

Lecture Notes in Computer Science

Edited by G. Goos, J. Hartmanis and J. van Leeuwen

1863

Springer

Larry Carter Jeanne Ferrante (Eds.)

Languages and Compilers for Parallel Computing

12th International Workshop, LCPC'99
La Jolla, CA, USA, August 4-6, 1999
Proceedings



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany
Juris Hartmanis, Cornell University, NY, USA
Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editors

Larry Carter
Jeanne Ferrante
University of California, San Diego
Department of Computer Science
9500 Gilman Drive, La Jolla, CA 92093-0114, USA
E-mail: {carter, ferrante}@cs.ucsd.edu

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Languages and compilers for parallel computing : 12th international workshop ; proceedings / LCPC '99, La Jolla, CA, USA, August 4 - 6, 1999. Larry Carter ; Jeanne Ferrante (ed.). - Berlin ; Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ; Paris ; Singapore ; Tokyo : Springer, 2000
(Lecture notes in computer science ; Vol. 1863)
ISBN 3-540-67858-1

CR Subject Classification (1998): D.3, D.1.3, F.1.2, B.2.1, C.2

ISSN 0302-9743

ISBN 3-540-67858-1 Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York
a member of BertelsmannSpringer Science+Business Media GmbH
© Springer-Verlag Berlin Heidelberg 2000
Printed in Germany

Typesetting: Camera-ready by author, data conversion by PTP-Berlin, Stefan Sossna
Printed on acid-free paper SPIN: 10722248 06/3142 5 4 3 2 1 0

Preface

In August 1999, the Twelfth Workshop on Languages and Compilers for Parallel Computing (LCPC) was hosted by the Hierarchical Tiling Research group from the Computer Science and Engineering Department at the University of California San Diego (UCSD). The workshop is an annual international forum for leading research groups to present their current research activities and the latest results. It has also been a place for researchers and practitioners to interact closely and exchange ideas about future directions. Among the topics of interest to the workshop are language features, code generation, debugging, optimization, communication and distributed shared memory libraries, distributed object systems, resource management systems, integration of compiler and runtime systems, irregular and dynamic applications, and performance evaluation. In 1999, the workshop was held at the International Relations/Pacific Studies Auditorium and the San Diego Supercomputer Center at UCSD. Seventy-seven researchers from Australia, England, France, Germany, Korea, Spain, and the United States attended the workshop, an increase of over 50% from 1998.

The program committee of LCPC 99 along with external reviewers as needed, were responsible for evaluating the submitted papers. Forty-eight papers were submitted, and of those, twenty-seven were selected to be presented as full papers at the workshop. In addition, thirteen submissions were presented as posters in a special poster session. Using feedback provided both before and after the presentations, all authors were given the opportunity to improve their papers before submitting the final versions contained in this volume. Short abstracts of the poster presentations are also included.

In addition to the paper and poster sessions, LCPC 99 also featured an invited talk by Burton Smith, Chief Scientist at Tera Computer (now renamed Cray) on “Optimization for the Tera MTA”. This talk gave an overview of the MTA architecture and the program transformations in the MTA compiler that allow it to take advantage of the MTA’s unique architectural characteristics. The home of the first Tera MTA is the San Diego Supercomputer Center, and a tour of SDSC, including the Tera MTA, was offered to all participants. We gratefully thank Burton Smith for his excellent presentation and for his full participation in the workshop.

Another feature of this year’s workshop was a panel session on Benchmarking organized by Keshav Pingali. This session grew out of a seminar on Tiling for Optimal Resource Allocation (hosted by the International Conference and Research Center for Computer Science at Schloss Dagstuhl in 1998) in which setting up a suite of benchmarks for locality that could be used by the general community was proposed. The panel, whose members also included Rudi Eigenmann, David Padua, and Sanjay Rajopadhye, presented a lively and diverse discussion on the merits of such a suite.

The organizers wish to acknowledge the San Diego Supercomputer Center and UCSD for their help. In particular, the conference was organized by Joann Pagan of UCSD Conference Services, with help from Nancy Jensen at SDSC. We especially wish to thank the software support staff, particularly Cindy Paloma, and graduate students Kang Su Gatlin, Karin Hogstedt, Beth Simon, and Michelle Mills Strout, all of the Computer Science and Engineering Department, for their excellent help. We also wish to acknowledge the great help of Chanathip Nampremre in editing and putting together this volume.

We would like to give special thanks to the LCPC09 program committee and the nameless external reviewers for their efforts in reviewing the submissions. Both the steering committee and the program committee helped with advice and suggestions on the organization of the workshop. Finally, we wish to thank all of the participants who helped to create a lively and constructive atmosphere of discussion, and the authors for sharing their significant research with us at LCPC 09.

May 2000

Larry Carter, Jeanne Ferrante
Program Chair
LCPC09

Organization

Program Committee

Program Chair: Larry Carter (University of California, San Diego, USA)
Jeanne Ferrante (University of California, San Diego, USA)

General Chair: Larry Carter (University of California, San Diego, USA)
Jeanne Ferrante (University of California, San Diego, USA)

Program Committee: Larry Carter (University of California, San Diego, USA)
Jeanne Ferrante (University of California, San Diego, USA)
Manish Gupta (IBM Research, USA)
Zhiyuan Li (Purdue University, USA)
Sam Midki (IBM Research, USA)
Jose Moreira (IBM Research, USA)
Jan Prins (University of North Carolina at Chapel Hill, USA)
Pen-Chung Yew (University Minnesota, USA)

Panel

Keshav Pingali (Cornell University, USA) (organizer)
Rudi Eigenmann (Purdue University, USA)
David Padua (University of Illinois at Urbana Champaign, USA)
Sanjay Rajopadhy (IRISA, France)

Steering Committee

Utpal Banerjee (Intel Corporation, USA)
Alex Nicolau (University of California, Irvine, USA)
David Gelernter (Yale University, USA)
David Padua (University of Illinois at Urbana Champaign, USA)

Sponsoring Institutions

San Diego Supercomputer Center, La Jolla, CA USA
University of California, San Diego, La Jolla, CA USA

Table of Contents

Java

High Performance Numerical Computing in Java: Language and Compiler Issues	1
<i>Pedro V. Artigas, Manish Gupta, Samuel P. Midkiff, and José E. Moreira</i>	

Instruction Scheduling in the Presence of Java's Runtime Exceptions	18
<i>Matthew Arnold, Michael Hsiao, Ulrich Kremer, and Barbara Ryder</i>	

Dependence Analysis for Java	35
<i>Craig Chambers, Igor Pechtchanski, Vivek Sarkar, Mauricio J. Serrano, and Harini Srinivasan</i>	

Low-Level Transformations A

Comprehensive Redundant Load Elimination for the IA-64 Architecture ..	53
<i>Youngfeng Wu and Yong-fong Lee</i>	

Minimum Register Instruction Scheduling: A New Approach for Dynamic Instruction Issue Processors	70
<i>R. Govindarajan, Chihong Zhang, and Guang R. Gao</i>	

Unroll-Based Copy Elimination for Enhanced Pipeline Scheduling	85
<i>Suhyun Kim, Soo-Mook Moon, Jinpyo Park, and HanSaem Yun</i>	

Data Distribution

A Linear Algebra Formulation for Optimising Replication in Data Parallel Programs	100
<i>Olav Beckmann and Paul H.J. Kelly</i>	

Accurate Data and Context Management in Message-Passing Programs ...	117
<i>Dhruva R. Chakrabarti and Prithviraj Banerjee</i>	

An Automatic Iteration/Data Distribution Method Based on Access Descriptors for DSM.	133
<i>Angeles G. Navarro and Emilio L. Zapata</i>	

High-Level Transformations

Inter-array Data Regrouping	149
<i>Chen Ding and Ken Kennedy</i>	

Iteration Space Slicing for Locality	164
<i>William Pugh and Evan Rosser</i>	

A Compiler Framework for Tiling Imperfectly-Nested Loops	185
<i>Yonghong Song and Zhiyuan Li</i>	

Models

Parallel Programming with Interacting Processes	201
<i>Peiyi Tang and Yoichi Muraoka</i>	

Application of the Polytope Model to Functional Programs	219
<i>Nils Ellmenreich, Christian Lengauer, and Martin Griebel</i>	

Multilingual Debugging Support for Data-Driven and Thread-Based Parallel Languages	236
<i>Parthasarathy Ramachandran and Laxmikant V. Kale'</i>	

Array Analysis

An Analytical Comparison of the I-Test and Omega Test	251
<i>David Niedzielski and Kleantes Psarris</i>	

The Access Region Test	271
<i>Jay Hoeflinger and Yunheung Paek</i>	

A Precise Fixpoint Reaching Definition Analysis for Arrays	286
<i>Jean-François Collard and Martin Griebel</i>	

Demand-Driven Interprocedural Array Property Analysis	303
<i>Yuan Lin and David Padua</i>	

Language Support

Language Support for Pipelining Wavefront Computations	318
<i>Bradford L. Chamberlain, E. Christopher Lewis, and Lawrence Snyder</i>	

The Data Mover: A Machine-Independent Abstraction for Managing Customized Data Motion	333
<i>Scott B. Baden and Stephen J. Fink</i>	

Optimization of Memory Usage Requirement for a Class of Loops Implementing Multi-dimensional Integrals	350
<i>Chi-Chung Lam, Daniel Cociorva, Gerald Baumgartner, and P. Sadayappan</i>	

Compiler Design and Cost Analysis

Compile-Time Based Performance Prediction	365
<i>Calin Cascaval, Luiz DeRose, David A. Padua, and Daniel A. Reed</i>	

Designing the Agassiz Compiler for Concurrent Multithreaded Architectures	380
<i>B. Zheng, J. Y. Tsai, B. Y. Zang, T. Chen, B. Huang, J. H. Li, Y. H. Ding, J. Liang, Y. Zhen, P. C. Yew, and C. Q. Zhu</i>	

The Scc Compiler: SWARing at MMX and 3DNow!	399
<i>Randall J. Fisher and Henry G. Dietz</i>	

Low-Level Transformation B

Loop Shifting for Loop Compaction	415
<i>Alain Darté and Guillaume Huard</i>	

Speculative Predication Across Arbitrary Interprocedural Control Flow . . .	432
<i>Hank G. Dietz</i>	

Posters

Porting an Ocean Code to MPI Using TSF	447
<i>F. Bodin, Y. Mével, S. Chauveau, and E. Rohou</i>	

A Geometric Semantics for Program Representation in the Polytope Model	451
<i>Brian J. d'Auriol</i>	

Compiler and Run-Time Support for Improving Locality in Scientific Codes	455
<i>Hwansoo Han, Gabriel Rivera, and Chau-Wen Tseng</i>	

Code Restructuring for Improving Real Time Response through Code Speed, Size Trade-offs on Limited Memory Embedded DSPs	459
<i>Vipin Jain, Siddharth Rele, Santosh Pande, and J. Ramanujam</i>	

Compiling for Speculative Architectures	464
<i>Seon Wook Kim and Rudolf Eigenmann</i>	

Symbolic Analysis in the PROMIS Compiler	468
<i>Nicholas Stavrakos, Steven Carroll, Hideki Saito, Constantine Polychronopoulos, and Alex Nicolau</i>	

Data I/O Minimization for Loops on Limited Onchip Memory Processors .	472
<i>Lei Wang and Santosh Pande</i>	

Time Skewing for Parallel Computers	477
<i>David Wonnacott</i>	

Run-Time Parallelization Optimization Techniques	481
<i>Hao Yu and Lawrence Rauchwerger</i>	

Thresholding for Work Distribution of Recursive, Multithreaded Functions	485
<i>Gary M. Zoppetti, Gagan Agrawal, and Lori L. Pollock</i>	

An Empirical Study of Function Pointers Using SPEC Benchmarks 490
Ben-Chung Cheng and Wen-mei W. Hwu

Data Driven Graph: A Parallel Program Model for Scheduling 494
V. D. Tran, L. Hluchy, and G. T. Nguyen

Author Index 499