Guest Editors' Introduction

Jean-Luc Gaudiot, *Fellow*, *IEEE*, and Fabrizio Lombardi, *Member*, *IEEE*

1 Introduction

This special issue of *IEEE Transactions on Computers* is devoted to Configurable Computing. Configurable computing has experienced a radical change throughout the last few years; from the early efforts to achieve "adaptiveness" under very restrictive conditions, configurable computing has emerged as a significant area of research and development for both the academic and industrial communities. This tremendous growth in the design and manufacturing of configurable systems is partly the result of advances in fundamental research and technology. Technological advances have allowed us to place on a single integrated circuit the computing power equivalent to early supercomputers and to completement that hardware with efficient software techniques; for example, a better understanding of the hardware/software issues involved has resulted in significant changes in the technology platform provided by field programmable gate arrays (FPGAs).

While the field of adaptive computing has been active for the last decade, the specific focus on configurable systems has taken on an increased importance due to the rapid expansion of the market for these products. Differently from the past, configurable systems are not used in a standalone configuration. Today's configurable systems are typically placed within a larger overall application, often performing real-time and high performance tasks. Example applications include digital signal processing (sensing, imaging, and filtering), computationally intensive tasks (arithmetic and embedded environments), personal communications devices (wireless phones), consumer electronics (printers, data storage), industrial control, and others.

There are many technical challenges facing the designers of configurable systems: these systems are truly complex in the integration of hardware and software resources. It is not appropriate to consider the various elements separately because the overall performance, modeling, and organization will depend on the interactions occurring between all of the system components. For example, the ability to meet a hard, real-time deadline will depend on the performance of the programmable hardware and on the software's use of that hardware, together with the timely response to the changing requirements of the applications. The ability to design a correctly functioning system depends on our ability to design and analyze software/hardware technologies and to properly account for the interactions between them in an embedded fashion.

Unfortunately, design environments and computer-aided design tools have not yet successfully integrated the technologies needed to design and analyze a complete configurable system. Existing techniques often depend on significant knowledge of the specific hardware/software implementations and tradeoffs. The increasing complexity of integrated circuits and the increasing use of FPGAs and ASICs hardware and supporting CAD software have meant that designers seldom have the detailed implementation knowledge needed to efficiently design these systems. Also, design approaches that depend on implementation specifics are not easily portable from one application to the next, so most reconfiguration techniques become both application and implementation dependent.

Design remains one of the most significant challenges for the configurable system community. In the early days, we assumed that software and hardware were two separate domains, each with its own set of problems. Now, the complexity of the system is such that its design must employ hardware and software in perfect synchronization and harmony to attain the quantum leap in performance demanded by many applications. In fact, the existing design paradigm is gradually making hardware and software indistinguishable such that integration takes a totally different role. The lack of sufficient analysis techniques and tools to support a correct modeling of configurable systems makes both their design and evaluation a formidable task. Communications links (including hardware and software components), real-time profiles, coordination and allocation of resources, and organization of application software provide an integrated hardware/software system with numerous possibilities. While the term configurable computing is becoming a reality, its complexity management is still a daunting task.

Configurable systems are also often very cost sensitive. The increase in cost of a few dollars is often crucial when the application involves the sale of millions of units. Cost sensitivity and high performance are often at odds with one another. Many of the elegant solutions for configurable systems have been rejected by industry because of the significant cost associated with them. Highly sophisticated architectures are a superb solution to a very difficult technical problem, but few applications can afford the features (such as cost and power) associated with these solutions. We must develop techniques that can provide efficient solutions to the practical configuration arena.

- *J.-L. Gaudiot is with the Department of Electrical Engineering-Systems, University of Southern California, Los Angeles, CA 90089-2563. E-mail: gaudiot@usc.edu.*
- *F. Lombardi is with the Department of Electrical and Computer Engineering, Northeastern University, Boston, MA 02115. E-mail: lombardi@ece.neu.edu.*

For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number 109958.

It is impossible to cover all of the technical issues in this one special issue, but we have selected seven papers that cover a spectrum of technologies and applications. The papers are authored by individuals from both industry and universities, and cover issues associated with hardware, software, design environments, partitioning and scheduling, and chip design.

Our first paper "Static and Dynamic Configurable Systems" authored by E. Sanchez, M. Sipper, J.-O. Haenni, J.-L. Beuchat, A. Stauffer, and A. Perez-Uribe reports on the general classification of configurable systems. Starting with the two common modes of configurability (static and dynamic), various systems which utilizes unique features are presented. These include FPGA-based back propagation neural networks, biological-based implementations capable of handling incomplete and/or modified specifications, and a novel configurable board that can be incorporated into an educational curriculum. This paper highlights the diversity and multiplicity of systems and applications to which configurable computing can be used. This is important for appreciating the paradigm involving programmable and configurable entities.

The second paper ("Partitioning Sequential Circuits on Dynamically Reconfigurable FPGAs" by D. Chang and M. Marek-Sadowska) examines the CAD topic of partitioning. This paper presents a new gate-level model that permits time-multiplexed computation on a dynamically reconfigurable FPGA. This is possible by utilizing a so-called enhanced force directed scheduling algorithm which partitions the original sequential circuit to minimize logic and communication overhead (subject to enhanced performance). The advantage of this approach is the capability to model communication between nonadjacent on-chip configurations and multiple levels of logic.

The issue of partitioning is also addressed in the third paper, "Temporal Partitioning and Scheduling Data Flow Graphs for Reconfigurable Computers" by K.M. Gajala Purna and D. Bhatia. This paper presents a technique for partitioning and scheduling data flow graphs for configurable computing. This is required to accommodate scalability by appropriately generating the temporal mappings and reusing the hardware. Both logic and temporal requirements are considered and examined such that a complete hardware environment can be realized through CAD tools for processing and synthesis. An intermediate format and a communication protocol for sequentially executing partitions in the time domain are also proposed. Its application to logic emulation is described.

The title of our fourth paper is "Dynamic Reconfiguration to Support Concurrent Applications," by J. Jean, K. Tomko, V. Yavagal, J. Shah, and R. Cook. The issue of multiple applications running concurrently is examined. This is made possible through the design of a FPGA resource manager to enable efficient management and pre-loading of the resources. The resource manager is developed based on application flow graphs; this relieves developers from the management of FPGA resources. Evaluation through extensive simulation supports the conclusion that the overhead associated with using the manager is compensated by a large speed-up in application execution due to concurrency.

The fifth paper, "Synthesis of Application Specific Instructions for Embedded DSP Software" by H. Choi, J.-S. Kim, C.-W. Yoon, I.-C. Park, S.H. Hwang, and C.-M. Kyung, describes a modified subset-sum-based approach for generating application specific instructions for digital signal processing. The issue of multicycle complex instructions is considered for use in an ASIP (application specific instruction processors). This approach results in enhanced parallelism, thus achieving higher performance in both theoretical and practical sized problems (such as those having multipaths). Code size is also reduced, while meeting timing constraints.

The sixth paper "A Scan-based Configurable, Programmable and Scalable Architecture for Sliding Window Based Operations" is authored by C. Thibeault and G. Begin. This practical architecture is applicable to many applications in which programmable signal processing with high bandwidth and real-time performance are required. In particular, the sliding-window operations are dealt with by using many examples involving convolutional encoding and finite impulse response filtering. This architecture is based on a loop of processing elements controlled by a set of circulating tokens which identify the processing of incoming data.

The last paper in this special issue, "The GRD Chip: Genetic Reconfiguration of DSPs for Neural Network Processing" by M. Murakawa, S. Yoshizawa, I. Kajitani, X. Yao, N. Kajihara, M. Iwata, and T. Higuchi, describes the application of genetic reconfiguration for DSP (GRD); this process is made possible by embedding it directly onto a chip. This chip is particularly suited to industrial applications in which time-varying problems and real-time constraints must be addressed. This chip does not require a host to realize autonomous configuration of hardware structures and on-line performance.

We sincerely hope that you enjoy this special issue. The topics covered in the papers are timely and important, and the authors have done an excellent job of presenting the material. We extend our sincere thanks to all the authors and reviewers. We also thank Prof. Jane Liu, former editor-in-chief of IEEE TC, for allowing us to create this special issue. Finally, a special thanks is due to all the staff for editing and assembling this issue. Please feel free to contact us if you have questions or comments.

Jean Luc Gaudiot, University of Southern California
Fabrizio Lombardi, Northeastern University

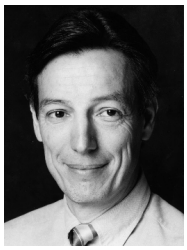

Jean-Luc Gaudiot received the Diplome d'Ingénieur from the Ecole Supérieure d'Ingénieurs en Electrotechnique et Electronique, Paris, France, in 1976, and the MS and PhD degrees in computer science from the University of California, Los Angeles, in 1977 and 1982, respectively. Dr. Gaudiot's experience includes microprocessor systems design at Teledyne Controls, Santa Monica, California (1979-1980) and research in innovative architectures for the TRW Technology Research Center, El Segundo, California (1980-1982). Since receiving his PhD in 1982, he has been a member of the faculty of the Department of Electrical Engineering-Systems at the University of Southern California, where he is currently a professor. His research interests include data-flow architectures, fault-tolerant multiprocessors, and implementation of artificial neural networks. In addition to his academic duties, he has consulted for several aerospace companies in the Southern California area.

Dr. Gaudiot is a member of the ACM, the ACM SIGARCH, the IFIP Working Group 10.3 (Parallel Processing), and a fellow of the IEEE. He was the general chairman of the 1992 International Symposium on Computer Architecture, the program chairman for the 1993 IFIP Working Conference on Architectures and Compilation Techniques for Fine and Medium Grain Parallelism, and the systems track program chairman for the 1993 IEEE Symposium on Parallel and Distributed Processing. He is the editor-in-chief of the *IEEE Transactions on Computers* and is currently an Advisory Board member of the IEEE Technical Committee on Computer Architecture.

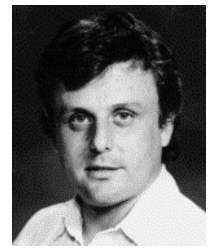

Fabrizio Lombardi graduated in 1977 from the University of Essex (UK) with a B.Sc. (Hons.) in electronic engineering. In 1977, he joined the Microwave Research Unit at University College London, where he received the Master's in microwaves and modern optics (1978), the Diploma in microwave engineering (1978), and the PhD from the University of London (1982). He is currently the chairperson of the Department of Electrical and Computer Engineering and holder of the International Test Conference (ITC) Endowed Professorship at Northeastern University, Boston. He was previously a faculty member at Texas Tech University, the University of Colorado-Boulder, and Texas A&M University. He has received many professional awards: the visiting fellowship at the British Columbia Advanced System Institute, University of Victoria, Canada (1988), twice the TEES Research Fellowship (1991-1992, 1997-1998) the Halliburton Professorship (1995), and an International Research Award from the Ministry of Science and Education of Japan (1993-1999).

Dr. Lombardi was the recipient of the 1985/86 Research Initiation Award from the IEEE/Engineering Foundation, a Silver Quill Award from Motorola-Austin (1996), and a distinguished visitor of the IEEE Computer Society for the period 1990-1993. He is an editor of the *IEEE Transactions on Computers*. He has been involved in organizing many international symposia, conferences, and workshops sponsored by organizations such as NATO and IEEE. His research interests are fault-tolerant computing, testing and design of digital systems, configurable computing, defect tolerance, and CAD VLSI. He has published extensively in these area and edited six books. He is a member of the IEEE.