Topic 8
Parallel Computer Architecture and Instruction-Level Parallelism

Stamatis Vassiliadis, Nikitas Dimopoulos, Jean-Francois Collard, and Arndt Bode
Topic Chairs

Parallel computer architecture and instruction-level parallelism are hot topics at Euro-Par conferences, since these techniques are present in most contemporary computing systems. At Euro-Par 2003, 18 papers were submitted to the topic, from which 1 distinguished, 4 regular and 4 short papers were accepted. The scope of this topic includes but is not limited to parallel computer architectures, processor architecture (architecture and micro architecture as well as compilation), the impact of emerging microprocessor architectures on parallel computer architectures, innovative memory designs to hide and reduce the excess latency, multi-threading, and impact of emerging applications on parallel computer architecture design.

The distinguished paper by George Almasi et al. gives an overview of the Blue Gene/L System Software Organization. With its 360 Teraflops of peak computing power and 65536 compute nodes it is a special purpose architecture well beyond the first entries of the TOP 500 list. The authors explain, how the system software will deal with this extreme amount of parallelism. The paper is therefore general for the future developments of parallel computer architectures.

Trace substitution is the topic of Hans Vandierendonck, Hans Logie and Koen De Bosschere from Gent University, Belgium. This technique is useful for wide-issue superscalar processors and their ability to react on branches in the instruction stream. The authors show that their new technique improves the fetch bandwidth of such processors.

Counteracting bank misprediction in sliced first-level caches again deals with possible cache misses and is presented by E.F. Torres, P. Ibanez, V. Vinals, J.M. Llaberia. The authors discuss techniques to reduce the bank misprediction penalty in future processors having sliced memory pipelines.

Juan Moure, Dolores Rexachs, Emilio Luque from the Computer Architecture and Operating Systems Group of the Universidad Autónoma de Barcelona present a mechanism to optimize a decoupled front-end architecture: the indexed fetch target buffer (iFTB). Once again, wide issue superscalar processors are addressed and an indexed variation of the fetch target buffer, a decoupled front-end architecture is presented to increase the fetch rate of such processors.

Seong-Won Lee from the University of Southern California at Los Angeles and Jean-Luc Gaudiot from the University of California at Irvine present the technique of clustered microarchitecuture simultaneous multithreading. SMT is

H. Kosch, L. Böszörményi, H. Hellwagner (Eds.): Euro-Par 2003, LNCS 2790, pp. 541–542, 2003.

favourable for future generation microprocessors, since it exploits instruction-level parallelism and thread-level parallelism. The clustered SMT architecture proposed by the authors significantly reduces power consumption without performance degradation.

Spiros Kalogeropulos from Sun proposes a global instruction scheduling technique addressing processors with moderate support for parallelism. The enhanced trace scheduler includes techniques for trace formation, a renaming scheme and a cost benefit analysis.

Masamichi Takagi and Kei Hiraki from the Department of Computer Science, University of Tokyo, propose compression in data caches with compressible field isolation for recursive data structures. The field array compression technique tends to alleviate the penalty for cache misses in program execution.

Hideyuki Miura et al. from the University of Tokyo discuss a compiler-assisted thread level control speculation for speculative multithreading execution. Two new techniques are presented.

The last paper of this topic again deals with caches in microprocessors. In this paper, Carles Aliagas et al. address the reduction of power to run the cache and show with their paper on "value compression to reduce power in data caches" the way to power reduction and reduction in die area without a performance penalty.