Editorial

High-Performance Reconfigurable Computing

Khaled Benkrid,¹ Esam El-Araby,² Miaoqing Huang,³ Kentaro Sano,⁴ and Thomas Steinke⁵

¹ School of Engineering, The University of Edinburgh, Edinburgh EH9 3JL, UK

² Electrical Engineering and Computer Science, The Catholic University of America, Washington, DC 20064, USA

³ Department of Computer Science and Computer Engineering, University of Arkansas, Fayetteville, AR 72701, USA

⁴ Graduate School of Information Sciences, Tohoku University, 6-6-01 Aramaki Aza Aoba, Sendai 980-8579, Japan

⁵ Zuse-Institut Berlin (ZIB), Takustraße 7, 14195 Berlin-Dahlem, Germany

Correspondence should be addressed to Khaled Benkrid, k.benkrid@ed.ac.uk

Received 28 February 2012; Accepted 28 February 2012

Copyright © 2012 Khaled Benkrid et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

This special issue presents some of the latest developments in the burgeoning area of high-performance reconfigurable computing (HPRC) which aims to harness the high-performance and low power of reconfigurable hardware in the forms of field programmable gate arrays (FPGAs) in high-performance computing (HPC) applications.

The issue starts with three widely popular HPC applications, namely, financial computing, bioinformatics and computational biology, and high-throughput data search. First, Starke et al. from the University of Kiel in Germany present the use of a massively parallel FPGA platform, called RIVYERA, in the high-performance and low-power optimization of investment strategies. The authors demonstrate an FPGA-based implementation of an investment strategy algorithm that considerably outperforms a single CPU node in terms of raw processing power and energy efficiency. Furthermore, it is shown that the implemented optimized investment strategy outperforms a buy-and-hold strategy. Then, Vanderbauwhede et al. from Glasgow University and the University of Massachusetts Lowell propose a design for the scoring part of a high-throughput real-time search application on FPGAs. The authors use a low-latency Bloom filter to realize high-performance information filtering. An analytical model of the application throughput is built around the Bloom filter. The experimental results on the Novo-G reconfigurable supercomputer demonstrate a 20× speedup compared with a software implementation on a 3.4 GHz Intel Core i7 processor. After that, Eusse et al. from the University of Brasilia and the Federal University of Mato Grosso do Sul in Brazil present an FPGA-accelerated

protein sequence analysis solution. The authors integrate the concept of divergence to the Viterbi algorithm used in the HMMER program suite for biological sequence alignment, in order to reduce the area of the score matrix in which the trace-back alignment is made. This technique leads to large speedups (182×) compared to nonaccelerated pure software processing.

The issue then presents a number of architectural concerns in the design of HPRC systems, namely: reconfigurable hardware architecture, communication network design, and arithmetic design. First, Wan et al. from the University of Illinois at Urbana Champaign and Magma Design Automation Inc. in the USA present a coarse-grained reconfigurable architecture (CGRA) with a Fast Data Relay (FDR) mechanism to enhance its performance. This is achieved through multicycle data transmission concurrent with computation, and effective reduction of communication traffic congestion. The authors also propose compiler techniques to efficiently utilize the FDR feature. The experimental results for various multimedia applications show that FDR combined with the new compiler delivers up to 29% and 21% higher performance than other CGRAs: ADRES and RCP, respectively. The following paper by Schmidt et al. from the University of Southern California and the University of North Carolina present an integrated on-chip/off-chip network with MPI-style point-to-point message, implemented on an all-FPGA computing cluster. The most salient differences between the network architecture presented in this paper and state-of-the-art Network-on-Chip (NoC) architectures is the use of a single full-crossbar switch. The results are different from

*Miaoqing Huang
Kentaro Sano
Thomas Steinke*

other efforts due to several reasons. First, the implementation target is the programmable logic of an FPGA. Second, most NoCs assume that a full crossbar is too expensive in terms of resources while within the context of an HPRC system the programmable logic resources are fungible. The authors justify their focus on the network performance by the fact that overall performance is limited by the bandwidth off the chip rather than by the mere number of compute cores on the chip. After that, El-Araby et al. from the Catholic University of America, Universidad Autonoma de Madrid, and the George Washington University present a technique for the acceleration of arbitrary-precision arithmetic on HPRC systems. Efficient support of arbitrary-precision arithmetic in very large science and engineering simulations is particularly important as numerical nonrobustness becomes increasingly an issue in such applications. Today's solutions for arbitrary-precision arithmetic are usually implemented in software and performance is significantly reduced as a result. In order to reduce this performance gap, the paper investigates the acceleration of arbitrary-precision arithmetic on HPRC systems.

The special issue ends with two papers which present reconfigurable hardware in HPC in the context of other computer technologies. First, Benkrid et al. from the University of Edinburgh, Scotland, and the University of Arizona, USA, present a comparative study of FPGAs, Graphics Processing Units (GPUs), IBM's Cell BE, and General Purpose Processors (GPPs) in the design and implementation of a biological sequence alignment application. Using speed, energy consumption, in addition to purchase and development costs, as comparison criteria, the authors argue that FPGAs are high-performance economic solutions for sequence alignment applications. In general, however, they argue that FPGAs need to achieve at least two orders of magnitude speedup compared to GPPs and one order of magnitude speedup compared to GPUs to justify their relatively longer development times and higher purchase costs. The following paper by Inta et al. from the Australian National University presents an off-the-shelf CPU, GPU, and FPGA heterogeneous computing platform, called Chimera, as a potential high-performance economic solution for certain HPC applications. Motivated by computational demands in the area of astronomy, the authors propose the Chimera platform as a viable alternative for many common computationally bound problems. Advantages and challenges of migrating applications to such heterogeneous platforms are discussed by using demonstrator applications such as Monte Carlo integration and normalized cross-correlation. The authors show that the most significant bottleneck in multidevice computational pipelines is the communications interconnect.

We hope that this special issue will serve as an introduction to those who have newly joined, or are interested in joining, the HPRC research community as well as provide specialists with a sample of the latest developments in this exciting research area.

*Khaled Benkrid
Esam El-Araby*