Reconfigurable High Performance Computing the Next Disruptive Innovation

Reiner Hartenstein
TU Kaiserslautern
http://hartenstein.de

Abstract.

Already years ago, Reconfigurable Computing (RC) has become mainstream in embedded systems, going practically into every application area. FPGAs (Field-Programmable Gate Arrays) are the fastest growing segment of the semiconductor market, currently with a market volume of about 6 billion US-Dollars. But also the much more area-efficient coarse-grained reconfigurability based on arrays of rDPUs (reconfigurable DataPath Units, featuring word level path width, such as 32 bits, for instance) is commercially available from a number of start-ups.

In all application segments of embedded systems a software to configware migration is an accelerator approach, often providing speed-up factors by one or two orders of magnitude, in fewer cases even up to three or four orders of magnitude. Migration to configware has also been used successfully for a substantial reduction of power dissipation. Because of the flexibility of its commodity platforms, configware avoids the high design cost and manufacturing cost of hardwired accelerators. This is illustrated by the rapidly increasing percentage of FPGA-based design starts compared to ASIC use. This presentation briefly surveys the area of reconfigurable platforms, their applications, algorithms, compilation environments and operating systems.

Reconfigurable platform use means a second RAM-based programming paradigm: structural programming instead of procedural programming. Configware code can be downloaded to the (hidden) configuration RAM - even over the internet. Why do supercomputing scenes hesitate to go reconfigurable? The challenges for 2005 - 2010 are [1]: new compilers, middleware and programming model changes, and others. Exploring the architectural space will be expensive [1], but is highly promising [2].

The talk briefly surveys differences to embedded systems: what algorithms to be implemented (there is a substantial overlap), what preferred languages and programming models, differences in scalability, compilation techniques and operating systems, what additional qualifications needed, and other differences. The talk discusses the impact on supercomputing: highly promising features vs. problems which stem from the software / configware chasm - a clash of cultures: instruction-stream-based versus processing without instruction fetch at run time. Reconfigurable-only is not feasible for supercomputing. But an instruction-stream-based-only roadmap will lead to stalled progress as well. A symbiosis of both cultures is viable, which has been successfully demonstrated by the area of embedded systems, as well as by an increasing number of publications reporting successful FPGA use to accelerate typical supercomputing applications.

References

[1] H. Simon (keynote): Progress in Supercomputing: the Top Three Breakthroughs of the Last 20 and the Top Three Challenges for the Next 20 years; Proc. ISC 2005, Heidelberg, Germany, June 22 - 24, 2005;

[2] R. Hartenstein (keynote): Software or Configware? About the Digital Divide of Computing; IPDPS 2004, Santa Fe, NM, USA, April 2004