International Journal of Parallel Programming, Vol. 24, No. 3, 1996

Part I: Special Issue on Parallel Architectures and Compilation Techniques

Welcome to part one of this special issue on Parallel Architectures and Compilation Techniques. The papers you will find here are a compendium of some (by no means the only) of the most relevant and strongest papers which were presented at the *Conference on Parallel Architectures and Compilation Techniques* (*PACT '95*). A word about *PACT '95*: it met in Cyprus, in June 1995, and it was the third in a series of meetings which were founded with a unique goal in mind: to provide a forum for the interaction of researchers working at the precise boundary between architecture and compilers for parallel machines.

Now back to this special issue: the papers it contains are a good sample of the topics covered by the conference as they range from analytical modeling to parallel compilation techniques, to the design of parallel architectures.

The first paper, "Using predicated execution to improve the performance of a dynamically scheduled machine with speculative execution" by P-Y. Chang, E. Hao, Y. Patt, and P. Chang, recognizes that two mechanisms for handling conditional branches, namely speculative execution and predicated execution, have very different characteristics and uses, and shows how a combination of the two can be used to advantage in wide-issue, deeply pipelined processors.

In the second paper, "A mean value analysis multiprocessor model incorporating superscalar processors and latency-tolerating techniques," D. Albonesi and I. Koren have demonstrated a Mean Value Analysis model to evaluate the tradeoffs involved in the design of multiprocessors whose nodes are individual superscalar processors.

How to generate code to enumerate the integer points of a polyhedron is the topic of the third paper, "A simple algorithm for the generation of efficient loop structures," by M. Cosnard and M. Loi. It is a known important

828/24/3-1

problem in the design of parallelizing compilers, and they demonstrate an algorithm, the Hierarchical Domain Descriptor (HDD), which provides for the efficient generation of loop structures.

In the second part of this special issue, "A partitioning-independent paradigm for nested data parallelism" by D. Engelhardt and A. Wendelborn attacks the problem of generalizing the data parallelism model of execution by an extension to nested and irregular data structures. Since the extension requires a new model of execution, they show how a multithreaded abstract machine can be used to express nested data parallelism. Their theory is corroborated by experimental observations made on a Thinking Machine CM-5.

The next paper, "A study of the EARTH-MANNA multithreaded system," by H. Hum, O. Maquelin, K. Theobald, X. Tian, G. Gao, and L. Hendren is a detailed description of the EARTH multithreaded multiprocessor system which has been designed at McGill University (Canada) and of its emulation on the MANNA 2.0, an existing multiprocessor. The results of the experiments point to the viability of the multithreaded model of execution both in terms of latency tolerance for multiprocessor environments and in terms of single processor performance.

In the paper, "Memory referencing behavior in compiler-parallelized applications," E. Torrie, M. Martonosi, M. Hall, and C-W. Tseng have concentrated on evaluating the influence of memory design on the performance of typical applications (such as can be found amongst SPEC or NAS, etc. benchmarks). They have shown that granularity was important (coarser means better), that wide cache lines would weaken the performance, and that an appropriate design of the compiler (for instance, by introducing advanced analyses and optimizations such as array privatization and interprocedural analysis) could reduce some of these detrimental effects. These results can have implications both for the architects and the compiler designers.

Finally, in the paper "An empirical evaluation of the CONVEX SPP-1000 hierarchical shared memory system," T. Sterling, D. Savarese, P. Merkey, and K. Olson evaluate the performance of the Convex SPP-1000 hierarchical cache coherency mechanisms. They demonstrate the high scalability of the implementation of the N-body problem, notorious for its complexity and difficulty of efficient parallel implementation.

I would like to thank all of the referees for this two part special issue as well as the PACT '95 Program Committee members for their efforts: their detailed comments have helped substantially improve the papers. I would also like to thank the authors for bearing with my demands and for dutifully revising the papers according to the referees' comments. All this work has resulted in this special issue which I hope you will find useful.

Jean-Luc Gaudiot
University of Southern California