

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, Lancaster, 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, Zurich, Switzerland

John C. Mitchell

Stanford University, Stanford, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Dortmund, Germany

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbrücken, Germany

More information about this series at <http://www.springer.com/series/7407>

Guang R. Gao · Depei Qian
Xinbo Gao · Barbara Chapman
Wenguang Chen (Eds.)

Network and Parallel Computing

13th IFIP WG 10.3 International Conference, NPC 2016
Xi'an, China, October 28–29, 2016
Proceedings

Editors

Guang R. Gao
University of Delaware
Newark, DE
USA

Barbara Chapman
Stony Brook University
Stony Brook, NY
USA

Depei Qian
Beihang University
Beijing
China

Wenguang Chen
Tsinghua University
Beijing
China

Xinbo Gao
Xidian University
Xi'an
China

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-47098-6 ISBN 978-3-319-47099-3 (eBook)
DOI 10.1007/978-3-319-47099-3

Library of Congress Control Number: 2016952885

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

© IFIP International Federation for Information Processing 2016

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.

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.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

These proceedings contain the papers presented at the 2016 IFIP International Conference on Network and Parallel Computing (NPC 2016), held in Xi'An, China, during October 28–29, 2016. The goal of the conference was to establish an international forum for engineers and scientists to present their ideas and experiences in network and parallel computing.

A total of 99 submissions were received in response to our Call for Papers. These papers originate from Australia, Asia (China, Japan), and North America (USA). Each submission was sent to at least three reviewers. Each paper was judged according to its originality, innovation, readability, and relevance to the expected audience. Based on the reviews received, a total of 19 papers were retained for inclusion in the proceedings. Among the 19 papers, 12 were accepted as full papers for presentation at the conference. We also accepted seven papers as short papers for a possible brief presentation at the conference. We accepted another ten papers for a poster session (but without proceedings). Thus, only 19 % of the total submissions could be included in the final program, and 29 % of the submitted work was proposed to be presented at the conference.

The topics tackled at this year's conference include resource management, in particular solid-state drives and other non volatile memory systems; resiliency and reliability; job and task scheduling for batch systems and big data frameworks; heterogeneous systems based on accelerators; data processing, in particular in the context of big data; and more fundamental algorithms and abstractions for parallel computing.

We wish to thank the contributions of the other members of the Organizing Committee. We thank the publicity chairs, Xiaofei Liao, Cho-Li Want, and Koji Inoue, for their hard work to publicize NPC 2016 under a very tight schedule. We are deeply grateful to the Program Committee members. The large number of submissions received and the diversified topics made this review process a particularly challenging one.

August 2016

Guang R. Gao
Depei Qian
Xinbo Gao
Barbara Chapman
Wenguang Chen

Organization

General Co-chairs

Guang Rong Gao	University of Delaware, USA
Xinbo Gao	Xidian University, China
Depei Qian	Beihang University, China

Organization Chair

Quan Wang	Xidian University, China
-----------	--------------------------

Program Co-chairs

Barbara Chapman	Stony Brook University, USA
Wenguang Chen	Tsinghua University, China

Publication Chair

Stephane Zuckerman	University of Delaware, USA
--------------------	-----------------------------

Local Arrangements Co-chair

Qiguang Miao	Xidian University, China
--------------	--------------------------

Publicity Chairs

Koji Inoue	Kyushu University, Japan
Xiaofei Liao	Huanzhong University of Science and Technology, China
Cho-Li Wang	University of Hong Kong, SAR China

Web Chair

Yining Quan	Xidian University, China
-------------	--------------------------

Steering Committee

Cheng Ding	University of Rochester, USA
Jack Dongarra	University of Tennessee, USA
Kemal Ebcioglu (Chair)	Global Supercomputing, USA
Guang Rong Gao	University of Delaware, USA

Jean-Luc Gaudiot	University of California, Irvine, USA
Tony Hey	Microsoft, USA
Hai Jin	Huanzhong University of Science and Technology, China
Guojie Li	Institute of Computing Technology, China
Yoichi Muraoka	Waseda University, Japan
Viktor Prasanna	University of Southern California, USA
Daniel Reed	University of Iowa, USA
Weisong Shi	Wayne State University, USA
Zhiwei Xu	Institute of Computing Technology, China

Program Committee

Abramson	University of Queensland, Australia
Hong An	University of Science and Technology of China, China
Pavan Balaji	Argonne National Lab, USA
Taisuke Boku	University of Tsukuba, Japan
Sunita Chandrasekaran	University of Delaware, USA
Barbara Chapman	Stony Brook University, USA
Wenguang Chen	Tsinghua University, China
Yurong Chen	Intel, China
Yeching Chung	National Tsinghua University, Taiwan
Yuefan Deng	Stony Brook University, USA
Zhihui Du	Tsinghua University, China
Robert Harrison	Stony Brook University, USA
Torsten Hoefler	ETH, Switzerland
Kise Kenji	Tokyo Institute of Technology, Japan
Keiji Kimura	Waseda University, Japan
Chao Li	Shanghai Jiao Tong University, China
Miron Livny	University of Wisconsin at Madison, USA
Yi Liu	Beihang University, China
Kai Lu	National University of Defense Technology, China
Yutong Lu	National University of Defense Technology, China
Yingwei Luo	Peking University, China
Xiaosong Ma	Qatar Computing Research Institute, Qatar
Philip Papadopoulos	University of California, San Diego, USA
Xuanhua Shi	Huazhong University of Science and Technology, China
Weiguo Wu	Xi'An Haotong University, China
Jingling Xue	University of New South Wales, Australia
Chao Yang	Institute of Software, Chinese Academy of Sciences, China
Jun Yao	Huawei, China
Li Zha	ICT, Chinese Academy of Sciences, China
Weihua Zhang	Fudan University, China
Yunquan Zhang	ICT, Chinese Academy of Sciences, China

Contents

Memory: Non-Volatile, Solid State Drives, Hybrid Systems

VIOS: A Variation-Aware I/O Scheduler for Flash-Based Storage Systems. . .	3
<i>Jinhua Cui, Weiguo Wu, Shiqiang Nie, Jianhang Huang, Zhuang Hu, Nianjun Zou, and Yinfeng Wang</i>	
Exploiting Cross-Layer Hotness Identification to Improve Flash Memory System Performance	17
<i>Jinhua Cui, Weiguo Wu, Shiqiang Nie, Jianhang Huang, Zhuang Hu, Nianjun Zou, and Yinfeng Wang</i>	
Efficient Management for Hybrid Memory in Managed Language Runtime	29
<i>Chenxi Wang, Ting Cao, John Zigman, Fang Lv, Yunquan Zhang, and Xiaobing Feng</i>	

Resilience and Reliability

Application-Based Coarse-Grained Incremental Checkpointing Based on Non-volatile Memory	45
<i>Zhan Shi, Kai Lu, Xiaoping Wang, Wenzhe Zhang, and Yiqi Wang</i>	
DASM: A Dynamic Adaptive Forward Assembly Area Method to Accelerate Restore Speed for Deduplication-Based Backup Systems	58
<i>Chao Tan, Luyu Li, Chentao Wu, and Jie Li</i>	

Scheduling and Load-Balancing

A Statistics Based Prediction Method for Rendering Application.	73
<i>Qian Li, Weiguo Wu, Long Xu, Jianhang Huang, and Mingxia Feng</i>	
IBB: Improved K-Resource Aware Backfill Balanced Scheduling for HTCondor	85
<i>Lan Liu, Zhongzhi Luan, Haozhan Wang, and Depei Qian</i>	
Multipath Load Balancing in SDN/OSPF Hybrid Network	93
<i>Xiangshan Sun, Zhiping Jia, Mengying Zhao, and Zhiyong Zhang</i>	

Heterogeneous Systems

A Study of Overflow Vulnerabilities on GPUs	103
<i>Bang Di, Jianhua Sun, and Hao Chen</i>	

Streaming Applications on Heterogeneous Platforms 116
Zhaokui Li, Jianbin Fang, Tao Tang, Xuhao Chen, and Canqun Yang

Data Processing and Big Data

DSS: A Scalable and Efficient Stratified Sampling Algorithm for
Large-Scale Datasets 133
Minne Li, Dongsheng Li, Siqi Shen, Zhaoning Zhang, and Xicheng Lu

A Fast and Better Hybrid Recommender System Based on Spark 147
*Jiali Wang, Hang Zhuang, Changlong Li, Hang Chen, Bo Xu,
Zhuocheng He, and Xuehai Zhou*

Discovering Trip Patterns from Incomplete Passenger Trajectories
for Inter-zonal Bus Line Planning 160
*Zhaoyang Wang, Beihong Jin, Fusang Zhang, Ruiyang Yang,
and Qiang Ji*

FCM: A Fine-Grained Crowdsourcing Model Based on Ontology in
Crowd-Sensing 172
Jian An, Ruobiao Wu, Lele Xiang, Xiaolin Gui, and Zhenlong Peng

QIM: Quantifying Hyperparameter Importance for Deep Learning 180
Dan Jia, Rui Wang, Chengzhong Xu, and Zhibin Yu

Algorithms and Computational Models

Toward a Parallel Turing Machine Model 191
Peng Qu, Jin Yan, and Guang R. Gao

On Determination of Balance Ratio for Some Tree Structures 205
Daxin Zhu, Tinran Wang, and Xiaodong Wang

Author Index 213