

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>

Xuanhua Shi · Hong An
Chao Wang · Mahmut Kandemir
Hai Jin (Eds.)

Network and Parallel Computing

14th IFIP WG 10.3 International Conference, NPC 2017
Hefei, China, October 20–21, 2017
Proceedings

Editors

Xuanhua Shi
Huazhong University of Science
and Technology
Wuhan
China

Hong An
University of Science and Technology
of China
Hefei
China

Chao Wang
University of Science and Technology
of China
Hefei
China

Mahmut Kandemir
Pennsylvania State University
University Park, PA
USA

Hai Jin
Huazhong University of Science
and Technology
Wuhan
China

ISSN 0302-9743

ISSN 1611-3349 (electronic)

Lecture Notes in Computer Science

ISBN 978-3-319-68209-9

ISBN 978-3-319-68210-5 (eBook)

DOI 10.1007/978-3-319-68210-5

Library of Congress Control Number: 2017954908

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

© IFIP International Federation for Information Processing 2017

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. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

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

Message from the NPC 2017 Chairs

Welcome to the proceedings of the 14th IFIP International Conference on Network and Parallel Computing (NPC 2017)!

The goal of the NPC conferences is to establish an international forum for engineers and scientists to present their excellent ideas and experiences in system fields of distributed and parallel computing.

High-performance computing and big data are two main areas where NPC 2017 provided a dynamic forum to explore, discuss, and debate state-of-the-art technology issues and challenges. High-performance computers and big-data systems are tied inextricably to the broader computing ecosystem and its designs and market adoption. They also highlight information security needs and economic competitiveness in ways that distinguish them from most other scientific instruments. We strongly believe that the stakes are high, and it is far beyond the boundaries of nations and continents, and should strongly encourage a broad international participation.

For their contribution toward a successful conference, we would like to thank the reviewers for the amazing service that they provided for the NPC conference to ensure that submissions (88 for the conference) received the consideration and attention that they deserve. As the NPC chairs, we especially appreciate the timely completion of reviews for the final selection phases. Each paper was reviewed by three to five expert reviewers ensuring that review details were considered in the decision-making of the final acceptance of 21 full papers (slightly below 24%), including 12 papers published as Special Issue papers of the *International Journal of Parallel Programming*, and nine papers published as LNCS proceedings. A number of strong papers that could not be accepted to the full papers track were considered for the short paper tracks. Finally, we selected 11 short papers (for an acceptance rate of 12.5%). These papers cover traditional areas of network and parallel computing, including parallel applications, distributed algorithms, parallel architectures, software environments, and distributed tools.

We sincerely appreciate the work and effort of the authors in preparing their submissions for review, and addressing the reviewers' comments before submitting the camera-ready copies of their accepted papers, and attending the conference to present and discuss their work.

We also want to thank every member of the NPC 2017 Organizing Committee and Steering Committee for their help in putting together such an exciting program. Finally, we thank all the attendees.

August 2017

Xuanhua Shi
Hong An
Chao Wang
Mahmut Kandemir
Hai Jin

Organization

Organizing Committee

General Co-chairs

Michael Gschwind IBM, USA
Weisong Shi Wayne State University, USA

Organization Chair

Guohua Cheng Hangzhou Jianpei Technology Co., Ltd., China

Program Co-chairs

Hong An University of Science and Technology of China, China
Mahmut Kandemir Pennsylvania State University, USA

Publications Chair

Xuanhua Shi Huazhong University of Science and Technology, China

Local Arrangements Chair

Yu Zhang University of Science and Technology of China, China

Publicity Co-chairs

Yong Chen Texas Tech University, USA
Guangzhong Sun University of Science and Technology of China, China
Keiji Kimura Waseda University, Japan
Stephane Zuckerman University of Delaware, USA

Registration Chair

Chao Wang University of Science and Technology of China, China

Web Chair

Hui Sun Anhui University, China

Steering Committee

Kemal Ebcioglu (Chair) Global Supercomputing USA
Hai Jin (Vice Chair) Huazhong University of Science and Technology, China
Chen Ding University of Rochester, USA
Jack Dongarra University of Tennessee, USA
Guangrong Gao University of Delaware, USA

Jean-Luc Gaudiot	University of California Irvine, USA
Tony Hey	Science and Technology Facilities Council, UK
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
Ninghui Sun	Institute of Computing Technology, China
Zhiwei Xu	Institute of Computing Technology, China

Program Committee

Meena Arunachalam	Intel, USA
Nachi Nachiappan	Apple, USA
Myoungsoo Jung	Yonsei University, Republic of Korea
Mahmut Kandemir	Pennsylvania State University, USA
David Abramson	University of Queensland, Australia
Pavan Balaji	Argonne National Laboratory, USA
Taisuke Boku	University of Tsukuba, Japan
Sunita Chandrasekaran	University of Delaware, USA
Barbara Chapman	Stony Brook University, USA
Robert Harrison	Brook University, USA
Miron Livny	University of Wisconsin at Madison, USA
Keiji Kimura	Waseda University, Japan
Kise Kenji	Tokyo Institute of Technology, Japan
Xiaosong Ma	Qatar Computing Research Institute
Yuefan Deng	Stony Brook University, USA
Hong An	University of Science and Technology of China, China
Wenguang Chen	Tsinghua University, China
Yeching Chung	National Tsinghua University, Taiwan
Chen Ding	University of Rochester, USA
Qing Yi	University of Colorado at Colorado Springs, USA
Zhihui Du	Tsinghua University, China
Xiaobing Feng	Institute of Computing Technology, CAS, China
Kai Lu	National University of Defense Technology, China
Yutong Lu	Sun Yat-sen University, China
Yun Xu	University of Science and Technology of China, China
Yifeng Chen	Peking University, China
Yingwei Luo	Peking University, China
Yong Chen	Texas Tech University, USA
Xuanhua Shi	Huazhong University of Science and Technology, China
Weiguo Wu	Xi'an Jiaotong University, China
Yungang Bao	Institute of Computing Technology, CAS, China
Weihua Zhang	Fudan University, China
Yunquan Zhang	Institute of Computing Technology, CAS, China
Li Shen	National University of Defense Technology, China

Chao Yang
Dongrui Fan
Chunyuan Zhang
Di Wu
Jinlei Jiang

Institute of Software, CAS, China
Institute of Computing Technology, CAS, China
National University of Defense Technology, China
Sun Yat-sen University, China
Tsinghua University, China

Contents

SCMKV: A Lightweight Log-Structured Key-Value Store on SCM.	1
<i>Zhenjie Wang, Linpeng Huang, and Yanmin Zhu</i>	
Adaptive Length Sliding Window-Based Network Coding for Energy Efficient Algorithm in MANETs.	13
<i>Baolin Sun, Chao Gui, Ying Song, and Hua Chen</i>	
A Virtual Machine Migration Algorithm Based on Group Selection in Cloud Data Center.	24
<i>Zhen Guo, Wenbin Yao, and Dongbin Wang</i>	
Two-Stage Job Scheduling Model Based on Revenues and Resources	37
<i>Yuliang Shi, Dong Liu, Jing Hu, and Jianlin Zhang</i>	
TCon: A Transparent Congestion Control Deployment Platform for Optimizing WAN Transfers.	49
<i>Yuxiang Zhang, Lin Cui, Fung Po Tso, Quanlong Guan, and Weijia Jia</i>	
Regional Congestion Mitigation in Lossless Datacenter Networks	62
<i>Xiaoli Liu, Fan Yang, Yanan Jin, Zhan Wang, Zheng Cao, and Ninghui Sun</i>	
A Fast and Accurate Way for API Network Construction Based on Semantic Similarity and Community Detection.	75
<i>Xi Yang and Jian Cao</i>	
Improving Branch Prediction for Thread Migration on Multi-core Architectures	87
<i>Tan Zhang, Chaobing Zhou, Libo Huang, Nong Xiao, and Sheng Ma</i>	
Optimizing OpenCL Implementation of Deep Convolutional Neural Network on FPGA	100
<i>Yuran Qiao, Junzhong Shen, Dafei Huang, Qianming Yang, Mei Wen, and Chunyuan Zhang</i>	
ACO-inspired ICN Routing Scheme with Density-Based Spatial Clustering	112
<i>Jianhui Lv, Xingwei Wang, and Min Huang</i>	

An Efficient Polarity Optimization Approach for Fixed Polarity Reed-Muller Logic Circuits Based on Novel Binary Differential Evolution Algorithm 118
Zhenxue He, Guangjun Qin, Limin Xiao, Fei Gu, Zhisheng Huo, Li Ruan, Haitao Wang, Longbing Zhang, Jianbin Liu, Shaobo Liu, and Xiang Wang

Stem: A Table-Based Congestion Control Framework for Virtualized Data Center Networks 122
Jie Wu, Binzhang Fu, and Mingyu Chen

SDAC: Porting Scientific Data to Spark RDDs 127
Tian Yang, Kenjiro Taura, and Liu Chao

Unified Access Layer with PostgreSQL FDW for Heterogeneous Databases . . . 131
Xuefei Wang, Ruohang Feng, Wei Dong, Xiaoqian Zhu, and Wenke Wang

Balancing Global and Local Fairness Allocation Model in Heterogeneous Data Center. 136
Bingyu Zhou, Guangjun Qin, Limin Xiao, Zhisheng Huo, Bing Wei, Jiulong Chang, Nan Zhou, Li Ruan, Haitao Wang, and Zipeng Wei

System Problem Detection by Mining Process Model from Console Logs . . . 140
Jian Li and Jian Cao

Quantifying the Isolation Characteristics in Container Environments 145
Chang Zhao, Yusen Wu, Zujie Ren, Weisong Shi, Yongjian Ren, and Jian Wan

CSAS: Cost-Based Storage Auto-Selection, a Fine Grained Storage Selection Mechanism for Spark. 150
Bo Wang, Jie Tang, Rui Zhang, and Zhimin Gu

An In-Depth Performance Analysis of Many-Integrated Core for Communication Efficient Heterogeneous Computing 155
Jie Zhang and Myoungsoo Jung

Author Index 161