

2019 IEEE/ACM International Symposium on Code Generation and Optimization (CGO 2019)

**Washington, DC, USA
16 – 20 February 2019**



**IEEE Catalog Number: CFP19CGO-POD
ISBN: 978-1-7281-1437-8**

**Copyright © 2019 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP19CGO-POD
ISBN (Print-On-Demand):	978-1-7281-1437-8
ISBN (Online):	978-1-7281-1436-1

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

Contents

Frontmatter

Message from the General Chair	iii
Message from the Program Chairs	v
CGO 2019 Organization	vi
Report from the Artifact Evaluation Committee	x
Sponsors	xi
ACM Student Research Competition	xii

Keynote

Rethinking Compilation in a Heterogeneous World (Keynote) Michael O’Boyle — <i>University of Edinburgh, UK</i>	1
--	---

Research Papers

Binary Optimization

BOLT: A Practical Binary Optimizer for Data Centers and Beyond Maksim Panchenko, Rafael Auler, Bill Nell, and Guilherme Ottoni — <i>Facebook, USA</i>	2
Janus: Statically-Driven and Profile-Guided Automatic Dynamic Binary Parallelisation Ruoyu Zhou and Timothy M. Jones — <i>University of Cambridge, UK</i>	15

Bugs and Security

Smokestack: Thwarting DOP Attacks with Runtime Stack Layout Randomization Misiker Tadesse Aga and Todd Austin — <i>University of Michigan, USA</i>	26
Automatic Equivalence Checking for Assembly Implementations of Cryptography Libraries Jay P. Lim and Santosh Nagarakatte — <i>Rutgers University, USA</i>	37
CSOD: Context-Sensitive Overflow Detection Hongyu Liu, Sam Silvestro, Xiaoyin Wang, Lide Duan, and Tongping Liu — <i>University of Texas at San Antonio, USA</i>	50
Reasoning about the Node.js Event Loop using Async Graphs Haiyang Sun, Daniele Bonetta, Filippo Schiavio, and Walter Binder — <i>USI Lugano, Switzerland; Oracle Labs, USA</i>	61

GPUs and Tensors

Automatic Generation of Warp-Level Primitives and Atomic Instructions for Fast and Portable Parallel Reduction on GPUs Simon Garcia De Gonzalo, Sitao Huang, Juan Gómez-Luna, Simon Hammond, Onur Mutlu, and Wen-mei Hwu — <i>University of Illinois at Urbana-Champaign, USA; ETH Zurich, Switzerland; Sandia National Laboratories, USA</i>	73
A Code Generator for High-Performance Tensor Contractions on GPUs Jinsung Kim, Aravind Sukumaran-Rajam, Vineeth Thumma, Sriram Krishnamoorthy, Ajay Panyala, Louis-Noël Pouchet, Atanas Rountev, and P. Sadayappan — <i>Ohio State University, USA; Pacific Northwest National Laboratory, USA; Colorado State University, USA</i>	85

Potpourri

Transforming Query Sequences for High-Throughput B+ Tree Processing on Many-Core Processors Ruiqin Tian, Junqiao Qiu, Zhijia Zhao, Xu Liu, and Bin Ren — <i>College of William and Mary, USA; University of California at Riverside, USA</i>	96
Quantifying and Reducing Execution Variance in STM via Model Driven Commit Optimization Girish Mururu, Ada Gavrilovska, and Santosh Pande — <i>Georgia Institute of Technology, USA</i>	109

White-Box Program Tuning	
Wen-Chuan Lee, Yingqi Liu, Peng Liu, Shiqing Ma, Hongjun Choi, Xiangyu Zhang, and Rajiv Gupta — <i>Purdue University, USA; University of California at Riverside, USA</i>	122
Generation of In-Bounds Inputs for Arrays in Memory-Unsafe Languages	
Marcus Rodrigues, Breno Guimarães, and Fernando Magno Quintão Pereira — <i>Federal University of Minas Gerais, Brazil</i>	136
Code Generation	
Function Merging by Sequence Alignment	
Rodrigo C. O. Rocha, Pavlos Petoumenos, Zheng Wang, Murray Cole, and Hugh Leather — <i>University of Edinburgh, UK; Lancaster University, UK</i>	149
An Optimization-Driven Incremental Inline Substitution Algorithm for Just-in-Time Compilers	
Aleksandar Prokopec, Gilles Duboscq, David Leopoldseider, and Thomas Würthinger — <i>Oracle Labs, Switzerland; JKU Linz, Austria</i>	164
Tensor Algebra Compilation with Workspaces	
Fredrik Kjolstad, Peter Ahrens, Shoaib Kamil, and Saman Amarasinghe — <i>Massachusetts Institute of Technology, USA; Adobe, USA</i>	180
Kernel Optimization	
Tiramisu: A Polyhedral Compiler for Expressing Fast and Portable Code	
Riyadh Baghdadi, Jessica Ray, Malek Ben Romdhane, Emanuele Del Sozzo, Abdurrahman Akkas, Yunming Zhang, Patricia Suriana, Shoaib Kamil, and Saman Amarasinghe — <i>Massachusetts Institute of Technology, USA; Politecnico di Milano, Italy; Google, USA; Adobe, USA</i>	193
Super-Node SLP: Optimized Vectorization for Code Sequences Containing Operators and Their Inverse Elements	
Vasileios Porpodas, Rodrigo C. O. Rocha, Evgueni Brevnov, Luís F. W. Góes, and Timothy Mattson — <i>Intel, USA; University of Edinburgh, UK; PUC-MG, Brazil</i>	206
Locus: A System and a Language for Program Optimization	
Thiago S. F. X. Teixeira, Corinne Ancourt, David Padua, and William Gropp — <i>University of Illinois at Urbana-Champaign, USA; MINES ParisTech, France</i>	217
GPUs	
Decoding CUDA Binary	
Ari B. Hayes, Fei Hua, Jin Huang, Yanhao Chen, and Eddy Z. Zhang — <i>Rutgers University, USA</i>	229
From Loop Fusion to Kernel Fusion: A Domain-Specific Approach to Locality Optimization	
Bo Qiao, Oliver Reiche, Frank Hannig, and Jürgen Teich — <i>University of Erlangen-Nuremberg, Germany</i>	242
IGC: The Open Source Intel Graphics Compiler	
Anupama Chandrasekhar, Gang Chen, Po-Yu Chen, Wei-Yu Chen, Junjie Gu, Peng Guo, Shruthi Hebbur Prasanna Kumar, Guei-Yuan Lueh, Pankaj Mistry, Wei Pan, Thomas Raoux, and Konrad Trifunovic — <i>Intel, USA; Intel, Poland</i>	254
Student Research Competition	
<hr/>	
Undergraduate	
Automatic Parallelization of Irregular x86-64 Loops	
Brandon Neth and Michelle Mills Strout — <i>University of Arizona, USA</i>	266
Graduate	
A Shared BTB Design for Multicore Systems	
Moumita Das, Ansuman Banerjee, and Bhaskar Sardar — <i>Jadavpur University, India; Indian Statistical Institute, India</i>	267
Optimizing RNA-RNA Interaction Computations	
Swetha Varadarajan — <i>Colorado State University, USA</i>	269
Code Generation from Formal Models for Automatic RTOS Portability	
Renata Martins Gomes and Marcel Baunach — <i>Graz University of Technology, Austria</i>	271
Understanding RDMA Behavior in NUMA Systems	
Jacob Nelson and Roberto Palmieri — <i>Lehigh University, USA</i>	273
Translating Traditional SIMD Instructions to Vector Length Agnostic Architectures	
Sheng-Yu Fu and Wei-Chung Hsu — <i>National Taiwan University, Taiwan</i>	275

Accelerating GPU Computing at Runtime with Binary Optimization	
Guangli Li, Lei Liu, and Xiaobing Feng — <i>Institute of Computing Technology at Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China</i>	276
Extending LLVM for Lightweight SPMD Vectorization: Using SIMD and Vector Instructions Easily from Any Language	
Robin Kruppe, Julian Oppermann, Lukas Sommer, and Andreas Koch — <i>TU Darmstadt, Germany</i>	278
Multi-target Compiler for the Deployment of Machine Learning Models	
Oscar Castro-Lopez and Ines F. Vega-Lopez — <i>Autonomous University of Sinaloa, Mexico</i>	280
A Tool for Performance Analysis of GPU-Accelerated Applications	
Keren Zhou and John Mellor-Crummey — <i>Rice University, USA</i>	282
Kernel Fusion/Decomposition for Automatic GPU-Offloading	
Alok Mishra, Martin Kong, and Barbara Chapman — <i>Stony Brook University, USA; Brookhaven National Laboratory, USA</i>	283
Translating CUDA to OpenCL for Hardware Generation using Neural Machine Translation	
Yonghae Kim and Hyesoon Kim — <i>Georgia Institute of Technology, USA</i>	285
Author Index	287