The Re-definition of Low Power Digital System Design - for slashing the Electricity Bill by Millions of Dollars

opening keynote, SBCCI 2006

Reiner Hartenstein, TU Kaiserslautern
http://hartenstein.de

Abstract. Reconfigurable Computing (RC) went mainstream in embedded systems already years ago. Overwhelmingly high hit rates by Google illustrate the pervasiveness of FPGA and RC applications. More recently RC rapidly obtains acceptance within scientific computing and even supercomputing. Cray and Silicon Graphics have launched products including FPGAs used as accelerators. Some application areas report speedup factors by up to three orders of magnitude (and sometimes even more). This is astonishing, since the clock frequency of FPGAs is substantially lower than the speed of microprocessors and the effective integration density on FPGAs is about four orders of magnitude behind Gordon Moore's law. Here algorithmic cleverness is the secret of success on the way to new frontiers in supercomputing. This is based on software to configware migration mechanisms, striving away from the traditional memory-cycle-hungry instruction-stream-based computing paradigm.

FPGA applications for scientific computing and supercomputing report similar high speedup factors. A side effect is the reduction of floor space requirements and of the electricity invoice amount by about one order of magnitude. This is important, since a major obstacle on the classical road map toward petaFLOP computing is the enormous electric power requirement of such monster machines, and the difficulty to find an electric power company being ready to deliver such dimensions.

This is a new aspect for those groups in Low Power Design community, traditionally aiming at designing nomadic devices running with a tiny battery as the only power supply. But in the new direction of high performance computing, technology and circuit level methodologies are not so much the key issues. Instead, algorithmic cleverness for software to configware migration is needed, which means a paradigm shift. Extending into the area of supercomputer architectures, Low Power Design needs to be re-defined to meet the goals of national and global significance: the energy consumption of our globe.

It's an important question, wether from the spectacular success stories in high performance computing we can also learn for designing energy-efficient servers. E. g., Google's annual electricity bill is 50,000,000 $, and 25% of the electricity consumption of Amsterdam goes into server farms. The US server market is expected to grow from 2.8 million units in 2005 to 4.9 million units in 2009. Mid July this year the US house of representatives approved a bill for studying the use of energy-efficient servers.

A main hurdle are educational deficits, also causing the configware / software chasm. Experts with a typical CS background usually need extensive support by EE type colleagues, coming along with communication problems. A taxonomy of algorithms for design space exploration in software to configware migration is completely missing. Scanning the more than 500 pages of the 2004 curriculum recommendations of the ACM / IEEE-CS / AIS Joint Task Force for Computing Curricula, we find zero encounters of the term „FPGA" and all its synonyms, also of RC. This is criminal. These recommendations miss the most important part of the IT-based job market. To-day, by far most cumulative MIPS equivalents running worldwide have been migrated to configware.