

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, Chennai, 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>

Feng Zhang · Jidong Zhai
Marc Snir · Hai Jin
Hironori Kasahara · Mateo Valero (Eds.)

Network and Parallel Computing

15th IFIP WG 10.3 International Conference, NPC 2018
Muroran, Japan, November 29 – December 1, 2018
Proceedings

Editors

Feng Zhang 
Renmin University of China
Beijing, China

Jidong Zhai
Tsinghua University
Beijing, China

Marc Snir
University of Illinois at Urbana-Champaign
Urbana, IL, USA

Hai Jin
Huazhong University of Science
and Technology
Wuhan, China

Hironori Kasahara
Waseda University
Shinjuku-ku, Japan

Mateo Valero
Barcelona Supercomputing Center
Barcelona, Spain

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-05676-6 ISBN 978-3-030-05677-3 (eBook)
<https://doi.org/10.1007/978-3-030-05677-3>

Library of Congress Control Number: 2018963718

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

© IFIP International Federation for Information Processing 2018

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.

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

Preface

These proceedings contain the papers presented at the 2018 IFIP International Conference on Network and Parallel Computing (NPC 2018), held in Muroran, Hokkaido, Japan, from November 29 to December 1, 2018. The goal of the conference is to establish an international forum for engineers and scientists to present their ideas and experiences in network and parallel computing.

A total of 72 submissions were received in response to our Call for Papers. These papers originate from Australia, Asia (China, Japan), Europe, 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, 22 full papers (about 30%), including 12 papers published as Special Issue papers of the *International Journal of Parallel Programming*, and ten papers published as LNCS proceedings were retained. A number of strong papers that could not be accepted to the full-paper track were considered for the short-paper tracks. Finally, we selected 12 short papers. These papers cover traditional areas of network and parallel computing, including parallel applications, distributed algorithms, parallel architectures, software environments, and distributed tools.

We share the view that, during the past decade, the tools and cultures of high-performance computing and big data analytics are diverging to the detriment of both, and the international community should find a unified path that can best serve the need of a broad spectrum of major application areas. Unlike other tools, which are limited to particular scientific domains, computational modeling and data analytics are applicable to all areas of science and engineering, as they breathe life into the underlying mathematics of scientific models.

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 2018 Organizing Committee and Steering Committee for their help in putting together such an exciting program. Finally, we thank all the attendees.

August 2018

Feng Zhang
Jidong Zhai
Marc Snir
Hai Jin
Hironori Kasahara
Mateo Valero

Organization

Organizing Committee

General Co-chairs

Hai Jin	Huazhong University of Science and Technology, China
Hironori Kasahara	Waseda University, Japan
Mateo Valero	Barcelona Supercomputing Center, Spain

Program Co-chairs

Marc Snir	University of Illinois at Urbana-Champaign, USA
Jidong Zhai	Tsinghua University, China

Publications Chair

Feng Zhang	Renmin University of China, China
------------	-----------------------------------

Local Arrangements Co-chairs

Mianxiong Dong	Muroran Institute of Technology, Japan
He Li	Muroran Institute of Technology, Japan

Publicity Chairs

Bingsheng He	National University of Singapore, Singapore
Keiji Kimura	Waseda University, Japan
Shuaiwen Leon Song	Pacific Northwest National Laboratory, USA
Yunquan Zhang	Institute of Computing Technology, China
Stéphane Zuckerman	University of Cergy-Pontoise, France

Web Chair

Yuyang Jin	Tsinghua University, China
------------	----------------------------

Advisory Committee

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
Guang R. 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

Yungang Bao	Institute of Computing Technology
Dehao Chen	Google
Wenguang Chen	Tsinghua University, China
Guoyang Chen	Alibaba Group US Inc., USA
Huimin Cui	Institute of Computing Technology
Yufei Ding	University of California, Santa Barbara, USA
Zhihui Du	Tsinghua University, China
Masato Edahiro	Nagoya University, Japan
Keiji Kimura	Waseda University, Japan
Ang Li	Pacific Northwest National Lab
Chao Li	Shanghai Jiao Tong University, China
Dong Li	University of California, Merced, USA
Yun Liang	Peking University, China
Weifeng Liu	Norwegian University of Science and Technology, Norway
Yingwei Luo	Peking University, China
Xiaosong Ma	Qatar Computing Research Institute, Qatar
Daniel A. Orozco	Google
Antoni Pop	University of Manchester, UK
Xuehai Qian	University of Southern California, USA
Lawrence Rauchwerger	Texas A&M University, USA
Bin Ren	College of William and Mary, USA
Jinglei Ren	Microsoft
Larry Rudolph	Two Sigma Investments
Li Shen	National University of Defense Technology, China
Xuanhua Shi	Huazhong University of Science and Technology, China
Yao Shi	Futurewei Technologies
GuangZhong Sun	University of Science and Technology of China, China
Shanjiang Tang	Tianjin University, China
Dingwen Tao	The University of Alabama, USA
Parimala Thulasiraman	University of Manitoba, Canada
Lei Wang	National University of Defense Technology, China
Zhaoguo Wang	Shanghai Jiao Tong University, China
Zheng Wang	Lancaster University, UK
Bo Wu	Colorado School of Mines, USA

Zhibin Yu	Shenzhen Institutes of Advanced Technology, Chinese Academy of Science, China
Weihua Zhang	Fudan University, China
Zhijia Zhao	University of California Riverside, USA
Amelie Chi Zhou	Shenzhen University, China
Stéphane Zuckerman	University of Cergy-Pontoise, France

Contents

CNLoc: Channel State Information Assisted Indoor WLAN Localization Using Nomadic Access Points	1
<i>Jiang Xiao, Huichuwu Li, He Li, and Hai Jin</i>	
ALOR: Adaptive Layout Optimization of Raft Groups for Heterogeneous Distributed Key-Value Stores	13
<i>Yangyang Wang, Yunpeng Chai, and Xin Wang</i>	
STrieGD: A Sampling Trie Indexed Compression Algorithm for Large-Scale Gene Data	27
<i>Yanzhen Gao, Xiaozhen Bao, Jing Xing, Zheng Wei, Jie Ma, and Peiheng Zhang</i>	
On Retargeting the AI Programming Framework to New Hardwares	39
<i>Jiacheng Zhao, Yisong Chang, Denghui Li, Chunwei Xia, Huimin Cui, Ke Zhang, and Xiaobing Feng</i>	
An Efficient Method for Determining Full Point-to-Point Latency of Arbitrary Indirect HPC Networks	52
<i>Chengchun Liu, Zhang Yang, Limin Xiao, Baicheng Yan, Zhihao Wang, and Hongyun Tian</i>	
KT-Store: A Key-Order and Write-Order Hybrid Key-Value Store with High Write and Range-Query Performance	64
<i>Haobo Wang, Yinliang Yue, Shuibing He, and Weiping Wang</i>	
GRAM: A GPU-Based Property Graph Traversal and Query for HPC Rich Metadata Management	77
<i>Wenke Li, Xuanhua Shi, Hong Huang, Peng Zhao, Hai Jin, Dong Dai, and Yong Chen</i>	
GPU-Accelerated Clique Tree Propagation for Pouch Latent Tree Models . . .	90
<i>Leonard K. M. Poon</i>	
HPC-SFI: System-Level Fault Injection for High Performance Computing Systems	103
<i>Yanqi Wang, Qi Zhang, Yi Liu, and Depei Qian</i>	
Data Fine-Pruning: A Simple Way to Accelerate Neural Network Training. . .	114
<i>Junyu Li, Ligang He, Shenyuan Ren, and Rui Mao</i>	

Balancing the QOS and Security in Dijkstra Algorithm by SDN Technology	126
<i>Zhao JinJing, Ling Pang, Xiaohui Kuang, and Rong Jin</i>	
Labeled Network Stack: A Co-designed Stack for Low Tail-Latency and High Concurrency in Datacenter Services.	132
<i>Wenli Zhang, Ke Liu, Hui Song, Lan Yu, and Mingyu Chen</i>	
A Deep Learning Approach for Network Anomaly Detection Based on AMF-LSTM	137
<i>Mingyi Zhu, Kejiang Ye, Yang Wang, and Cheng-Zhong Xu</i>	
FSObserver: A Performance Measurement and Monitoring Tool for Distributed Storage Systems	142
<i>Xiao Zhang, Lanxin Kong, Shunyi Zhu, Zhanhuai Li, and Xiaonan Zhao</i>	
vGroupier: Optimizing the Performance of Parallel Jobs in Xen by Increasing Synchronous Execution of Virtual Machines	148
<i>Peng Jiang, Ligang He, Shen yuan Ren, Junyu Li, and Yuhua Cui</i>	
Systolic Array Based Accelerator and Algorithm Mapping for Deep Learning Algorithms	153
<i>Zhijie Yang, Lei Wang, Dong Ding, Xiangyu Zhang, Yu Deng, Shiming Li, and Qiang Dou</i>	
A Fine-Grained Performance Bottleneck Analysis Method for HDFS.	159
<i>Yi Liu, Yunchun Li, Honggang Zhou, Jingyi Zhang, Hailong Yang, and Wei Li</i>	
Mimir+: An Optimized Framework of MapReduce on Heterogeneous High-Performance Computing System	164
<i>Nan Hu, Zhiguang Chen, Yunfei Du, and Yutong Lu</i>	
DLIR: An Intermediate Representation for Deep Learning Processors	169
<i>Huiying Lan and Zidong Du</i>	
GPU Memory Management Solution Supporting Incomplete Pages	174
<i>Li Shen, Shiqing Zhang, Yaohua Yang, and Zhiying Wang</i>	
Leveraging Subgraph Extraction for Performance Portable Programming Frameworks on DL Accelerators	179
<i>Xiao Zhang, Huiying Lan, and Tian Zhi</i>	
An Intelligent Parking Scheduling Algorithm Based on Traffic and Driver Behavior Predictions	185
<i>Jiazao Lin, Shi-Yong Chen, Chih-Yung Chang, and Guilin Chen</i>	
Author Index	191