

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

George Almási Călin Caşcaval
Peng Wu (Eds.)

Languages and Compilers for Parallel Computing

19th International Workshop, LCPC 2006
New Orleans, LA, USA, November 2-4, 2006
Revised Papers

Volume Editors

George Almási
Călin Cașcaval
Peng Wu

IBM Research Division
Thomas J. Watson Research Center
Yorktown Heights, New York 10598
E-mail: {gheorghe, cascaval, pengwu}@us.ibm.com

Library of Congress Control Number: 2007926757

CR Subject Classification (1998): D.3, D.1.3, F.1.2, B.2.1, C.2.4, C.2, E.1, D.4

LNCS Sublibrary: SL 1 – Theoretical Computer Science and General Issues

ISSN 0302-9743
ISBN-10 3-540-72520-2 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-72520-6 Springer 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. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2007
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12063598 06/3180 5 4 3 2 1 0

Preface

The 19th Workshop on Languages and Compilers for Parallel Computing was held in November 2006 in New Orleans, Louisiana USA. More than 40 researchers from around the world gathered together to present their latest results and to exchange ideas on topics ranging from parallel programming models, code generation, compilation techniques, parallel data structure and parallel execution models, to register allocation and memory management in parallel environments.

Out of the 49 paper submissions, the Program Committee, with the help of external reviewers, selected 24 papers for presentation at the workshop. Each paper had at least three reviews and was extensively discussed in the committee meeting. The papers were presented in 30-minute sessions at the workshop. One of the selected papers, while still included in the proceedings, was not presented because of an unfortunate visa problem that prevented the authors from attending the workshop.

We were fortunate to have two outstanding keynote addresses at LCPC 2006, both from UC Berkeley. Kathy Yelick presented “Compilation Techniques for Partitioned Global Address Space Languages.” In this keynote she discussed the issues in developing programming models for large-scale parallel machines and clusters, and how PGAS languages compare to languages emerging from the DARPA HPCS program. She also presented compiler analysis and optimization techniques developed in the context of UPC and Titanium source-to-source compilers for parallel program and communication optimizations.

David Patterson’s keynote focused on the “Berkeley View: A New Framework and a New Platform for Parallel Research.” He summarized trends in architecture design and application development and he discussed how these will affect the process of developing system software for parallel machines, including compilers and libraries. He also presented the Research Accelerator for Multiple Processors (RAMP), an effort to develop a flexible, scalable and economical FPGA-based platform for parallel architecture and programming systems research. Summaries and slides of the keynotes and the program are available from the workshop Web site <http://www.lcpcworkshop.org>.

The success of the LCPC 2006 workshop would not have been possible without help from many people. We would like to thank the Program Committee members for their time and effort in reviewing papers. We wish to thank Gerald Baumgartner, J. Ramanujam, and P. Sadayappan for being wonderful hosts. The LCPC Steering Committee, especially David Padua, provided continuous support and encouragement. And finally, we would like to thank all the authors who submitted papers to LCPC 2006.

March 2007

Gheorghe Almási
Călin Caşcaval
Peng Wu

Organization

Steering Committee

Utpal Banerjee	Intel Corporation
David Gelernter	Yale University
Alex Nicolau	University of California, Irvine
David Padua	University of Illinois, Urbana-Champaign

Organizing Committee

Program Co-chairs	Gheorghe Almási, IBM Research Călin Cașcaval, IBM Research Peng Wu, IBM Research
Local Co-chairs	Gerald Baumgartner, Louisiana State University J. Ramanujam, Louisiana State University P. Sadayappan, Ohio State University

Program Committee

Vikram Adve	University of Illinois at Urbana-Champaign
Gheorghe Almási	IBM Research
Eduard Ayguad	Universitat de Politècnica de Catalunya
Gerald Baumgartner	Louisiana State University
Călin Cașcaval	IBM Research
Rudolf Eigenmann	Purdue University
Maria-Jesus Garzaran	University of Illinois at Urbana-Champaign
Zhiyuan Li	Purdue University
Sam Midkiff	Purdue University
Paul Petersen	Intel Corp.
J. Ramanujam	Louisiana State University
P. Sadayappan	Ohio State University
Peng Wu	IBM Research

Table of Contents

Keynote I

Compilation Techniques for Partitioned Global Address Space Languages	1
<i>Kathy Yelick</i>	

Session 1: Programming Models

Can Transactions Enhance Parallel Programs?	2
<i>Troy A. Johnson, Sang-Ik Lee, Seung-Jai Min, and Rudolf Eigenmann</i>	
Design and Use of htalib – A Library for Hierarchically Tiled Arrays ...	17
<i>Ganesh Bikshandi, Jia Guo, Christoph von Praun, Gabriel Tanase, Basilio B. Fraguera, María J. Garzarán, David Padua, and Lawrence Rauchwerger</i>	
SP@CE - An SP-Based Programming Model for Consumer Electronics Streaming Applications	33
<i>Ana Lucia Varbanescu, Maik Nijhuis, Arturo González-Escribano, Henk Sips, Herbert Bos, and Henri Bal</i>	

Session 2: Code Generation

Data Pipeline Optimization for Shared Memory Multiple-SIMD Architecture	49
<i>Weihua Zhang, Tao Bao, Binyu Zang, and Chuanqi Zhu</i>	
Dependence-Based Code Generation for a CELL Processor	64
<i>Yuan Zhao and Ken Kennedy</i>	
Expression and Loop Libraries for High-Performance Code Synthesis ...	80
<i>Christopher Mueller and Andrew Lumsdaine</i>	
Applying Code Specialization to FFT Libraries for Integral Parameters	96
<i>Minhaj Ahmad Khan and Henri-Pierre Charles</i>	

Session 3: Parallelism

A Characterization of Shared Data Access Patterns in UPC Programs	111
<i>Christopher Barton, Călin Cașcaval, and José Nelson Amaral</i>	

Exploiting Speculative Thread-Level Parallelism in Data Compression Applications	126
<i>Shengyue Wang, Antonia Zhai, and Pen-Chung Yew</i>	

On Control Signals for Multi-Dimensional Time	141
<i>DaeGon Kim, Gautam, and S. Rajopadhye</i>	

Keynote II

The Berkeley View: A New Framework and a New Platform for Parallel Research	156
<i>David Patterson</i>	

Session 4: Compilation Techniques

An Effective Heuristic for Simple Offset Assignment with Variable Coalescing	158
<i>Hassan Salamy and J. Ramanujam</i>	

Iterative Compilation with Kernel Exploration	173
<i>D. Barthou, S. Donadio, A. Duchateau, W. Jalby, and E. Courtois</i>	

Quantifying Uncertainty in Points-To Relations	190
<i>Constantino G. Ribeiro and Marcelo Cintra</i>	

Session 5: Data Structures

Cache Behavior Modelling for Codes Involving Banded Matrices	205
<i>Diego Andrade, Basilio B. Fraguela, and Ramón Doallo</i>	

Tree-Traversal Orientation Analysis	220
<i>Kevin Andrusky, Stephen Curial, and José Nelson Amaral</i>	

UTS: An Unbalanced Tree Search Benchmark	235
<i>Stephen Olivier, Jun Huan, Jinze Liu, Jan Prins, James Dinan, P. Sadayappan, and Chau-Wen Tseng</i>	

Session 6: Register Allocation

Copy Propagation Optimizations for VLIW DSP Processors with Distributed Register Files	251
<i>Chung-Ju Wu, Sheng-Yuan Chen, and Jenq-Kuen Lee</i>	

Optimal Bitwise Register Allocation Using Integer Linear Programming	267
<i>Rajkishore Barik, Christian Grothoff, Rahul Gupta, Vinayaka Pandit, and Raghavendra Udupa</i>	

Register Allocation: What Does the NP-Completeness Proof of Chaitin et al. Really Prove? Or Revisiting Register Allocation: Why and How . . .	283
<i>Florent Bouchez, Alain Darté, Christophe Guillon, and Fabrice Rastello</i>	

Session 7: Memory Management

Custom Memory Allocation for Free	299
<i>Alin Jula and Lawrence Rauchwerger</i>	
Optimizing the Use of Static Buffers for DMA on a CELL Chip	314
<i>Tong Chen, Zehra Sura, Kathryn O'Brien, and John K. O'Brien</i>	
Runtime Address Space Computation for SDSM Systems	330
<i>Jairo Balart, Marc González, Xavier Martorell, Eduard Ayguadé, and Jesús Labarta</i>	
A Static Heap Analysis for Shape and Connectivity: Unified Memory Analysis: The Base Framework	345
<i>Mark Marron, Deepak Kapur, Darko Stefanovic, and Manuel Hermenegildo</i>	
Author Index	365