IEEE/ACM

2023 INTERNATIONAL CONFERENCE ON COMPUTER-AIDED DESIGN

ICCAD

42nd Edition

- WELCOME MESSAGE
- TABLE OF CONTENTS
- TECHNICAL PAPERS
- AUTHOR INDEX

IEEE/ACM ICCAD 2023

CONFERENCE PROCEEDINGS

SPONSORS AND ORGANIZERS

IEEE | CAS | CEDA IEEE Council on Electronic Design Automation | acm Association for Computing Machinery | sigda

ISBN: 979-8-3503-2225-5
Part Number: CFP23CAD-ART

Technical Support
Conference Catalysts
Phone: +1 352 872 5544
cdyer@conferencecatalysts.com

2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)Proceedings

Additional copies may be ordered from:

IEEE Service Center
445 Hoes Lane
Piscataway, NJ 08855-1331 USA

+1 800 678 IEEE (+1 800 678-4333)
+1 732 981 1393
+1 732 981 9667 (FAX)
email: customer-service@ieee.org

IEEE Catalog Number: CFP23Q10-ART
ISBN: 979-8-3503-1559-2

cadence®

ACADEMIC NETWORK

Are You Ready to Create Tomorrow's Technology, Today?

Explore career opportunities at Cadence

Make Your Mark

Cadence is a pivotal leader in electronic systems design, building upon more than 30 years of computational software expertise. The company applies its underlying Intelligent System Design strategy to deliver software, hardware, and IP that turn design concepts into reality.

Cadence customers are the world's most innovative companies, delivering extraordinary products from chips to systems, chemicals to drugs, and specification to manufacturing for the most dynamic market applications, including hyperscale computing, 5G communications, automotive, mobile, aerospace, consumer, industrial, and life sciences. Join our teams that are playing a valuable role in creating the technologies that modern life depends on. Apply today at www.cadence.com/cadence/careers.

Academic Network

The Cadence® Academic Network delivers Intelligent System Design™ technology, training, and programs to universities and innovators in the global academic community. Learn more about how you can join fellow experts in the Academic Network at www.cadence.com/site/academicnetwork.

FORTUNE 100 BEST COMPANIES TO WORK FOR® 2023

SYNOPSYS® | ACADEMIC & RESEARCH ALLIANCES

Bridging academia and industry to create new opportunities for innovation and collaboration

Calling all Academics!

Dive into a world of opportunity with Synopsys Academic & Research Alliances

From Sunday—Wednesday, our team will be on site, engaging in various exciting activities:

- Catch us at the ACM Student Research Competition, where innovation takes center stage
- Visit our table at the EDA job fair to explore promising career opportunities
- Engage with us for an enlightening career panel on Wednesday to learn what's in store after graduation

Join us for an enriching discussion on our various programs dedicated to enhancing your educational experience.

Scan the QR code for more information or email us at **SARA-NA@synopsys.com** to schedule a time to sync up.

Table of Contents

General Information

Registration

Location: Lobby, 2nd Floor

Sunday, October 29	7:30 – 17:00
Monday, October 30	7:00 – 19:00
Tuesday, October 31	7:00 – 19:00
Wednesday, November 1	7:00 – 18:00
Thursday, November 2	7:00 – 16:00

Whova

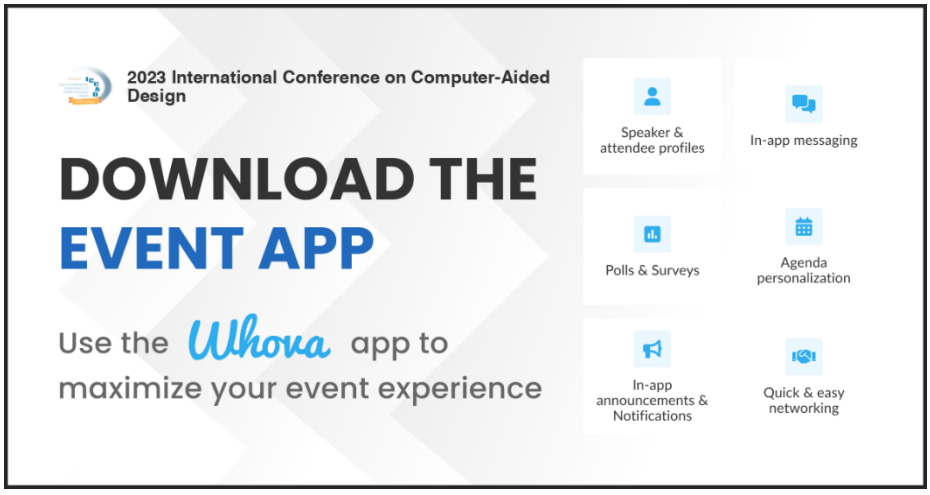

Access the virtual platform to build your schedule, watch presentation videos, view posters, download slides, and connect with other attendees!
Click Here to log in or create your account.

Speaker's Breakfast

Location: Gallery Ballroom III (Please attend the day of your presentation!)

Monday, October 30	7:00 – 8:00
Tuesday, October 31	7:00 – 8:00
Wednesday, November 1	7:00 – 8:00

ICCAD Social Media Platforms

Conference website: https://iccad.com/
LinkedIn: https://www.linkedin.com/company/9289426

Venue Map & Information

Hyatt Regency San Francisco Downtown Soma

Address: 50 Third Street, San Francisco, California 94103 United States
Phone: +1 415 974 6400

Parking at the Hotel

Valet parking with In-Out privileges is available for guests at the San Francisco Downtown SoMa Hotel. Please note that vehicles have a 6’ 10” height restriction and that EV charging stations are available.

- Valet Parking: $75 /Daily
- Oversized vehicles: $85 /Daily

Venue Map

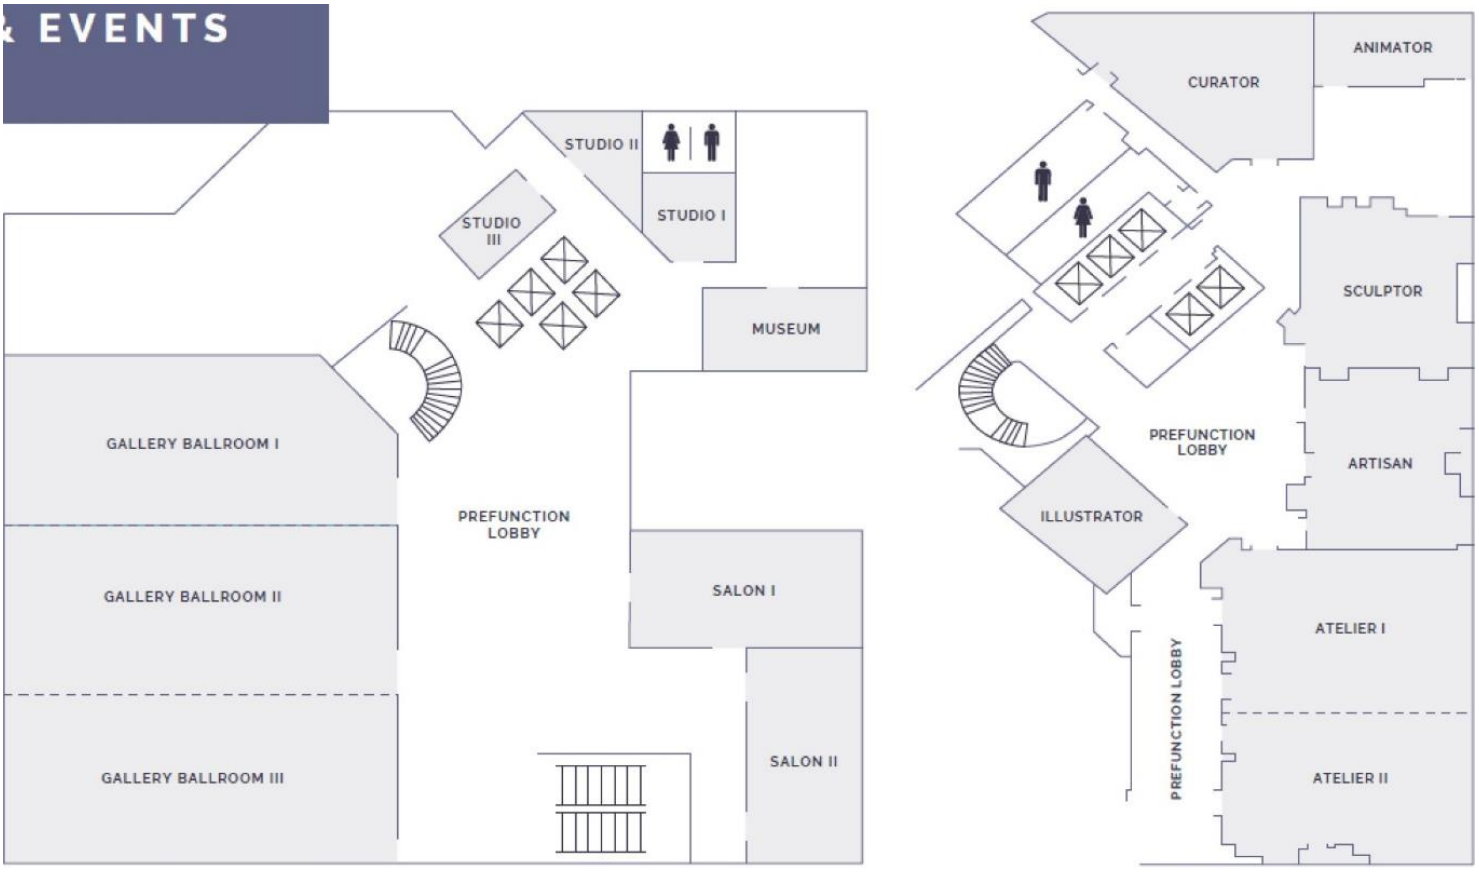

Welcome to ICCAD 2023

Welcome to the 42nd International Conference on Computer-Aided Design!

This year marks our return to a fully in-person ICCAD after the COVID-19 pandemic. We are thrilled to resume an event with personal interactions and extensive networking, while retaining some good practice adopted during the pandemic era of having a virtual platform that participants and authors can connect remotely.

Jointly sponsored by IEEE and ACM, ICCAD has always been the premier venue to explore forefront technology challenges in electronic design automation, present leading-edge R&D solutions, and identify future research directions. In this era of AI, the ICCAD scope has been extensively expanded to encompass exciting AI technology in EDA and hardware design, and to address automation challenges in emergent technology and paradigms.

This is yet another record year of blessings for ICCAD! We received 750 final submissions, organized into 18 tracks and reviewed by 349 outstanding international TPC members from both academia and industry. The TPC meetings were conducted online without compromising the quality of the double-blind review process. Finally, 172 papers were accepted, marking a competitive acceptance rate of 22.9%. The Tracks were organized into 39 sessions with interleaving poster sessions. We also received a record high number of 29 special session and tutorial proposal submissions. Eventually, 9 special sessions and 2 embedded tutorials were accepted, with 28 invited papers included in the proceedings. These special sessions and tutorials perfectly complemented the regular sessions, forming a strong technical ICCAD 2023 program!

We are delighted to host several distinguished keynote speakers. The Monday morning keynote on “EDA in the Age of Deep Learning” will be given by Professor Bill Dally, Nvidia Chief Scientist and SVP Research and Stanford Professor. On Tuesday, IEEE CEDA will host its annual ICCAD Luncheon Distinguished Lecture by Professor Margaret Martonosi from Princeton University with the title “Mind the Gap: Challenges and Opportunities in Closing the Algorithms-to-Devices Gap in Quantum Computing.” On Wednesday morning, Dr. Vamsi Boppana, SVP of AI in AMD, will present the keynote on “AI and EDA: Powering the Next Frontier of Design.” Last but not least, ACM SIGDA will host the keynote in its annual SIGDA Award Dinner on Tuesday by Professor Ron Rohrer, Professor Emeritus from Carnegie Mellon University, on “Lessons Learned in 60 years in EDA (& Other Industries)”. You will undoubtedly find all these keynotes transforming and insightful.

Additionally, we have five co-located workshops on a spectrum of attractive topics. One of them, “Top Picks in Hardware & Embedded Security”, will be held on Wednesday, while the others, “Fast ML for Science (Fast ML),” “System Level Interconnect Path Finding (SLIP)”, “Zero Trust Hardware Architectures” and “Workshop on Sustainable Hardware Security (SUSHI)” will be on Thursday. While some are long-time staples of ICCAD, the “Fast ML” workshop tests the waters for the first time. They all have exciting programs, and we hope you will stay one more day and join some of

them. There will also be a forum on “VLSI Education Community” to be held on Thursday to address the undergraduate and graduate education needs for the semiconductor workforce.

In keeping with its storied tradition, ICCAD remains a home to strong student-oriented activities: the SIGDA CADathlon, ACM Student Research Competition, SIGDA/CEDA Job Fair, ACM/IEEE Contest for Machine Learning on Hardware, and the IC/CAD Contests. This is also the year we launched the inaugural ICCAD Student Scholar Program supported by academic and industrial sponsors, through which we support many students traveling to ICCAD for various student activities and research experiences. ICCAD is committed to creating and providing more opportunities and support to young scholars and students and to promoting diversity.

Once again, ICCAD promises to be an ultimate destination for those working on the cutting-edge EDA research. We do hope you will be able to join us. We are grateful to our ICCAD 2023 sponsors, supporters, and volunteers for making this year’s conference a great success. The members of the executive committee, the technical program committee, and numerous volunteers have spent an enormous effort to prepare an outstanding program for this year. We do hope you will enjoy the conference!

Have a wonderful ICCAD 2023!

General Chair of ICCAD 2023
Evangeline Young
The Chinese University of Hong Kong

Awards Committees

2023 IEEE/ACM William J. McCalla ICCAD Best Paper Award Selection Committee

Takaishi Sato (Kyoto University)
Tei-Wei Guo (National Taiwan University)
Hussam Amrouch (Technical University of Munich)
Tim Kwang-Ting CHENG (Hong Kong University)
Krishnendu Chakrabarty (Arizona State University)
Renu Mehra (Synopsys)

2023 William J. McCalla ICCAD Ten Year Retrospective Most Influential Paper Award Selection Committee

David Z. Pan (University of Texas at Austin)
Lukas Sekanina (Brno University of Technology)
Jie Han (University of Alberta)
Gang Qu (University of Maryland)
Iris Bahar (Colorado School of Mines)

2023 Award Nominations & Winners

IEEE/ACM William J. McCalla ICCAD Best Paper Award Nominations

Front End:

Session AI-Tools 1: 63: Robust GNN-based Representation Learning for HLS

Atefeh Sohrabizadeh (University of California)
Yunsheng Bai (University of California)
Yizhou Sun (University of California)
Jason Cong (University of California)

Session AI-Sys 3: 594: Fluid Batching: Exit-Aware Preemptive Serving of Early-Exit Neural Networks on Edge NPUs

Alexandros Kouris (Samsung AI)
Stylianos Venieris (Samsung AI)
Stefanos Laskaridis (Samsung AI)
Nicholas Lane (University of Cambridge)

Session Security-Alg 2: 717: Striving for Both Quality and Speed: Logic Synthesis for Practical Garbled Circuits

Mingfei Yu (Integrated Systems Laboratory, EPFL)
Giovanni De Micheli (Integrated Systems Laboratory, EPFL)

Back End:

Session AI-Sys 2: 120: Improving Realistic Worst-Case Performance of NVCiM DNN Accelerators through Training with Right-Censored Gaussian Noise

Zheyu Yan (University of Notre Dame)
Yifan Qin (University of Notre Dame)
Wujie Wen (North Carolina State University)
X. Sharon Hu (University of Notre Dame)
Yiyu Shi (University of Notre Dame)

Session HLS 2: 403: MapBuf: Simultaneous Technology Mapping and Buffer Insertion for HLS Performance Optimization

Hanyu Wang (ETH Zurich)
Carmine Rizzi (ETH Zurich)
Lana Josipovic (ETH Zurich)

Session DFM1: 577: OPT: Optimal Proposal Transfer for Efficient Yield Optimization for Analog and SRAM Circuits

Yanfang Liu (Beihang University)
Guohao Dai (Shenzhen University)
Yuanqing Cheng (Beihang University)
Wang Kang (Beihang University)
Wei Xing (Beihang University)

Reduction and IR-drop Compensations Techniques for Reliable Neuromorphic Computing Systems

Beiye Liu (University of Pittsburgh)
Hai Li (University of Pittsburgh)
Yiran Chen (University of Pittsburgh)
Xin Li (Carnegie Mellon University)
Tingwen Huang (Texas A&M University)
Qing Wu (Air Force Research Laboratory)
Mark Barnell (Air Force Research Laboratory)

Published in 2014 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), pp. 63-70. IEEE, 2014.

This award is given to the paper judged to be the most influential on research and industry practice in computer-aided design of integrated circuits over the ten years since its original appearance at ICCAD.

2023 IEEE CEDA Ernest S. Kuh Early Career Award

Pierluigi Nuzzo (University of Southern California)

"For outstanding contribution to high-assurance design of cyber-physical systems using contract-based methodology."

2023 IEEE CEDA Outstanding Service Recognition

Tulika Mitra (National University of Singapore)

"For outstanding service to the EDA community as ICCAD General Chair in 2022"

2023 IEEE ICCAD Inaugural Best Reviewer Award

Stephan Held (University of Bonn)

Mahmut Kandemir (Penn State University)

Yibo Lin (Peking University)

Georgios Zervakis (University of Patras)

"In recognition of the outstanding service to our community through an exemplary commitment to high-quality reviews"

Keynote Speaker - Bill Dally

Monday, October 30, 8:30 – 9:30

Gallery I & II Ballroom

Bill Dally
NVIDIA & Stanford
NVIDIA Chief Scientist, SVP Research and Stanford Professor

Bill Dally joined NVIDIA in January 2009 as chief scientist, after spending 12 years at Stanford University, where he was chairman of the computer science department. Dally and his Stanford team developed the system architecture, network architecture, signaling, routing, and synchronization technology that is found in most large parallel computers today. Dally was previously at the Massachusetts Institute of Technology from 1986 to 1997, where he and his team built the J-Machine and the M-Machine, experimental parallel computer systems that pioneered the separation of mechanism from programming models and demonstrated very low overhead synchronization and communication mechanisms. From 1983 to 1986, he was at the California Institute of Technology (Cal Tech), where he designed the MOSSIM Simulation Engine and the Torus Routing chip, which pioneered "wormhole" routing and virtual-channel flow control. He is a member of the National Academy of Engineering, a Fellow of the American Academy of Arts & Sciences, a Fellow of the IEEE and the ACM, and has received the ACM Eckert-Mauchly Award, the IEEE Seymour Cray Award, and the ACM Maurice Wilkes award. He has published over 250 papers, holds over 120 issued patents, and is an author of four textbooks. Dally received a bachelor's degree in Electrical Engineering from Virginia Tech, a master's in Electrical Engineering from Stanford University, and a Ph.D. in Computer Science from Cal Tech. He was a cofounder of Velio Communications and Stream Processors.

EDA in the Age of Deep Learning

Deep Learning is affecting a change in EDA tools. Deep networks can accelerate analysis tools. Reinforcement learning and generative models can increase designer productivity and produce better designs. This talk will survey the application of deep learning to EDA tools and give an outlook for the future.

Keynote Speaker - Margaret Martonosi

Tuesday, October 31, 13:00 – 14:00

Gallery I & II Ballroom

Margaret Martonosi
Princeton University
Professor of Computer Science

Margaret Martonosi is the H.T. Adams '35 Professor of Computer Science at Princeton University, where she has been on the faculty since 1994. Martonosi's research focuses on hardware-software interface approaches in both classical and quantum computing systems. Martonosi is an elected member of the US National Academy of Engineering and the American Academy of Arts and Sciences. In 2021, she received computer architecture's highest honor, the ACM/IEEE Eckert-Mauchly Award, for contributions to the design, modeling, and verification of power-efficient computer architecture. She is a Fellow of IEEE and ACM. Her work with others to co-found the ACM CARES movement was recognized by the Computing Research Association's 2020 Distinguished Service Award.

Mind the Gap: Challenges and Opportunities in Closing the Algorithms-to-Devices Gap in Quantum Computing

From its initial proposal, Quantum Computing (QC) has had captivating potential, and scientists have worked on advancing toward that potential. With well-known algorithms as motivation, and increasingly capable hardware devices, QC has now reached an interesting and important inflection point. The Algorithms-to-Devices gap in QC refers to the orders of magnitude difference between the quantity and quality of resources needed by QC algorithms, and what has been successfully built today. What is needed now are computer scientists and engineers to develop the crucial intermediate tool flows, abstraction layers, and programming languages that will help QC systems close this gap and reach practical quantum advantage. My talk will offer some recent results from my group about new applications, architecture, and design synthesis approaches for bridging the gap. More broadly, I will advocate for the role that computer scientists and engineers must play for QC to reach its full potential.

Keynote Speaker - Ron Rohrer

Tuesday, October 31, 19:30 – 20:30

Gallery II & III Ballroom

Ron Rohrer
Carnegie Mellon University
Professor Emeritus - Electrical and Computer Engineering

Ron Rohrer is Founder and Chairman of Alto Technologies, Inc. and University Professor Emeritus of Electrical and Computer Engineering at Carnegie Mellon University. Ron received the B.S. (1960) degree from Massachusetts Institute of Technology (MIT) and the M.S. (1961) and Ph.D. (1963) degrees from the University of California (UC), Berkeley. A noted professor, tool innovator and entrepreneur, Ron is recognized as an early developer of circuit simulation, integrated circuit interconnect reduction and delay calculation. He is credited with driving Electronic Design Automation (EDA) tools into broader industry use. In 1968, in a graduate course he taught at UC, Berkeley, he oversaw the production of an integrated circuit simulation program similar to the FairCirc program he had developed earlier at Fairchild Semiconductor, a forerunner to SPICE (Simulation Program with Integrated Circuit Emphasis). Based on earlier Adjoint Circuit simulation work, in 1971, he developed the foundation of what became the industry standard technique for simulation of analog and RF IC noise. In 1982, while at GE, he launched the logic synthesis work that led to the Socrates logic synthesis tool that led to the formulation of Synopsys. In 1989 he was inducted into the National Academy of Engineering for contributions to circuit simulation that have enabled deep submicron IC design. He has authored five textbooks and over 100 technical papers. His many awards include the IEEE 1993 Education Medal “for innovation in bringing electrical engineering practice into the classroom and merging academic research with industrial need”, the 1996 NEC Computer and Communication Prize, a worldwide honor for pioneering contributions in electronics; and the 2002 Phil Kaufman award, presented by the EDA Consortium. Ron is also the SIGDA Pioneer Award recipient for 2023.

Lessons learned in 60 years in EDA (& other industries)

Ron Rohrer has held both academic and industrial positions during his career of over sixty years. He has been fortunate to have seen and, in some cases, to have impacted the EDA field since the 1960s. During that period, he's observed that while some things have changed, some have stayed the same and some may be hindering more rapid progress. The talk will cover experiences, observations and insights garnered from academic and industry perspectives.

Keynote Speaker - Vamsi Boppana

Wednesday, November 1, 8:00 – 9:00

Gallery I & II Ballroom

Vamsi Boppana
AMD
Senior Vice President of Artificial Intelligence

Vamsi Boppana is the senior vice president of the Artificial Intelligence Group at AMD. He is responsible for the AI strategy across AMD, including driving the AI roadmap across client, edge, and cloud for the company's AI engines and AI software and ecosystem efforts. In his previous role, he was senior vice president of the Central Products Group (CPG), responsible for developing and marketing Xilinx's Adaptive and AI product portfolio. He also served as executive sponsor for the Xilinx integration into AMD. In his prior roles, Boppana led the development of industry-leading products such as Versal™ and Zynq™ UltraScale™+ MPSoC. Before joining the company in 2008, he held engineering management roles at Open-Silicon and Zenasis Technologies, a company he co-founded. Boppana earned a Bachelor of Technology, with honors, in computer science and engineering from the Indian Institute of Technology, Kharagpur, and Master of Science and doctorate degrees in electrical engineering from the University of Illinois at Urbana-Champaign. The author and co-author of more than 40 technical papers, he has received six patents for his design.

AI and EDA: Powering the Next Frontier of Design

Over the past several decades, EDA innovation has been foundational to design progress. Rapid advances in design have, in turn, created computational capabilities that are now powering the AI revolution. In this talk, we discuss key advances that have brought us to this point of inflection, review state-of-the-art AI platform capabilities, and look at the exciting road ahead with AI advances powering EDA innovation that enables the next frontier of design.

CADathlon 2023: Sunday, October 29

The **CADathlon** is a challenging, all-day programming competition focusing on practical problems at the forefront of Computer-Aided Design and Electronic Design Automation in particular. The contest emphasizes the knowledge of algorithmic techniques for CAD applications, problem-solving and programming skills, as well as teamwork. As the "**Olympic games of EDA**," the contest brings together the best and the brightest of the next generation of CAD professionals. It gives academia and the industry a unique perspective on challenging problems and rising stars, and it also helps attract top graduate students to the EDA field.

The contest is open to **two-person teams** of **undergraduate**/graduate students specializing in CAD and currently full-time enrolled in a Ph.D. granting institution in any country. Students are selected based on their academic backgrounds and their relevant EDA programming experiences. Partial or full travel grants are provided to qualifying students. CADathlon competition consists of *six* problems in the following areas:

- Circuit Design & Analysis
- Physical Design & Design for Manufacturability
- Logic & High-Level Synthesis
- System Design & Analysis
- Functional Verification & Testing
- Future technologies (Bio-EDA, Security, AI, etc.)

More specific information about the problems and relevant research papers will be released on the Internet **one week prior** to the competition. The writers and judges that construct and review the problems are experts in EDA from both academia and industry. At the contest, students will be given the problem statements and example test data, but they will not have the judges' test data. Solutions will be judged on correctness and efficiency. Where appropriate, partial credit might be given.

The team that earns the highest score is declared the winner. In addition to handsome trophies, the first place and the second-place teams receive cash awards, and the contest winners will be announced at the ICCAD conference.

Organization Committee:

Chair: Andy, Yu-Guang Chen (National Central University)
Co-chair: Jeff (Jun) Zhang (Arizona State University)
Co-chair: Zahra Ghodsi (Purdue University)
Co-chair: Pei-Yu (Billy) Lee (Synopsys)

Monday, October 30: Program Grid

	Sculptor	Artisan	Atelier I	Atelier II	
7:00 - 8:00	Speakers Breakfast (Gallery Ballroom III)				
8:00 - 9:30	Opening & Keynote: Bill Dally (Gallery Ballroom I & II)				
9:30 - 10:00	Coffee Break (Gallery Ballroom III)				
10:00 - 11:30	Reconf-LP 3: Approximate Computing	SS1: AI for Sign-Off Showdown: Unveiling the Cost & Effect of Artificial Intelligence	Analog 1: Advancements in Machine Learning and Modeling Techniques for RF Circuits a Power Grid Networks	Timing 2: Routing and Routability Enhancement	
11:30 - 12:30	Lunch (Gallery Ballroom III)				
12:30 - 14:00	Security-Alg 2: Efficiency for Security	SS2: Towards Generative AI for EDA: Datasets, Benchmarks and Infrastructures	NewComp 1: Navigating the Future: A Voyage from Beyond CMOS to Beyond von Neumann	AI-Sys 5: Machine Learning on Emerging Platforms	SRC Poster Session (Gallery III)
14:00 - 15:30	Security-Arch 1: Microarchitectural security: to speculate or not?	SS3: Sustainable AI Training at the Large and Tiny Scales	Nano 1: Emerging technologies for computation in memory	AI-Tools 2: Alternative approaches to effective and efficient training of neural networks	
15:30 - 16:00	Coffee Break (Gallery Ballroom III)				
16:00 - 17:30	PD 1: Routability and Congestion Prediction	SS4: Next-Generation Computing Paradigm for Next-Generation (6G) Wireless	Bio 1: Neuromorphic Wonders: Bridging Minds and Machines	SysDesign 1: Designing Chiplet-Based Systems	
17:30 - 19:00	PD 2: Advanced Floorplanning and Partitioning	CPS 2: Fast and Trustworthy Embedded Systems	LLM-Aided Design Panel	SysDesign 2: System Level Design and Exploration	
19:00 - 20:30		ACM Student Research Competition (SRC) @ ICCAD			
20:30 - 22:00	TPC Reception (Gallery III & Gallery Foyer) *Invite Only Event*				

Tuesday, October 31: Program Grid

	Sculptor	Artisan	Atelier I	Atelier II
7:00 - 8:00	Speakers Breakfast (Gallery Ballroom III)			
8:00 - 10:00	HSL 1: Novel ideas in logic synthesis	SS5: Frontiers in Edge AI: Technology, Algorithms, and Emerging Trends	Security-Alg 1: IP and IC Trust: Reflections and Advancement	AI-Tools 1: Towards graph-learning assisted electronic design automation (EDA)
10:00 - 10:30	Coffee Break (Gallery Ballroom III)			
10:30 - 12:30	Test 1: Simulation and Emulation	Tutorial: Introduction to Hybrid Quantum-Classical Programming Using C++ Quantum Extension	Reconf-LP 1: "Hot" Microarchitectures: Modeling and Design	AI-Sys 1: Efficient Accelerator Design
12:30 - 14:00	Keynote Lunch: Margaret Martonosi *Tickets Required Limited Access: First Come, First Served* (Gallery Ballroom I & II)			
14:00 - 15:30	Security-Arch 2: How Do We Patch Our Neural Nets (Against Information Leakage)?	NewComp2: Quantum Circuitry Unleashed: Innovations in Simulation, Synthesis, and Optimization	Security-Arch 3: Microelectronics Security: What Lies Ahead?	Analog 2: Novel Frameworks and Methodologies for Optimizing Analog/Mixed-Signal Circuits
15:30 - 16:00	Coffee Break (Gallery Ballroom III)			
16:00 - 18:00	Timing 1: Analyzing and Optimizing for PPA, Timing and Reliability	SS7: 2023 CAD Contest at ICCAD	Contest for Machine Learning on Hardware	AI-Sys 2: Design Tools and Performance Optimization for DNN Acceleration
18:00 - 19:00	IEEE/ACM EDA Job Fair & SIDGDA Dinner *Tickets Required For Dinner Limited Access: First Come, First Served* (Gallery Ballroom I, II, & III)			

Wednesday, November 1: Program Grid

	Sculptor	Artisan	Atelier I	Atelier II	Curator
7:00 - 8:00	Speakers Breakfast (Gallery Ballroom I & II)				
8:00 - 9:00	Keynote: Vamsi Boppana (Gallery Ballroom III)				
9:00 - 10:30	NewComp 3: Frontiers in Quantum Computing: Novel Algorithms and Beyond	SS8: Accelerating EDA Algorithms with Heterogeneous Parallelism	PD 3: The State-of-the-Art Placement	SysDesign 3: Designing Memory and Storage Systems	Top Picks in Hardware and Embedded Security Workshop
10:30 - 11:00	Coffee Break (Gallery Ballroom I & II)				
11:00 - 13:00	HLS 2: Technology Mapping Revived	Tutorial: RapidWright: Unleashing the Full Power of FPGA Technology with Domain-Specific Tooling	AI-Tools 3: Anywhere Anytime – Full Stack AI Deployment	AI-Sys 3: Efficient and Fair Machine Learning on the Edge	Top Picks in Hardware and Embedded Security Workshop
13:00 - 14:00	Lunch (Gallery Ballroom I & II)				
14:00 - 16:00	Test 2: Advances in Verification	SS10: In-Sensor AI Computing Towards Next Generation Autonomous Edge Intelligence	AI-Tools 4: Adaptation, Acceleration, Automation – Tripple-A AI Computing	AI-Sys 4: Software/Hardware Co-design of AI Accelerator	Top Picks in Hardware and Embedded Security Workshop
16:00 - 16:30	Coffee Break (Gallery Ballroom I & II)				
16:30 - 18:00	DFM 1: Revolutionizing Design for Reliability and Manufacturability	SS11: From Arithmetic Components to Side Channel Attacks: A Deep Dive in Post Quantum Hardware	Reconf-LP 2: Modern Reconfigurable Architectures and Design Tools	CPS 1: Efficient Machine Learning for Real-World Applications	Top Picks in Hardware and Embedded Security Workshop
18:30 - 19:30	Synopsys Panel - Engineering your Future: Exploring Diverse Career Paths After Graduation (Gallery Ballroom I & II)				

Thursday, November 2: Program Grid

	Sculptor	Artisan	Salon I	Salon II	Atelier I & II
7:30 - 8:00	Breakfast (3rd Floor Foyer)				
8:00 - 10:00	Fast ML for Science	System Level Interconnect Path Finding	VLSI Education Community	Zero Trust Hardware Architectures	SUSHI
10:00 - 10:30	Coffee Break (3rd Floor Foyer)				
10:30 - 12:00	Fast ML for Science	System Level Interconnect Path Finding	VLSI Education Community	Zero Trust Hardware Architectures	SUSHI
12:00 - 13:00	Lunch (Buffet: 3rd Floor Foyer - Meals to be taken back to breakout rooms)				
13:00 - 15:00	Fast ML for Science	System Level Interconnect Path Finding		Zero Trust Hardware Architectures	SUSHI
15:00 - 15:30	Coffee Break (3rd Floor Foyer)				
15:30 - 17:00	Fast ML for Science	System Level Interconnect Path Finding		Zero Trust Hardware Architectures	SUSHI

Technical Program: Monday, October 30, 2023

8:00 – 9:30, Gallery Ballroom I & II
Opening Ceremony & Keynote: *Bill Dally*

10:00 – 11:30, Room: Sculptor
Reconf-LP 3: Approximate Computing
Session Chairs: Weikang Qian (Shanghai Jiao Tong University)
Vojtech Mrazek (BRNO University of Technology)

730: Bespoke Approximation of Multiplication-Accumulation and Activation Targeting Printed Multilayer Perceptrons
Florentia Afentaki (Karlsruhe Institute of Technology)
Gurol Saglam (Karlsruhe Institute of Technology)
Argyris Kokkinis (Aristotle University of Thessaloniki)
Kostas Siozios (Aristotle University of Thessaloniki)
Georgios Zervakis (University of Patras)
Mehdi Tahoori (Karlsruhe Institute of Technology)

749: DASALS: Differentiable Architecture Search-driven Approximate Logic Synthesis
Xuan Wang (UM-STJTU Joint Inst.)
Zheyu Yan (University of Notre Dame)
Chang Meng (UM-STJTU Joint Inst.)
Yiyu Shi (University of Notre Dame)
Weikang Qian (Shanghai Jiao Tong University)

409: Constant Coefficient Multipliers Using Self-Similarity-Based Hybrid Binary-Unary Computing
Alireza Khataei (University of Minnesota)
Kia Bazargan (University of Minnesota)

603: Xel-FPGAs: An End-to-End Automated Exploration Framework for Approximate Accelerators in FPGA-Based Systems
Bharath Srinivas Prabakaran (TU Wien)
Vojtech Mrazek (TU Wien)
Zdenek Vasicek (Brno University of Technology)
Lukas Sekanina (Brno University of Technology)
Muhammad Shafique (NYUAD)

10:00 – 11:30, Room: Artisan
SS1: AI for Sign-Off Showdown: Unveiling the Cost & Effect of Artificial Intelligence
Session Chairs: Yibo Lin (Peking University)
Cheng Zhuo (Zhejiang University)

42: Unleashing the Potential of Machine Learning: Harnessing the Dynamics of Supply Noise for Timing Sign-Off
Yufei Chen (Zhejiang University)
Xiao Dong (Zhejiang University)
Wei-Kai Shih (Synopsys Inc.)
Cheng Zhuo (Zhejiang University)

52: Solving Fine-grained Static 3DIC Thermal with ML Thermal Solver Enhanced with Decay Curve Characterization
Haiyang He (Ansys Inc)
Norman Chang (Ansys Inc)
Jie Yang (Ansys Inc)
Akhilesh Kumar (Ansys Inc)
Wenbo Xia (Ansys Inc)
Lang Lin (Ansys Inc)
Rishikesh Ranade (Ansys Inc)

57.1: The Inevitability of AI Infusion into Design Closure and Signoff
Jiang Hu (Texas A&M University)
Andrew Kahng (UC San Diego)

Industry Trends for Power and Thermal Signoff in Advanced Nodes and Packaging
Zhiyu Zeng (Cadence Design Systems)

10:00 – 11:30, Room: Atelier I
Analog 1: Advancements in Machine Learning and Modeling Techniques for RF Circuits a Power Grid Networks
Session Chairs: Zheng Zhang (UCSB)
Liang Chen (Shanghai Univ)

413: FuNToM: Functional Modeling of RF Circuits Using a Neural Network Assisted Two-Port Analysis Method
Morteza Fayazi (University of Michigan)
Morteza Tavakoli Taba (University of Michigan)
Amirata Tabatabavakili (University of Michigan)
Ehsan Afshari (University of Michigan)
Ronald Dreslinski (University of Michigan)

886: One-Dimensional Deep Image Prior for Curve Fitting of S-Parameters from Electromagnetic Solvers
Sriram Ravula (University of Texas)
Varun Gorti (University of Texas)
Bo Deng (University of Texas)
Swagato Chakraborty (Siemens)
James Pingenot (Siemens)
Bhyrav Mutnury (Dell)
Doug Wallace (Dell)
Doug Winterberg (Dell)
Adam Klivans (University of Texas)
Alexandros G. Dimakis (University of Texas)

455: Accuracy-Preserving Reduction of Sparsified Reduced Power Grids with A Multilevel Node Aggregation Scheme
Zhiqiang Liu (Tsinghua University)
Wenjian Yu (Tsinghua University)

393: Multi-Task Evolutionary to PVT Knowledge Transfer for Analog Integrated Circuit Optimization
Jintao Li (University of Electronic Science and Technology of China)
Haochang Zhi (Southeast University)
Weiwei Shan (Southeast University)
Yongfu Li (Shanghai Jiao Tong University)
Yanhan Zeng (Guangzhou University)
Yun Li (University of Electronic Science and Technology of China)

10:00 – 11:30, Room: Atelier II
Timing 2: Routing and Routability Enhancement
Session Chair: Biying Xu (Cadence)

371: Towards Timing-Driven Routing: An Efficient Learning Based Geometric Approach
Liying Yang (University of Science and Technology of China)
Guowei Sun (University of Science and Technology of China)
Hu Ding (University of Science and Technology of China)

612: NeuroEscape: Ordered Escape Routing via Monte-Carlo Tree Search and Neural Network
Zhiyang Chen (Tsinghua University)
Tsung-Yi Ho (The Chinese University of Hong Kong)
Ulf Schlichtmann (Technical University of Munich)
Datao Chen (Huawei Technologies Co.)
Mingyu Liu (Huawei Technologies Co.)
Hailong Yao (University of Science and Technology Beijing)
Xia Yin (Tsinghua University)

642: Delay-Matching Routing for Advanced Packages
Chun An Lee (Realtek)
Wen-Hao Liu (National Tsing Hua University)
Gary Lin (Cadence)
Tsung-Yi Ho (University of Hong Kong)

853: Power Distribution Network Optimization Using HLA-GCN for Routability Enhancement
Younggwang Jung (KAIST)
Youngsoo Shin (KAIST)
Soyoon Choi (KAIST)
Daijoon Hyun (Sejong University)

12:30 – 14:00, Room: Sculptor
Security-Alg 2: Efficiency for Security
Session Chair: Siddarth Garg (New York University)

717: Striving for Both Quality and Speed: Logic Synthesis for Practical Garbled Circuits
Mingfei Yu (Integrated Systems Laboratory, EPFL)
Giovanni De Micheli (Integrated Systems Laboratory, EPFL)

363: EMSim+: Accelerating Electromagnetic Security Evaluation with Generative Adversarial Network
Ya Gao (Tianjin University)
Haocheng Ma (Tianjin University)
Jindi Kong (Tianjin University)
Jiaji He (Tianjin University)
Yiqiang Zhao (Tianjin University)
Yier Jin (Tianjin University)

942: Hyperdimensional Computing as a Rescue for Efficient Privacy-Preserving Machine Learning-as-a-Service
Jaewoo Park (UNIST)
Chenghao Quan (UNIST)
Hyungon Moon (UNIST)
Jongeun Lee (UNIST)

924: PSOFuzz: Fuzzing Processors with Particle Swarm Optimization
Chen Chen (Texas A&M University)
Vasudev Gohil (Texas A&M University)
Rahul Kande (Texas A&M University)
Ahmad-Reza Sadeghi (Texas A&M University)
Jeyavijayan Rajendran (Texas A&M University)

Denotes 2023 Best Paper Nominee

12:30 – 14:00, Room: Artisan
SS2: Towards Generative AI for EDA: Datasets, Benchmarks and Infrastructures
Session Chair: Haoyu Yang (NVIDIA)

31.1: CircuitOps: An ML Infrastructure Enabling Generative AI for VLSI Circuit Optimization
Rongjian Liang (NVIDIA)
Anthony Agnesina (NVIDIA)
Geraldo Pradipta (NVIDIA)
Haoxing Ren (NVIDIA)
Vidya A. Chhabria (Arizona State University)

35: VerilogEval: Evaluating Large Language Models for Verilog Code Generation
Mingjie Liu (NVIDIA)
Nathaniel Pinckney (NVIDIA)
Brucek Khailany (NVIDIA)
Haoxing Ren (NVIDIA)

51: Verilog-to-PyG -- A Framework for Graph Learning and Augmentation on RTL Designs
Yingjie Li (University of Maryland)
Mingju Liu (University of Maryland)
Cunxi Yu (University of Maryland)
Alan Mishchenko (UC Berkeley)

55: Towards the Imagenets of ML4EDA
Animesh Basak Chowdhury (New York University)
Shailja Thakur (New York University)
Ramesh Karri (New York University)
Siddharth Garg (New York University)
Hammond Pearce (University of South Wales)

12:30 – 14:00, Room: Atelier I
NewComp 1: Navigating the Future: A Voyage from Beyond CMOS to Beyond von Neumann
Session Chair: Saptadeep Pal (Auradine)

298: PBA: Percentile-Based Level Allocation for Multiple-Bits-Per-Cell RRAM
Anjiang Wei (Stanford University)
Akash Levy (Stanford University)
Pu Yi (Stanford University)
Robert Radway (Stanford University)
Priyanka Raina (Stanford University)
Subhasish Mitra (Stanford University)
Sara Achour (Stanford University)

349: TL-nvSRAM-CIM: Ultra-High-Density Three-Level ReRAM-Assisted Computing-in-nvSRAM with DC-Power Free Restore and Ternary MAC Operations
Dengfeng Wang (Shanghai Jia Tong University)
Liukai Xu (Shanghai Jia Tong University)
Songyuan Liu (Shanghai Jia Tong University)
Zhi Li (Shanghai Jia Tong University)
Weifeng He (Shanghai Jia Tong University)
Xueqing Li (Tsinghua University)
Yiming Chen (Tsinghua University)

1037: VECOM: Variation-Resilient Encoding and Offset Compensation Schemes for Reliable ReRAM-Based DNN Accelerator
Jewoo Jang (Yonsei University)
Joon-Sung Yang (Yonsei University)
Thai Hoang Nguyen (Sungkyunkwan University)

131: Automated Synthesis for In-Memory Computing
Muhammad Rashedul Haq Rashed (University of Central Florida)
Sven Thijssen (University of Central Florida)
Rickard Ewetz (University of Central Florida)
Sumit Jha (Florida International University)

12:30 – 14:00, Room: Atelier II
AI-Sys 5: Machine Learning on Emerging Platforms
Session Chair: Marina Zapater (HEIG-VD)
Jiaqi Gu (ASU)

26: Reliable Hyperdimensional Reasoning on Unreliable Emerging Technologies
Hamza Errahmouni Barkam (University of California Irvine)
Sanggeon Yun (University of California Irvine)
Andrew Ding (University of California Irvine)
Hanning Chen (University of California Irvine)
Mohsen Imani (University of California Irvine)
Paul Genssler (University of Stuttgart)
Albi Mema (University of Stuttgart)
Hussam Amrouch (University of Stuttgart)
George Michelogiannakis (Lawrence Berkeley National Laboratory)

615: NearUni: Near-Unitary Training for Efficient Optical Neural Networks
Amro Eldebiky (Technical University of Munich)
Bing Li (Technical University of Munich)
Grace Li Zhang (TU Darmstadt)

590: Kernel Shape Control for Row-Efficient Convolution on Processing-In-Memory Arrays
Johnny Rhe (Sungkyunkwan University)
Kang Eun Jeon (Sungkyunkwan University)
Joo Chan Lee (Sungkyunkwan University)
Seongmoon Jeong (Sungkyunkwan University)
Jong Hwan Ko (Sungkyunkwan University)

62: FIONA: Photonic-Electronic Co-Simulation Framework and Transferable Prototyping for Photonic Accelerator
Yinyi LIU (The Hong Kong University of Science and Technology)
Bohan HU (The Hong Kong University of Science and Technology)
Zhenguo LIU (The Hong Kong University of Science and Technology)
Peiyu CHEN (The Hong Kong University of Science and Technology)
Linfeng DU (The Hong Kong University of Science and Technology)
Jiaqi LIU (The Hong Kong University of Science and Technology)
Xianbin LI (The Hong Kong University of Science and Technology)
Wei ZHANG (The Hong Kong University of Science and Technology)
Jiang XU (The Hong Kong University of Science and Technology)

14:00 – 15:30, Room: Sculptor
Security-Arch 1: Microarchitectural Security: To Speculate or Not?
Session Chairs: Ramya Jayaram Masti (Ampere Computing)
Anupam Chattopadhyay (Nanyang Technological University)

231: BeKnight: Guarding against Information Leakage in Speculatively Updated Branch Predictor

Md Hafizul Islam Chowdhuryy (UCF)
Zhenkai Zhang (Clemson University)
Fan Yao (UCF)

16: HidFix: Efficient Mitigation of Cache-based Spectre Attacks through Hidden Rollbacks

Arash Pashrashid (National University of Singapore)
Ali Hajiabadi (National University of Singapore)
Trevor E. Carlson (National University of Singapore)

638: Exploration and Exploitation of Hidden PMU Events

Yihao Yang (Beijing Ministry of Education)
Pengfei Qiu (Beijing Ministry of Education)
Chunlu Wang (Beijing Ministry of Education)
Yu Jin (Beijing Ministry of Education)
Qiang Gao (Beijing Ministry of Education)
Xiaoyong Li (Beijing Ministry of Education)
Dongsheng Wang (Zhongguancun Lab)
Gang Qu (University of Maryland)

226: Secure-by-Construction Design Methodology for CPUs: Implementing Secure Speculation on the RTL'

Tobias Jauch (RPTU)
Alex Wezel (RPTU)
Dominik Stoffel (RPTU)
Wolfgang Kunz (RPTU)
Philipp Schmitz (RPTU)
Sayak Ray, Jason M. Fung (Intel)
Mohammad Rahmani Fadiheh (Stanford University)
Christopher W. Fletcher (University of Illinois)

14:00 – 15:30, Room: Artisan
SS3: Sustainable AI Training at the Large and Tiny Scales
Session Chairs: Xiaofan Zhang (Google)
Peipei Zhou (University of Pittsburgh)

Hardware-aware Sparsity: Accurate & Efficient Foundation Model Training
Beidi Chen (Meta FAIR and CMU)

Federated Learning at Scale: Efficiency and Sustainability
Ang Li (Qualcomm and University of Maryland at College Park)

DeepZero: Scaling Up Zeroth-order Optimization for Deep Model Training
Sijia Liu (Michigan State University & MIT-IBM Watson AI Lab)

Sustainable Training via Tensor Optimization
Zheng Zhang (UCSB)

14:00 – 15:30, Room: Atelier I
Nano 1: Emerging Technologies for Computation in Memory
Session Chair: Mehdi B. Tahoori (Karlsruhe Institute of Technology)

573: Lowering Latency of Embedded Memory by Exploiting In-Cell Victim Cache Hierarchy Based on Emerging Multi-Level Memory Devices

Juejian Wu (Tsinghua University)
Tianyu Liao (Tsinghua University)
Taixin Li (Tsinghua University)
Yongpan Liu (Tsinghua University)
Huazhong Yang (Tsinghua University)
Xueqing Li (Tsinghua University)
Yixin Xu (Penn State University)
Vijaykrishnan Narayanan (Penn State University)

130: SEE-MCAM: A Scalable Multi-bit FeFET Content Addressable Memory for Energy Efficient Associative Search

Shengxi Shou (Zhejiang University)
Xunzhao Yin (Zhejiang University)
Jianyi Yang (Zhejiang University)
Cheng Zhuo (Zhejiang University)
Sanggeon Yun (University of California Irvine)
Mohsen Imani (University of California Irvine)
Che-Kai Liu (Georgia Institute of Technology)
Zishen Wan (Georgia Institute of Technology)
Kai Ni (Georgia Institute of Technology)
X. Sharon Hu (University of Notre Dame)

421: Accelerating Polynomial Modular Multiplication with Crossbar-Based Compute-in-Memory

Mengyuan Li (University of Notre Dame)
Haoran Geng (University of Notre Dame)
Michael Niemier (University of Notre Dame)
X. Sharon Hu (University of Notre Dame)

928: LIORAT: NN Layer I/O Range Training for Area/Energy-Efficient Low-Bit A/D Conversion System Design in Error-Tolerant Computation-in-Memory

Ayumu Yamada (University of Tokyo)
Naoko Misawa (University of Tokyo)
Chihiro Matsui (University of Tokyo)
Ken Takeuchi (University of Tokyo)

14:00 – 15:30, Room: Atelier II
AI-Tools 2: Alternative Approaches to Effective and Efficient Training of Neural Networks
Session Chairs: Ashutosh Dhar (Nvidia)
Chen Zhang (Shanghai Jiao Tong University)

214: Towards Effective Training of Robust Spiking Recurrent Neural Networks under General Input Noise via Provable Analysis
Wendong Zheng (Sun Yat-sen University)
Yu Zhou (Sun Yat-sen University)
Gang Chen (Sun Yat-sen University)
Zonghua Gu (Sun Yat-sen University)
Kai Huang (Sun Yat-sen University)

229: Analog or Digital In-memory Computing? Benchmarking through Quantitative Modeling
Jiacong Sun (MICAS)
Pouya Houshmand (MICAS)
Marian Verhelst (MICAS)

367: Falcon: Accelerating Homomorphically Encrypted Convolutions for Efficient Private Mobile Networks Inference
Tianshi Xu (Peking University)
Meng Li (Peking University)
Runsheng Wang (Peking University)
Ru Huang (Peking University)

411: Rapid-INR: Storage Efficient CPU-free DNN Training Using Implicit Neural Representation
Hanqiu Chen (Georgia Institute of Technology)
Hang Yang (Georgia Institute of Technology)
Stephen BR Fitzmeyer (Georgia Institute of Technology)
Cong Hao (Georgia Institute of Technology)

16:00 – 17:30, Room: Sculptor
PD 1: Routability and Congestion Prediction
Session Chair: Haocheng Li (Synopsys)

114: Routability Prediction and Optimization Using Explainable AI
Seonghyeon Park (POSTECH)
Daeyeon Kim (POSTECH)
Seongbin Kwon (POSTECH)
Seokhyeong Kang (POSTECH)

137: Routablity-driven Orientation-aware Analytical Placement for System in Package
Jai-Ming Lin (National Cheng Kung University)
Tsung-Chun Tsai (National Cheng Kung University)
Rui-Ting Shen (National Cheng Kung University)

510: Lay-Net: Grafting Netlist Knowledge on Layout-Based Congestion Prediction
Su Zheng (Chinese University of Hong Kong)
Peng XU (Chinese University of Hong Kong)
Siting Liu (Chinese University of Hong Kong)
Bei Yu (Chinese University of Hong Kong)
Martin Wong (Chinese University of Hong Kong)
Lancheng Zou (Wuhan University)

601: ClusterNet: Routing Congestion Prediction and Optimization using Netlist Clustering and Graph Neural Networks
Kyungjun Min (University of Science and Technology Pohang)
Seongbin Kwon (University of Science and Technology Pohang)
Sung-Yun Lee (University of Science and Technology Pohang)
Dohun Kim (University of Science and Technology Pohang)
Sunghye Park (University of Science and Technology Pohang)
Seokhyeong Kang (University of Science and Technology Pohang)

16:00 – 17:30, Room: Artisan
SS4: Next-Generation Computing Paradigm for Next-Generation (6G) Wireless
Session Chairs: Bo Yuan (Rutgers University)
Jiang Hu (Texas A&M University)

40.1: Accelerating Next-G Wireless Communications with FPGA-based AI Accelerators
Chunxiao Lin (Virginia Tech)
Muhammad Farhan Azmine (Virginia Tech)
Cindy Yang Yi (Virginia Tech)

41: MU-MIMO Detection Using Oscillator Ising Machines
Shreesha Sreedhara (UC Berkeley)
Jaijeet Roychowdhury (UC Berkeley)
Joachim Wabnig (Nokia Bella Labs)
Pavan Koteshwar Srinath (Nokia Bella Labs)

In-Memory Computing for Belief Propagation-based Wireless DSP
Y. Gong (Rutgers Univ.)
Y. Sui (Rutgers Univ.)
Bo Yuan (Rutgers Univ.)
A. Sridharan (Arizona State Univ.)
F. Zhang (Arizona State Univ.)
D. Fan (Arizona State Univ.)

NVidia Roadmap to AI-Infused 6G: Requirement, Architecture & Key Technologies
Chris Dick (NVIDIA)

16:00 – 17:30, Room: Atelier I
Bio 1: Neuromorphic Wonders: Bridging Minds and Machines
Session Chairs: Debjit Pal (UIC)

621: ARMM: Adaptive Reliability Quantification Model of Microfluidic Designs and Its Graph-Transformer-Based Implementation

Siyuan Liang (University of Hong Kong)
Tsung-Yi Ho (University of Hong Kong)
Meng Lian (Technical University of Munich)
Mengchu Li (Technical University of Munich)
Tsun-Ming Tseng (Technical University of Munich)
Ulf Schlichtmann (Technical University of Munich)

339: A Novel and Efficient Block-Based Programming for ReRAM-Based Neuromorphic Computing

Wei-Lun Chen (National Cheng Kung University)
Fang-Yi Gu (National Cheng Kung University)
Ing-Chao Lin (National Cheng Kung University)
Grace Li Zhang (TU Darmstradt)
Bing Li (Technical University of Munich)
Ulf Schlichtmann (Technical University of Munich)

745: Multi-Objective Architecture Search and Optimization for Heterogeneous Neuromorphic Architecture

Juseong Park (POSTECH)
Yongwon Shin (POSTECH)
Hyojin Sung (POSTECH)

402: Power-Aware Training for Energy-Efficient Printed Neuromorphic Circuits

Haibin Zhao (Karlsruhe Institute of Technology)
Priyanjana Pal (Karlsruhe Institute of Technology)
Mehdi Tahoori (Karlsruhe Institute of Technology)
Michael Beigl (Karlsruhe Institute of Technology)
Michael Hefenbrock (RevoAT GmbH)

16:00 – 17:30, Room: Atelier II
SysDesign 1: Designing Chiplet-Based Systems
Session Chair: Hao Zheng (UCF)

5: ARIES: Accelerating Distributed Training in Chiplet-based Systems via Flexible Interconnects
Lingxiang Yin (UCF)
Amir Ghazizadeh Ahsaei (UCF)
Hao Zheng (UCF)
Ahmed Louri (George Washington University)

927: Monad: Towards Cost-effective Specialization for Chiplet-based Spatial Accelerators
Xiaochen Hao (Peking University)
Zijian Ding (Peking University)
Yuan Wang (Peking University)
Yun (Eric) Liang (Peking University)
Jieming Yin (Nanjing University)

560: RONet: Scaling GPU System with Silicon Photonic Chiplet
Chengeng Li (The Hong Kong University of Science and Technology)
Fan Jiang (The Hong Kong University of Science and Technology)
Shixi Chen (The Hong Kong University of Science and Technology)
Xianbin LI (The Hong Kong University of Science and Technology)
Yinyi LIU (The Hong Kong University of Science and Technology)
Lin Chen (The Hong Kong University of Science and Technology)
Xiao LI (The Hong Kong University of Science and Technology)
Jiang Xu (The Hong Kong University of Science and Technology)

365: Thermally-aware Multi-core Chiplet Stacking
Gaurav Kothari (State University of New York)
Kanad Ghose (State University of New York)

17:30 – 19:00, Room: Sculptor
PD 2: Advanced Floorplanning and Partitioning
Session Chair: Tsung-Yi Ho (UHK)

267: Handling Orientation and Aspect Ratio of Modules in Electrostatics-based Large Scale Fixed-Outline Floorplanning

Xingquan Li (Peng Cheng Laboratory)
Bei Yu (The Chinese University of Hong Kong)
Fuxing Huang (Fuzhou University)
Duanxiang Liu (Fuzhou University)
Wenxing Zhu (Fuzhou University)

280: Floorplanning for Embedded Multi-die Interconnect Bridge Packages

Chung-Chia Lee (National Taiwan University)
Yao-Wen Chang (National Taiwan University)

219: iPL-3D: A Novel Bilevel Programming Model for Die-to-Die Placement

Xueyan Zhao (Peng Cheng Laboratory)
Shijian Chen (Peng Cheng Laboratory)
Yihang Qiu (Peng Cheng Laboratory)
Jiangkao Li (Peng Cheng Laboratory)
Zhipeng Huang (Peng Cheng Laboratory)
Biwei Xie (Peng Cheng Laboratory)
Xingquan Li (Peng Cheng Laboratory)
Yungang Bao (Peng Cheng Laboratory)

843: An Open-Source Constraints-Driven General Partitioning Multi-Tool for VLSI Physical Design

Ismail Bustany (AMD)
Grigor Gasparyan (AMD)
Andrew Kahng (University of California San Diego)
Bodhisatta Pramanik (University of California San Diego)
Zhiang Wang (University of California San Diego)
Ioannis Koutis (New Jersey Institute of Technology)

17:30 – 19:00, Room: Artisan
CPS 2: Fast and Trustworthy Embedded Systems
Session Chair: Caiwen Ding (University of Connecticut)

61： Data Recomputation for Multithreaded Applications
Gulsum Gudukbay Akbulut (Pennsylvania State University)
Mahmut T. Kandemir (Pennsylvania State University)
Mustafa Karakoy (Hanyang University)
Wonil Choi (Tubitak- Bilgem)

40：HAPIC: a Scalable, Lightweight and Reactive Cache for Persistent-Memory-based Index
Chih-Ting Lo (National Taiwan University)
Yun-Chih Chen (National Taiwan University)
Yuan-Hao Chang (National Taiwan University)
Tei-Wei Kuo (National Taiwan University)

39：DOMINO: Domain-Invariant Hyperdimensional Classification for Multi-Sensor Time Series Data
Junyao Wang (University of California Irvine)
Luke Chen (University of California Irvine)
Mohammad Al Faruque (University of California Irvine)

721： PARseL: Towards a Verified Root-of-Trust over seL4
Ivan De Oliveira Nunes (RIT)
Norrathep Rattanavipanon (Prince of Songkla Univ.)
Seoyeon Hwang (UCI)
Sashidhar Jakkamsetti (UCI)
Gene Tsudik (UCI)

LLM-Aided Design Panel

17:30 – 19:00, Room: Atelier I

Organizer: Deming Chen (University of Illinois at Urbana-Champaign)

Anna Goldie (Google/Stanford)

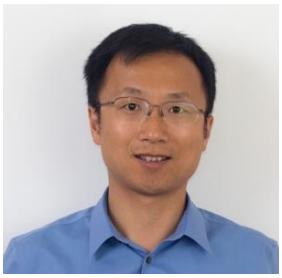

Mark Ren (NVIDIA)

Bei Yu (Chinese University of Hong Kong)

Ramesh Karri (NYU)

Ruchir Puri (IBM)

Moderator: Andrew Kahng (UCSD)

17:30 – 19:00, Room: Atelier II
SysDesign 2: System Level Design and Exploration
Session Chair: Hao Zheng (UCF)

680: Path-based Processing using In-Memory Systolic Arrays for Accelerating Data-Intensive Applications

Sumit Jha (FIU)
Hao Zheng (UCF)
Sven Thijssen (UCF)
Muhammad Rashedul Haq Rashed (UCF)
Rickard Ewetz (UCF)

87: PANDA: Architecture-Level Power Evaluation by Unifying Analytical and Machine Learning Solutions

Shiyu Li (Duke University)
Guanglei Zhou (Duke University)
Jingyu Pan (Duke University)
Chen-Chia Chang (Duke University)
Yiran Chen (Duke University)
Qijun Zhang (Hong Kong University of Science and Technology)
Zhiyao Xie (Hong Kong University of Science and Technology)

266: A Transfer Learning Framework for High-accurate Cross-workload Design Space Exploration of CPU

Duo Wang (Chinese Academy of Sciences)
Mingyu Yan (Chinese Academy of Sciences)
Yihan Teng (Chinese Academy of Sciences)
Dengke Han (Chinese Academy of Sciences)
Haoran Dang (Chinese Academy of Sciences)
Xiaochun Ye (Chinese Academy of Sciences)
Dongrui Fan (Chinese Academy of Sciences)

281: A General Wavelength-Routed Optical Networks-on-Chip Model with Applications to Provably Good Customized and Fault-Tolerant Topology Designs

Yan-Lin Chen (National Taiwan University)
Wei-Che Tseng (National Taiwan University)
Wei-Yao Kao (National Taiwan University)
Yao-Wen Chang (National Taiwan University)

ACM Student Research Competition (SRC) @ ICCAD

19:00 – 20:30, Room: Artisan

The ACM Student Research Competition is an internationally recognized venue enabling undergraduate and graduate students who are ACM members to:

1. Experience the research world — this is a first for many undergraduates!
2. Share research results and exchange ideas with other students, judges, and conference attendees.
3. Rub shoulders with academic and industry luminaries.
4. Understand the practical applications of their research.
5. Perfect their communication skills.
6. Receive prizes and gain recognition from ACM and the greater computing community.

The ACM Special Interest Group on Design Automation (ACM SIGDA) is organizing such an event in conjunction with the International Conference on Computer-Aided Design (ICCAD). Authors of accepted submissions will receive a grant from ACM SIGDA to attend the event at ICCAD. The event consists of several rounds, as described at http://src.acm.org/ and http://www.acm.org/student-research-competition, where you can find more details on student eligibility and timeline.

2023 Co-Chairs:

Debjit Pal (The University of Illinois at Chicago)
Michael Zuzak (Rochester Institute of Technology)

Technical Program: Tuesday, October 31, 2023

8:00 – 10:00, Room: Sculptor
HSL 1: Novel Ideas in Logic Synthesis
Session Chair: Cunxi Yu (University of Maryland)

204: Fast Exact NPN Classification with Influence-aided Canonical Form
Yonghe Zhang (Shenzhen University)
Liwei Ni (Peng Cheng Laboratory)
Shenggen Zheng (Peng Cheng Laboratory)
Jiaxi Zhang (Peking University)
Guojie Luo (Peking University)
Huawei Li (Chinese Academy of Sciences)

292: LIM-GEN: A Data-guided Framework for Automated Generation of Heterogeneous Logic-in-Memory Architecture
Libo Shen (Chinese Academy of Sciences)
Boyu Long (Chinese Academy of Sciences)
Rui Liu (Chinese Academy of Sciences)
Xiaoyu Zhang (Chinese Academy of Sciences)
Yinhe Han (Chinese Academy of Sciences)
Xiaoming Chen (Chinese Academy of Sciences)

331: EasySO: Exploration-enhanced Reinforcement Learning for Logic Synthesis Sequence Optimization and a Comprehensive RL Environment
Junjie Ye (Huawei Noah's Ark Lab)
Mingxuan Yuan (Huawei Noah's Ark Lab)
Jianye Hao (Huawei Noah's Ark Lab)
Jianyong Yuan (Shanghai Jiao Tong University)
Peiyu Wang (Shanghai Jiao Tong University)
Junchi Yan (Shanghai Jiao Tong University)

437: MiniTNtk: An Exact Synthesis-based Method for Minimizing Transistor Network
Weihua Xiao (Shanghai Jiao Tong University)
Shanshan Han (Shanghai Jiao Tong University)
Yue Yang (Shanghai Jiao Tong University)
Shaoz Yang (Shanghai Jiao Tong University)
Cheng Zheng (Shanghai Jiao Tong University)
Jingsong Chen (Huawei Technologies Co., Ltd.)
Tingyuan Liang (Huawei Technologies Co., Ltd.)
Lei Li (Huawei Technologies Co., Ltd.)
Weikang Qian (Shanghai Jiao Tong University)

482: WolFEx: Word-Level Function Extraction and Simplification from Gate-Level Arithmetic Circuits

Kuo-Wei Ho (National Taiwan University)
Shao-Ting Chung (National Taiwan University)
Tian-Fu Chen (National Taiwan University)
Yu-Wei Fan (National Taiwan University)
Che Cheng (National Taiwan University)
Cheng-Han Liu (National Taiwan University)
Jie-Hong Roland Jiang (National Taiwan University)

8:00 – 10:00, Room: Artisan
SS5: Frontiers in Edge AI: Technology, Algorithms, and Emerging Trends
Session Chair: Xiaoxuan Yang (Duke University)

Towards the Efficiency, Heterogeneity, and Robustness of Edge AI
Bokyung Kim (Duke University)
Zhixu Du (Duke University)
Jingwei Sun (Duke University)
Yiran Chen (Duke University)

Ultra-Efficient Edge AI using FeFET-based Monolithic 3D Integration
Shubham Kumar (STAR)
Yogesh Chauhan (Indian Institute of Technology)
Hussam Amrouch (TUM)

Algorithm/Hardware Codesign for Few-Shot Learning at the Edge
Ann Franchesca Laguna (De La Salle University)
Mohammad Mehdi Sharifi (Notre Dame)
Liu Liu, Andrew Hennessee (Notre Dame)
Clayton O'Dell (Notre Dame)
Ian O'Connor (Notre Dame)
Michael Niemier (Notre Dame)
X. Sharon Hu (Notre Dame)
Dayane Reis (USF)

Hyperdimensional Computing for Resilient Edge Learning
Hamza Errahmouni Barkam (University of California Irvine)
SungHeon Jeong (University of California Irvine)
Sanggeon Yun (University of California Irvine)
Calvin Yeung (University of California Irvine)
Zhuowen Zou (University of California Irvine)
Mohsen Imani (University of California Irvine)
Xun Jiao (Vilanova University)
Narayan Srinivasa (Intel Labs)

Interpretable and Robust Reasoning on Edge
Hanning Chen (UC Irvine)
Prathyush Poduval (UC Irvine)
Haleh Alimohamadi (UCLA)
Hugo Latapie (CISCO)
Mohsen Imani (UC Irvine)

8:00 – 10:00, Room: Atelier I
Security-Alg 1: IP and IC Trust: Reflections and Advancement
Session Chairs: Amin Rezaei (California State University)
Satwik Patnaik (University of Delaware)

974: Reflections on trusting TrustHUB
Christian Krieg (ICT)

286: Automated Hardware Trojan Detection at LUT Using Explainable Graph Neural Networks
Lingjuan Wu (Northwestern Polytechnical University)
Xuelin Zhang (Northwestern Polytechnical University)
Han Li (Northwestern Polytechnical University)
Hao Su (Huazhong Agricultural University)
Yu Tai (Huazhong Agricultural University)
Wei Hu (Huazhong Agricultural University)

727: Risk-Aware and Explainable Framework for Ensuring Guaranteed Coverage in Evolving Hardware Trojan Detection
Rahul Vishwakarma (Long Beach University)
Amin Rezaei (Long Beach University)

713: PDNSig: Identifying Multi-Tenant Cloud FPGAs with Power Distribution Network-based Signatures
Huifeng Zhu (Washington University STL)
Xuan Zhang (Washington University STL)
Weidong Cao (George Washington University)

319: An Anti-removal-attack Hardware Watermarking Method based on Polymorphic Gates
Yongliang Chen (Peking University)
Xiaole Cui (Peking University)
Pengyuan Yang (Peking University)
Gang Qu (University of Maryland)

8:00 – 10:00, Room: Atelier II
AI-Tools 1: Towards Graph-Learning Assisted Electronic Design Automation (EDA)
Session Chair: Debjit Sinha (Google)

63: Robust GNN-based Representation Learning for HLS
Atefeh Sohrabizadeh (University of California LA)
Yunsheng Bai (University of California LA)
Yizhou Sun (University of California LA)
Jason Cong (University of California LA)

186: Accelerating Exact Combinatorial Optimization via RL-based Initialization -- A Case Study in Scheduling
Jiaqi Yin (University of Utah)
Cunxi Yu (University of Utah)

353: Memory-aware Scheduling for Complex Wired Networks with Iterative Graph Optimization
Shuzhang Zhong (Peking University)
Meng Li (Peking University)
Yun (Eric) Liang (Peking University)
Runsheng Wang (Peking University)
Ru Huang (Peking University)

428: GraPhSyM: Graph Physical Synthesis Model
Ahmed Agiza (Brown University)
Sherief Reda (Brown University)
Rajarshi Roy (NVIDIA)
Teodor-Dumitru Ene (NVIDIA)
Saad Godil (NVIDIA)
Bryan Catanzaro (NVIDIA)

498: GRAFT: Graph-assisted Reinforcement learning for Automated SSD Firmware Testing
Gyohun Jeong (Samsung)
Mingyu Pi (Samsung)
Hyukil Kwon (Samsung)
Hakyoung Lim (Samsung)
Eungchae Kim (Samsung)
Sunghee Lee (Samsung)
Yoon Hyeok Lee (Samsung)
Youngmin Oh (Samsung)
Bosun Hwang (Samsung)

Denotes 2023 Best Paper Nominee

10:30 – 12:30, Room: Sculptor
Test 1: Simulation and Emulation
Session Chair: Vidya Chhabria (ASU)

10: Fast and Scalable Gate-level Simulation in Massively Parallel Systems
Haichuan Hu (Nanchang University)
Zichen Xu (Nanchang University)
Yuhao Wang (Nanchang University)
Fangming Liu (Pengcheng Lab)

161: SurgeFuzz: Surge-Aware Directed Fuzzing for CPU Designs
Yuichi Sugiyam (University of Tokyo)
Reoma Matsuo (University of Tokyo)
Ryota Shioya (University of Tokyo)

743: TaintFuzzer: SoC Security Verification using Taint Inference-enabled Fuzzing
Muhammad Monir Hossain (University of Florida)
Nusrat Farzana Dipu (University of Florida)
Kimia Zamiri Azar (University of Florida)
Fahim Rahman (University of Florida)
Farimah Farahmandi (University of Florida)
Mark Tehranipoor (University of Florida)

528: Checkpoint Placement for Systematic Fault-Injection Campaigns
Christian Dietrich (Hamburg University)
Matthias Mnich (Hamburg University)
Tim-Marek Thomas (Leibnix University)

475: Sphinx: A Hybrid Boolean Processor-FPGA Hardware Emulation System
Ruiyao Pu (Fudan University)
Yiwei Sun (Fudan University)
Xuan Zeng (Fudan University)
Fan Yang (Fudan University)
Li Shang (Fudan University)
Pei-Hsin Ho (UniVista Industrial Software Group)

10:30 – 12:30, Room: Artisan
Tutorial: Introduction to Hybrid Quantum-Classical Programming Using C++ Quantum Extension
Session Chairs: Xin-Chuan (Ryan) Wu (Intel)
Shavindra Premaratne (Intel)
Kevin Rasch (Intel)

10:30 – 12:30, Room: Atelier I
Reconf-LP 1: "Hot" Microarchitectures: Modeling and Design
Session Chairs: Georgios Zervakis (University of Patras)
Aoyang Zhang (Tsinghua University)

1020: ACOR: On the Design of Energy-Efficient Autocorrelation for Emerging Edge Applications
Charalampos Eleftheriadis (Queen's University)
Georgios Karakonstantis (Queen's University)

235: Fast Full-Chip Parametric Thermal Analysis Based on Enhanced Physics Enforced Neural Networks
Liang Chen (University of California Riverside)
Jincong Lu (University of California Riverside)
Wentian Jin (University of California Riverside)
Sheldon Tan (University of California Riverside)

734: DiCA: A Hardware-Software Co-Design for Differential Check-Pointing in Intermittently Powered Devices
Antonio Joia Neto (Rochester Institute of Technology)
Adam Caulfield (Rochester Institute of Technology)
Christabelle Alvares (Rochester Institute of Technology)
Ivan De Oliveira Nunes (Rochester Institute of Technology)

752: IT-DSE: Invariant Risk Minimized Transfer Microarchitecture Design Space Exploration
Ziyang Yu (Chinese University of Hong Kong)
Chen BAI (Chinese University of Hong Kong)
Bei Yu (Chinese University of Hong Kong)
Martin Wong (Chinese University of Hong Kong)
Shoubo Hu (Huawei Noah's Ark Lab)
Ran Chen (Huawei Noah's Ark Lab)
Mingxuan Yuan (Huawei Noah's Ark Lab)
Taohai He (HiSilicon)

420: Real-time Thermal Map Estimation for AMD Multi-Core CPUs using Transformer
Jincong Lu (University of California Riverside)
Jinwei Zhang (University of California Riverside)
Sheldon Tan (University of California Riverside)

10:30 – 12:30, Room: Atelier II
AI-Sys 1: Efficient Accelerator Design
Session Chairs: Hyoukjou Kwon (UC-Irvine)
Ziyun Li (Meta)

511: An Energy-Efficient 3D Point Cloud Neural Network Accelerator With Efficient Filter Pruning, MLP Fusion, and Dual-Stream Sampling

Changchun Zhou (Peking University)
Yuzhe Fu (Peking University)
Min Liu (Peking University)
Siyuan Qiu (Peking University)
Ge Li (Peking University)
Hailong Jiao (Peking University)
Yifan He (Reconova Technologies)

258: BOOST: Block Minifloat-Based On-Device CNN Training Accelerator with Transfer Learning

Chuliang Guo (Zhejiang University)
Binglei Lou (University of Sydney)
Xueyuan Liu (University of Sydney)
David Boland (University of Sydney)
Philip H.W. Leong (University of Sydney)
Cheng Zhuo (Key Lab od CS&AUS)

96: SpOctA: A 3D Sparse Convolution Accelerator with Octree-Encoding-Based Map Search and Inherent Sparsity-Aware Processing

Dongxu Lyu (Shanghai Jiao Tong University)
Zhenyu Li (Shanghai Jiao Tong University)
Yuzhou Chen (Shanghai Jiao Tong University)
Jinming Zhang (Shanghai Jiao Tong University)
Ningyi Xu (Shanghai Jiao Tong University)
Guanghui He (Shanghai Jiao Tong University)

422: RNA-ViT: Reduced-Dimension Approximate Normalized Attention Vision Transformers for Latency Efficient Private Inference

Dake Chen (USC)
Yuke Zhang (USC)
Peter Beerel (USC)
Chenghao Li (USC)
Souvik Kundu (Intel labs)

256: SAGA: Sparsity-Agnostic Graph Convolutional Network Acceleration with Near-optimal Workload Balance

Sanjay Gandham (UCF)
Lingxiang Yin (UCF)
Hao Zheng (UCF)
Mingjie Lin (UCF)

12:30 – 14:00, Room: Gallery II Ballroom
Keynote Lunch: *Margaret Martonosi (Princeton)*
Tickets Required Limited Access: First Come, First Served

14:00 – 15:30, Room: Sculptor
Security-Arch 2: How Do We Patch Our Neural Nets (Against Information Leakage)?
Session Chairs: Nimisha Limaye (Synopsys)
Jeremy Blackstone (Howard University)

697: Side Channel-assisted Inference Attack on Machine Learning-based ECG Classification
Chongzhou Fang (University of California Davis)
Ning Miao (University of California Davis)
Houman Homayoun (University of California Davis)
Jialin Liu (Temple University)
Han Wang (Temple University)

631: THE-V: Verifiable Privacy-Preserving Neural Network via Trusted Homomorphic Execution
Yuntao Wei (Beihang University)
Xueyan Wang (Beihang University)
Song Bian (Beihang University)
Weisheng Zhao (Beihang University)
Yier Jin (University of Science and Technology China)

141: Deep-learning Model Extraction through Software-based Power Side-channel
Xiang Zhang (Northeastern University)
Aidong Adam Ding (Northeastern University)
Yunsi Fei (Northeastern University)

228: SystemC Model of Power Side-Channel Attacks Against AI Accelerators: Superstition or not?
Andrija Neskovic (Lubeck University)
Saleh Mulhem (Lubeck University)
Alexander Treff (Lubeck University)
Rainer Buchty (Lubeck University)
Thomas Eisenbarth (Lubeck University)
Mladen Berekovic (Lubeck University)

14:00 – 15:30, Room: Artisan
NewComp2: Quantum Circuitry Unleashed: Innovations in Simulation, Synthesis, and Optimization
Session Chair: Tsung-Wei Huang (UW-Madison)

132: Full State Quantum Circuit Simulation Beyond Memory Limit
Yilun Zhao (Chinese Academy of Sciences)
Yu Chen (Chinese Academy of Sciences)
Ying Wang (Chinese Academy of Sciences)
Kaiyan Chang (Chinese Academy of Sciences)
Yinhe Han (Chinese Academy of Sciences)
Bingmeng Wang (Capital Normal University)
Bing Li (Capital Normal University)
He Li (Southeast University)

224: Optimizing LUT-based Quantum Circuit Synthesis using Relative Phase Boolean Operations
David Lawrence (Ritsumeikan University)
Bantug Clarino (Ritsumeikan University)
Naoya Asada (Ritsumeikan University)
Shigeru Yamashita (Ritsumeikan University)

225: Optimal Layout Synthesis for Quantum Circuits as Classical Planning
Irfansha Shaik (Aarhus University)
Jaco Van de Pol (Aarhus University)

352: Single-Qubit Gates Matter for Optimizing Quantum Circuit Depth in Qubit Mapping
Sanjiang Li (University of Sydney)
Ky Dan Nguyen (University of Sydney)
Zachary Clare (University of Sydney)
Yuan Feng (University of Sydney)

14:00 – 15:30, Room: Atelier I
Security-Arch 3: Microelectronics Security: What Lies Ahead?
Session Chairs: Patanjali SLPSK (UF)

692: Protection Against Physical Attacks Through Self-Destructive Polymorphic Latch

Andrew John Cannon (University of Florida)
Tasnuva Farheen (University of Florida)
Sourav Roy (University of Florida)
Domenic Forte (University of Florida)
Shahin Tajik (Worcester Polytechnic Institute)

75: SAM: A Scalable Accelerator for Number Theoretic Transform Using Multi-Dimensional Decomposition

Cheng Wang (Xi'an Jiatong University)
Mingyu Gao (Tsinghua University)

792: KyberMat: Efficient Hardware Accelerator for Matrix-Vector Multiplication in CRYSTALS-Kyber Scheme via NTT and Polyphase Decomposition

Weihang Tan (University of Minnesota)
Keshab Parhi (University of Minnesota)
Yingjie Lao (Clemson University)

540: CRYSTALS-Dilithium on RISC-V Processor: Lightweight Secure Boot using Post Quantum Digital Signature

Naina Gupta (NTU)
Arpan Jati (NTU)
Anupam Chattopadhyay (NTU)

14:00 – 15:30, Atelier II
Analog 2: Novel Frameworks and Methodologies for Optimizing Analog/Mixed-Signal Circuits
Session Chair: Sheldon Tan (UCR)

494: Multi-Product Optimization for 3D Heterogeneous Integration with D2W Bonding
Kai-Yuan Chao (Huawei Technology)
Zhen Zhuang (Chinese University of Hong)
Bei Yu (Chinese University of Hong)
Tsung-Yi Ho (Chinese University of Hong)
Martin Wong (Chinese University of Hong)

307: Distributionally Robust Circuit Design Optimization under Variation Shifts
Yifan Pan (University of California Santa Barbara)
Zichang He (University of California Santa Barbara)
Zheng Zhang (University of California Santa Barbara)
Nanlin Guo (Fudan University)

683: Practical Layout-Aware Analog/Mixed-Signal Design Automation with Bayesian Neural Networks
Ahmet Budak (University of Texas)
Keren Zhu (University of Texas)
David Z. Pan (University of Texas)

154: Design and Optimization of Low-Dropout Voltage Regulator Using Relational Graph Neural Network and Reinforcement Learning in Open-Source SKY130 Process
Zonghao Li (University of Toronto)
Anthony Chan Carusone (University of Toronto)

16:00 – 18:00, Room: Sculptor
Timing 1: Analyzing and Optimizing for PPA, Timing and Reliability
Session Chairs: Masanori Hashimoto (Kyoto University)

57: MasterRTL: A Pre-Synthesis PPA Estimation Framework for Any RTL Design
Wenji Fang (Hong Kong University of Science and Technology)
Yao Lu (Hong Kong University of Science and Technology)
Shang Liu (Hong Kong University of Science and Technology)
Qijun Zhang (Hong Kong University of Science and Technology)
Zhiyao Xie (Hong Kong University of Science and Technology)
Hongce Zhang (Hong Kong University of Science and Technology)
Lisa Wu Wills (Duke University)
Ceyu Xu (Duke University)

486: Local Layout Effect-aware Static Timing Analysis by use of a New Sensitivity-based Library
Juyeon Kim (Samsung)
Cheoljun Bae (Samsung)
Yoobeom Kim (Samsung)
Jae Hoon Kim (Samsung)
Hyun-Seung (Samsung)
Seo Changho Han (Kumoh National Institute of Technology)

847: Risk Propagation Aware Vector Profiling for High Coverage Dynamic IR-drop Analysis
Yihan Wen (Beijing University)
Juan Li (Beijing University)
Xiaoyi Wang (Beijing University)

276: Design and Technology Co-optimization for Useful Skew Scheduling on Multi-bit Flip-flops
Suwan Kim (Seoul National University)
Taewhan Kim (Seoul National University)

468: READ: Reliability-Enhanced Accelerator Dataflow Optimization using Critical Input Pattern Reduction
Zuodong Zhang (Peking University)
Renjie Wei (Peking University)
Meng Li (Peking University)
Yibo Lin (Peking University)
Runsheng Wang (Peking University)
Ru Huang (Peking University)

16:00 – 18:00, Room: Artisan
SS7: 2023 CAD Contest at ICCAD
Session Chairs: Chun-Yao Wang) Andy Yu-Guang Chen, (National Central University)

48: Overview of 2023 CAD Contest at ICCAD
Takashi Sato (Kyoto University)
Chun-Yao Wang (National Tsing Hua University)
Yu-Guang Chen (National Central University)
Tsung-Wei Huang (University of Wisconsin)

39.1: 2023 ICCAD CAD Contest Problem A: Multi-bit Large-scale Boolean Matching
Chung-Han Chou (Cadence)
Chih-Jen Hsu (Cadence)
Chi-An Wu (Cadence)
Kuan-Hua Tu (Cadence)
Kei-Yong Khoo (Cadence)

32: 2023 ICCAD CAD Contest Problem B: 3D Placement with Macros
Kai-Shun Hu (Synopsys)
Hao-Yu Chi (Synopsys)
I-Jye Lin (Synopsys)
Yi-Hsuan Wu (Synopsys)
Wei-Hsu Chen (Synopsys)
Yi-Ting Hsieh (Synopsys)

53: 2023 ICCAD CAD Contest Problem C: Static IR Drop Estimation Using Machine Learning
Gana Surya Prakash Kadagala (Arizona State University)
Vidya A. Chhabria (Arizona State University)

56: IEEE CEDA DATC: Emerging Foundations in IC Physical Design and ML-CAD Research
Jinwook Jung (UCSD)
Andrew Kahng (UCSD)
Sayak Kundu (UCSD)
Zhiang Wang (UCSD)
Dooseok Yoon (UCSD)

Contest for Machine Learning on Hardware

16:00 – 18:00, Room: Atelier I

ACM/IEEE Contest for Machine Learning on Hardware at ICCAD is a challenging research and development competition, focusing on real-world problems that require the implementation of machine learning algorithms on both conventional hardware platforms as well as emerging hardware platforms. It is open to teams world-wide. The top three teams of each track will be announced during 2023 International Conference on Computer-Aided Design (ICCAD). Winners will be invited to give presentations on their methods at the conference.

Given the great interest received last year, this year's contest is separated into two tracks (tinyML contest on microcontrollers and quantum computers challenge) to accommodate researchers with diverse backgrounds.

1. **Track 1** is on **Conventional Platforms**, please refer to the 2023 TinyML Design Contest at ICCAD.
2. **Track 2** is on **Emerging Platforms**, please refer to the 2023 Quantum Computing for Drug Discovery Challenge at ICCAD.

The top winners in each track will receive up to $2,500 USD.
For details, please visit: https://mlhardwarecontest.github.io/Contest-MachineLearning-On-Hardware-2023/index.html

Contest Chairs:
Track 1: Dawei Li (South-Central Minzu University)
Track 2: Zhiding Liang (University of Notre Dame)

Contest Sponsors

16:00 – 18:00, Room: Atelier II
AI-Sys 2: Design Tools and Performance Optimization for DNN Acceleration
Session Chairs: Bingzhe Li (UT Dallas)
Syed Shakib Sarwar (Meta)

120： Improving Realistic Worst-Case Performance of NVCiM DNN Accelerators through Training with Right-Censored Gaussian Noise

Zheyu Yan (University of Notre Dame)
Yifan Qin (University of Notre Dame)
Yiyu Shi (University of Notre Dame)
X. Sharon Hu (University of Notre Dame)
Wujie Wen (NC State University)

451： Runtime Row/Column Activation Pruning for ReRAM-based Processing-in-Memory DNN Accelerators

Xikun Jiang (Shandong University)
Zhaoyan Shen (Shandong University)
Zhiping Jia (Shandong University)
Lei Ju (Shandong University)
Zhiyong Zhang (Shandong University)
Dongxiao Yu (Shandong University)
Siqing Sun (Cloud Inspur)
Ping Yin (Cloud Inspur)

343: An Open Source Mixed-Precision Neural Network Accelerator Design Framework for FPGAs

Erjing Luo (Chinese Academy of Sciences)
Haitong Huang (Chinese Academy of Sciences)
Cheng Liu (Chinese Academy of Sciences)
Guoyu Li (Chinese Academy of Sciences)
Xiaowei Li (Chinese Academy of Sciences)
Ying Wang (Chinese Academy of Sciences)
Huawei Li (Chinese Academy of Sciences)
Bing Yang (Harbin University)

Denotes 2023 Best Paper Nominee

746: Accel-GCN: High-Performance GPU Accelerator Design for Graph Convolution Networks

Xi Xie (University of Connecticut)
Hongwu Peng (University of Connecticut)
Amit Hasan (University of Connecticut)
Shaoyi Huang (University of Connecticut)
Jiahui Zhao (University of Connecticut)
Haowen Fang (University of Connecticut)
Wei Zhang (University of Connecticut)
Caiwen Ding (University of Connecticut)
Omer Khan (University of Connecticut)
Tong Geng (University of Rochester)

308：MirrorNet: A TEE-Friendly Framework for Secure On-device DNN Inference

Ziyu Liu (Northeastern University)
Yukui Luo (Northeastern University)
Shijin Duan (Northeastern University)
Tong Zhou (Northeastern University)
Xiaolin Xu (Northeastern University)

IEEE/ACM EDA Job Fair

18:00 – 19:00, Room: Gallery II & III

ACM SIGDA and IEEE CEDA will be holding their annual (2nd) EDA job fair at the 2023 International Conference on Computer-Aided Design (ICCAD) conference, in-person, on Tuesday October 31, 2023 (6:00 PM—9:00 PM PDT) in San Francisco, CA. The job fair is open to all attendees of the ICCAD conference. Attendees are encouraged to submit a CV or résumé for circulation beforehand.

The EDA job fair is a place for students and professionals looking for internships or jobs to meet with representatives from companies and academia. In the beginning of the event, one representative from each organization has the opportunity to introduce the organization and its job opening(s).

Afterward, all attendees have a chance to mingle. If you are representing a company, research organization or university and would like to participate in the job fair, please send an email with your contact information to sigdajobfair@gmail.com. In addition, all job fair participants must register (at least one-day) for the ICCAD.

Participating Companies

- **Synopsys**
- **Cadence**
- **NVIDIA Research**
- **The Chinese University of Hong Kong**
- **And more!**

Submit Your CV

As part of the job fair, attendees are encouraged to submit their CV to all of the participating companies before the job fair. This will help companies identify and possibly schedule an onsite interview with you when you attend. Please clearly indicate on your CV or résumé your availability and type of job interest, along with contact information! By uploading your CV, you are agreeing to be contacted by participating and sponsor companies about recruiting and job openings.

Organizers

Chair: Vidya Chhabria (Arizona State University)
Co-Chair: Jeff (Jun) Zhang (Arizona State University)

8:00 – 9:00, Room: Gallery I & II
Keynote: *Vamsi Boppana (AMD)*

9:00 – 10:30, Room: Sculptor
NewComp 3: Frontiers in Quantum Computing: Novel Algorithms and Beyond
Session Chair: Hussam Amrouch (TUM)

178: Effective and Efficient Qubit Mapper
Hao Fu (University of Science and Technology China)
Mingzheng Zhu (University of Science and Technology China)
Jun Wu (University of Science and Technology China)
Wei Xie (University of Science and Technology China)
Zhaofeng Su (University of Science and Technology China)
Xiang-yang Li (University of Science and Technology China)

446: Exact Logic Synthesis for Reversible Quantum-Flux-Parametron Logic
Olivia Chen (Tokyo City University)
Nobuyuki Yoshikawa (Yokohama National University)
Rongliang Fu (Chinese University of Hong Kong)
Tsung-Yi Ho (Chinese University of Hong Kong)

449: DLPlace: A Delay-Line Clocking-based Placement Framework for AQFP Circuits
Olivia Chen (Tokyo City University)
Nobuyuki Yoshikawa (Yokohama National University)
Bei Yu (Chinese University of Hong Kong)
Rongliang Fu (Chinese University of Hong Kong)
Tsung-Yi Ho (Chinese University of Hong Kong)

487: QPulseLib: Accelerating the Pulse Generation of Quantum Circuit with Reusable Patterns
Wuwei Tian (Zhejiang University)
Xinghui Jia (Zhejiang University)
Siwei Tan (Zhejiang University)
Jianwei Yin (Zhejiang University)
Liqiang Lu (Zhejiang University)
Zixuan Song (Zhejiang University)

9:00 – 10:30, Room: Artisan
SS8: Accelerating EDA Algorithms with Heterogeneous Parallelism
Session Chair: Mark Ren (NVIDIA)

43: Programming Dynamic Task Parallelism for Heterogeneous EDA Algorithms

Cheng-Hsiang Chiu (University of Wisconsin Madison)
Dian-Lun Lin (University of Wisconsin Madison)
Tsung-Wei Huang (University of Wisconsin Madison)

46: Accelerating Routability and Timing Optimization with Open-Source AI4EDA Dataset CircuitNet and Heterogeneous Platforms

Xun Jiang (Peking University)
Zizheng Guo (Peking University)
Zhuomin Chai (Peking University)
Yuxiang Zhao (Peking University)
Yibo Lin (Peking University)
Runsheng Wang (Peking University)
Ru Huan (Peking University)

Macro Placement with GPU-accelerated Placer and Automated Parameter Tuning

Anthony Agnesina (NVidia Research)

Heterogenous Acceleration for Design Rule Checking

Bei Yu (Chinese University of Hong Kong)

9:00 – 10:30, Room: Atelier I
PD 3: The State-of-the-Art Placement
Session Chair: Nima Karimpour Darav (AMD)

138: HyPlace-3D: A Hybrid Placement Approach for 3D ICs Using Space Transformation Technique

Jai-Ming Lin (National Cheng Kung University)
Yu-Chien Lin (National Cheng Kung University)
Hsuan Kung (National Cheng Kung University)
Wei-Yuan Lin (National Cheng Kung University)

324: Clock Tree Aware Global Placement for Low Power

Jinghao Ding (Southwest University)
Linhao Lu (Southwest University)
Zhaoqi Fu (Southwest University)
Jie Ma (Southwest University)
Yuanrui Qi (Southwest University)
Mengshi Gong (Southwest University)
Wenxin Yu (Southwest University)

354: Stronger Mixed-Size Placement Backbone Considering Second-Order Information

Yifan Chen (Peking University)
Zaiwen Wen (Peking University)
Yun (Eric) Liang (Peking University)
Yibo Lin (Peking University)

381: Systolic Array Placement on FPGAs

Hailiang Hu (Texas A&M University)
Donghao Fang (Texas A&M University)
Jiang Hu (Texas A&M University)
Bo Yuan (Rutgers)
Wuxi Li (AMD)

9:00 – 10:30, Room: Atelier II
SysDesign 3: Designing Memory and Storage Systems
Session Chair: Xueqing Li (Tsinghua University)

1018: Exploring Error Bits for Memory Failure Prediction: An In-Depth Correlative Study
Wengui Zhang (Huawei)
Jorge Cardoso (Huawei)
Qiao Yu (University of Berlin)
Odej Kao (University of Berlin)

912: SSDe: FPGA-based SSD Express Emulation Framework
Yizhen Lu (University of Illinois Urbana)
Luyang Yu (University of Illinois Urbana)
Deming Chen (University of Illinois Urbana)

1031: HF-Dedupe: Hierarchical Fingerprint Scheme for High Efficiency Data Deduplication on Flash-based Storage Systems
Yun-Shan Hsieh (Academia Sinica)
Yuan-Hao Chang (National Taiwan University)
Yen-Ting Chen (National Taiwan University)
Yu-Pei Liang (National Chung Cheng University
Po-Chun Huang (Taipei Tech)
Kai-Ting Weng (National Tsing Hua University)
Wei-Kuan Shih (National Tsing Hua University)

11:00 – 13:00, Room: Sculptor
HLS 2: Technology Mapping Revived
Session Chair: Ing-Chao Lin (National Cheng Kung University)

403: MapBuf: Simultaneous Technology Mapping and Buffer Insertion for HLS Performance Optimization
Hanyu Wang (ETH)
Carmine Rizzi (ETH)
Lana Josipovic (ETH)

213: EffiSyn: Efficient Logic Synthesis with Dynamic Scoring and Pruning
Jiantang Zhang (Huawei Noah's Ark Lab)
Shuang Wen (Huawei Noah's Ark Lab)
Weihua Sheng (Huawei Noah's Ark Lab)
Yu Huang (Huawei Noah's Ark Lab)
Xing Li (Hisilicon)
Lei Chen (Hisilicon)
Mingxuan Yuan (Hisilicon)

566: EasyMap: Improving Technology Mapping via Exploration-Enhanced Heuristics and Adaptive Sequencing
Junjie Ye (Huawei Noah's Ark Lab)
Xing Li (Huawei Noah's Ark Lab)
Lei Chen (Huawei Noah's Ark Lab)
Mingxuan Yuan (Huawei Noah's Ark Lab)
Jianye Hao (Huawei Noah's Ark Lab)
Anqi Lu (Shanghai Jiao Tong University)
Peiyu Wang (Shanghai Jiao Tong University)
Junchi Yan (Shanghai Jiao Tong University)

593: AlphaSyn: Logic Synthesis Optimization with Efficient Monte Carlo Tree Search
Zehua Pei (Chinese University of Hong Kong)
Zhuolun He (Chinese University of Hong Kong)
Guojin Chen (Chinese University of Hong Kong)
Bei Yu (Chinese University of Hong Kong)
Keren Zhu (Shanghai AI Lab)
Fangzhou Liu (Shanghai AI Lab)
Haisheng Zheng (Shanghai AI Lab)

691: Technology Mapping Using Multi-output Library Cells
Alessandro Tempia Calvino (EPFL)
Giovanni De Micheli (EPFL)

Denotes 2023 Best Paper Nominee

11:00 – 13:00, Room: Artisan
Tutorial: RapidWright: Unleashing the Full Power of FPGA Technology with Domain-Specific Tooling
Session Chairs: Chris Lavin (AMD)
Eddie Hung (AMD)

11:00 – 13:00, Room: Atelier I
AI-Tools 3: Anywhere Anytime – Full Stack AI Deployment
Session Chairs: Jihye Kwon (Cadence)
Uday Mallappa (Intel)

614: DeepGate2: Functionality-Aware Circuit Representation Learning
Zhengyuan Shi (Chinese University of Hong Kong)
Qiang Xu (Chinese University of Hong Kong)
Sadaf Khan (Chinese University of Hong Kong)
Min Li (Chinese University of Hong Kong)
Yi Liu (Chinese University of Hong Kong)
Junhua Huang (Huawei Noah's Ark Lab)
Hui-Ling Zhen (Huawei Noah's Ark Lab)
Mingxuan Yuan (Huawei Noah's Ark Lab)
Zhufei Chu (Ningbo University)
Hongyang Pan (Ningbo University)

744: DiviML: A Module-based Heuristic for Mapping Neural Networks onto Heterogeneous Platforms
Yassine Ghannane (Cornell University)
Mohamed Abdelfattah (Cornell University)

985: PRIMO: A Full-Stack Processing-in-DRAM Emulation Framework for Machine Learning Workloads
Jaehoon Heo (KAIST)
Joo-Young Kim (KAIST)
Sangjin Choi (KAIST)
Sungwoong Yune (KAIST)
Junghoon Kim (KAIST)
Hyojin Sung (KAIST)
Youngjin Kwon (KAIST)
Yongwon Shin (POSTECH)

967: ARES: A Mapping Framework of DNNs towards Diverse PIMs with General Abstractions

Xiuping Cui (Peking University)
Size Zheng (Peking University)
Tianyu Jia (Peking University)
Le Ye (Peking University)
Yun (Eric) Liang (Peking University)

537: Learned Formal Proof Strengthening for Efficient Hardware Verification

Minwoo Kang (UC Berkeley)
Azade Nova (Google)
Eshan Singh (Google)
Geetheeka Sherronn Bathini (Google)
Yuriy Viktorov (Google)

11:00 – 13:00, Room: Atelier II
AI-Sys 3: Efficient and Fair Machine Learning on the Edge
Session Chairs: Peipei Zhou (University of Pittsburgh)

594: Fluid Batching: Exit-Aware Preemptive Serving of Early-Exit Neural Networks on Edge NPUs

Alexandros Kouris (Samsung AI)
Stylianos Venieris (Samsung AI)
Stefanos Laskaridis (Samsung AI)
Nicholas Lane (University of Cambridge)

708: Edge-MoE: Memory-Efficient Multi-Task Vision Transformer Architecture with Task-level Sparsity via Mixture-of-Experts

Hanxue Liang (Georgia Institute of Technology)
Zhiwen Fan (Georgia Institute of Technology)
Zhangyang Wang (Georgia Institute of Technology)
Rishov Sarkar (University of Texas)
Cong "Callie" Hao (University of Texas)

646: FET-OPU: A Flexible and Efficient FPGA-based Overlay Processor for Transformer Networks

Yueyin Bai (Fudan University)
Hao Zhou (Fudan University)
Keqing Zhao (Fudan University)
Hongji Wang (Fudan University)
Jianli Chen (Fudan University)
Jun Yu (Fudan University)
Kun Wang (Fudan University)

636: PP-Transformer: Enable Efficient Deployment of Transformers through Pattern Pruning

Jialin Cao (Fudan University)
Xuanda Lin (Fudan University)
Manting Zhang (Fudan University)
Kejia Shi (Fudan University)
Jun Yu (Fudan University)
Kun Wang (Fudan University)

Denotes 2023 Best Paper Nominee

956: Fast and Fair Medical AI on the Edge through Neural Architecture Search for Hybrid Vision Models

Changdi Yang (Northeastern University)
Peiyan Dong (Northeastern University)
Yanyu Li (Northeastern University)
Zhenglun Kong (Northeastern University)
Pinrui Yu (Northeastern University)
Yanzhi Wang (Northeastern University)
Xue Lin Yi Sheng (George Mason University)
Lei Yang (George Mason University)

14:00 – 16:00, Room: Sculptor
Test 2: Advances in Verification
Session Chairs: Aman Arora (ASU)
Seetal Potluri (University of Albany)

185: Integrating Exact Simulation into Sweeping for Datapath Combinational Equivalence Checking

Zhihan Chen (Chinese Academy of Sciences)
Xindi Zhang (Chinese Academy of Sciences)
Yuhang Qian (Chinese Academy of Sciences)
Shaowei Ca (Chinese Academy of Sciences)
Qiang Xu (Chinese University of Hong Kong)

151: Verification of Flow-Based Computing Systems using Bounded Model Checking

Suraj Singireddy (University of Texas)
Sumit Jha (FIU)
Muhammad Rashedul Haq Rashed (UCF)
Sven Thijssen (UCF)
Rickard Ewetz (UCF)

844: Efficient Formal Verification and Debugging of Arithmetic Divider Circuits

Jiteshri Dasari (University of Massachusetts)
Maciej Ciesielski (University of Massachusetts)

779: Automatic Inductive Invariant Generation for Scalable Dataflow Circuit Verification

Jiahui Xu (ETH)
Lana Josipovic (ETH)

562: Accurate Hybrid Delay Models for Dynamic Timing Analysis

Arman Ferdowsi (TU Wien)
Ulrich Schmid (TU Wien)
Josef Salzmann (TU Wien)

14:00 – 16:00, Room: Artisan
SS10: In-Sensor AI Computing Towards Next Generation Autonomous Edge Intelligence
Session Chair: Yiyu Shi (University of Notre Dame)

Neural Sensors: AI driven In-Pixel Compute for Compressed Readout Bandwidth
Haley So (Stanford University)
Julien Martel (Stanford University)
Laurie Bose (University of Manchester)
Piotr Dudek (University of Manchester)
Gordon Wetzstein (Stanford University)

On-Silicon In-Sensor Computing
Donhee Ham (Harvard University)

The Path and Challenges to Bionic Vision as the Next "Moore's Law" Won't Be Defined by Transistor Density
Charbel Rizk (Oculi)

38: Learned In-Sensor Visual Computing: From Compression to Eventification
Yu Feng (University of Rochester)
Yuhao Zhu (University of Rochester)
Tianrui Ma (Washington University STL)
Adith Boloor (Washington University STL)
Xuan Zhang (Washington University STL)

50: In-Sensor Radio Frequency Computing for Energy-Efficient Intelligent Radar
Yang Sui (Rutgers University)
Minning Zhu (Rutgers University)
Lingyi Huang (Rutgers University)
Chung-Tse Michael Wu (Rutgers University)
Bo Yuan (Rutgers University)

14:00 – 16:00, Room: Atelier I
AI-Tools 4: Adaptation, Acceleration, Automation – Tripple-A AI Computing
Session Chairs: Kishor Kunal (NVIDIA)

508: Automatic Kernel Generation for Large Language Models on Deep Learning Accelerators
Fuyu Wang (Sun Yat-Sen University)
Minghua Shen (Sun Yat-Sen University)

619: SATformer: Transformer-Based UNSAT Core Learning
Junhua Huang (Noah's Ark Lab)
Hui-Ling Zhen (Noah's Ark Lab)
Mingxuan Yuan (Noah's Ark Lab)
Zhengyuan Shi (Chinese University of Hong Kong)
Min Li (Chinese University of Hong Kong)
Yi Liu (Chinese University of Hong Kong)
Sadaf Khan (Chinese University of Hong Kong)
Qiang Xu (Chinese University of Hong Kong)

700: GPT4AIGChip: Towards Next-Generation AI Accelerator Design Automation via Large Language Models
Yonggan Fu (Georgia Institute of Technology)
Yongan Zhang (Georgia Institute of Technology)
Zhongzhi Yu (Georgia Institute of Technology)
Sixu Li (Georgia Institute of Technology)
Zhifan Ye (Georgia Institute of Technology)
Chaojian Li (Georgia Institute of Technology)
Cheng Wan (Georgia Institute of Technology)
Yingyan (Celine) Lin (Georgia Institute of Technology)

823: 3DNN-Xplorer: A Machine Learning Framework for Design Space Exploration of Heterogeneous 3D DNN Accelerators
Gauthaman Murali (Georgia Institute of Technology)
Aditya Iyer (Georgia Institute of Technology)
Navneeth Ravichandran (Georgia Institute of Technology)
Sung Kyu Lim (Georgia Institute of Technology)

1006: Klotski: DNN Model Orchestration Framework for Dataflow Architecture Accelerators
Chen BAI (Chinese University of Hong Kong)
Xuechao Wei (Chinese University of Hong Kong)
Youwei Zhuo (Chinese University of Hong Kong)
Yi Cai (Chinese University of Hong Kong)
Hongzhong Zheng (Chinese University of Hong Kong)
Bei Yu (Chinese University of Hong Kong)
Yuan Xie (Chinese University of Hong Kong)

14:00 – 16:00, Room: Atelier II
AI-Sys 4: Software/Hardware Co-design of AI Accelerator
Session Chairs: Cheng Tan (Northeastern)

955: SOLE: Hardware-Software Co-design of Softmax and LayerNorm for Efficient Transformer Inference

Wenxun Wang (Tsinghua University)
Shuchang Zhou (MEGVII)
Wenyu Sun (Tsinghua University)
Peiqin Sun (MEGVII)
Yongpan Liu (Tsinghua University)

1002: TSTC: Two-level Sparsity Tensor Core Enabling both Algorithm Flexibility and Hardware Efficiency

Jun Liu (Shanghai Jiao Tong University)
Guohao Dai (Shanghai Jiao Tong University)
Hao Xia (Shanghai Jiao Tong University)
Lidong Guo (Tsinghua University)
Xiangsheng Shi (Tsinghua University)
Jiaming Xu (Xidian University)
Huazhong Yang (Tsinghua University)
Yu Wang (Tsinghua University)

993: A Point Transformer Accelerator with Fine-Grained Pipelines and Distribution-Aware Dynamic FPS

Yaoxiu Lian (Shanghai Jiao Tong University)
Xinhao Yang (Tsinghua University)
Ke Hong (Tsinghua University)
Yu Wang (Tsinghua University)
Guohao Dai (Shanghai Jiao Tong University)
Ningyi Xu (Shanghai Jiao Tong University)

695: INR-Arch: A Dataflow Architecture and Compiler for Arbitrary-Order Gradient Computations in Implicit Neural Representation Processing

Stefan Abi-Karam (Georgia Institute of Technology)
Rishov Sarkar (Georgia Institute of Technology)
Dejia Xu (University of Texas)
Zhiwen Fan (University of Texas)
Zhangyang Wang (University of Texas)
Cong "Callie" Hao (Georgia Institute of Technology)

685: Efficient Sampling and Grouping Acceleration for Point Cloud Deep Learning via Single Coordinate Comparison

Hyunsung Yoon (Pohang University of Science and Technology)
Jae-Joon Kim (Seoul National University)

16:30 – 18:00, Room: Sculptor
DFM 1: Revolutionizing Design for Reliability and Manufacturability
Session Chair: Takashi Soto (Kyoto University)

767: Frequency-Domain Transient Electromigration Analysis Using Circuit Theory
Vidya A. Chhabria (ASU)
Nestor Evmorfopoulos (University of Thessaly)
Mohammad Abdullah Al Shohel (University of Minnesota)
Sachin S. Sapatnekar (University of Minnesota)

577: Frequency-Domain Transient Electromigration Analysis Using Circuit Theory
Guohao Dai (Shenzhen University)
Yanfang Liu (Beihang University)
Yuanqing Cheng (Beihang University)
Wang Kang (Beihang University)
Wei Xing (Beihang University)

505: Spacing Cost-aware Optimal and Efficient Mixed-Cell-Height Detailed Placement for DFM Considerations
Da-Wei Huang (National Taiwan University of Science and Technology)
Ying-Jie Jiang (National Taiwan University of Science and Technology)
Shao-Yun Fang (National Taiwan University of Science and Technology)

737: An Adversarial Active Sampling-based Data Augmentation Framework for AI-Assisted Lithography Modeling
Mingjie Liu (NVIDIA)
Haoyu Yang (NVIDIA)
Brucek Khailany (NVIDIA)
Haoxing Ren (NVIDIA)

434: PostPINN-EM: Fast Post-Voiding Electromigration Analysis Using Two-Stage Physics-Informed Neural Networks
Subed Lamichhane (University of California)
Wentian Jin (University of California)
Liang Chen (University of California)
Mohammadamir Kavousi (University of California)
Sheldon Tan (University of California)

Denotes 2023 Best Paper Nominee

16:30 – 18:00, Room: Artisan

SS11: From Arithmetic Components to Side Channel Attacks: A Deep Dive in Post Quantum Hardware

Session Chairs: Francesco Regazzoni (University of Amsterdam)
Apostolos Fournaris (ISI – ATHENA)

45: Instruction Set Extensions for Post-Quantum Cryptography

Marco Brohet (University of Amsterdam)
Felipe Valencia (Universita della Svizzera italiana)
Francesco Regazzoni (CSEM)

47: Machine Learning based Blind Side-Channel Attacks on PQC-based KEMs - A Case Study of Kyber KEM

PRASANNA RAVI (Nanyang Tech University)
Dirmanto Jap (Nanyang Tech University)
Shivam Bhasin (Nanyang Tech University)
Anupam Chattopadhyay (Nanyang Tech University)

49: A Scalable Hardware/Software Co-design Approach for Efficient Polynomial Multiplication

Lóránt Meszlényi (University of Passau)
Irem Keskinkurt-Paksoy (Middle East Tech University)
Ayesha Khalid (Queen's University Belfast)
Tolga Yalçın (Qualcomm)
Elif Bilge Kavun (University of Passau)

54: Dilithium Hardware-Accelerated Application using OpenCL-based High-Level Synthesis

Alexander Islam El-Kady (University of Patras)
Vassilis Paliuras (University of Patras)
Apostolos Fournaris (ATHENA)

A Scalable Hardware/Software Co-design Approach for Efficient Polynomial Multiplication

Lóránt Meszlényi
İrem Keskinkurt-Paksoy
Elif Bilge Kavun (University of Passau)
Ayesha Khalid
Tolga Yalcin

16:30 – 18:00, Room: Atelier I
Reconf-LP 2: Modern Reconfigurable Architectures and Design Tools
Session Chair: Eli Bozorgzadeh (UCI)

197: FLEX : Introducing FLEXible Execution on CGRA with Spatio-Temporal Vector Dataflow

Thilini Bandara (National University of Singapore)
Dan Wu (National University of Singapore)
Rohan Juneja (National University of Singapore)
Dhananjaya Wijerathne (National University of Singapore) (National University of Singapore)
Tulika Mitra (National University of Singapore)
Li-Shiuan Peh (National University of Singapore)

329: Meltrix: A RRAM-based Polymorphic Architecture Enhanced by Function Synthesis

Boyu Long (Chinese Academy of Sciences)
Libo Shen (Chinese Academy of Sciences)
Xiaoyu Zhang (Chinese Academy of Sciences)
Yinhe Han (Chinese Academy of Sciences)
Xiaoming Chen (Chinese Academy of Sciences)
Xian-He Sun (Illinois Institute of Technology)

493: VecPAC: A Vectorizable and Precision-Aware CGRA

Cheng Tan (Google)
Antonino Tumeo (Pacific Northwest National Lab)
Gabriel Weisz (Microsoft)
Steve Reinhardt (Microsoft)
Deepak Patil (ASU)
Jeff Zhang (ASU)

834: AIM: Accelerating Arbitrary-precision Integer Multiplication on Heterogeneous Reconfigurable Computing Platform Versal ACAP

Jiaqi Yin (University of Maryland)
Cunxi Yu (University of Maryland)
Alex Jones (University of Maryland)
Zhuoping Yang (University of Pittsburgh)
Jinming Zhuang (University of Pittsburgh)
Peipei Zhou (University of Pittsburgh)

979: IIBLAST: Speeding Up Commercial FPGA Routing by Decoupling and Mitigating the Intra-CLB Bottleneck

Chirag Ravishankar (AMD)
Dinesh Gaitonde (AMD)
Shashwat Shrivastava (EPFL)
Stefan Nikolic (EPFL)
Mirjana Stojilovic (EPFL)

16:30 – 18:00, Room: Atelier II
CPS 1: Efficient Machine Learning for Real-World Applications
Session Chair: Meng Li (Peking University)

1030: TRAIN: A Reinforcement Learning Based Timing-Aware Neural Inference on Intermittent Systems

Shu-Ting Cheng (National Cheng Kung University)
Chia-Heng Tu (National Cheng Kung University)
Wen Sheng Lim (National Taiwan University)
Yuan-Hao Chang (National Taiwan University)

4: HyperNode: An Efficient Node Classification Framework Using Hyperdimensional Computing

Yichi Chen (Tianjin University)
Haomin Li (Shanghai Jiao University)
Fangxin Liu (Shanghai Jiao University)
Li Jiang (Shanghai Jiao University)

239: Brain-inspired Trustworthy Hyperdimensional Computing with Efficient Uncertainty Quantification

Mohsen Imani (University of California Irvine)
Yang Ni (University of California Irvine)
Hanning Chen (University of California Irvine)
Zhuowen Zou (University of California Irvine)
Prathyush Poduval (University of Maryland)
Pietro Mercati (Intel Labs)

328: EDS-SLAM: An Energy-efficient Accelerator for Real-time Dense Stereo SLAM with Learned Feature Matching

Qian Huang (Sun Yat-sen University)
Gaoxing Shang (Sun Yat-sen University)
Yu Zhang (Sun Yat-sen University)
Gang Chen (Sun Yat-sen University)

1032: MOC: Multi-Objective Mobile CPU-GPU Co-optimization for Power-efficient DNN Inference

Yushu Wu (Northeastern University)
Chao Wu (Northeastern University)
Yifan Gong (Northeastern University)
Zheng Zhan (Northeastern University)
Yanzhi Wang (Northeastern University)
Yanyu Li (Northeastern University)
Qi Wang (Northeastern University)
Geng Yuan (University of Georgia)

Synopsys Panel - Engineering your Future: Exploring Diverse Career Paths After Graduation

18:30 – 19:30, Room: Gallery I & II

Industry Sponsored Session

Engineering your Future: Exploring Diverse Career Paths After Graduation

In this engaging panel discussion, our panelists from Synopsys will share their own experience on the various career paths one can pursue with an engineering degree. Join us to discover the creative possibilities beyond the engineer's path.

Panelists

Harish Balan (Global Ambassador, Academic & Research Alliances)
Ron Duncan (Sr. Staff Program Manager, Chief Innovation Office)
Rob Aitken (Distinguished Architect, Product & Business Management)
Shela Aboud (Sr. Staff Technical Product Mgmt., Sr Staff, Product & Business Management)

Workshops: Wednesday, November 1

9:00 – 19:00, Room: Curator
Top Picks in Hardware and Embedded Security Workshop

Top Picks recognizes the best of the best in hardware security, spanning the gamut from hardware to microarchitecture to embedded systems.

Top Picks will be selected from hardware security papers that have appeared in leading conferences/journals, including but not limited to top security (e.g., IEEE S&P, CCS), architecture (e.g., ISCA, ASPLOS), CAD (e.g., ICCAD, DAC, DATE), and hardware security (e.g., CHES) venues.

After submission, papers will undergo a down select, whereby "shortlisted" papers will be invited to the Top Picks workshop, co-located with ICCAD 2023. An author(s) of each shortlisted paper is required to present the paper in-person at the workshop. A subset of these will be selected as Top Picks. Select papers are then invited for submission to a special journal issue at IEEE Design & Test.

The in-person Top Picks workshop is on November 1st, 2023

Workshop Agenda is available here: https://www.ieee-hsttc.org/tp-program-2023/

For More Information Please Visit: https://www.ieee-hsttc.org/top-picks-2023/

Workshops: Thursday, November 2

8:00 – 12:00, Room: Salon I | Virtual
VLSI Education Community

New attention is being paid to the semiconductor industry and its needs for workforce development at all levels. VLSI education is an important aspect of that effort. This meeting will provide a forum for all interested parties to discuss undergraduate and graduate VLSI education and how to address this workforce development requirement. Both academic and industry participants are welcome. The meeting will provide a Zoom feed as well as the in-person meeting at ICCAD to allow for the broadest possible participation.

Join Zoom Meeting: https://unl.zoom.us/j/9018725169
Meeting ID: 901 872 5169

8:00 – 17:30, Room: Sculptor
Fast ML for Science

This workshop aims to address emerging challenges and explore innovative solutions in the field of computer-aided design (CAD) for integrated circuits and systems for ultra-low latency and high bandwidth scientific applications. The workshop builds on the ideas laid out in the "Applications and Techniques for Fast Machine Learning in Science" white paper and the corresponding Fast Machine Learning for Science conference series (2023 edition). This workshop at ICCAD 2023 aims to bring domains together and forge new connections with the CAD community.

Scientific applications across particle physics, astrophysics, material sciences, quantum information sciences, fusion energy (and beyond!) utilize data acquisition and in situ processing systems which require very low latency and high data bandwidth custom processing elements and real-time control modules. Integrating data reduction and control applications with real-time machine learning algorithms can enable significant breakthroughs in the sciences. We will bring together researchers, practitioners, and industry experts to exchange ideas, share applications, and discuss the latest advancements in CAD methodologies, algorithms, and tools.

8:00 – 17:30, Room: Artisan
The 25th ACM/IEEE International Workshops on System-Level Interconnect Pathfinding (SLIP)

The 2023 ACM/IEEE International Workshop on System-Level Interconnect Pathfinding (SLIP) is the 25th edition of the Workshop. SLIP, co-located with ICCAD 2023, will bring together researchers and practitioners who have a shared interest in the challenges and futures of system-level interconnect, coming from wide-ranging backgrounds that span system, application, design and technology.

The technical goal of the workshop is to:

1. Identify fundamental problem
2. Foster new pathfinding of design, analysis, and optimization of system-level interconnects with emphasis on system-level interconnect modeling and pathfinding, DTCO-enhanced interconnect fabrics, memory and processor communication links, novel dataflow mapping for machine learning, 2.5/3D architectures, and new fabrics for the beyond-Moore era.

8:00 – 17:30, Room: Atelier I & II
SUSHI: Sustainable Hardware Security: An Interactive Workshop

The desire for digital sovereignty, recent global semiconductor shortages, and geopolitical interests are driving forces behind various worldwide initiatives to strengthen semiconductor technology and manufacturing on a national and regional basis. In this context, hardware security will play a vital role since hardware is at the heart of all computing systems, and insecure hardware will put critical systems and our society at risk.

However, in recent years, we are observing the discovery of a growing number of hardware design and implementation vulnerabilities that could be exploited by unprivileged software, leading to potential exposure of sensitive data or compromise of whole computing systems. This new attack paradigm casts a long shadow on decades of research on system security and disrupts the traditional threat models that have mainly focused on software-only vulnerabilities and often assume that the underlying hardware is behaving correctly and is trustworthy. Unfortunately, existing solutions are often ad-hoc, limited, inefficient, and address only specific problems.

The main goal of this workshop is to bring together international researchers and experts from academia, industry, and government to exchange knowledge and explore new ideas and research directions for tackling the challenges related but not limited to security-by-design for hardware, scalable assurance methodologies for hardware security and resilience, and security-aware electronic design automation that pave the way for establishing sustainable security for computing platforms.

8:00 – 17:30, Room: Salon II
Workshop on Zero Trust Hardware Architectures (ZTHA)

In recent times, there has been a major push and urgency to adopt the zero-trust model for cybersecurity. The zero-trust model is based on the principle of “never trust, always verify” and is aimed at eliminating all implicit trust in a system. While adopting a zero-trust model for network security generally involves authenticating, authorizing and continuously validating the credentials of users in a network, these measures alone are not enough to create a true zero trust-based architecture. The underlying hardware needs to be trusted and secured as well. Thus, novel approaches for building zero trust architectures, from systems all the way down to silicon, is one of the big challenges for next generation hardware system design.

Traditionally, research on establishing trust and security in hardware has primarily focused on the host and its associated memory subsystems. These include principles of trusted execution environments, silicon roots of trust, Trusted Platform Modules, encryption at rest etc. However, in modern system architectures such as edge/cloud computing, composable systems and chiplet based integrated circuits, the realm of trust needs to be extended beyond the host system and incorporate many hardware devices and IPs. In view of threats such as compromised supply chain integrity, counterfeit chips, hardware trojan implants, malicious firmware, malware, etc., it is important to establish trust in hardware components and to communicate trust between different components of a system. The different kinds of communication could range from that between different IPs inside a SoC, between a host and its attached peripherals, as well as between chiplets inside a multi-chip module. Trust also needs to be established and revoked in a dynamic manner, with the ability to handle large number of subcomponents in the design. Thus, a new set of protocols that can work to establish trust and security in these new types of system architectures has become necessary. While some of these protocols are being developed as industry and government standards, large-scale effort is required to bring them to adoption. It is equally important to develop open source and verifiable hardware designs that can bring security without compromising system parameters such as performance or functionality. The focus of this workshop will be on all aspects of security and trust required to create zero-trust hardware architectures for heterogeneous computing systems.

2023 ICCAD Executive Committee

General Chair
Evangeline Young
The Chinese University of Hong Kong

Past Chair
Tulika Mitra
National University of Singapore

Program Chair
Jinjun Xiong
University of Buffalo

Vice Program Chair
Robert Wille
Johannes Kepler University Linz

Tutorial & Special Session Chair
Deming Chen
University of Illinois

Workshop Chair
Ismail S. K. Bustany
AMD

Asian Representative
Wei Zhang
Hong Kong University of Science and Technology

European Representative
Ulf Schlichtmann
Technical University of Munich

ACM SIGDA Representative
Yiyu Shi
University of Notre Dame

CEDA Representative
Tsung-Yi Ho
The Chinese University of Hong Kong

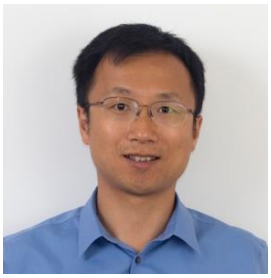

Industry Representative
Mark Ren
NVIDIA

2023 Patrons & Sponsors

Special Thanks To Our

Platinum Patrons

cadence®

SYNOPSYS®

ACADEMIC & RESEARCH ALLIANCES

Gold Patrons

FUTUREWEI Technologies

Silver Patrons

AMD

EMPYREAN

Conference Sponsors

IEEE/ACM

2024 INTERNATIONAL CONFERENCE ON COMPUTER-AIDED DESIGN

ICCAD

43rd Edition

CALL FOR PAPERS

DEADLINES

April 28, 2024
Abstract Technical Submission

May 5, 2024
Paper Technical Submission

May 5, 2024
Proposal for Workshops, Tutorials, Special Sessions, Panels

CONTACT

General Chair

Jinjun Xiong
jinjun@buffalo.edu

Program Chair

Robert Wille
robert.wille@tum.de

Jointly sponsored by IEEE and ACM, ICCAD is the premier forum to explore new challenges, present leading-edge innovative solutions, and identify emerging technologies in the electronic design automation research areas. ICCAD covers the full range of CAD topics – from device and circuit level up through system level, as well as post-CMOS design. ICCAD has a long-standing tradition of producing cutting-edge, innovative technical program for attendees.

COVERED TOPICS

Original technical submissions on, but not limited to, the following topics are invited:

- System-Level CAD
- Synthesis, Verification, Physical Design, Analysis, Simulation, and Modeling
- CAD For Emerging Technologies, Paradigms

Paper submissions must be made through the online submission system at the ICCAD website. More details on the covered topics and the submission process will be provided soon.

CALL FOR PROPOSALS

In addition to presentations of technical submissions, ICCAD will also cover

- Workshops
- Tutorials
- Special Sessions
- Panels
- Keynotes

Corresponding proposals for the above will be greatly appreciated. A call for proposals will be distributed soon.

More information is available at www.iccad.com.

Author Index