additional manuscripts are published in *IEEE Transactions on Computers* focused on resilience, interconnect design, low-energy reconfigurable devices, and application-specific fabrics. Articles appearing in this Special Section were also invited to prepare companion abstract videos, which if available will appear on the Transactions on Computers portal page at <http://www.computer.org/web/tc>.

The utilization of emerging devices within reconfigurable fabrics is addressed through several approaches in this Special Section. The article “A High-performance FPGA Architecture Using One-Level RRAM-based Multiplexers” by Xifan Tang, Giovanni De Micheli, and Pierre-Emmanuel Gaillardon shows how Resistive RAM technology can be successfully applied to the routing of FPGA architectures. In particular, it is shown that RRAM-based routing multiplexers, can be employed to achieve significantly reduced delay compared to state-of-art SRAM-based implementations. Based on the specific features of RRAM multiplexers, the paper introduces FPGA architecture optimizations that significantly improve both Area-Delay Product and Delay-Power Product. Meanwhile, the manuscript “Energy Efficient Reconfigurable Threshold Logic Circuit with Spintronic Devices” by Zhezhi He and Deliang Fan develops an energy-efficient approach by intrinsically-supporting a threshold-based logic model. It focuses on reconfigurable threshold logic gates whereby a spintronic weight device enables the design of Spintronic Threshold Logic Gates (STLGs) within a reconfigurable fabric. This paper shows that the proposed STLG-based reconfigurable logic could provide important improvements with respect to state-of-the-art, including much lower Energy-Delay Product (EDP) compared with CMOS based FPGAs.

Another article addressing emerging technologies, titled “Inkjet-Configurable Gate Arrays” by Jordi Carrabina, Mohammad Mashayekhi, Jofre Pallarès, and Lluís Terés, illustrates how gate array configurations could be applicable to inkjet-printed circuitry. The authors show that using inkjets for the configuration layer of a regular gate array structure allows them to combine low cost of inkjet technology with the higher yield achieved via flexibility of a post-manufacturing reconfiguration phase.

Interconnect innovations for reconfigurable fabrics are showcased within two manuscripts. The article “Design and Applications for Embedded Networks-on-Chip on FPGAs” by Mohamed S. Abdelfattah, Andrew Bitar, and Vaughn Betz proposes the inclusion of a packet-based network-on-chip in FPGA architectures to facilitate on-chip communication. Through several case studies including an Ethernet switch and a packet processor, the authors show that the use of an embedded network-on-chip supports increases in clock frequency while saving on long wires, which are scarce resources in FPGAs. Meanwhile, the manuscript “SAT-Based Formulation for Logical Capacity Evaluation of VIA-Configurable Structured ASIC” by Felipe Marranghello, Vinicius Dal Bem, André Reis, and Renato Ribas examines a different angle of reconfigurable fabrics from the perspective of a new Via-Configurable Structured ASIC (VCSA) design. It proposes a satisfiability formulation for determining the logical capacity of a VCSA block template or fabric, thus looking at the optimal configuration of a regular fabric at manufacturing time. The proposed SAT formulation improves the performance over previous solutions while eliminating multiple iterations.

New approaches to fault-handling are addressed at various levels of design abstraction. The article “Hierarchical Strategies for Efficient Fault Recovery on the Reconfigurable PAnDA Device” by Martin A. Trefzer, David M. R. Lawson, Simon J. Bale, James A. Walker, and Andy M. Tyrrell focuses on fault-tolerance and recovery in reconfigurable fabrics. The authors

discuss two different strategies, a stochastic and a deterministic one, that are able to repair circuits under random faults injected at random locations on the PAnDA device. PAnDA realizes a programmable analog and digital array that features fine-grained hierarchical partial reconfiguration. Fault-tolerance continues to be the main topic in the paper “Fault-tolerant FPGA with Column-based Redundancy and Power Gating Using RRAM” by Kibum Lee and S. Simon Wong. The approach achieves fault-tolerance by inserting redundant columns into the logic fabric and storing fault locators in non-volatile RRAM monolithically integrated on top of the CMOS circuit. Additionally, the fabric's columns are power-gated which not only helps minimize leakage power but also provides redundancy in the power network. With regards to fault-tolerance through configuration diversity, the article “Aging Resilience and Fault Tolerance in Runtime Reconfigurable Architectures” by Hongyan Zhang, Lars Bauer, Michael A. Koche, Eric Schneider, Hans-Joachim Wunderlich, and Jörg Henkel, strives for achieving fault-tolerance and aging resilience by creating different configurations for runtime-reconfigurable accelerators. The configurations differ in their use of logic blocks. A corresponding placement algorithm utilizes the diversity in the configurations to mask faulty logic blocks and to balance the execution-induced stress for logic blocks.

Considering aspects of hardware security, the article “Binary Decision Diagram Assisted Modeling of FPGA-based Physically Unclonable Function by Genetic Programming” by Rajat Subhra, Ratan Rahul Jeldi, Indrasish Saha, and Jimson Mathew looks at a new method for creating computational models for large Boolean functions for which only a small fraction of the truth table is initially known. The authors apply genetic programming on logic representations in the form of reduced ordered binary decision diagrams. The importance of this work stems from the fact that creating such computational models corresponds to the problem of finding acceptable physically unclonable functions for FPGAs.

Approaches to low-energy reconfigurable computing span from configurable elements to non-Boolean computing approaches. An FPGA architecture employing hybrid logic blocks combining small-input look-up tables with reconfigurable, dedicated gate functions of a few inputs is discussed in the manuscript “PEAF: A Power-Efficient Architecture for SRAM-Based FPGAs Using Reconfigurable Hard Logic Design in Dark Silicon Era” by Zahra Ebrahimi, Behnam Khaleghi, and Hossein Asadi. Power gating is used to selectively turn off the dedicated functions and the look-up tables, which results in a reduction in static power and the power-delay product. On the other hand, IBM's TrueNorth processor adopts an alternative computational paradigm within a configurable platform for implementing deep neural networks at very low power budgets. In their article titled “Always-on Speech Recognition using TrueNorth, a Reconfigurable, Neurosynaptic Processor,” Wei-Yu Tsai, Davis R. Barch, Andrew S. Cassidy, Michael V. DeBole, Alexander Andreopoulos, Bryan L. Jackson, Myron D. Flickner, John Arthur, Dharmendra S. Modha, Jack Sampson, and Vijaykrishnan Narayanan demonstrate the use of this processor for implementing an audio processing pipeline. For one part of the pipeline, i.e. an audio feature extractor, they explore trade-

offs in accuracy, power and performance by customizing TrueNorth's neuromorphic network structures.

The guest editors thank all of the reviewers for their valued time, expertise, and constructive feedback in their reviews. We also thank all of the authors for their submissions and their accommodation of publication deadlines and constraints. Finally, we are indebted to the Editor-in-Chief of *IEEE Transactions on Computers*, Dr. Paolo Montuschi, and the Editor-in-Chief of *IEEE Transactions on Emerging Topics in Computing*, Dr. Fabrizio Lombardi, who have collectively made this joint Special Section possible.

Sincerely,

Ronald F. DeMara  
Marco Platzner  
Marco Ottavi  
Guest Editors



**Ronald F. DeMara** (S'87-M'93-SM'05) received the Ph.D. degree in Computer Engineering from the University of Southern California in 1992. Since 1993, he has been a full-time faculty member at the University of Central Florida where he is a Professor of Electrical and Computer Engineering, and joint faculty of Computer Science, and has served as Associate Chair, ECE Graduate Coordinator, and Computer Engineering

Program Coordinator. His research interests are in computer architecture with emphasis on reconfigurable logic devices, evolvable hardware, and emerging devices, on which he has published approximately 200 articles and holds one patent. He is a Senior Member of IEEE and has served on the Editorial Boards of *IEEE Transactions on VLSI Systems*, *Journal of Circuits, Systems, and Computers*, the journal *Microprocessors and Microsystems*, and as Associate Guest Editor of *ACM Transactions on Embedded Computing Systems*, as well as a Keynote Speaker of the *International Conference on Reconfigurable Computing and FPGAs (ReConFig)* and the *Reconfigurable Architectures Workshop (RAW)*. He is currently an Associate Editor of *IEEE Transactions on Computers* and serves on various IEEE conference program committees, including *ISVLSI* and *SSCI*. He received *IEEE's Joseph M. Bidebach Outstanding Engineering Educator Award* in 2008.



**Marco Platzner** (M'95-SM'10) is a Professor of Computer Engineering at Paderborn University. Previously, he held research positions at the Computer Engineering and Networks Lab at ETH Zurich, Switzerland, the Computer Systems Lab at Stanford University, USA, the GMD - Research Center for Information Technology in Sankt Augustin, Germany, and the Graz University of Technology, Austria. Marco Platzner holds diploma and PhD degrees in Telematics (Graz University of Technology, 1991 and 1996), and a

"Habilitation" degree for the area hardware-software codesign (ETH Zurich, 2002). His research interests are in reconfigurable computing, hardware-software codesign, and parallel architectures, on which he has published more than 150 papers. He is a senior member of the IEEE, a member of the ACM, a member of the boards of the Paderborn Center for Parallel Computing and the Paderborn Institute of Advanced Studies in Computer Science and Engineering, a faculty member of the International Graduate School Dynamic Intelligent Systems of Paderborn University, and of the Advanced Learning and Research Institute at Universita' della Svizzera Italiana (USI), in Lugano. He is currently Associate Editor for the *ACM Transactions on Reconfigurable Technology and Systems*, the *Hindawi International Journal on Reconfigurable Computing*, and the *Springer International Journal on Design Automation for Embedded Systems*.



**Marco Ottavi** (M'04-SM'10) received the Ph.D. degree in telecommunications and microelectronics engineering from the University of Rome Tor Vergata Italy in 2004. Since 2009 he has been with the same University as a recipient of a "rientro dei cervelli" fellowship, awarded by the Italian Ministry of University and Research, and, from 2014, as an Associate Professor. Previously, he was

a Senior Design Engineer with AMD in Boxborough, (MA) USA, and he held Postdoctoral positions with Sandia National Laboratories in Albuquerque, (NM) USA, and with the ECE Department of Northeastern University in Boston, (MA) USA. His research interests include VLSI yield and reliability modeling, test, design for testability, fault-tolerant architectures, and online testing and design of nanoscale circuits and systems. In these fields, he has published about 100 articles on archival journals and peer-reviewed conferences. From December 2011 to November 2015 he was the Chair of COST Action IC1103 "Manufacturable and Dependable Multicore Architectures at Nanoscale" (MEDIAN). He serves as a member of several conference program committees and as a reviewer for various IEEE journals and highly cited conferences. He is an Associate Editor of *IEEE Transactions on Emerging Topics in Computing* and a Senior Member of IEEE.