MLCAD

2nd ACM/IEEE Workshop on Machine Learning for CAD (MLCAD)

16-20 November, 2020

Sponsors & Supporters:

Keynote "MLCAD Today and Tomorrow: Learning, Optimization and Scaling"

7am PST, 10am EST, 4pm CET, 8:30pm India, 11pm Asia, midnight Japan

Andrew B. Kahng
CSE and ECE Departments,
UC San Diego, USA

Abstract

The scaling imperative challenges us to always do better and faster, with less resources. The semiconductor industry has looked to machine learning (ML) as a design-based lever for scaling that will reduce design costs and design schedules while improving quality of results. As a result, in recent years Machine Learning for CAD (MLCAD) has dominated the conversation at leading conferences. Numerous ML-based enhancements and their benefits have been highlighted by EDA vendors and their customers. With this as backdrop, this talk will offer some thoughts on future directions for MLCAD.

First, MLCAD lies on a road to "self-driving IC design tools and flows" that make design ideation and design space exploration both accurate and accessible. Eventually, MLCAD (ML for CAD) will lead us to MLDA (ML-enabled Design Automation). But for this to happen, researchers and practitioners will need to deliver (i) human-quality prediction, evaluation and decision-making with no humans; (ii) design tools and flows that never require iteration and never fail; (iii) modeling of design processes that continually improves; and more.

Second, the trajectory of MLCAD will need to keep three concepts in foreground: Learning, Optimization and Scaling. "Learning" seems obvious from "ML", but it brings open questions about data and models, ranging from statistics to standards, sharing and openness. "Optimization" is the essence of CAD, and brings open questions about both synergies and boundaries with learning. "Scaling" is how practical realization of Learning and Optimization will satisfy the needs of design within an ever-tighter box of compute, schedule and other resources. Finally, there is a meta-question of how the MLCAD community will itself learn, optimize, and scale.

Speaker Bio:

Andrew B. Kahng is Distinguished Professor of CSE and ECE and holder of the endowed chair in high-performance computing at UC San Diego. He was visiting scientist at Cadence (1995-97) and founder/CTO at Blaze DFM (2004–06). He is coauthor of 3 books and over 500 journal and conference papers, holds 35 issued U.S. patents, and is a fellow of ACM and IEEE. He served as general chair of DAC, ISPD and other conferences, and from 2000-2016 as international chair/co-chair of the ITRS Design and System Drivers working groups. He currently serves as PI of "OpenROAD" https://theopenroadproject.org/, a $15M U.S. DARPA project targeting open-source, autonomous ("no human in the loop") tools for IC implementation.

Session 1: DNN for CAD

8am PST, 11am EST, 5pm CET, 9:30pm India, midnight Asia, 1am (Tue) Japan
Session Chair: Hussam Amrouch – *University of Stuttgart*

1.1 An Adaptive Analytic FPGA Placement Framework based on Deep-Learning

Abeer Alhyari, Ahmed Shamli, Timothy Martin, Shawki Areibi, Gary Grewal – *University of Guelph*

1.2 Design Rule Checking with a CNN Based Feature Extractor

Luis Francisco, Tanmay Lagare, Arpit Jain, Somal Chaudhary, Madhura Kulkarni, Divya Sardana, W. Rhett Davis, Paul Franzon – *North Carolina State University*

1.3 Using DNNs and Smart Sampling for Coverage Closure Acceleration

Raviv Gal, Avi Ziv – *IBM Reserach Lab, Haifa, Israel*
Eldad Haber – *UBC*

1.4 R^2AD: Randomization and Reconstructor based Adversarial Defense on Deep Neural Network

Marzieh Ashrafiamiri – *University of California Irvine*
Sai Manoj Pudukotai Dinakarrao – *George Mason University*
Amir Hosein Afandizadeh Zargari, Minjun Seo, Fadi Kurdahi – *University of California Irvine*
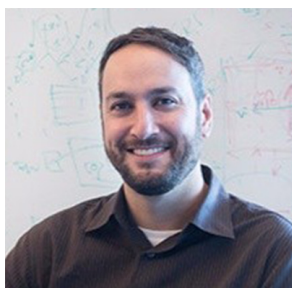Houman Homayoun – *University of California Davis*

1.5 DAVE: Deriving Automatically Verilog from English

Hammond Pearce, Benjamin Tan, Ramesh Karri – *New York University*

Plenary "Accelerating Chip Design with Machine Learning"

9:15am PST, 12:15pm EST, 6:15pm CET, 10:45pm India, 01:15am (Tue) Asia, 02:15am (Tue) Japan

Brucek Khailany
NVIDIA
Austin, TX, USA

Abstract

As Moore's law has provided an exponential increase in chip transistor density, the unique features we can now include in large chips are no longer predominantly limited by area constraints. Instead, new capabilities are increasingly limited by the engineering effort associated with digital design, verification, and implementation. As applications demand more performance and energy efficiency from specialization in the post-Moore's-law era, we expect required complexity and design effort to increase.

Historically, these challenges have been met through levels of abstraction and automation. Over the last few decades, Electronic Design Automation (EDA) algorithms and methodologies were developed for all aspects of chip design - design verification and simulation, logic synthesis, place-and-route, and timing and physical signoff analysis. With each increase in automation, total work per chip has increased, but more work has also been offloaded from manual effort to software. Just as machine learning (ML) has transformed software in many domains, we expect advancements in ML will also transform EDA software and as a result, chip design workflows.

In this talk, we highlight work from our research group and the community applying ML to various chip design prediction tasks. We show how deep convolutional neural networks and graph-based neural networks can be used in the areas of automatic design space exploration, power analysis, VLSI physical design, and analog design. We also present a future vision of an AI-assisted chip design workflow to automate optimization tasks. In this future vision, GPU acceleration, neuralnetwork predictors, and deep reinforcement learning techniques combine to automate VLSI design and optimization.

Speaker Bio:

Brucek Khailany joined NVIDIA in 2009 and is the Director of the ASIC and VLSI Research group. He leads research into innovative design methodologies for IC development, ML and GPU assisted EDA, and energy efficient ML accelerators. Over 10 years at NVIDIA, he has contributed to many projects in research and product groups spanning computer architecture and VLSI design. Previously, Dr. Khailany was a Co-Founder and Principal Architect at Stream Processors, Inc where he led R&D related to parallel processor architectures. At Stanford University, he led the VLSI implementation of the Imagine processor, which introduced the concepts of stream processing and partitioned register organizations. He received his PhD in Electrical Engineering from Stanford University and BSE degrees in Electrical and Computer Engineering from the University of Michigan. He is a Senior Member of the IEEE.

Tuesday, November 17, 2020

Keynote "SoC Design Automation with ML – It's Time for Research"

7am PST, 10am EST, 4pm CET, 8:30pm India, 11pm Asia, midnight Japan

Wolfgang Ecker
Infineon Technologies
Munich, Germany

Abstract

The AI-hype started a few years ago, with advances in object recognition. Soon the EDA research community made proposals on applying AI in EDA and all major players announced new AI-based tools at DAC 2018. Unfortunately, few new AI-based EDA-tools made it to productive use today. This talk analyses general challenges of AI in EDA, outlines promising use cases, and motivates more AI research in EDA: More HI (=Human Intelligence) is needed to make AI successful in EDA.

Speaker Bio:

Wolfgang Ecker is Distinguished Engineer at Infineon and Adjunct Professor at Technical University of Munich. He is (co-)author of over 200 papers on modeling and design automation, received 5 best paper awards, was granted with the German EDA achievement award. He is member of Acatech, the German Academy of Science and Engineering. Wolfgang Ecker leads the Infineon Deep Learning internal think tank. In addition, he is member of the AI commission of inquiry of the German Government.

Session 2: Design Methodology and Optimization

8am PST, 11am EST, 5pm CET, 9:30pm India, midnight Asia, 1am (Wed) Japan
Session Chair: Helen Hai Li – *Duke University*

2.1 Cost Optimization at Early Stages of Design Using Deep Reinforcement Learning (best paper award nominee)

Lorenzo Servadei – *Infineon Technologies AG*
Jiapeng Zheng – *Technical University of Munich*
José Antonio Arjona – *Johannes Kepler University Linz*
Michael Werner, Volkan Esen – *Infineon Technologies AG*
Sepp Hochreiter – *Johannes Kepler University Linz*
Wolfgang Ecker – *Infineon Technologies AG*
Robert Wille – *Johannes Kepler University Linz*

2.2 F-LEMMA: Fast Learning-based Energy Management for Multi/Many-core Processors (best paper award nominee)

An Zou, Karthik Garimella – *Washington University in St. Louis*
Benjamin Lee – *University of Pennsylvania*
Christopher Gill, Xuan Zhang – *Washington University in St. Louis*

2.3 CALT: A Classification with Adaptive Labeling Thresholds for Analog Design

Zhengfeng Wu, Ioannis Savidis – *Drexel University*

2.4 Decision Making in Synthesis cross Technologies using LSTMs and Transfer Learning

Cunxi Yu – *University of Utah*
Wang Zhou – *IBM T.J Watson Research Center*

2.5 Application of Machine Learning to Quantum VLSI Placement

George Li, Isaac Turtletaub, Mohannad Ibrahim, Paul Franzon – *North Carolina State University*

Plenary "From Tuning to Learning: Why the FPGA physical design flow offers a compelling case for ML?"

9:15am PST, 12:15pm EST, 6:15pm CET, 10:45pm India, 01:15am (Wed) Asia, 02:15am (Wed) Japan
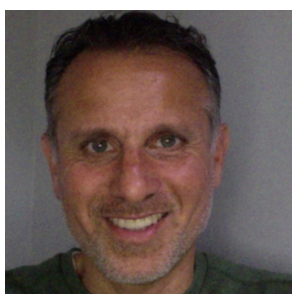
Ismail Bustany
Xilinx
USA

Abstract

ML is particularly suited for the FPGA Physical design (PD) flow since each device family generation innately provides a rich platform for device/design feature data harvesting: (1) A vast amount of device architecture-specific interconnect/layout fabric data and (2) significant amount of large design suite data from and from broad set of application domains. These bode well for developing robust predictive ML models. Furthermore, the long lifespan of these device families affords a favorable ROI. In this talk, we will highlight some data harvesting and ML solutions we have developed in Xilinx' Vivado PD flow and share some initial results. These include a strategy recommendation framework for design closure, design classification for computational resource allocation, device characteristics modeling, and routing congestion estimation. Furthermore, we will outline potential MLCAD opportunities in trend identification, algorithm parameter optimization, and reinforcement learning paradigms where we foresee potential collaborations with the academic community.

Speaker Bio:

Ismail Bustany is a Distinguished Engineer at Xilinx, where he works on physical design, MLCAD, and sparse computation hardware acceleration. He has served on the technical programming committees for the ISPD, the ISQED, and DAC. He was the 2019 ISPD general chair. He currently serves on the organizing committees for ICCAD and SLIP. He organized the 2014 and 2015 ISPD detailed routing-driven placement contests and co-organized the 2017 ICCAD detailed placement contest. His research interests include physical design, computationally efficient optimization algorithms, MLCAD, sparse matrix computations, and hardware acceleration. He earned his M.S. and Ph.D. in electrical engineering from UC Berkeley.

Keynote "Data-driven CAD or Algorithm-Driven CAD: Competitors or Collaborators?"

3pm (Tue) PST, 6pm (Tue) EST, midnight (Wed) CET, 4:30am (Wed) India, 7am (Wed) Asia, 8am (Wed) Japan

Rajeev Jain
Qualcomm
USA

Abstract

As with any new technology there has been a lot of research in applying machine learning to anything and everything. In this talk we will focus on the gaps between algorithm-based CAD and the needs of the designer, why these gaps exist, and where specifically data-driven techniques can help to close these gaps and compliment algorithm-based CAD. We will survey both the opportunities and technical challenges with the data-driven approach, specifically by formulating the ML problems to close the gaps in design automation. As technology scaling reaches its limits and the gains in performance, power and area (PPA) from technology scaling get limited, the pressure on designers and design tools to close the PPA gap through design will increase significantly, creating opportunities for data-driven CAD to disrupt CAD flows and SoC design methodologies.

Speaker Bio:

Rajeev Jain received his B.Tech in EE from IIT Delhi in 1978 and Ph.D in EE from K.U. Leuven in 1985. He started his career in IC design in 1980 at Siemens (later Infineon) where he worked on the first DSP subscriber line codec IC (in 10 micron NMOS) for digital exchanges. More recently he leads an ML team at Qualcomm for designing SoCs in deep-submicron CMOS. He has done research in IC design as a Group Leader at IMEC, Post-Doc at UC Berkeley, Professor at UCLA (where he is currently Professor Emeritus). He has been a technology consultant for Communications IC design at Conexant and Atheros. Since 2011 he has been at Qualcomm where he has led the development of various machine learning technologies for mobile SoCs. In 1999 he was elected IEEE Fellow for his contributions to Computer-Aided Design Tools for Signal Processing Circuits.

** Attention: the third day of MLCAD 2020 runs on a different time schedule as the other days, to better accommodate Asia participants **

Session 3: ML for Reliability Improvement

4pm (Tue) PST, 7pm (Tue) EST, 1am (Wed) CET, 5:30am (Wed) India, 8am (Wed) Asia, 9am (Wed) Japan
Session Chair: Bei Yu – *The Chinese University of Hong Kong*

3.1 Data-Driven Fast Electrostatics and TDDB Aging Analysis

Shaoyi Peng, Wentian Jin – *University of California, Riverside*
Liang Chen – *Shanghai Jiaotong University*
Sheldon Tan – *University of California, Riverside*

3.2 HAT-DRL: Hotspot-Aware Task Mapping for Lifetime Improvement of Multi-core System using Deep Reinforcement Learning

Jinwei Zhang, Sheriff Sadiqbatcha, Yuanqi Gao, Michael O'Dea, Nanpeng Yu, Sheldon Tan – *University of California Riverside*

3.3 Can Wear-Aware Memory Allocation be Intelligent?

Christian Hakert, Kuan-Hsun Chen, Jian-Jia Chen – *TU Dortmund*

3.4 An Enhanced Machine Learning Model for Adaptive Monte Carlo Yield Analysis

Richard Kimmel, Tong Li, David Winston – *IBM*

3.5 Towards NN-based Online Estimation of the Full-Chip Temperature and the Rate of Temperature Change

Martin Rapp – *Karlsruhe Institute of Technology*
Omar Elfatairy – *German University in Cairo*
Marilyn Wolf – *University of Nebraska – Lincoln*
Jörg Henkel – *Karlsruhe Institute of Technology*
Hussam Amrouch – *University of Stuttgart*

** Attention: the third day of MLCAD 2020 runs on a different time schedule as the other days, to better accommodate Asia participants **

Plenary "Design Challenges on post Moore's Law Era"

5:15pm (Tue) PST, 8:15pm (Tue) EST, 2:15am (Wed) CET, 6:45am (Wed) India, 9:15am (Wed) Asia, 10:15am (Wed) Japan

Matthew Leung
Director
Huawei Hong Kong Research Center
Hong Kong

Abstract

IC companies nowadays are busy struggling between increasing challenges in deep submicron process and, in the same time, more stringent time-to-market cycle to entertain the more demanding consumers. As a result, engineers have to turn to more holistic optimizations across software, architecture, micro-architecture, circuit design and physical implementations. The increase in complexity also demands for high level of automation and help from design tools. We shall look into some of the solutions that we are exploring to cope with the situation.

Speaker Bio:

Mr. Matthew Leung serves as the director of Huawei Hong Kong Research Center, with a current focus in the development of hardware, software and algorithm for artificial intelligence. Prior to that, he served as the director and a founding member of HiSilicon (a subsidiary of Huawei) Hong Kong R&D Center. His expertise and experience lies in the fields of VLSI design for advanced communication chipsets, microprocessors and artificial intelligence. Mr. Leung received his BSc and MSc degrees of Electrical Engineering in University of Michigan and Stanford University respectively.

** Attention: the third day of MLCAD 2020 runs on a different time schedule as the other days, to better accommodate Asia participants **

Thursday, November 19, 2020

Keynote "Machine Learning in EDA: Opportunities and Challenges"

7am PST, 10am EST, 4pm CET, 8:30pm India, 11pm Asia, midnight Japan

Elias Fallon
Engineering Group Director
Cadence Design Systems
USA

Abstract

Electronic Design Automation software has delivered semiconductor design productivity improvements for decades. The next leap in productivity will come from the addition of machine learning techniques to the toolbox of computational software capabilities employed by EDA developers. Recent research and development into machine learning for EDA point to clear patterns for how it impacts EDA tools, flows, and design challenges. This research has also illustrated some of the challenges that will come with production deployment of machine learning techniques into EDA tools and flows. This talk will detail patterns observed in ML for EDA development, as well as discussing challenges with productization of ML for EDA developments and the opportunities that it presents for researchers.

Speaker Bio:

Elias Fallon is currently Engineering Group Director at Cadence Design Systems, a leading Electronic Design Automation company. He has been involved in EDA for more than 20 years from the founding of Neolinear, Inc, which was acquired by Cadence in 2004. Elias was co-Primary Investigator on the MAGESTIC project, funded by DARPA to investigate the application of Machine Learning to EDA for Package/PCB and Analog IC. Elias also leads an innovation incubation team within the Custom IC R&D group as well as other traditional EDA product teams. Beyond his work developing electronic design automation tools, he has led software quality improvement initiatives within Cadence, partnering with the Carnegie Mellon Software Engineering Institute. Elias graduated from Carnegie Mellon University with an M.S. and B.S. in Electrical and Computer Engineering. Elias, his wife and two children live north of Pittsburgh, PA.

https://www.linkedin.com/in/elias-fallon/

Session 4: Intelligent Modeling

8am PST, 11am EST, 5pm CET, 9:30pm India, midnight Asia, 1am (Fri) Japan
Session Chair: Helen Hai Li – *Duke University*

4.1 Track-Assignment Detailed Routing Using Attention-based Policy Model With Supervision

Haiguang Liao, Qingyi Dong – *Carnegie Mellon University*
Weiyi Qi, Elias Fallon – *Cadence Design Systems*
Levent Kara – *Carnegie Mellon University*

4.2 Compact Models for Initial MOSFET Sizing based on Higher-order Artificial Neural Networks

Husni Habal, Dobroslav Tsonev – *Infineon Technologies AG*
Matthias Schweikardt – *Hochschule Reutlingen*

4.3 An Efficient and Flexible Learning Framework for Dynamic Power and Thermal Co-Management

Yuan Cao, Tianhao Shen, Li Zhang, Xunzhao Yin, Cheng Zhuo – *Zhejiang University*

4.4 Partial Sharing Neural Networks for Multi-Target Regression on Power and Performance of Embedded Memories

Felix Last – *Technical University of Munich (TUM) & Intel Deutschland GmbH*
Ulf Schlichtmann – *Technical University of Munich (TUM)*

4.5 Explaining and Interpreting Machine Learning CAD Decisions: An IC Testing Case Study

Prashanth Krishnamurthy, Animesh Basak Chowdhury, Benjamin Tan, Farshad Khorrami, Ramesh Karri – *New York University*

Plenary "Machine-Learning Enabled Next-Generation Physical Design – An EDA Perspective"

9:15am PST, 12:15pm EST, 6:15pm CET, 10:45pm India, 1:15am (Fri) Asia, 2:15am (Fri) Japan

Vishal Khandelwal
Synopsys
USA

Abstract

Physical design is an ensemble of NP-complete problems that P&R tools attempt to solve in (pseudo) linear time. Advanced process nodes and complex signoff requirements bring in new physical and timing constraints into the implementation flow, making it harder for physical design algorithms to deliver industry-leading power, performance, area (PPA), without giving up design turn-around-time. The relentless pursuit for low-power high-performance designs is putting constant pressure to limit any over-design, creating an acute need to have better models/predictions and advanced analytics to drive implementation flows. Given the advancements in supervised and reinforcement learning, combined with the availability of large-scale compute, Machine Learning (ML) has the potential to become a disruptive paradigm change for EDA tools. In this talk, I would like to share some of the challenges and opportunities for innovation in next-generation physical design using ML.

Speaker Bio:

Vishal leads the physical optimization team for the Digital Implementation products at Synopsys. He has 15 years of R&D experience in building state-of-the-art optimization engines and P&R flows targeting advanced-node low-power high-performance designs. More recently, he has been looking at bringing machine-learning paradigms into digital implementation tools to improve power, performance, area and productivity. Vishal has a B.Tech. from IIT-Kanpur and a Ph.D. from University of Maryland, College Park. He has co-authored several patents and over 20 IEEE/ACM publications.

Panel "ML for CAD - Where is the Treasure Hiding?"

7am PST, 10am EST, 4pm CET, 8:30pm India, 11pm Asia, midnight Japan

The panelists:

- David Z. Pan – *The University of Texas at Austin*

- Mark Ren – *Nvidia Research*

- Manish Pandey – *Synopsys*

- Marilyn Wolf – *University of Nebraska - Lincoln*

- Avi Ziv – *IBM Research - Haifa, Israel*

Moderator: Raviv Gal – *IBM Research - Haifa, Israel*

Session 5: ML for Systems

8:30am PST, 11:30am EST, 5:30pm CET, 10pm India, 12:30am Asia, 1:30am (Sat) Japan
Session Chair: Hussam Amrouch – *University of Stuttgart*

5.1 Transfer Learning for Design-Space Exploration with High-Level Synthesis (best paper award nominee)

Jihye Kwon, Luca Carloni – *Columbia University*

5.2 Using Machine Learning Clustering To Find Large Coverage Holes

Raviv Gal, Giora Simchoni, Avi Ziv – *IBM Research - Haifa*

5.3 Exploring Logic Optimizations with Reinforcement Learning and Graph Convolutional Network

Keren Zhu, Mingjie Liu, Hao Chen – *The University of Texas at Austin*
Zheng Zhao – *Synopsys*
David Z. Pan – *The University of Texas at Austin*

5.4 AdaPool: Multi-Armed Bandits for Adaptive Virology Screening on Cyber-Physical Digital-Microfluidic Biochips

Mohamed Ibrahim – *Intel Corporation*

5.5 Automatic compiler optimization on embedded software through k-means clustering

Michael Werner, Lorenzo Servadei – *Infineon Technologies AG*
Robert Wille – *Johannes Kepler University Linz*
Wolfgang Ecker – *Infineon Technologies AG*

5.6 Footprint Classification of Electric Components on Printed Circuit Boards

Yun-Jie Ni – *National Tsing Hua University*
Yan-Jhih Wang – *Footprintku. Inc.*
Tsung-Yi Ho – *National Tsing Hua University*