Travel Report from ASLPOS IX and PACS'00

http://foothill.lcs.mit.edu/asplos2k/
http://dynamo.ecn.purdue.edu/~pacs00/

November 21, 2000

Flavius Gruian
Lund Institute of Technology
<Flavius.Gruian@cs.lth.se>

The 9th conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS IX) took place in Cambridge, Massachusetts, from November 13 to 15. Two full-day workshops were held one day before the conference, namely the 2nd Workshop on Intelligent Memory Systems and the Workshop on Power Aware Computing Systems (PACS'00). The purpose of my visit was, at first, to present a paper (*System-Level Design Methods for Low-Energy Architectures Containing Variable Voltage Processors*) at the latter. Second, an event such as ASPLOS is a good place to get the latest news in microarchitectural decisions, and not only.

The trip started with a pleasant surprise in the first day, when I realized that there were many participants interested in my work (Currently we are planning a longer term collaboration with one of the research groups I came in contact there). The following day begun to unwind in a hectic but fruitful pace with presentations involving MEMS (micro-electromechanical memory systems, which seem to be the next big step in the storage technology), FLASH, and network design. One presentation from Berkeley, on networked sensors (*System Architecture Directions for Networked Sensors*, by J.Hill & al.) was particularly interesting. The speaker sent working sensors into the audience, while showing on-screen how the network reconfigured itself as sensor were moving around or switched on/off.

The high-point of the next day was the session on Wild-and-Crazy Ideas, where most were centered on the low fabrication cost microelectronics will exhibit in the near future. More precisely, instead of Design-Test-Build paradigm, we'll see a shift towards a Build-Measure-Design approach. Pretty much as in FPGAs today, but at a higher level of reconfigurability. Although off the RT community interests, there was another "visionary" presentation from Berkeley's J.Kubiatowicz, *OceanStore: An Architecture for Global-Scale Persistent Storage* (Take a look at http://oceanstore.cs.berkeley.edu).

Finally, the last day revolved around speeding-up computations using application characteristic behavior. Intel and HP presented a joint paper on IA-64 (*OS and Compiler Considerations in the Design of the IA-64 Architecture*), which supports both data and control speculation, but leaves the compiler and the OS to decide when and how these should be pursued. Another interesting concept, connected to speculative execution, came from the North Carolina State University and says "two redundant programs combined run faster than either can alone." In principle, using a two-way microprocessor chip, two instances of the same program run simultaneously with a slight delay - one gathering profile information and skipping some code, the other patching up if things go wrong.

As a conclusion, there were presented a whole bunch of architectural features, very much application and data specific, that can lead to performance improvements. Some of these features are just improvements of older approaches, (e.g. using a LRU algorithm for cache-line power down) others were entirely new concepts. The impact of these on the RT-community is the need for more complex time estimators, because of their highly dynamic behavior.

One of the eye-striking facts at the conference was the large number of participants from industry (Sun, Intel, Transmeta, Compaq, HP-labs). Many of the papers presented work carried out during internships at IBM, Intel, etc. (Internship is something missing from the swedish industry, unfortunately.) Furthermore, not only were two dinners sponsored by HP and Compaq, but there was quite a lot of active recruiting going on all the time.

All in all, it was a nice experience....And now I can also say “I’ve been to Harvard.”