John von Neumann Institute for Computing

Parallel Computing with FPGAs – Concepts and Applications

Erik H. D'Hollander, Dirk Stroobandt, Abdellah Touhafi

published in

Parallel Computing: Architectures, Algorithms and Applications ,
C. Bischof, M. Bücker, P. Gibbon, G.R. Joubert, T. Lippert, B. Mohr, F. Peters (Eds.),
John von Neumann Institute for Computing, Jülich,
NIC Series, Vol. **38**, ISBN 978-3-9810843-4-4, pp. 739-740, 2007.

http://www.fz-juelich.de/nic-series/volume38

Parallel Computing with FPGAs - Concepts and Applications

Erik H. D'Hollander[1], **Dirk Stroobandt**[1], **and Abdellah Touhafi**[2]

[1] Ghent University
Electronics and Information Systems Department
Parallel Information Systems, B-9000 Ghent, Belgium
E-mail: {Erik.DHollander, Dirk.Stroobandt}@UGent.be

[2] Brussels University Association, Department IWT
B-1070 Brussels, Belgium
E-mail: atouhafi@info.vub.ac.be

The Mini-Symposium "Parallel computing with FPGAs" aimed at exploring the many ways in which field programmable gate arrays can be arranged into high-performance computing blocks. Examples include high-speed operations obtained by sheer parallelism, numerical algorithms mapped into hardware, co-processing time critical sections and the development of powerful programming environments for hardware software co-design.

Introduction

The idea to launch a mini-symposium on parallel computing with FPGAs, was inspired by the need to explore the huge performance potential which can be tapped from the tiny computing blocks called field programmable gate arrays. Their features are so flexible and reconfigurable that they are capable of massively parallel operations, explicitly tailored to the problem at hand. That said, there have been a lot of paradigms to put FPGAs at work in a high performance computing environment. We are all beginning to see the new and exciting possibilities of reconfigurable computing. Because a new idea is as good as its successful application, we found that it is in the tradition of the ParCo parallel computing conferences to focus on application oriented solutions. This has led to seven interesting papers, which have been presented in this symposium.

Contributions

The paper *Parallel Computing with Low-Cost FPGAs - A Framework for COPACOBANA* by Tim Güneysu, Christoph Paar, Jan Pelzl, Gerd Pfeiffer, Manfred Schimmler and Christian Schleiffer, describes a novel extensible framework of clusters of FPGAs, geared at high-performance computing. A communication library is used to configure up to 120 FPGAs simultaneously and the system is operated from a host computer. Applications in the area of cryptanalysis show the potential of the system when compared to high-cost alternatives.

The following paper *Accelerating the Cube Cut Problem with an FPGA-augmented Compute Cluster* by Tobias Schumacher, Enno Lübbers, Paul Kaufmann and Marco Platzner, employ FPGAs to speed up the bit-operations of a compute intensive exhaustive

search problem. Using parallel compute nodes each equipped with FPGAs, the authors obtain speedups of respectively 27 and 105 on 1 and 4 processors, compared to the same computations done on 1 and 4 processors without FPGAs.

A Cache Subsystem Supporting Run-time Reconfigurability by Fabian Nowak, Rainer Buchty and Wolfgang Karl, addresses performance from a cache point of view. Instead of optimizing code and data locality for a particular type of cache, a reconfigurable system is presented which optimizes the cache for a particular type of locality. During the execution the cache can be adapted to the characteristics of the different program phases. The idea is to change the associativity, the number of lines and the replacement strategy, without flushing the cache.

In *A Brain Derived Vision System Accelerated by FPGAs* by Jeff Furlong, Andrew Felch, Jayram Moorkanikara Nageswaran, Nikil Dutt, Alex Nicolau, Alex Veidenbaum, Ashok Chandrashekar and Richard Granger, a highly parallel neural model is used to overcome the limited parallelism in most programs due to sequential code. Using a winner-take-all competition between competing neurons, a massively parallel FPGA system is put in place which is able to outperform a general-purpose CPU by more than an order of magnitude.

Accelerating digital signal processing by FPGAs is studied in *Programmable Architectures for Realtime Music Decompression* by Martin Botteck, Holger Blume, Jörg von Livonius, Martin Neuenhahn and Tobias G. Noll. The paper analyzes the efficiency gained by using FPGAs for decoding MP3 streams. A pipelined implementation of the decoder gives good results, but at the same time it is shown that the power consumption of the solution is too heavy for portable devices.

High-level programming of FPGAs is a subject of *The HARWEST High Level Synthesis Flow to Design an FPGA-Based Special-Purpose Architecture to Simulate the 3D Ising Model* by Alessandro Marongiu and Paolo Palazzari. ANSI type C programs are analyzed and converted into a control and data flow graph, which are further converted into a data path and a control finite state machine. This approach is applied to the 3-D Ising spin glass model.

Web crawlers have the huge task to correlate and rank web pages. This application is dominated by a sparse matrix times vector multiplication. In the paper *Towards an FPGA Solver for the PageRank Eigenvector Problem* by Séamas McGettrick, Dermot Geraghty and Ciarán McElroy , it is shown that after pipelining and parallelizing, the computations can be mapped onto an FPGA accelerator. Reordering the data structure of the matrix allows the accelerator to outperform the PC, even when the FPGA clock is about 10 times slower.

Conclusion

FPGAs offer a number of paradigms to speed up calculations in a hardware software co-design environment. Creativity and innovation is needed to exploit all avenues and select promising and effective solutions. Trade-offs are necessary between competing goals such as portability, power consumption, performance and communication. This mini-symposium has shown that in these various areas successful ideas and implementations are obtainable. However, much work remains to be done to integrate these efforts into a framework unifying FPGAs with high-performance parallel computing.