

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

Alfred Kobsa

University of California, Irvine, CA, 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

TU Dortmund University, Germany

Madhu Sudan

Microsoft Research, Cambridge, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbruecken, Germany

Chenggang Wu Albert Cohen (Eds.)

Advanced Parallel Processing Technologies

10th International Symposium, APPT 2013
Stockholm, Sweden, August 27-28, 2013
Revised Selected Papers

Volume Editors

Chenggang Wu

Chinese Academy of Sciences, Institute of Computing Technology

State Key Laboratory of Computer Architecture

No. 6 Kexueyuan South Road, Haidian District, 100190 Beijing, China

E-mail: wucg@ict.ac.cn

Albert Cohen

INRIA and École Normale Supérieure

Département d'Informatique

45 rue d'Ulm, 75005 Paris, France

E-mail: albert.cohen@inria.fr

ISSN 0302-9743

e-ISSN 1611-3349

ISBN 978-3-642-45292-5

e-ISBN 978-3-642-45293-2

DOI 10.1007/978-3-642-45293-2

Springer Heidelberg New York Dordrecht London

Library of Congress Control Number: 2013955010

CR Subject Classification (1998): C.1.4, I.3.1, D.2, D.1.3, C.2, J.1, H.5, D.4

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

© Springer-Verlag Berlin Heidelberg 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

With the continuity of Moore's law in the multicore era, the success of cloud computing, and the emerging heterogeneous systems, parallelism pervades almost every domain of computing and information processing. This creates grand challenges to computer architectures and systems, and puts enormous pressure on the design of a new generation of tools and programming methods to harness these systems. These challenges formed the core theme of APPT 2013. Our two-day technical program of APPT 2013 provided an excellent venue capturing some of the state of the art and practice in parallel architecture, parallel software, concurrent and distributed systems, cloud computing, with a highlight on computing systems for big data applications.

We believe this biennial event provides a forum for the presentation of this community's research efforts and exchanging viewpoints. We would like to express our thankfulness to all the colleagues who submitted papers and congratulate those whose papers were accepted. As an event that has taken place for 18 years, APPT aims at providing a high-quality program for all attendees. In all, 62 papers were submitted, a 50% increase over the last conference. We accepted 30 papers: 18 as an oral presentation and 12 as a poster presentation. The regular paper acceptance rate is 29%.

Most submissions were reviewed by three Program Committee(PC) members. An online PC meeting was held during June 9–14. Consensus was reached for each submission.

We would like to thank the authors for submitting their fine work to APPT 2013, and we would also like to show our sincere appreciation to this year's dream-team PC. The 18 PC members did an excellent job in returning high-quality reviews in time and engaging in a constructive online discussion.

We would also like to thank the general chairs (Prof. Yong Dou and Mats Brorsson), and the local organization, publicity and publication chairs for making APPT 2013 possible. Finally, the contributions from our sponsors and supporters were invaluable: We would like to thank the China Computer Federation, the Technical Committee on Computer Architecture of the China Computer Federation, the Royal Institute of Technology in Stockholm, the National Laboratory for Parallel and Distributed Processing, and the State Key Laboratory of High Performance Computing. Our thanks also goes to Springer for its assistance in putting the proceedings together.

August 2013

Albert Cohen
Chenggang Wu

Organization

APPT 2013 was organized by the China Computer Federation.

General Chairs

| | |
|---------------|---|
| Yong Dou | National University of Defense Technology, China |
| Mats Brorsson | KTH Royal Institute of Technology, Sweden |

Organization Chairs

| | |
|--------------|---|
| Zhonghai Lu | KTH Royal Institute of Technology, Sweden |
| Xingwei Wang | Northeastern University, China |

Program Chairs

| | |
|--------------|--|
| Chenggang Wu | Institute of Computing Technology, China |
| Albert Cohen | Inria, France |

Program Committee

| | |
|----------------------|--|
| David Black-Schaffer | Uppsala University, Sweden |
| Haibo Chen | Shanghai Jiaotong University, China |
| Yunji Chen | Institute of Computing Technology, China |
| Chen Ding | University of Rochester, USA |
| Lieven Eeckhout | Ghent University, Belgium |
| Robby Findler | Northwestern University, USA |
| Georgi Gaydadjiev | University of Gothenburg, Sweden |
| R. Govindarajan | Indian Institute of Science, India |
| Wei-Chung Hsu | National Chiao Tung University Hsinchu, Taiwan, China |
| Wolfgang Karl | Karlsruhe Institute of Technology, Germany |
| Lawrence Rauchwerger | Texas A&M University, USA |
| Xipeng Shen | University of William and Mary, USA |
| Olivier Temam | INRIA, France |
| Zhenjiang Wang | Institute of Computing Technology, China |
| Youfeng Wu | Intel, USA |
| Pen-Chung Yew | University of Minnesota, USA |
| Qing Yi | University of Colorado at Colorado, USA |
| Yunlong Zhao | Harbin Institute of Technology, China |

Publicity Chairs

Gongxuan Zhang

Nanjing University of Science and Technology,
China

Hans Vandierendonck

Queen's University Belfast, Queen's University
Belfast, UK

Publication Chairs

Junjie Wu

National University of Defense Technology,
China

Zhicai Shi

Shanghai University of Engineering Science,
China

Sponsor

China Computer Federation China

Supporters

Technical Committee on Computer Architecture of CCF China

KTH Royal Institute of Technology Sweden

Springer-Verlag Germany

National Key Laboratory of Parallel and Distributed Processing China

State Key Laboratory of High Performance Computing China

Table of Contents

| | |
|---|-----|
| Inference and Declaration of Independence in Task-Parallel Programs . . . <i>Foivos S. Zakkak, Dimitrios Chasapis, Polyvios Pratikakis, Angelos Bilas, and Dimitrios S. Nikolopoulos</i> | 1 |
| BDDT: Block-Level Dynamic Dependence Analysis for Task-Based Parallelism <i>George Tzenakis, Angelos Papatriantafyllou, Hans Vandierendonck, Polyvios Pratikakis, and Dimitrios S. Nikolopoulos</i> | 17 |
| A User-Level NUMA-Aware Scheduler for Optimizing Virtual Machine Performance <i>Yuxia Cheng, Wenzhi Chen, Xiao Chen, Bin Xu, and Shaoyu Zhang</i> | 32 |
| Towards RTOS: A Preemptive Kernel Basing on Barrelfish <i>Jingwei Yang, Xiang Long, Xukun Shen, Lei Wang, Shuaitao Feng, and Siyao Zheng</i> | 47 |
| Interrupt Modeling and Verification for Embedded Systems Based on Time Petri Nets <i>Gang Hou, Kuanjiu Zhou, Junwang Chang, Rui Li, and Mingchu Li</i> | 62 |
| Pruning False Positives of Static Data-Race Detection via Thread Specialization <i>Chen Chen, Kai Lu, Xiaoping Wang, Xu Zhou, and Li Fang</i> | 77 |
| G-Paradex: GPU-Based Parallel Indexing for Fast Data Deduplication <i>Bin Lin, Xiangke Liao, Shanshan Li, Yufeng Wang, He Huang, and Ling Wen</i> | 91 |
| Research on SPH Parallel Acceleration Strategies for Multi-GPU Platform <i>Lei Hu, Xukun Shen, and Xiang Long</i> | 104 |
| Binarization-Based Human Detection for Compact FPGA Implementation <i>Shuai Xie, Yibin Li, Zhiping Jia, and Lei Ju</i> | 119 |
| HPACS: A High Privacy and Availability Cloud Storage Platform with Matrix Encryption <i>Yanzhang He, Xiaohong Jiang, Kejiang Ye, Ran Ma, and Xiang Li</i> | 132 |

| | |
|--|-----|
| ECAM: An Efficient Cache Management Strategy for Address Mappings in Flash Translation Layer | 146 |
| <i>Xuchao Xie, Qiong Li, Dengping Wei, Zhenlong Song, and Liquan Xiao</i> | |
| A Performance Study of Software Prefetching for Tracing Garbage Collectors | 160 |
| <i>Hao Wu, Zhenzhou Ji, Suzia Zhu, and Zhigang Chen</i> | |
| Adaptive Implementation Selection in the SkePU Skeleton Programming Library | 170 |
| <i>Usman Dastgeer, Lu Li, and Christoph Kessler</i> | |
| Automatic Skeleton-Based Compilation through Integration with an Algorithm Classification | 184 |
| <i>Cedric Nugteren, Pieter Custers, and Henk Corporaal</i> | |
| Optimizing Program Performance via Similarity, Using a Feature-Agnostic Approach | 199 |
| <i>Rosario Cammarota, Laleh Aghababae Beni, Alexandru Nicolau, and Alexander V. Veidenbaum</i> | |
| Scalable NIC Architecture to Support Offloading of Large Scale MPI Barrier | 214 |
| <i>Shaogang Wang, Weixia Xu, Dan Wu, Zhengbin Pang, and Pingjing Lu</i> | |
| Hierarchical Clustering Routing Protocol Based on Optimal Load Balancing in Wireless Sensor Networks | 227 |
| <i>Tianshu Wang, Gongxuan Zhang, Xichen Yang, and Ahmadreza Vajdi</i> | |
| An Efficient Parallel Mechanism for Highly-Debuggable Multicore Simulator | 241 |
| <i>Xiaochun Ye, Dongrui Fan, Da Wang, Fenglong Song, Hao Zhang, and Zhimin Tang</i> | |
| Data Access Type Aware Replacement Policy for Cache Clustering Organization of Chip Multiprocessors | 254 |
| <i>Chongmin Li, Dongsheng Wang, Haixia Wang, Guohong Li, and Yibo Xue</i> | |
| Agent-Based Credibility Protection Model for Decentralized Network Computing Environment | 269 |
| <i>Xiaolong Xu, Qun Tu, and Xinheng Wang</i> | |
| A Vectorized K-Means Algorithm for Intel Many Integrated Core Architecture | 277 |
| <i>Fuhui Wu, Qingbo Wu, Yusong Tan, Lifeng Wei, Lisong Shao, and Long Gao</i> | |

| | |
|---|-----|
| Towards Eliminating Memory Virtualization Overhead..... | 295 |
| <i>Xiaolin Wang, Lingmei Weng, Zhenlin Wang, and Yingwei Luo</i> | |
| An Improved FPGAs-Based Loop Pipeline Scheduling Algorithm for Reconfigurable Compiler | 307 |
| <i>Zhenhua Guo, Yanxia Wu, Guoyin Zhang, and Tianxiang Sui</i> | |
| ACF: Networks-on-Chip Deadlock Recovery with Accurate Detection and Elastic Credit | 319 |
| <i>Nan Wu, Yuran Qiao, Mei Wen, and Chunyuan Zhang</i> | |
| An Auction and League Championship Algorithm Based Resource Allocation Mechanism for Distributed Cloud | 334 |
| <i>Jiajia Sun, Xingwei Wang, Keqin Li, Chuan Wu, Min Huang, and Xueyi Wang</i> | |
| Accelerating Software Model Checking Based on Program Backbone.... | 347 |
| <i>Kuanjiu Zhou, Jiawei Yong, Xiaolong Wang, Longtao Ren, Gang Hou, and Junwang Chang</i> | |
| A Cloud Computing System for Snore Signals Processing | 359 |
| <i>Jian Guo, Kun Qian, Zhaomeng Zhu, Gongxuan Zhang, and Huijie Xu</i> | |
| Research on Optimum Checkpoint Interval for Hybrid Fault Tolerance | 367 |
| <i>Lei Zhu, Jianhua Gu, Yunlan Wang, and Tianhai Zhao</i> | |
| Programming Real-Time Image Processing for Manycores in a High-Level Language | 381 |
| <i>Essayas Gebrewahid, Zain-ul-Abdin, Bertil Svensson, Veronica Gaspes, Bruno Jego, Bruno Lavigueur, and Mathieu Robart</i> | |
| Self-adaptive Retransmission for Network Coding with TCP | 396 |
| <i>Chunqing Wu, Hongyun Zhang, Wanrong Yu, Zhenqian Feng, and Xiaofeng Hu</i> | |
| Author Index | 409 |