Special Issue on Reconfigurable Systems: Foundations

By JAMES LYKE

CHRISTOS G. CHRISTODOULOU

ALONZO VERA

ART H. EDWARDS

I. INTRODUCTORY REMARKS

This issue of Proceedings of the IEEE is the first of two special issues that will provide an in-depth treatment of the subject of reconfigurable systems as enabled through advancements in the electrical engineering field, as well as the approaches we use to manage complexity in systems having millions of fixed and programmable elements. It is our view that the frenetic evolution of the field programmable gate array (FPGA) as a preeminent digital platform will not only continue for the foreseeable future, but will serve as an inspiration/template for other functional domains to include analog, radio-frequency, photonic, and even those not strictly electrical in nature to include thermal, fluidic, and mechanical. This issue will serve as an immersion into the subject of reconfigurable systems, both for the purposes of gaining familiarity with essential theoretical elements, the key developments in the field, and the potential future directions (not just the next few years, but decades beyond the present).

This issue is the first of two special issues that will provide an in-depth treatment of the subject of reconfigurable systems as enabled through advancements in the electrical engineering field.

Our future will depend on reconfigurability. To some degree, it already does. Every memory pattern in our computers, cameras, and phones is based on harnessing a rudimentary form of reconfigurability—the binary storage of electronic charge. The power of binary reconfigurability has given us the von Neumann (stored program) computer, mass storage, and the FPGA, along applications such as software-defined radio and software-defined networking. But we have only scratched the surface of a vast frontier of possibilities in reconfigurability. The implications of this frontier could change our society in ways as profoundly as has the "digital revolution."

In the two dedicated special issues of the Proceedings of the IEEE, we explore the fundamental concepts behind reconfigurable systems, from the earliest foundations of computation to architectures of modern FPGAs. We will outline the principles through which reconfigurable systems can be understood, applied, and extended.

Digital Object Identifier: 10.1109/JPROC.2015.2399071

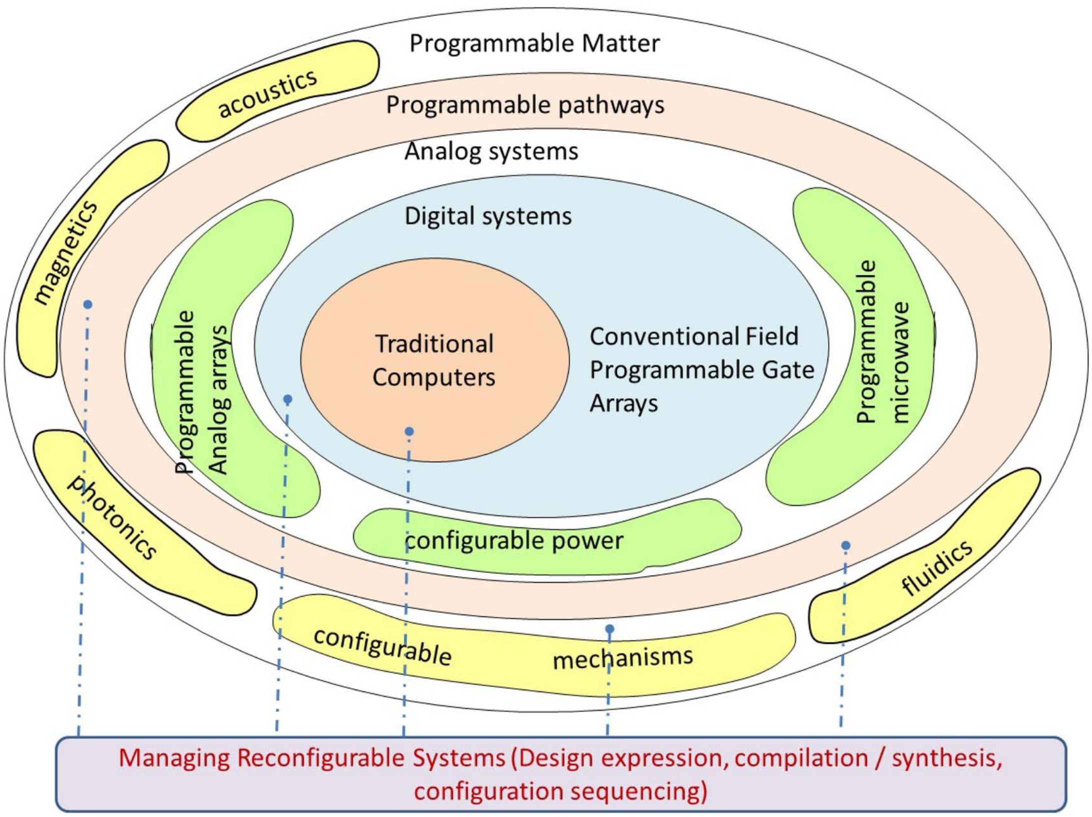

Fig. 1. *A Venn diagram of prospective reconfigurable systems.*

Understanding reconfigurability involves understanding what the building blocks might be and how these blocks can be composed and managed on very large scales to form programmable devices, circuits, and systems. The beginnings of modern reconfigurable systems trace their origins to early stored program computers, cellular automata concepts, and programmable logic arrays. The building blocks, originally bits of logic, memory, and switches to interconnect them, have grown in sophistication over time from dozens of primitive elements to millions. Contemporary reconfigurable systems constitute complex fabrics, properly called "systems on a chip" having a variety of functional units that in some cases are beginning to include analog and mixed-signal functions. Empowered mostly through Moore's Law, their sophistication level has grown well beyond the ability for unaided humans to harness them. In most cases, they are managed through automated synthesis tools and design languages to serve ubiquitously in an ever-widening range of consumer, industrial, and military applications.

The flexibility and power of reconfigurable devices have given rise to a number of interesting and useful applications, and we hear the term "software-defined" being applied as an adjective to an increasing number of things: radios, networks, storage, data centers, and the list continues. Rather than locking the circuitry patterns of particular radio designs with fixed arrangements of primitive circuit elements, it is now possible to render much of the circuitry in software form, leading to dramatically more capable protocols and dynamic architectures. In many ways, this revolution, which is going by names such as "cognitive radio," is even today only in its earliest stages. Similarly, new advances in software-defined networking are building upon FPGA performance and flexibility, further empowering the significant trends in cloud computing, benefitting not just from virtualizing machines, but also the networks that they connect to.

But backing the ability to deliver on "software-defined anything" requires extending reconfigurability in new domains. We can imagine extending the ideas of FPGAs into other realms (as suggested in the notional Fig. 1 Venn diagram), to include fine-grain analog and microwave circuitry, antennas, photonics, pathways and mechanisms, and ultimately, programmable matter. What will be the FPGAs of the future? Fusing the influences of materials, nanoscale processing and engineering, and systems engineering is expected to lead to breakthrough possibilities of reconfigurability. These include systems that are more profoundly flexible and adaptive, systems capable of self-organizing and self-healing in response to the whims of designers and users. We consider issues, such as tools for achieving effective productivity in these new systems. For examples, if we program computers with C++ and Java and FPGAs with VHDL and Verilog, what are the new concepts of language and reconfigurable expression in programmable matter?

II. PART I: FOUNDATIONS

The three themes that we will address in the first special issue on reconfigurable systems (the second issue is planned for May 2015) include digital reconfigurability, tools for reconfigurable systems, and configurable radio technologies. The first paper of this special issue, "Introduction to Reconfigurable Systems," developed by the guest editorial team, will attempt to provide a broad overarching framework for what we believe to be an emerging discipline of reconfigurable systems. We discuss some of the key defining qualities of reconfigurable systems and examine two broad facets: designing *for* reconfigurability and designing *with* reconfigurability. In the former case, we are concerned with the disciplines of creating reconfigurable versions of traditional devices and systems, which involves engineering soft-defined features and establishing architectures to manage these engineered degrees of freedom. In the second case, we consider how one would take such systems and exploit their intrinsic flexibility. Ultimately, the designers include, not just the platform developers, but the users themselves.

The next four papers focus on the field programmable gate array and reconfigurable computation. The FPGA, entering now its fourth decade of existence, is the quintessentially flexible digital system, with the ability to shape at will, pools of memory, logic, and interconnects to create virtually any digital system—even entire computers—on a chip. FPGAs went from a curiosity to a fact of life over the years to become one the most ubiquitous components in the microelectronics industry today. As such, it serves as an archetype and reference model by which other reconfigurable systems can be patterned. In fact, it may well be that the FPGA in its own evolution eventually becomes the penultimate reconfigurable system by absorbing additional functions, features, and even entire domains (such as analog, microwave, power, and optical) into an overall framework.

In "Three Ages of FPGAs" (the second paper of the special issue), Xilinx's Steve Trimberger reflects on how technology evolution has driven the design of FPGAs—reconfigurable digital technology—through three epochs: the age of invention, the age of expansion, and the age of accumulation, showing how reconfigurability moved from an "overhead" cost to an economic advantage as costs and demands shifted with technology. The history of the FPGA is a fascinating story, but the future will be even more exciting.

The next two papers, "Reconfigurable Computing Architectures" and "Fundamental Underpinnings of Reconfigurable Computing Architectures," consider how reconfigurable digital technologies have blurred boundaries of hardware and software. Reconfigurable computing offers a broader interpretation of architectural design space, in which reconfigurability, through flexibility and dynamicism, may offer additional benefits to gain performance and/or improve the energy efficiency to a growing class of real-world problems. In the first of these two papers, Tessier *et al.* provide an overview of the broad body-of-knowledge developed in the field, while in the second, DeHon presents a unified framework for understanding the advantages and tradeoffs in reconfigurable computing architectures.

In the next paper of this issue, "High-Reliability FPGA-Based Systems: Space, High-Energy Physics, and Beyond," Wirthlin highlights some of the challenges in creating reliable reconfigurable systems. The space environment provide special challenges, since the tiny patterns of charge that underlie all digital functions can be disrupted by energetic particles and cosmic rays that naturally occur in space. We believe this paper will provide readers the opportunity to reflect on some potential challenges in creating more robust reconfigurable systems.

The role of tools for managing reconfigurability will be as important as the technology itself. We address this topic in a dedicated paper on high-level synthesis (HLS) for reconfigurable computing (and discuss later a second paper on tools for software radio), but the subject is also featured in a number of other articles in the special issue. In "High-Level Language Tools for Reconfigurable Computing," Windh *et al.* provide a focused survey of five tools (Xilinx Vivaldo, Altera OpenCL, Bluespec System Verilog, UC Riverside's Optimizing Compiler for Configurable Computing, and University of Toronto's LegUp) for enhanced productivity code development for FPGAs.

While software radio can be viewed as a "killer application" for reconfigurability, the challenges of harnessing the full potential of software radio will be daunting. Four papers will address the relation between reconfigurable systems technologies and radio applications. In "Software-Defined Radio: Bridging the Analog-Digital Divide," Machado and Wyglinski describe the evolution of software radio. The next two papers span reconfigurability from the RF front end through the signal chain. Costantine *et al.*, in the paper "Reconfigurable Antennas: Design and Applications," discuss design, manufacture, and control concepts for antennas that employ programmable features. Rais-Zadeh *et al.* examine potential technologies in their paper ("Reconfigurable Radios: A Possible Solution to Reduce Entry Costs in Wireless Phones") that could be used to implement fully reconfigurable radios and also evaluate the ability of such radios to reduce spectrum acquisition entry costs for new competitors in wireless service.

We reprise the subject of tools again, this time focused on the radio domain, in a paper by Reed *et al.*, titled "Software Frameworks for SDR." Their paper discusses how recent advances in software frameworks (in particular, REDHAWK) allow for more reconfigurability and gluing together of reconfigurable components to make a system.

III. PART 2: ADVANCED APPLICATIONS AND TECHNOLOGIES FOR RECONFIGURABLE SYSTEMS

In the second special issue (later this year), we will examine other important applications of reconfigurability in another potential "killer app" domain: software-definable networking (wired and wireless). We hope to illuminate other extensions in the performance of digital reconfigurability. We will also focus on more advanced concepts, namely those pertaining to analog systems, photonics, metamaterials, with some commentary on the possibilities of programmable matter. ■

ABOUT THE GUEST EDITORS

James Lyke (Senior Member, IEEE) received the B.S. degree in electrical engineering at the University of Tennessee, Knoxville, TN, USA in 1984, the M.S. degree in electrical engineering at the Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, USA in 1989, and the Ph.D. degree in electrical engineering from University of New Mexico, Albuquerque, NM, USA in 2004.

He was in active duty miltary service with the U.S. Air Force from 1984 through 1995. Since 1990, he has supported the Air Force Research Laboratory (AFRL), Space Vehicles Directorate (AFRL/RV), Kirtland Air Force Base, NM, USA, including its precursor organizations (Weapons Laboratory, 1990–1991, and Phillips Laboratory, 1991–1998), in a number of capacities. He is currently technical advisor to the AFRL Space Electronics Branch (Space Vehicles Directorate) and an AFRL Fellow since 2008. He has lead over one hundred in-house and contract research efforts involving two- and three-dimensional advanced packaging, radiation-hardened microelectronics, and scalable, reconfigurable computational and systems architectures, with recent emphasis on modularity and the rapid formation of complex systems. He has authored over 100 publications (journal and conference papers, book chapters, and technical reports), receiving four best paper awards, and he has been awarded 11 U.S. patents.

Dr. Lyke is an Associate Fellow of the American Institute of Aeronautics and Astronautics (AIAA) and serves on the AIAA Computer Systems Technical Committee. He was selected as recipient of the Federal Laboratory Consortium award for Excellence in Technology Transfer in 1992, and for the U.S. Air Force Science and Engineering Award in Exploratory and Advanced Technology Development in 1997 and 2000.

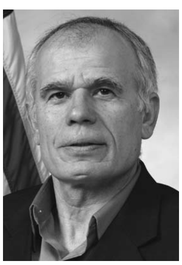

Christos G. Christodoulou (Fellow, IEEE) received the Ph.D. degree in electrical engineering from North Carolina State University, Raleigh, NC, USA in 1985.

From 1985 to 1998, he served as a faculty member in the University of Central Florida, Orlando. In 1999, he joined the faculty of the Electrical and Computer Engineering Department of the University of New Mexico (UNM), Albuquerque, NM, USA, where he served as the Chair of the Department from 1999 to 2005. He is a member of Commission B of URSI and a Distinguished Professor at UNM. Currently, he is the Associate Dean of Research for the School of Engineering at UNM, and the director for COSMIAC (Configurable Space Microsystems Innovations & Applications Center) at UNM. He was appointed as an IEEE AP-S Distinguished Lecturer (2007–2010) and served as an associate editor for the IEEE Transactions on Antennas and Propagation for six years, as a guest editor for a special issue on "Applications of Neural Networks in Electromagnetics" in the Applied Computational Electromagnetics Society (ACES) journal, the co-editor of a the IEEE Antennas and Propagation Special issue on "Synthesis and Optimization Techniques in Electromagnetics and Antenna System Design" (March 2007), and the guest editor for a Special issue on "Antenna Systems and Propagation for Cognitive Radio" in 2014. He has also given keynote and invited talks in several places in Asia, Europe, Latin America and the United States. He served as the major advisor for 25 Ph.D. candidates and 70 M.S. students. He published about 500 papers in journals and conferences, has 14 book chapters, and has coauthored six books. His research interests are in the areas of modeling of electromagnetic systems, cognitive radio, machine learning in electromagnetics, high-power microwave antennas, and reconfigurable antennas for cognitive radio.

Dr. Christodoulou is the recipient of the 2010 IEEE John Krauss Antenna Award for his work on reconfigurable fractal antennas using MEMS switches, the Lawton-Ellis Award, and the Gardner Zemke Professorship at the University of New Mexico.

Alonzo Vera (Senior Member, IEEE) is currently a member of the technical staff at the Configurable Space Microsystems Innovations Applications Center (COSMIAC) at the University of New Mexico, Albuquerque, NM, USA, where he and his group work on the design of small spacecraft for a broad range of missions. At COSMIAC, Dr. Vera has focussed on embedded system design for aerospace applications, dynamic partial reconfiguration applications using FPGAs, radiation effects mitigation techniques for FPGA-based systems and custom ASICs as well as radiation effects testing. He has numerous publications and has also been an invited lecturer on digital design and digital signal processing using FPGAs in numerous countries. His current areas of interest are reconfigurable computing, embedded system design for aerospace applications and digital signal processing.

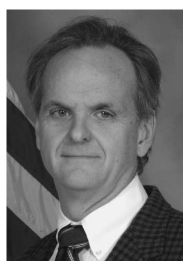

Art H. Edwards is a senior research physicist in the Space Vehicles Directorate of the Air Force Research Laboratory. He received all of his degrees from Lehigh University, including the Ph.D. degree in physics in 1981, with an emphasis in solid-state theory. He has worked in industry (Westinghouse Electric), academia (University of North Carolina at Charlotte), and government (Army and Air Force Research Laboratories). He has over 60 publications in several areas including electronic structure of defects in insulators and semiconductors, device physics, and physics of memristor materials and mechanisms.