From the Guest Editors

It is with great pleasure that we introduce this two-part special issue on Instruction-Level Parallel Processing. With the entire computer industry continuing to turn out hardware and software products based on instruction-level parallel processing technology, this is a particularly timely special issue on the subject.

The papers that appear in this two-part special issue were chosen based on the recommendations of the Program Committee of the *27th Annual ACM/IEEE International Symposium on Microarchitecture*. Six outstanding papers from the symposium were selected, and the authors invited to revise the papers to meet journal quality standards. All the authors responded with enthusiasm and hard work, and the result is the distinguished collection of six papers presented in this two-part special issue.

In Part I, Rau's paper, "Iterative Modulo Scheduling," describes a practical algorithm for module scheduling that is capable of dealing with realistic program and machine constraints. The author clearly defines the algorithm, characterizes its output schedule quality, and shows its computation complexity.

The second paper in Part I (Volume 24, Number 1), "Parallelization of Control Recurrences for ILP Processors" by Schlansker *et al.* focuses on the problem of speeding up the execution of loops containing control recurrences and data dependent exits. The authors show that their technique, when combined with data dependence height reduction techniques, provides a comprehensive approach to accelerating loops with conditional exits.

In Part II (Volume 24, Number 2), "Minimizing Register Requirements of Modulo Schedule via Optimum Stage Scheduling" by Eichenberger *et al.* presents a technique to minimize the register requirements of a given modulo reservation table. The authors show that their approach can result in a significant reduction in register usage compared to a register insensitive modulo scheduler.

The paper by Chang *et al.*, "Branch Classification: A New Mechanism for Improving Branch Predictor Performance," introduces a method of

classifying branches such that a hybrid branch predictor can be used to significantly increase predication accuracy.

“Evaluating the Effect of Predicated Execution on Branch Prediction” by Tyson and Farrens investigates the sources of mispredictions for a variety of branch predictors and the effect of predicated execution on these poorly predicted branches.

The final paper, “Hardware-Based Profiling: An Effective Technique for Profile-driven Optimization” by Conte *et al.*, introduces the concept of using existing hardware features in current high performance processors to conduct very low overhead program execution profiling. Such low overhead profiling can enable a much wider usage of profile-driven optimization on large production programs.

We would especially like to thank all the authors for their extensive efforts to incorporate the (sometimes substantial) modifications suggested by the reviewers. It is the teamwork of the authors and the reviewers that make this two-part special issue a particularly worthwhile contribution to the field of instruction-level parallel processing.

Guest Editors

Matthew Farrens, University of California, Davis

Wen-mei Hwu, University of Illinois at Urbana-Champaign