

Founding Editors

Gerhard Goos

Karlsruhe Institute of Technology, Karlsruhe, Germany

Juris Hartmanis

Cornell University, Ithaca, NY, USA

Editorial Board Members

Elisa Bertino

Purdue University, West Lafayette, IN, USA

Wen Gao

Peking University, Beijing, China

Bernhard Steffen 

TU Dortmund University, Dortmund, Germany

Gerhard Woeginger 

RWTH Aachen, Aachen, Germany

Moti Yung

Columbia University, New York, NY, USA

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

Santosh Pande · Vivek Sarkar (Eds.)

Languages and Compilers for Parallel Computing

32nd International Workshop, LCPC 2019
Atlanta, GA, USA, October 22–24, 2019
Revised Selected Papers

Editors

Santosh Pande
Georgia Institute of Technology
Atlanta, GA, USA

Vivek Sarkar
Georgia Institute of Technology
Atlanta, GA, USA

ISSN 0302-9743

ISSN 1611-3349 (electronic)

Lecture Notes in Computer Science

ISBN 978-3-030-72788-8

ISBN 978-3-030-72789-5 (eBook)

<https://doi.org/10.1007/978-3-030-72789-5>

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

© Springer Nature Switzerland AG 2021

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, expressed 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

The **32nd Workshop on Languages and Compilers for Parallel Computing (LCPC)** was held during **October 22–24, 2019 at Georgia Institute of Technology, Atlanta**. LCPC is the longest-running workshop in the intersection of programming systems and parallel computing. Since its inception, LCPC has been a leading venue for disseminating cutting-edge research on all aspects of parallel programming systems – including programming models, languages, compilers, runtime systems, and tools – for a diverse range of application domains and computing platforms. The scope of LCPC 2019 spanned a wide variety of topics related to parallel computing, including current domains and platforms, such as scientific computing, batch/streaming/real-time data analytics, machine learning, cognitive computing, heterogeneous/reconfigurable computing, mobile computing, cloud computing, IoT; as well as forward-looking domains such as analog and quantum computing. Bold new ideas in programming systems are critical for the success of computing platforms in the Moore’s Law end-game and post-Moore eras. LCPC 2019 offered a highly interactive forum for the dissemination of innovative research contributions as well as in-depth discussions of novel and emerging ideas. As in past years, LCPC 2019 brought together researchers from academia, national labs, and industry with the aim of creating and strengthening research collaborations. LCPC 2019 also hosted a co-located event on October 21, 2019 – a special workshop on MLIR for HPC organized by Prof. Mary Hall.

We were fortunate to have a diverse set of expert program committee (PC) members, spanning junior and senior researchers, women and under-represented groups, and researchers from different cross sections of the community, including academia, industry, and national labs. The details of the PC can be found at <https://lcpc19.cc.gatech.edu/committee>. Two types of papers were invited for submission:

- **Full papers** on innovative and original research that describe new research contributions. 15 pages limit.
- **Short papers** on preliminary results, surveys, demonstrations, or vision for future research. 8 pages limit.

The review process was guided by novelty; reviewers were given explicit instructions to look for novel, intriguing ideas in the submissions. The inclusion of papers that propose new ideas – a new problem, a new research topic, radical insight into an existing topic, surprising results, etc. – was one of the key goals in devising the workshop program. Another important consideration was whether the paper could provoke interesting discussions during the workshop. The PC ensured that each paper received at least 4 reviews. Each paper was discussed in an online PC meeting, and the final selection of the papers was based on the degree of novelty and quality. Regardless of acceptance or rejection, authors were provided with detailed feedback. Out of 17 submissions, 8 were accepted as regular papers and 3 were accepted as short papers. As in past LCPC workshops, a two-tier revision process was followed for the accepted

papers: first, the authors were asked to incorporate the reviewers' feedback and prepare a pre-proceedings version of the paper, which was made available at the workshop. Next, the authors were asked to incorporate the feedback received during the workshop and prepare a final camera-ready version, which is included in these proceedings.

Apart from the accepted papers, the workshop program included two exciting keynote talks, one on resource-efficient quantum computing by Prof. Fred Chong, University of Chicago, and another on the role of compilers in supporting modern cryptography-based computations by Dr. Madan Musuvathi from Microsoft Research. The program also included eight invited talks and a panel that covered a variety of frontiers and emerging topics in the field. The full program of the workshop can be found at <https://lcpc19.cc.gatech.edu/program>. The eight invited talks were:

1. 'HPVM: Performance, Portability, Programmability and Approximation for Heterogeneous Parallel Systems' by Vikram Adve, UIUC
2. 'Optimization in the Sparse Tensor Algebra Compiler' by Fredrik Kjolstad, MIT
3. 'Optimizing Data Movement and Achieving Performance Portability with Brick Data Layouts' by Mary Hall, University of Utah
4. 'DPC++: A Direct Programming Model for Diverse Architectures' by Xinmin Tian, Intel
5. 'Model-driven multi-level tiling for tensor contractions' by P. (Saday) Sadayappan, University of Utah
6. 'A framework for compositional transformations of recursions and loops' by Milind Kulkarni, Purdue University
7. 'A Lightweight Polyhedral Abstraction for MLIR' by Albert Cohen, Google
8. 'Filling in the gaps between applications and OpenMP for in-node programming on exascale systems' by Oscar Hernandez, Oak Ridge National Laboratory

The social part of the workshop included an excursion to Georgia Aquarium, the largest aquarium in the United States and in the Western Hemisphere. It is home to whale sharks, beluga whales, California sea lions, bottlenose dolphins, and manta rays that reside in 10 million gallons of fresh and saltwater (workshop photos are available at: <https://lcpc19.cc.gatech.edu/photos>).

A total of 48 students, faculty, and researchers participated in LCPC 2019. They contributed to interesting discussions in a close-knit setting during the workshop sessions, breaks, lunches, excursion, and dinner, while enjoying beautiful fall weather in Atlanta.

The LCPC 2019 organizers would like to thank our sponsors, Futurewei and Intel, for their generous support. We are grateful to the keynote and invited speakers, the authors, and all the participants for making LCPC 2019 a great success. We would especially like to thank Sri Raj Paul (Publicity and Web Chair) and Francella Tonge (Administrative Support), without whose hard work the workshop would not have been possible. We would also like to thank all the student volunteers. We thank Georgia Tech and GT Hotel for the use of their conference facilities and Springer-Verlag for providing publication support for the workshop proceedings through their LNCS series. Finally, we would like to thank the LCPC steering committee for providing us the

opportunity to organize LCPC 2019, and for their guidance and support throughout the entire process.

Sincerely,

Santosh Pande
Vivek Sarkar

Organization

Workshop Chairs

Santosh Pande	Georgia Tech
Vivek Sarkar	Georgia Tech

Steering Committee

Rudolf Eigenmann	U. Delaware
Alex Nicolau	UC Irvine
David Padua	UIUC
Lawrence Rauchwerger	TAMU, UIUC
Vivek Sarkar	Georgia Tech

Keynote/Panel Chair

Oscar Hernandez	Oak Ridge National Labs
-----------------	-------------------------

Program Committee

Milind Chabbi	Scalable Machines Research, Uber
Tiago Cogumbreiro	UMass Boston
Damian Dechev	University of Central Florida
Chen Ding	University of Rochester
Yufei Ding	UC Santa Barbara
Mary Hall	University of Utah
Akihiro Hayashi	Rice University
Changhee Jung	Purdue University
Martin Kong	University of Oklahoma
Tushar Kumar	Google
Seyong Lee	Oak Ridge National Laboratory
Zhiyuan Li	Purdue University
Lawrence Rauchwerger	TAMU, UIUC
Hari Sundar	University of Utah
Rishi Surendran	Xilinx Inc

Publicity/Web Chair

Sri Raj Paul	Georgia Tech
--------------	--------------

Contents

Performance of Static and Dynamic Task Scheduling for Real-Time Engine Control System on Embedded Multicore Processor	1
<i>Yoshitake Oki, Hiroki Mikami, Hikaru Nishida, Dan Umeda, Keiji Kimura, and Hironori Kasahara</i>	
PostSLP: Cross-Region Vectorization of Fully or Partially Vectorized Code	15
<i>Vasileios Porpodas and Pushkar Ratnalikar</i>	
FLARE: Flexibly Sharing Commodity GPUs to Enforce QoS and Improve Utilization	32
<i>Wei Han, Daniel Mawhirter, Bo Wu, Lin Ma, and Chen Tian</i>	
Foundations of Consistency Types for a Higher-Order Distributed Language.	49
<i>Xin Zhao and Philipp Haller</i>	
Common Subexpression Convergence: A New Code Optimization for SIMT Processors	64
<i>Sana Damani and Vivek Sarkar</i>	
Using Performance Event Profiles to Deduce an Execution Model of MATLAB with Just-In-Time Compilation	74
<i>Patryk Kiepas, Corinne Ancourt, Claude Tadonki, and Jarosław Koźlak</i>	
CLAM: Compiler Leasing of Accelerator Memory	89
<i>Dong Chen, Chen Ding, and Dorin Patru</i>	
Abstractions for Polyhedral Topology-Aware Tasking [Position Paper]	98
<i>Martin Kong</i>	
SWIRL ++ : Evaluating Performance Models to Guide Code Transformation in Convolutional Neural Networks	108
<i>Tharindu R. Patabandi, Anand Venkat, Rajkishore Barik, and Mary Hall</i>	
A Structured Grid Solver with Polyhedral+Dataflow Representation	127
<i>Eddie C. Davis, Catherine R. M. Olschanowsky, and Brian Van Straalen</i>	

CubeGen: Code Generation for Accelerated GEMM-Based Convolution
with Tiling 147
 Amarin Phaosawasdi, Christopher Rodrigues, Long Chen, and Peng Wu

Author Index 165