Editorial

Nader Bagherzadeh

Department of Electrical Engineering and Computer Science,
University of California,
Irvine, CA 92697-2625, USA
E-mail: nader@uci.edu

Biographical notes: Nader Bagherzadeh received his BS and MS in Electrical Engineering and a PhD degree in Computer Engineering from the University of Texas at Austin in 1977, 1979 and 1987, respectively. From 1980 to 1984, he was with AT&T Bell Laboratories in Holmdel New Jersey. Since 1988, he has been with the Department of Electrical Engineering and Computer Science at the University of California, Irvine. He was the Chair of his department from 1998 to 2002. His research interests are in low-power and embedded digital signal processing, computer architecture, computer graphics and VLSI design.

The advance computer architecture arena has gone through several changes in the past 30 years. One of the most recent changes is the notion of reconfigurable computer systems. With the advent of Field Programmable Gate Array (FPGA) the idea of having a 'soft' hardware took a new meaning. Since the first introduction of FPGAs, computer architects have been busy improving this technology by introducing new and more efficient configurable blocks, provide better design tools and more efficient interconnects to meet the tight power and performance requirement. Additionally, researchers have embarked on yet a new generation of reconfigurable architectures that are coarse grained as compared to FPGAs. These new architectures provide an alternative approach to FPGAs by fixing the data path and the functionality of configurable blocks throughout the design, delivering close to hardwired designs (ASICs) performance and power consumption.

One of the key issues for the success of reconfigurable high performance architectures is how to design the software and algorithms to fully take advantage of these machines. This has always been a challenge for new and exotic architectures in the past, and reconfigurable designs are not immune from it either. The desirable goal has always been to provide a seamless path to take existing sequential code written in C or other popular languages and provide efficient parallel codes for these architectures. We have a long way to go before achieving this goal but the success of these architectures depends on it. The researchers have realised the importance of software tools and parallel algorithms in order for these high performance architectures become ubiquitous.

With the increasing reliance on multicore architectures to address the computing needs of our next generation systems, scalability of high performance designs has become an important topic of discussion. Current multicore designs rely heavily on a bus central design approach that has been the main stream design concept for decades. Unfortunately, Systems-on-Chip (SoC) with hundreds of cores will not be able to scale effectively if one relies on a bus-based design. The communication among cores becomes a bottleneck and the technology parameters for submicron fabrications alleviate the problem further because long wires delay become the dominant factor when compared to the gate delay. It is projected that in 10 years the difference between gate delay and wire delay will be more than three orders of magnitude, suggesting that having shorter wires such as the ones advocated for Network-on-Chip (NoC) a promising approach as compared to bus-based designs.

For this Special Issue of the *International Journal of High Performance Systems Architecture* we have selected seven top quality papers that were presented this year at the ITNG conference in Las Vegas. All of these papers directly or indirectly address the notion of reconfigurable and scalable high performance architectures. We are grateful to the authors for preparing an extended version of their conference paper for this special issue. Next, we will briefly introduce these papers.

In the paper titled: A flexible processor for the characteristic 3 η_T pairing, the hardware implementation of the η_T pairing on a super-singular elliptic curve of characteristic 3 is described. All characteristic 3 operations required for the computation of the pairing are outlined in detail. The required extension field operations can be performed in terms of subfield operations, many of which can be computed in parallel using hardware. The hardware architectures required for pairing computation are also introduced. An efficient and reconfigurable processor utilising these hardware architectures is presented and discussed.

In the paper titled: Reconfiguration support for vector operations, a programmable vector processor and its FPGA implementation are presented. As a specific benchmark for this vector processor, *W*-matrix sparse solver for linear equations is utilised. Using the FPGA implementation, actual power matrices with up to 1723 nodes can be realised with less than 1.1 ms. A performance comparison with ordinary PC shows

that the vector processor is competitive for such computation-intensive problems.

In the paper titled: Design of a router for network-on-chip, a new NoC architecture with a minimal adaptive router and associated packet protocols are introduced. By using their cycle-accurate System-C simulator and comparing with other routing algorithms, its performance is shown to be competitive. Finally, the prototype design of this NoC architecture demonstrates its hardware feasibility for multicore architecture interconnection network.

In the paper titled: Efficient finite field processor for GF (2^{163}) and its VLSI implementation, a high performance architecture elliptic curve processor is proposed. The modified Bit-Parallel Word-Serial (BPWS) finite field multiplication algorithm and its pipelined multiplier architecture with a throughput of one multiplication of $N+1$ cycles is discussed. Also by exploring parallelism at the instruction level and supporting separate hardware modules for multiplication, squaring and addition, up to three arithmetic operations can be executed in parallel. Finally, the processor is implemented on TSMC 0.18 μm CMOS technology and reported to be able to perform a scalar multiplication over GF (2^{163}) in 62 μs.

In the paper titled: An efficient fault tolerant mechanism to deal with permanent and transient failures in a network on Chip, a comprehensive fault tolerant mechanism for packet-based NoCs is proposed to handle packet losses or corruption due to transient faults. Also a dynamic routing mechanism is presented to tolerate permanent link and/or router failures on chip. By simulation, the proposed mechanism is shown to provide a graceful degradation of performance in the presence of transient and permanent faults.

In the paper titled: Compact FPGA-based systolic array architecture suitable for vision systems, the architecture for highly computation intensive block searching for a video codec application is presented. This architecture which utilises efficient memory accesses and parallelism achieves a peak performance of 9 GOPS.

In the paper titled: SPA resistant elliptic curve cryptosystem using addition chains, a reconfigurable architecture for cryptographic system is presented. The design uses a Simple Power Analysis (SPA) technique which is shown to be faster than double-and-add method.