

Zhiru Zhang

School of Electrical and Computer Engineering, Cornell University
office: 320 Rhodes Hall, Ithaca, NY 14853 **phone:** 607-255-5954
e-mail: zhiruz@cornell.edu **web:** <https://www.csl.cornell.edu/~zhiruz>

EDUCATION

University of California, Los Angeles, Los Angeles, CA

Ph.D., Computer Science, 2007 (with Highest Distinction)

- Thesis Topic: Behavior-Level Scheduling and Planning for Nanometer IC Designs
- Advisor: Professor Jason Cong

M.S., Computer Science, 2003

Peking University, Beijing, China

B.S., Computer Science, 2001 (with Highest Distinction)

PROFESSIONAL EXPERIENCE

Cornell University, Ithaca, NY

Associate Professor of Electrical and Computer Engineering Jul. 2018–present

Graduate Field Member of Computer Science Apr. 2013–present

Assistant Professor of Electrical and Computer Engineering Aug. 2012–Jun. 2018

Xilinx Inc., San Jose, CA

Jan. 2011–Aug. 2012

Software Development Manager

AutoESL Design Technologies Inc., Los Angeles, CA

May 2006–Jan. 2011

Co-founder, Director of Research and Development

AWARDS AND HONORS

- Elevation to IEEE Fellow, 2023.
- FPGA'04 Paper on ASIP Synthesis Inducted to TCFPGA Hall of Fame (Class of 2023), 2023.
- Michael Tien'72 Excellence in Teaching Award, College of Engineering, Cornell University, 2022.
- Best Paper Award, 30th International Symposium on Field-Programmable Gate Arrays, 2022.
- DAC'06 Paper on SDC Scheduling Inducted to TCFPGA Hall of Fame (Class of 2022), 2022.
- Best Paper Award Nomination, 58th Design Automation Conference (5 out of 916 submissions), 2021.
- Best Paper Award, 29th International Symposium on Field-Programmable Gate Arrays (1 out of 135 submissions), 2021.
- Best Paper Award Nomination, 29th International Symposium on Field-Programmable Gate Arrays (3 out of 135 submissions), 2021.
- IEEE HSTTC Top Picks in Hardware and Embedded Security, 2020.
- Facebook Research Award, 2020.
- Best Paper Award, 27th International Symposium on Field-Programmable Gate Arrays (1 out of 161 submissions), 2019.
- The Ruth and Joel Spira Award for Excellence in Teaching, 2018.
- Google Faculty Research Award, 2018.
- DAC Under-40 Innovators Award (named one of the top-five innovators), Design Automation Conference, 2018.
- Best Paper Award (short paper category), 26th International Symposium on Field-Programmable Custom Computing Machines, 2018.

- Rising Professional Achievement Award, UCLA Henry Samueli School of Engineering and Applied Science, 2018.
- Best Paper Nomination, 26th International Symposium on Field-Programmable Gate Arrays (4 out of 116 submissions), 2018.
- Best Paper Nomination, 25th International Symposium on Field-Programmable Gate Arrays (4 out of 101 submissions), 2017.
- Michael Tien '72 Excellence in Teaching Award, College of Engineering, Cornell University, 2016.
- DARPA Young Faculty Award (YFA), 2015.
- IEEE CEDA Ernest S. Kuh Early Career Award, 2015.
- NSF CAREER Award, National Science Foundation, 2015.
- Best Paper Award, ACM Transactions on Design Automation of Electronic Systems, 2012.
- Ross Freeman Award for Technical Innovation, Xilinx, 2012.
- Best Paper Nomination, 28th International Conference on Computer-Aided Design (14 out of 438 submissions), 2009.
- Outstanding Ph.D. Award, UCLA Computer Science Department, 2007.
- Phi Tau Phi Scholarship Award, Phi Tau Phi Scholastic Honor Society of America, 2005.

TEACHING EXPERIENCE

- **ECE/ENGRD 2300** – Digital Logic and Computer Organization
Spring 2022 (67 students), Spring 2021 (68 students), Spring 2018 (74 students), Spring 2017 (92 students), Spring 2016 (62 students), Spring 2015 (58 students), Spring 2014 (63 students)
- **ECE 5775** – High-Level Digital Design Automation
Fall 2022 (44 students), Fall 2018 (34 students), Fall 2017 (20 students), Fall 2016 (30 students), Fall 2015 (36 students), Fall 2014 (33 students), Fall 2013 (17 students)
- **ECE 6980** – Special Topics on Hardware Acceleration of Deep Learning
Spring 2023 (9 students)
- **ECE 5997** – Special Topics on Hardware Accelerator Design and Automation
Fall 2021 (9 students)
- **ENGRG 1050** – Freshman Engineering Seminar
Fall 2022 (20 students), Fall 2017 (21 students), Fall 2013 (18 students)
- **ECE 5950** – Special Topics on High-Level Digital Design Automation
Spring 2013 (14 students)

EXTERNAL PROFESSIONAL ACTIVITIES

Conference Organizing Committee Member

- The 31st ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA), *Program Chair*, 2023.
- The 30th IEEE International Symposium On Field-Programmable Custom Computing Machines (FCCM), *General Chair*, 2022.
- The 30th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA), *Publications Chair*, 2022.
- The 30th International Conference on Application-specific Systems, Architectures and Processors (ASAP), *General Chair*, 2019.
- The 29th International Conference on Application-specific Systems, Architectures and Processors (ASAP), *Technical Program Co-Chair*, 2018.
- The 28th International Conference on Application-specific Systems, Architectures and Processors (ASAP), *Publicity Chair*, 2017.
- The 25th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA), *Publicity Chair*, 2017.
- Electronic System Level Synthesis Conference (ESLSyn), *Technical Program Chair*, 2015.

Conference Technical Program Committee Member

- Design Automation Conference (DAC) 2016-2018, 2022 (track chair).
- Design, Automation, and Test in Europe (DATE) 2014, 2020, 2021.
- International Symposium On Field-Programmable Custom Computing Machines (FCCM) 2019-present.
- International Symposium on Field-Programmable Gate Arrays (FPGA) 2015-present.
- The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), 2022.
- Conference on Neural Information Processing Systems (NeurIPS), 2020.
- International Workshop on Logic & Synthesis (IWLS) 2018-2020.
- International Symposium on Design, Automation and Test (VLSI-DAT), 2021.
- International Conference on Application-specific Systems, Architectures and Processors (ASAP) 2017-2019.
- Asia and South Pacific Design Automation Conference (ASP-DAC) 2013-2015, 2017-2019 (track chair).
- International Conference on Computer-Aided Design (ICCAD) 2015-2017.
- International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS) 2015-2016.
- International Conference on Computing Frontiers (CF) 2016.
- International Conference on Computer Design (ICCD) 2015.
- International Conference on Field-Programmable Logic and Applications (FPL) 2016.
- International Symposium on Low Power Electronics and Design (ISLPED) 2016.
- IEEE International Symposium on Circuits and Systems (ISCAS) 2011-2015.
- International Conference on Very Large Scale Integration (VLSI-SOC) 2011, 2015.
- Electronic System Level Synthesis Conference (ESLSyn) 2012-2015.
- International Symposium on Application Accelerators in High Performance Computing (SAAHPC) 2009-2012.

Journal Editor

- IEEE Design and Test, *Associate Editor*, 2022-present.
- ACM Transactions on Reconfigurable Technology and Systems (TRETs), *Associate Editor*, 2019-present.
- IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), *Associate Editor*, 2019-present.
- IEEE Circuits and Systems Magazine (CAS-M), Special Issue on FPGA-Based Computing, *Lead Guest Editor*, 2020-2021.
- IEEE Design and Test, Special Issue on Machine Intelligence at the Edge, *Guest Editor*, 2019.
- Integration, the VLSI Journal, ASP-DAC 2019 Special Issue, *Guest Editor*, 2019.
- Springer Journal of Signal Processing Systems (JSPS), *Associate Editor*, 2017-present.
- Springer Journal on Design Automation for Embedded Systems (DAEM), *Associate Editor*, 2016-2019.
- Journal of Electrical and Computer Engineering, Special Issue on ESL Design Methodology (JECE), *Guest Editor*, 2011-2012.

Journal Reviewer

- ACM Transactions on Design Automation of Electronic Systems (TODAES)
- ACM Transactions on Architecture and Code Optimization (TACO)
- ACM Transactions on Reconfigurable Technology and Systems (TRETs)
- IEEE Transactions on Computers (TC)
- IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)
- IEEE Transactions on Signal Processing (T-SP)
- IEEE Computer Architecture Letters (CAL)
- IEEE Transactions on Very Large Scale Integration Systems (TVLSI)

- IET Computers & Digital Techniques (IET-CDT)
- IEEE Design & Test Magazine
- Journal of Parallel and Distributed Computing (JPDC)

Professional Society Memberships

- IEEE Fellow, 2023-
- IEEE Senior Member, 2018-2022
- ACM Member, 2015-present
- ACM SIGDA Member, 2012-present

PUBLICATIONS

Conference Publications

- E. S. Alcorta, A. Gerstlauer, C. Deng, Q. Sun, Z. Zhang, C. Xu, L. W. Wills, D. S. Lopera, W. Ecker, S. Garg, and J. Hu, Machine Learning for Embedded System Design. *International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, Sep. 2023. (Invited Paper)
- N. Lazarev, T. Ji, A. Kalia, D. Kim, I. Marinos, F. Y. Yan, C. Delimitrou, Z. Zhang, and A. Akella, Resilient Baseband Processing in Virtualized RANs with Slingshot, *Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, Sep. 2023.
- E. Ustun, C. Yu, and Z. Zhang. Equality Saturation for Datapath Synthesis: A Pathway to Pareto Optimality. *Design Automation Conference (DAC)*, Jul. 2023. (Invited Perspective Paper)
- Z. Zhang, M. Hofmann, and A. Butt. A Case for Open EDA Verticals. *International Symposium on Physical Design (ISPD)*, Mar. 2023. (Invited Perspective Paper)
- C. Deng, X. Li, Z. Feng, and Z. Zhang. GARNET: Reduced-Rank Topology Learning for Robust and Scalable Graph Neural Networks. *Learning on Graphs Conference (LoG)*, Dec. 2022. (Spotlight Paper)
- T. Yu, Y. Zhang, Z. Zhang, and C. De Sa. Understanding Hyperdimensional Computing for Parallel Single-Pass Learning. *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)*, Nov./Dec. 2022.
- W. Hua, M. Umar, Z. Zhang, and G. E. Suh. GuardNN: Secure Accelerator Architecture for Privacy-Preserving Deep Learning. *Design Automation Conference (DAC)*, Jul. 2022.
- D. Pal, Y.-H. Lai, S. Xiang, N. Zhang, H. Chen, J. Casas, P. Cocchini, Z. Yang, Jin Yang, L.-N. Pouchet, and Z. Zhang. Accelerator Design with Decoupled Hardware Customizations: Benefits and Challenges. *Design Automation Conference (DAC)*, Jul. 2022. (Invited Paper)
- Y. Zhang, Z. Zhang, and L. Lew. PokeBNN: A Binary Pursuit of Lightweight Accuracy. *The Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2022.
- M. Umar, W. Hua, Z. Zhang, and G. E. Suh. SoftVN: Efficient Memory Protection via Software-Provided Version Numbers. *International Symposium on Computer Architecture (ISCA)*, Jun. 2022.
- W. Hua, M. Umar, Z. Zhang, and G. E. Suh. MGX: Near-Zero Overhead Memory Protection for Data-Intensive Accelerators. *International Symposium on Computer Architecture (ISCA)*, Jun. 2022.
- E. Ustun, I. San, J. Yin, C. Yu, and Z. Zhang. IMpress: Large Integer Multiplication Expression Rewriting for FPGA HLS. *International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, May 2022.
- S. Xiang, Y.-H. Lai, Y. Zhou, H. Chen, N. Zhang, D. Pal, and Z. Zhang. HeteroFlow: An Accelerator Programming Model with Decoupled Data Placement for Software-Defined FPGAs. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb./Mar. 2022.
- Y. Du, Y. Hu, and Z. Zhang. High-Performance Sparse Linear Algebra on HBM-

Equipped FPGAs Using HLS: A Case Study on SpMV. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb./Mar. 2022.

- L. Guo, P. Maidee, Y. Zhou, C. Lavin, J. Wang, Y. Chi, W. Qiao, A. Kaviani, Z. Zhang, and J. Cong. RapidStream: Parallel Physical Implementation of FPGA HLS Designs. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb./Mar. 2022. **(Best Paper Award)**
- W. Hua, Y. Zhang, C. Guo, Z. Zhang, and G. E. Suh. BulletTrain: Accelerating Robust Neural Network Training via Boundary Example Mining. *Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS)*, Dec. 2021.
- Y. Zhou, H. Wang, J. Yin, and Z. Zhang. Distilling Arbitration Logic from Traces using Machine Learning: A Case Study on NoC. *Design Automation Conference (DAC)*, Dec. 2021. **(Best Paper Nominee)**
- W. Cheng, C. Deng, Z. Zhao, Y. Cai, Z. Zhang, and Z. Feng. SPADE: A Spectral Method for Black-Box Adversarial Robustness Evaluation. *International Conference on Machine Learning (ICML)*, Jul. 2021.
- N. Lazarev, S. Xiang, N. Adit, Z. Zhang, and C. Delimitrou. Dagger: Efficient and Fast RPCs in Cloud Microservices with Near-Memory Reconfigurable NICs. *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Apr. 2021. **(IEEE Micro Top Picks Honorable Mention)**
- Y. Zhang, J. Pan, X. Liu, H. Chen, and Z. Zhang. FracBNN: Accurate and FPGA-Efficient Binary Neural Networks with Fractional Activations. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2021. **(Best Paper Nominee)**
- L. Guo, Y. Chi, J. Wang, J. Lau, W. Qiao, E. Ustun, Z. Zhang, and J. Cong. Auto-Bridge: Coupling Coarse-Grained Floorplanning and Pipelining for High-Frequency HLS Design on Multi-Die FPGAs. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2021. **(Best Paper Award)**
- J. Jiao, D. Pal, C. Deng, and Z. Zhang. GLAIVE: Graph Learning Assisted Instruction Vulnerability Estimation. *Design, Automation and Test in Europe Conference (DATE)*, Feb. 2021.
- Y. Hu, Z. Ye, M. Wang, J. Yu, D. Zheng, M. Li, Z. Zhang, Z. Zhang, and Y. Wang. FeatGraph: A Flexible and Efficient Backend for Graph Neural Network Systems. *International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Nov. 2020.
- N. Srivastava, H. Jin, J. Liu, D. Albonesi, and Z. Zhang. MatRaptor: A Sparse-Sparse Matrix Multiplication Accelerator Based on Row-Wise Product. *International Symposium on Microarchitecture (MICRO)*, Oct. 2020.
- L. Guo, J. Lau, Y. Chi, J. Wang, C. Yu, Z. Chen, Z. Zhang, and J. Cong. Analysis and Optimization of the Implicit Broadcasts in FPGA HLS to Improve Maximum Frequency. *Design Automation Conference (DAC)*, Jul. 2020.
- E. Singh, F. Lonsing, S. Chattopadhyay, M. Strange, P. Wei, X. Zhang, Y. Zhou, J. Cong, D. Chen, Z. Zhang, P. Raina, C. Barrett, and S. Mitra. A-QED Verification of Hardware Accelerators. *Design Automation Conference (DAC)*, Jul. 2020.
- R. Nigam, S. Atapattu, S. Thomas, Z. Li, T. Bauer, Y. Ye, A. Koti, A. Sampson, and Z. Zhang. Predictable Design of FPGA Accelerators using Time-Sensitive Affine Types. *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Jun. 2020. **(Selected into MIT PL Review 2023)**
- C. Deng, Z. Zhao, Y. Wang, Z. Zhang, and Z. Feng. GraphZoom: A Multi-Level Spectral Approach for Accurate and Scalable Graph Embedding. *International Conference on Learning Representations (ICLR)*, Apr. 2020. **(Oral)**
- Y. Zhang, R. Zhao, W. Hua, N. Xu, G. E. Suh, and Z. Zhang. Precision Gating: Improving Neural Network Efficiency with Dynamic Dual-Precision Activations. *International Conference on Learning Representations (ICLR)*, Apr. 2020.

- N. Srivastava, H. Jin, S. Smith, H. Rong, D. Albonesi, and Z. Zhang. Tensaurus: A Versatile Accelerator for Mixed Sparse-Dense Tensor Computations. *International Symposium on High-Performance Computer Architecture (HPCA)*, Feb. 2020.
- W. Hua, Y. Zhou, C. De Sa, Z. Zhang, and G. E. Suh. Channel Gating Neural Networks. *Thirty-third Conference on Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
- W. Hua, Y. Zhou, C. De Sa, Z. Zhang, and G. E. Suh. Boosting the Performance of CNN Accelerators with Dynamic Fine-Grained Channel Gating. *International Symposium on Microarchitecture (MICRO)*, Oct. 2019.
- R. Zhao, Y. Hu, J. Dotzel, C. De Sa, and Z. Zhang. Building Efficient Deep Neural Networks with Unitary Group Convolutions. *The Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2019.
- R. Zhao, Y. Hu, J. Dotzel, C. De Sa, and Z. Zhang. Improving Neural Network Quantization without Retraining using Outlier Channel Splitting. *International Conference on Machine Learning (ICML)*, Jun. 2019.
- A. Rovinski, C. Zhao, K. Al-Hawaj, P. Gao, S. Xie, C. Torng, S. Davidson, A. Amarnath, L. Vega, B. Veluri, A. Rao, T. Ajayi, J. Puskar, S. Dai, R. Zhao, D. Richmond, Z. Zhang, I. Galton, C. Batten, M. Taylor, and R. Dreslinski. A 1.4 GHz 695 Giga RISC-V Inst/s 496-core Manycore Processor with Mesh On-Chip Network and an All-Digital Synthesized PLL in 16nm CMOS. *Symposium on VLSI Circuits (VLSI)*, Jun. 2019.
- Y. Zhou, H. Ren, Y. Zhang, B. Keller, B. Khailany, and Z. Zhang. PRIMAL: Power Inference using Machine Learning. *Design Automation Conference (DAC)*, Jun. 2019.
- C. Yu and Z. Zhang. Painting on Placement: Forecasting Routing Congestion using Conditional Generative Adversarial Nets. *Design Automation Conference (DAC)*, Jun. 2019.
- Z. Jiang, H. Jin, G. E. Suh, and Z. Zhang. Designing Secure Cryptographic Accelerators with Information Flow Enforcement: A Case Study on AES. *Design Automation Conference (DAC)*, Jun. 2019.
- S. Dai and Z. Zhang. Improving Scalability of Exact Modulo Scheduling with Specialized Conflict-Driven Learning. *Design Automation Conference (DAC)*, Jun. 2019.
- G. Liu, J. Primmer, and Z. Zhang. Rapid Generation of High-Quality RISC-V Processors from Functional Instruction Set Specifications. *Design Automation Conference (DAC)*, Jun. 2019.
- N. Srivastava, H. Rong, P. Barua, G. Feng, H. Cao, Z. Zhang, D. Albonesi, V. Sarkar, W. Chen, P. Petersen, G. Lowney, A. Herr, C. Hughes, T. Mattson, and P. Dubey. Productively Generating High-Performance Spatial Hardware for Dense Tensor Computations. *International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, Apr./May 2019.
- E. Ustun, S. Xiang, J. Gui, C. Yu, and Z. Zhang. LAMDA: Learning-Assisted Multi-Stage Autotuning for FPGA Design Closure. *International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, Apr./May 2019.
- Y.-H. Lai, Y. Chi, Y. Hu, J. Wang, C. H. Yu, Y. Zhou, J. Cong, and Z. Zhang. HeteroCL: A Multi-Paradigm Programming Infrastructure for Software-Defined Reconfigurable Computing. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2019. **(Best Paper Award)**
- Z. Jiang, S. Dai, G. E. Suh, and Z. Zhang. High-Level Synthesis with Timing-Sensitive Information Flow Enforcement. *International Conference On Computer Aided Design (ICCAD)*, Nov. 2018.
- W. Hua, Z. Zhang, and G. E. Suh. Reverse Engineering Convolutional Neural Networks Through Side-channel Information Leaks. *Design Automation Conference (DAC)*, Jun. 2018.

- S. Dai, Y. Zhou, H. Zhang, E. Ustun, E. F.Y. Young, and Z. Zhang. Fast and Accurate Estimation of Quality of Results in High-Level Synthesis with Machine Learning. *International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, May 2018. (**Best Paper Award**, short paper category)
- S. Dai, G. Liu, and Z. Zhang. A Scalable Approach to Exact Resource-Constrained Scheduling Based on a Joint SDC and SAT Formulation. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2018. (**Best Paper Nominee**)
- Y. Zhou, U. Gupta, S. Dai, R. Zhao, N. Srivastava, H. Jin, J. Featherston, Y.-H. Lai, G. Liu, G. Velasquez, W. Wang, and Z. Zhang. Rosetta: A Realistic High-Level Synthesis Benchmark Suite for Software Programmable FPGAs. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2018.
- G. Liu and Z. Zhang. Statistically Certified Approximate Logic Synthesis. *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2017.
- S. Dai, G. Liu, R. Zhao, and Z. Zhang. Enabling Adaptive Loop Pipelining in High-Level Synthesis. *52nd Annual Asilomar Conference on Signals, Systems, and Computers*, Oct. 2017. (Invited Paper)
- T. Ajayi, K. Al-Hawaj, A. Amarnath, S. Dai, S. Davidson, P. Gao, G. Liu, A. Lotfi, J. Puscar, A. Rao, A. Rovinski, L. Salem, N. Sun, C. Torng, L. Vega, B. Veluri, X. Wang, S. Xie, C. Zhao, R. Zhao, C. Batten, R. Dreslinski, I. Galton, R. Gupta, P. Mercier, M. Srivastava, M. Taylor, and Z. Zhang. Celerity: An Open-Source RISC-V Tiered Accelerator Fabric. *ACM/IEEE Symposium on High-Performance Chips (HOTCHIPS)*, Aug. 2017.
- J.H. Lin, T. Xing, R. Zhao, Z. Zhang, M. Srivastava, Z. Tu, and R. Gupta. Binarized Neural Networks with Separable Filters for Efficient Hardware Acceleration. *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Jul. 2017.
- E. Bartz, J. Chaves, Y. Gershtein, E. Halkiadakis, M. Hildreth, S. Kyriacou, K. Lannon, A. Lefeld, A. Ryd, L. Skinnari, R. Stone, C. Strohman, Z. Tao, B. Winer, P. Wittich, Z. Zhang, and M. Zientek. FPGA-based Real-time Charged Particle Trajectory Reconstruction at the Large Hadron Collider. *International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, May 2017.
- G. Liu and Z. Zhang. A Parallelized Iterative Improvement Approach to Area Optimization for LUT-Based Technology Mapping. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2017. (**Best Paper Nominee**)
- Y. Zhou, K. Al-Hawaj, and Z. Zhang. A New Approach to Automatic Memory Banking using Trace-Based Address Mining. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2017.
- R. Zhao, W. Song, W. Zhang, T. Xing, J.-H. Lin, M. Srivastava, R. Gupta, and Z. Zhang. Accelerating Binarized Convolutional Neural Networks with Software-Programmable FPGAs. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2017.
- C. Xu, G. Liu, R. Zhao, S. Yang, G. Luo, and Z. Zhang. A Parallel Bandit-Based Approach for Autotuning FPGA Compilation. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2017.
- S. Dai, R. Zhao, G. Liu, S. Srinath, U. Gupta, C. Batten, and Z. Zhang. Dynamic Hazard Resolution for Pipelining Irregular Loops in High-Level Synthesis. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2017.
- N. Srivastava, S. Dai, R. Manohar, and Z. Zhang. Accelerating Face Detection on Programmable SoC Using C-Based Synthesis. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2017.
- R. Zhao, G. Liu, S. Srinath, C. Batten, and Z. Zhang. Improving High-Level Synthesis with Decoupled Data Structure Optimization. *Design Automation Conference (DAC)*, Jun. 2016.
- A. Majumdar, Z. Zhang, and D. Albonesi. Characterizing the Benefits and Limi-

- tations of Smart Building Meeting Room Scheduling. *International Conference on Cyber-Physical Systems (ICCPS)*, Apr. 2016.
- D. Chen, J. Cong, S. Gurumani, W-M. Hwu, K. Rupnow, and Z. Zhang. System Synthesis and Automated Verification: Design Demands for IoT Devices. *Sensors to Cloud Architectures Workshop (SCAW)*, Mar. 2016.
 - F. Koushanfar, A. Mirhoseini, G. Qu, and Z. Zhang. DA Systemization of Knowledge: A Catalog of Prior Forward-Looking Initiatives. *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2015. (Invited Paper)
 - M. Tan, G. Liu, R. Zhao, S. Dai, and Z. Zhang. ElasticFlow: A Complexity-Effective Approach for Pipelining Irregular Loop Nests. *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2015.
 - G. Liu and Z. Zhang. A Reconfigurable Analog Substrate for Highly Efficient Maximum Flow Computation. *Design Automation Conference (DAC)*, Jun. 2015.
 - R. Zhao, M. Tan, S. Dai, and Z. Zhang. Area-Efficient Pipelining for FPGA-Targeted High-Level Synthesis. *Design Automation Conference (DAC)*, Jun. 2015.
 - M. Tan, S. Dai, U. Gupta, and Z. Zhang. Mapping-Aware Constrained Scheduling for LUT-Based FPGAs. *International Symposium on Field-Programmable Gate Arrays (FPGA)*, Feb. 2015.
 - S. Srinath, B. Ilbeyi, M. Tan, G. Liu, Z. Zhang, and C. Batten. Architectural Specialization for Inter-Iteration Loop Dependence Patterns. *International Symposium on Microarchitecture (MICRO)*, Dec. 2014.
 - M. Tan, B. Liu, S. Dai, and Z. Zhang. Multithreaded Pipeline Synthesis for Data-Parallel Kernels. *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2014.
 - G. Liu, Y. Tao, M. Tan, and Z. Zhang. CASA: Correlation-Aware Speculative Adders. *International Symposium on Low Power Electronics and Design (ISLPED)*, Aug. 2014.
 - S. Dai, M. Tan, K. Hao, and Z. Zhang. Flushing-Enabled Loop Pipelining for High-Level Synthesis. *Design Automation Conference (DAC)*, Jun. 2014.
 - Z. Zhang and B. Liu. SDC-Based Modulo Scheduling for Pipeline Synthesis. *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2013.
 - Z. Zhang and D. Chen. Challenges and Opportunities of ESL Design Automation. *IEEE International Conference on Solid-State and Integrated Circuit Technology*, Oct. 2012. (Invited Paper)
 - J. Zhang, Z. Zhang, S. Zhou, M. Tan, X. Liu, X. Cheng, and J. Cong. Bit-Level Optimization for High-Level Synthesis and FPGA-Based Acceleration. *International Symposium on FPGAs (FPGA)*, Feb. 2010.
 - J. Cong, B. Liu, and Z. Zhang. Scheduling with Soft Constraints. *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2009. (**Best Paper Nominee**)
 - J. Cong, B. Liu, and Z. Zhang. Behavior-Level Observability Don't-Cares and Application to Low-Power Behavioral Synthesis. *International Symposium on Low Power Electronics and Design (ISLPED)*, Aug. 2009.
 - J. Cong, K. Gururaj, B. Liu, C. Liu, Z. Zhang, S. Zhou, and Y. Zou. Evaluation of Static Analysis Techniques for Fixed-Point Precision Optimization. *IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM)*, Apr. 2009.
 - W. Jiang, Z. Zhang, M. Potkonjak, and J. Cong. Scheduling with Integer Time Budgeting for Low-Power Optimization. *Asia and South Pacific Design Automation Conference (ASP-DAC)*, Jan. 2008.
 - C.T. Hsieh, J. Cong, S.C. Chang, and Z. Zhang. Behavioral Synthesis with Activating Unused Flip-Flops for Reducing Glitch Power in FPGA. *Asia and South Pacific Design Automation Conference (ASP-DAC)*, Jan. 2008.
 - D. Chen, J. Cong, Y. Fan, and Z. Zhang. High-Level Power Estimation and Low-Power Design Space Exploration for FPGAs. *Asia and South Pacific Design Au-*

- tomation Conference (ASP-DAC), Jan. 2007.
- J. Cong, Y. Fan, G. Han, W. Jiang, and Z. Zhang. Platform-Based Behavior-Level and System-Level Synthesis. *IEEE International SOC Conference (SOCC)*, Sep. 2006. (Invited Paper)
 - J. Cong and Z. Zhang. An Efficient and Versatile Scheduling Algorithm Based on SDC Formulation. *Design Automation Conference (DAC)*, Jul. 2006. (**TCFPGA Hall of Fame – Class of 2022**)
 - J. Cong, Y. Fan, G. Han, W. Jiang, and Z. Zhang. Behavior and Communication Co-Optimization for Systems with Sequential Communication Media. *Design Automation Conference (DAC)*, Jul. 2006.
 - J. Cong, G. Han, and Z. Zhang. Architecture and Compilation for Data Bandwidth Improvement in Configurable Embedded Processors. *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2005.
 - J. Cong, Y. Fan, G. Han, A. Jagannathan, G. Reinman, and Z. Zhang. Instruction Set Extension with Shadow Registers for Configurable Processors. *International Symposium on FPGAs (FPGA)*, Feb. 2005.
 - J. Cong, Y. Fan, G. Han, Y. Lin, J. Xu, Z. Zhang, and X. Cheng. Bitwidth-Aware Scheduling and Binding in High-Level Synthesis. *Asia and South Pacific Design Automation Conference (ASP-DAC)*, Jan. 2005.
 - J. Cong, Y. Fan, and Z. Zhang. Architecture-Level Synthesis for Automatic Interconnect Pipelining. *Design Automation Conference (DAC)*, Jun. 2004.
 - J. Cong, Y. Fan, G. Han, and Z. Zhang. Application-Specific Instruction Generation for Configurable Processor Architectures. *International Symposium on FPGAs (FPGA)*, Feb. 2004. (**TCFPGA Hall of Fame – Class of 2023**)
 - J. Cong, Y. Fan, G. Han, X. Yang, and Z. Zhang. Architectural Synthesis Integrated with Global Placement for Multi-Cycle Communication. *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2003.
 - Z. Zhang, Y. Fan, M. Potkonjak, and J. Cong. Gradual Relaxation Technique with Application to Behavioral Synthesis. *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2003.
 - J. Cong, Y. Fan, X. Yang, and Z. Zhang. Architecture and Synthesis for Multi-Cycle Communication. *International Symposium on Physical Design (ISPD)*, Apr. 2003. (Invited Paper)

Journal Publications

- S. K. Venkataramanaiah, J. Meng, H. Suh, I. Yeo, J. Saikia, S. K. Cherupally, Y. Zhang, Z. Zhang, and J. Seo. A 28nm 8-bit Floating-Point Tensor Core based CNN Training Processor with Dynamic Activation/Weight Sparsification. *IEEE Journal of Solid-State Circuits (JSSC)*, Jul. 2023.
- L. Guo, P. Maidee, Y. Zhou, C. Lavin, E. Hung, W. Li, J. Lau, W. Qiao, Y. Chi, L. Song, Y. Xiao, A. Kaviani, Z. Zhang, and J. Cong. RapidStream 2.0: Automated Parallel Implementation of Latency Insensitive FPGA Designs Through Partial Reconfiguration. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, Apr. 2023.
- J. Liu, Z. Zhao, Z. Ding, B. Brock, H. Rong, and Z. Zhang. An Intermediate Language for General Sparse Format Customization. *IEEE Computer Architecture Letters (CAL)*, Mar. 2023.
- J. Cong, J. Lau, G. Liu, S. Neuendorffer, P. Pan, K. Vissers, and Z. Zhang. FPGA HLS Today: Successes, Challenges, and Opportunities. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, Dec. 2022. (**Keynote Paper**)
- W. Hua, Z. Zhang, and G. E. Suh. Reverse Engineering CNN Models using Side-Channel Attacks. *IEEE Design & Test*, Aug. 2022. (**IEEE HSTTC Top Picks in Hardware and Embedded Security**)

- L. Cheng, P. Pan, Z. Zhao, K. Ranjan, J. Weber, B. Veluri, S. Ehsani, M. Ruttenberg, D. Jung, P. Ivanov, D. Richmond, M. Taylor, Z. Zhang, and C. Batten. A Tensor Processing Framework for CPU-Manycore Heterogeneous Systems. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, Jun. 2022.
- Y.-H. Lai, E. Ustun, S. Xiang, Z. Fang, H. Rong, and Z. Zhang. Programming and Synthesis for Software-Defined FPGA Acceleration: Status and Future Prospects. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, Dec. 2021. (Invited Paper)
- C. Hao, J. Dotzel, J. Xiong, L. Benini, Z. Zhang, and D. Chen. Enabling Design Methodologies and Future Trends for Edge AI: Specialization and Co-Design. *IEEE Design & Test*, Aug. 2021.
- N. Lazarev, N. Adit, S. Xiang, Z. Zhang, and C. Delimitrou. Dagger: Towards Efficient RPCs in Cloud Microservices with Near-Memory Reconfigurable NICs. *IEEE Computer Architecture Letters (CAL)*, Aug. 2020.
- A. Rovinski, C. Zhao, K. Al-Hawaj, P. Gao, S. Xie, C. Torng, S. Davidson, A. Amarnath, L. Vega, B. Veluri, A. Rao, T. Ajayi, J. Puscar, S. Dai, R. Zhao, D. Richmond, Z. Zhang, I. Galton, C. Batten, M. Taylor, and R. Dreslinski. Evaluating Celerity: A 16nm 695 Giga-RISC-V Instructions/s Manycore Processor with Synthesizable PLL, *IEEE Solid-State Circuits Letters (SSC-L)*, Dec. 2019.
- G. Liu and Z. Zhang. PIMap: A Flexible Framework for Improving LUT-Based Technology Mapping via Parallelized Iterative Optimization. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, Jan. 2019.
- S. Davidson, S. Xie, C. Torng, K. Al-Hawai, A. Rovinski, T. Ajayi, L. Vega, C. Zhao, R. Zhao, S. Dai, A. Amarnath, B. Veluri, P. Gao, A. Rao, G. Liu, R. Gupta, Z. Zhang, R. Dreslinski, C. Batten, and M. Taylor. The Celerity Open-Source 511-Core RISC-V Tiered Accelerator Fabric: Fast Architectures and Design Methodologies for Fast Chips. *IEEE Micro*, Apr. 2018.
- G. Liu, M. Tan, S. Dai, R. Zhao, and Z. Zhang. Architecture and Synthesis for Area-Efficient Pipelining of Irregular Loop Nests. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, Feb. 2017.
- D. Chen, J. Cong, S. Gurumani, W.-M. Hwu, K. Rupnow, and Z. Zhang. Platform Choices and Design Demands for IoT Platforms: Cost, Power and Performance Tradeoffs. *IET Cyber-Physical Systems: Theory & Applications (IET-CPS)*, Nov. 2016.
- Z. Zhang, D. Chen, S. Dai, and K. Campbell. High-Level Synthesis for Low-Power Design. in *IPSS Transactions on System LSI Design Methodology (T-SLDM)*, Feb. 2015. (Invited Paper)
- J. Cong, B. Liu, S. Neuendorffer, J. Noguera, K. Vissers, and Z. Zhang. High-Level Synthesis for FPGAs: From Prototyping to Deployment. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*. 30(4):473–491. Apr. 2011. **(Keynote Paper)**
- J. Cong, B. Liu, R. Majumdar, and Z. Zhang. Behavior-Level Observability Analysis and Operation Gating in Low-Power Behavioral Synthesis. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 16(1):1–29. Nov. 2010. **(Best Paper Award)**
- J. Cong, G. Han, and Z. Zhang. Architecture and Compiler Optimization for Data Bandwidth Improvement in Configurable Processors. *IEEE Transaction on Very Large Scale Integration Systems (TVLSI)*, 14(9):986–997. Sep. 2006.
- J. Cong, Y. Fan, G. Han, X. Yang, and Z. Zhang. Architecture and Synthesis for On-Chip Multicycle Communication. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 23(4):550–564. Apr. 2004.

Book Chapters

- D. Pal, C. Deng, E. Ustun, C. Yu, and Z. Zhang. Machine Learning for Agile FPGA Design. *Machine Learning Applications in Electronic Design Automation* ed. H. Ren and J. Hu, Springer, 2022.
- N. Srivastava, G. Liu, Y.-H. Lai, and Z. Zhang. FPGA-Specific Compilers. *Handbook of Computer Architecture* ed. A. Chattopadhyay, Springer, 2022.
- Z. Zhang, Y. Fan, W. Jiang, G. Han, C. Yang, and J. Cong. AutoPilot: A Platform-Based ESL Synthesis System. *High-Level Synthesis: From Algorithm to Digital Circuit* ed. P. Coussy and A. Morawiec, Springer Publishers, 2008.

PATENTS

- *Soft Constraints in Scheduling*. U.S. Patent 8296710, (co-invented with B. Liu and J. Cong,) issued Oct. 2012.
- *Performance Visualization System*. U.S. Patent 7784008, (co-invented with M. Hutten and D. Karchmer,) issued Aug. 2010.

INVITED TALKS

- *High-Level Synthesis: Now is the time or will it continue to remain just a promising technology?*. DAC'23 Panel on HLS, San Francisco, CA, Jul. 2023.
- *Equality Saturation for Datapath Synthesis: A Pathway to Pareto Optimality* (Lightning Talk). DAC'23 Technical Session on "Breaking News from RTL Technology-Independent Optimization", San Francisco, CA, Jul. 2023.
- *Co-Design of Binarized Deep Learning*. CSTIC'23 Symposium IX, Virtual, Jun. 2023. (Keynote)
- *A Case for Open EDA Verticals*. ISPD'23 Panel on "EDA for Domain Specific Computing", Virtual, Mar. 2023.
- *Accelerator Synthesis for Agile Hardware Specialization: A New Dawn*. Amazon Tech Talk, Virtual, Dec. 2022.
- *Accelerator Design with Decoupled Hardware Customizations: Benefits and Challenges*. DAC Special Session on "New Perspective of HLS", San Francisco, CA, Jul. 2022.
- *A Pursuit of Efficient and Accurate Binary Neural Networks*. NSF A3D3 Institute, Virtual, May 2022.
- *Near-Memory Hardware Specialization for Fast and Efficient Sparse Processing*. DATE DATA-DREAM Workshop, Virtual, Mar. 2022.
- *Pruning on the Fly: Efficient Deep Learning with Adaptive Fine-Grained Optimizations*. Rutgers REFAI Seminar, Nov. 2021.
- *Co-Design for Low-Bitwidth Neural Networks with Dynamic Quantization*. ICCAD Halo Workshop, Virtual, Nov. 2021.
- *Faster, Slimmer, Smarter: Machine Learning for Agile Hardware Specialization*. MLCAD Workshop, Sep. 2021. (Plenary Talk)
- *Synthesis and Programming for Agile Hardware Specialization*. Qualcomm, Virtual, Jan 2021.
- *Accelerator Synthesis for Agile Hardware Specialization: A New Dawn*. ECE Colloquium, University of Minnesota, Virtual, Nov. 2020.
- *Design and Design Automation for Joint ML-Hardware Specialization*. Google Brain, Virtual, Nov. 2020.
- *Design and Design Automation for Joint ML-Hardware Specialization*. Samsung Forum, Virtual, Oct. 2020.
- *Accelerator Synthesis for Agile Hardware Specialization: A New Dawn*. Google Brain, Virtual, Aug. 2020.
- *Accelerator Synthesis for Agile Hardware Specialization: A New Dawn*. IWLS'20 Workshop, Virtual, Jul. 2020. (Keynote)
- *Dynamic Pruning of Neural Networks for Efficient Deep Learning*. DAC'20 ROAD4NN Workshop, Virtual, Jul. 2020.
- *Algorithm-Accelerator Co-Design for Neural Network Specialization*. VLSID'20, Ben-

- galuru, India, Jan. 2020.
- *Algorithm-Accelerator Co-Design for Neural Network Specialization*. NeurIPS'19 EMC2 Workshop, Vancouver, Canada, Dec. 2019.
 - *Efficient Machine Learning in the Age of Hardware Specialization*. Simon Fraser University, Vancouver, Canada, Dec. 2019.
 - *Building FPGA-Targeted Accelerators with HeteroCL*. TVM and Deep Learning Compilation Conference, University of Washington, Seattle, WA, Dec. 2019.
 - *Algorithm-Accelerator Co-Design for Deep Learning Specialization*. University of Washington, Seattle, WA, Dec. 2019.
 - *Efficient Machine Learning in the Age of Hardware Specialization*. ECE Seminars, UCLA, Los Angeles, CA, Nov. 2019.
 - *The Smart and The Quick: Efficient Machine Learning Greets Agile Hardware Design*. EE Seminars, Technion - Israel Institute of Technology, Haifa, Israel, Oct. 2019.
 - *Algorithm-Accelerator Co-Design for Neural Network Specialization*. Special ECE Seminar, University of Pittsburgh, Pittsburgh, PA, Sep. 2019.
 - *The Smart and The Quick: Efficient Machine Learning Greets Agile Hardware Design*. CALCM Seminars, Carnegie Mellon University, Pittsburgh, PA, Sep. 2019.
 - *Designing Application Specific AI Processors*. DAC'19 Tutorial, Las Vegas, NV, Jun. 2019.
 - *Deep Neural Nets with Small Footprint: Algorithms and Hardware Specialization*. Research Seminar, Xilinx, San Jose, CA, Nov. 2018.
 - *Deep Neural Networks with Small Footprint: Towards Near-Data Intelligence*. CRISP Center Research Review, University of Virginia, Charlottesville, VA, Oct. 2018.
 - *Co-Design for Deep Learning on the Edge: Algorithms and Hardware Specialization*. CTO Talk Series, MediaTek, Hsinchu, Taiwan, Jul. 2018.
 - *High-Level Synthesis for Accelerator-Rich Computing*. Pennsylvania State University, State College, PA, Mar. 2018.
 - *Towards Multi-Paradigm Programming for Software-Defined FPGAs*. Center of Domain-Specific Computing Symposium (CDSC) at UCLA, Los Angeles, CA, Mar. 2018.
 - *Accelerating Deep Neural Networks on FPGAs*. ASP-DAC'18 Tutorial, Jeju, Korea, Jan. 2018.
 - *Hardware Specialization for Quantized Neural Networks*, ICCAD'17 AMS Workshop, Irvine, CA, Nov. 2017.
 - *Adaptive Loop Pipelining in High-Level Synthesis*, Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, Oct 2017.
 - *Design Automation for Software-Programmable FPGAs*, University of Toronto, Toronto, Ontario, Canada, Oct. 2017.
 - *High-Level Synthesis for Accelerator-Rich Computing*, Columbia University, New York City, NY, Oct. 2017.
 - *High-Level Synthesis for Accelerator-Rich Computing*, Massachusetts Institute of Technology, Cambridge, MA, Oct. 2017.
 - *Mapping-aware Logic Synthesis with Parallelized Stochastic Optimization*, EPFL Workshop on Logic Synthesis and Emerging Technologies, EPFL, Lausanne, Switzerland, Sep. 2017.
 - *Accelerating Deep Neural Networks on FPGAs*, Summer School on Deep Learning On-Chip, Politecnico di Torino, Torino, Italy, Sep. 2017.
 - *Design Automation for Software-Programmable FPGAs*, University of Illinois at Urbana-Champaign, Champaign, IL, Aug. 2017.
 - *Design Automation for Software-Programmable FPGAs*, University of Texas at Austin, TX, Jun. 2017.
 - *Design Automation for Software-Programmable FPGAs*, Imperial College London, London, UK, May. 2017.

- *Design Automation for Software-Programmable FPGAs*, Princeton University, Princeton, NJ, May. 2017.
- *Design Automation for Software-Programmable FPGAs*, UC Berkeley, Berkeley, CA, Apr. 2017.
- *Design Automation for Software-Programmable FPGAs*, Stanford University, Stanford, CA, Apr. 2017.
- *Enabling Software-Defined Reconfigurable Computing*, ICC Distinguished Lecturer Series, Michigan Technological University, Houghton, MI, Dec. 2016.
- *Productive High-Level Programming for Deep Learning Acceleration on FPGAs*, Machine Learning on FPGAs Summit, Intel, San Jose, CA, Oct. 2016.
- *Rapid Hardware Specialization with HLS: Glass Half Full*, MICRO'16 Tutorial on "Rapid Exploration of Accelerator-rich Architectures: Automation from Concept to Prototyping", Taipei, Taiwan, Oct. 2016.
- *Design Automation for Software-Defined Reconfigurable Computing*, National Taiwan University, Taipei, Taiwan, Oct. 2016.
- *Complexity-Effective Loop Specialization*. Institute of Microelectronics, Tsinghua University, Beijing, China, Dec. 2015.
- *Intelligent High-Level Synthesis for Effort-Less Hardware Specialization*. ICCAD'15 Workshop on "Towards Efficient Computing in the Dark Silicon Era", Austin, TX, Nov. 2015.
- *DA Systemization of Knowledge: A Catalog of Prior Forward-Looking Initiatives*. ICCAD'15 Special Session on "From EDA to DA: Can We Evolve Beyond Our E-Roots?", Austin, TX, Nov. 2015.
- *Programming Software-Defined FPGAs: Progress and Roadblocks*. SEAK 2015: DAC Workshop on Suite of Embedded Applications and Kernels, San Francisco, CA, Jun. 2015.
- *Synthesis with Higher Abstractions for Effort-Less FPGA Programming*. CALCM Seminars, Carnegie Mellon University, Pittsburgh, PA, Feb. 2015.
- *Towards Reconfigurable Computing for Software Programmers*. Institute of Microelectronics, Tsinghua University, Beijing, China, Dec. 2014.
- *Architecture and Synthesis for Effort-Less Hardware Specialization*. Center for Energy-Efficient Computing and Applications, Peking University, Beijing, China, Dec. 2014.
- *Mapping-Aware Scheduling for LUT-Based FPGAs*. Research Seminar, Xilinx, San Jose, CA, Nov. 2014.
- *Synthesis with Higher Abstractions for Software Programmable FPGAs*. Distinguished Lecturer Series, USC Information Sciences Institute (ISI), Arlington, VA, Aug. 2014.
- *High-Level Synthesis for Low-Power Design*. ASP-DAC'14 Tutorial, Singapore, Jan. 2014.
- *High-Level Synthesis for FPGAs: From Prototyping to Deployment*. The Third Workshop on the Intersections of Computer Architecture and Reconfigurable Logic (CARL 2013), Davis, CA, Dec. 2013.
- *Intelligent Modulo Scheduling for Pipeline Synthesis*. Research Seminar, Xilinx, San Jose, CA, Nov. 2013.
- *Scheduling Algorithms for High-Quality High-Level Synthesis*. Research seminar, Electronic Engineering Department, Tsinghua University, Beijing, China, Dec. 2012.
- *Scheduling Algorithms for High-Quality High-Level Synthesis*. DA PIC seminar, IBM T.J. Watson Research Center, Yorktown Heights, NY, Dec. 2012.
- *C-Based Hardware Synthesis with AutoPilot*. Center of Domain-Specific Computing Symposium (CDSC) at UCLA, Los Angeles, CA, Feb. 2010.
- *Domain-Specific Reconfigurable Computing with C-Based Synthesis*. Invited Tutorial at IEEE International SOC Conference (SOCC), New Port Beach, CA, Sep. 2008.
- *C/C++/SystemC to RTL Synthesis for FPGA*. Invited Tutorial at Reconfigurable System Summer Institute (RSSI) at UIUC, Champaign, IL, Jul. 2008.