International Journal of Parallel Programming, Vol. 24, No. 6, 1996

Guest Editors' Introduction

This special issue contains enchanced versions of five papers selected from those presented at the *Eight International Workshop on Languages and Compilers for Parallel Computing*, which was held at Columbus, Ohio, on August 10–12, 1995.

The paper "Array Data Flow Analysis for Load-Store Optimizations in Fine-Grain Architectures," by Bodík and Gupta, presents an analysis framework to solve various dataflow problems in the context of array load-store optimizations. An algorithm is developed for the optimal placement of load-store operations for array references in loops, which minimizes the dynamic number of memory accesses caused by the array references.

In their paper "Interprocedural Array Region Analyses," Creusillet and Irigoin describe an approach to exact array region analysis. A general linear framework is used for interprocedural analysis, which handles array reshapes. The intra- and inter-procedural propagation of array regions is implemented in their interprocedural parallelizer PIPS.

Ghiya and Hendren present a paper on "Connection Analysis: A Practical Interprocedural Heap Analysis for C." The technique is designed to analyze programs that allocate many disjoint objects such as dynamically allocated arrays on the heap. The presented method has been implemented in the McCAT system. Experimental results compare the accuracy of connection analysis with a conservative estimate based on points-to analysis.

Kelly *et al.* present a paper on the "Transitive Closure of Infinite Graphs and Its Applications." Predicates with parameterized affine constraints are used as a closed-form description of integer tuple relations. It is shown that the transitive closure of such relations is not computable in the general case. Algorithms producing exact results for commonly occurring cases, and upper of lower bounds for other cases are developed. Applications of the transitive closure algorithm are provided.

The paper "Efficient Distribution Analysis via Graph Contraction," by Sheffler *et al.*, addresses the problem of efficient automatic determination of optimal array distributions. Because it is a difficult combinatorial optimization problem, the authors propose the use of graph contraction to reduce

828/24/6-1

the size of the distribution problem, which can be formulated as a graph labeling problem. Experimental results are provided that demonstrate the effectiveness of the approach.

We offer our sincere thanks to the referees of this special issue for providing timely reviews of the papers, and to the authors for working with a tight time schedule for the special issue.

Chua-Huang Huang and P. Sadayappan
Guest Editors for the Special Issue