sig da Special Interest Group Design Automation
The resource for EDA Professionals

1 August 2020, Vol. 50, No. 8
Online archive: http://www.sigda.org/publications/newsletter

Comments from the Editors

Dear ACM/SIGDA members,

We are excited to present to you August e-newsletter. We encourage you to invite your students and colleagues to be a part of the SIGDA newsletter. The newsletter covers a wide range of information from the upcoming conferences and hot research topics to technical news and activities from our community. Get involved and contact us if you want to contribute an article or announcement.

The newsletter is evolving, let us know what you think.

Happy reading!

Debjit Sinha, Keni Qiu, Editors-in-Chief, SIGDA E-News

To renew your ACM SIGDA membership, please visit http://www.acm.org/renew or call between the hours of 8:30am to 4:30pm EST at +1-212-626-0500 (Global), or

1-800-342-6626 (US and Canada). For any questions, contact acmhelp@acm.org

''SIGDA E-News Editorial Board:''

'''Debjit Sinha''', E-Newsletter co Editor-in-Chief

'''Keni Qiu''', E-Newsletter co Editor-in-Chief

'''Xiang Chen''', E-Newsletter Associate Editor for SIGDA News column

'''Yanzhi Wang''', E-Newsletter Associate Editor for SIGDA Local chapter news column

'''Pingqiang Zhou''', E-Newsletter Associate Editor for SIGDA Awards column

'''Wanli Chang''', E-Newsletter Associate Editor for SIGDA What is column

'''Xun Jiao''', E-Newsletter Associate Editor for SIGDA What is column

'''Jayita Das''', E-Newsletter Associate Editor for SIGDA Funding opportunities column

'''Qinru Qiu''', E-Newsletter Associate Editor for SIGDA Live column

'''Yiyu Shi''', E-Newsletter Associate Editor for SIGDA Live column

'''Rajsaktish Sankaranarayanan''', E-Newsletter Associate Editor for SIGDA Researcher spotlight column

'''Xin Zhao''', E-Newsletter Associate Editor for SIGDA Paper submission deadline column

'''Ying Wang''', E-Newsletter Associate Editor for SIGDA Technical activities column

Back to Contents

SIGDA News

(1) "The U.S.-China Tech Reckoning Has Arrived"
[https://www.eetimes.com/the-u-s-china-tech-reckoning-has-arrived/]

When all is said and done, it may turn out that horseshoe bats dwelling in the caves of China’ s Yunnan province helped reboot the U.S. semiconductor industry.

(2) "Covid Economy: How Damaged Are We?"
[https://www.eetimes.com/the-u-s-china-tech-reckoning-has-arrived/]

Six months since the pandemic touched all corners of the world is a good time to assess the depth and width of the damage done. EE Times sat down with two Yole Développement experts, analyst Eric Mounier and Guillaume Assogba, an economist who oversees Yole’ s macroeconomics research.

(3) "Tiny Alchip Uses AI Strength to Win 7nm Capacity from TSMC"
[https://www.eetimes.com/tiny-alchip-uses-ai-strength-to-win-7nm-capacity-from-ts...]

Alchip Technologies is using its AI and high-performance computing strength as a chip design house to win capacity from Taiwan Semiconductor Manufacturing Co. (TSMC) at the most advanced 7nm and 5nm nodes.

(4) "SMIC: Advanced Process Technologies and Gov’ t Funding"

[https://www.eetimes.com/smic-advanced-process-technologies-and-govt-funding-part...]

Semiconductor Manufacturing International Corp. (SMIC) has begun producing 14nm chips and has joined the relatively small club of semiconductor makers that can build finFETs. The company is on the verge of a stock offering that could reap in excess of $7 billion to keep investing in its business.

(5) "DDR5 Spec Published: High Capacities Meet Extreme Speeds"
[https://www.eetimes.com/ddr5-spec-published-high-capacities-meet-extreme-speeds/]

JEDEC, a semiconductor engineering organization that sets standards for dynamic random-access memory (DRAM), published the finalized JESD79-5 DDR5 specification last week. The new type of memory doubles per-pin data transfer rate over its predecessor increases the capacity of memory devices by four times, lowers operating voltage, and introduces several methods to improve DRAM reliability at thin nodes — as well as reduces power consumption.

(6) "Photonic Tensor Cores Boost Machine Learning Capacity for Optical Feeds and 5G"
[https://www.eetimes.com/photonic-tensor-cores-boost-machine-learning-capacity-fo...]

A new approach to performing neural network computations for machine learning using photonic tensor cores instead of graphics processing units (GPUs) suggests 2-3 orders higher performance can be achieved for processing optical data feeds. It also indicates that photonic processors have the potential to augment electronic systems and may perform exceptionally well in network-edge devices in 5G networks.

(7) "Edging Towards a 5G Future: Mobile Edge Computing"
[https://www.eetimes.com/edging-towards-a-5g-future-mobile-edge-computing/]

Two organizations that have been working on the development of Multi-Access Edge Computing (MEC) in the mobile space are expected to announce early next month specifications for the technology.

(8) "Intel’ s Stumble Signals AMD’ s Gain"
[https://www.eetimes.com/intels-stumble-signals-amds-gain/]

AMD’ s gain in market share against Intel is more likely, as we reported earlier, given Intel’ s yet another delay on its technology roadmap announced last week.

(9) "Insight of GaN and SiC Market"
[https://www.eetimes.com/photonic-tensor-cores-boost-machine-learning-capacity-fo...]

Power electronics has taken an interesting road with the adoption of GaN and SiC. Yole Développement (Yole) estimated a general view of these wide bandgap materials. While silicon is still dominating the market, GaN and SiC devices are already more efficient solutions in some applications.

(10) "NXP Customizes AI Compiler for MCU Products"
[https://www.eetimes.com/nxp-customizes-ai-compiler-for-mcu-products/]

In a sign that machine learning techniques are fast gaining adoption on embedded platforms, NXP announced that it has created a customized implementation of Glow for microcontrollers (MCU), including some of its i.MX RT family. Glow is a neural network compiler that optimizes neural networks for specific target hardware.

Back to Contents

"What is" Column

What is Timing Speculative DNN Acceleration?

Contributing author: Jeff (Jun) Zhang,
Postdoctoral Fellow,
John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, MA

Timing speculation is a promising approach to improve the processor energy efficiency. It is based on the empirical observation that critical path delays are rarely manifested during program execution. There is a considerable body of work on timing speculation based low-power digital logic and system [1, 2, 3, 4, 5]. The idea behind timing speculation is to optimistically execute a chip at a lower voltage without scaling the frequency, thus reducing energy at the expense of higher propagation delay and consequently occasional timing errors [1].

The timing error rate of a given circuit depends on the workload (since combinational logic delay is highly data-dependent) and on factors such as process variability and temperature [6]. Traditionally, timing errors can either be allowed to propagate if the application is itself error-tolerant [7, 8], or detected using so-called Razor flip-flops and corrected via safe re-execution [1].
Reducing the supply voltage results in a quadratic saving in Joules per Operation. Under timing speculation, the circuit still operates at its nominal frequency. If the extra operations (i.e., re-execution due to timing errors) are small, we get large net energy savings.

For power-hungry DNN acceleration, timing speculation opens a new dimension for energy efficiency. However, directly applying timing speculation on DNN accelerators is challenging. Prior works show that timing error propagation (TEP) causes DNN classification accuracy to drop sharply for timing error rates as low as 0.1% [9, 10]. Detailed post-synthesis timing simulation shows that timing errors frequently flip high-order bits of multiply-and-accumulate (MAC) outputs (i.e., the partial sums), resulting in large computational errors even at low timing error rates, limiting opportunities for voltage underscaling [10]. Timing error detection and recovery (TED) technique offered by Razor flip-flops can guarantee correct computation against timing errors. For example, Whatmough et al. show that TED works on a small DNN accelerator with only 8 loosely coupled MACs [11]. However, for systolic array (SA) based DNN accelerators where thousands of MAC units are highly synchronized, one timing-error related stall/re-execution in any MAC unit will disturb the precise schedule of the whole array. Stalling and re-executing the entire SA even if a single MAC is faulty results in a high re-execution rate and therefore significant performance degradation [10].

Recent work has proposed a new timing speculation technique for SA like DNN accelerator, referred to as TE-Drop [10, 12], that significantly outperforms the two aforementioned techniques and enables up to 57% energy savings with a negligible drop (< 1%) in image classification accuracy [12]. TE-Drop still instruments the SA with Razor flip-flops to detect timing errors, but recovers from timing errors by stealing a clock cycle from its successor MAC unit and bypassing that successor MAC' s partial sum update. Equivalently, the effect of TE-Drop on DNN execution can be viewed as randomly dropping a fraction of connections between input and output neurons, where the randomness arises from timing errors. Interestingly, randomly dropping connections during training is a commonly used technique, referred to as dropconnect [13], that prevents overfitting and helps reduce generalization error. Moreover, since timing errors vary significantly across DNN layers, dynamically adjust the voltage level for each layer of DNN can yield even greater energy savings at iso-accuracy [10, 12].

Inspired by TE-Drop, several works have achieved even greater energy efficiency and timing error resilience for systolic arrays. GreenTPU identifies the patterns in the erroneous activation sequences and thus prevents those patterns in later executions [14]. Choi et al. performs the sensitivity analysis on the weight-to-MAC mapping strategies and design heterogenous MACs for aggressive voltage underscaling [15]. Wu et al. propose weight distribution and error-aware quantization of DNNs to reduce timing errors without the need of hardware protection [16]. Li et al. use multiple voltage islands to have finer-grained control on the voltage for DNN execution [17].

To study the timing error behavior at the early design stage, traditional post-synthesis gate-level timing simulations are too time-consuming. Therefore, a line of work is trying to build fast and accurate timing simulation framework. For example, recent work by Jiao et al. [18] uses machine learning (ML) techniques to estimate timing errors for microprocessor function units; Zhang et al. [19] propose to train DNN based timing error prediction for deep learning accelerators. Fast simulations offered by these tools facilitate future architectural explorations for timing speculative DNN acceleration.

In summary, timing speculation is a promising approach for DNNs that pushes hardware to its limit. Timing speculative DNN accelerators are more energy-efficient, and timing error-resilient. With more intelligent IoT devices deployed in real-world under dynamic environmental conditions and uncertain operating variations, more interests and efforts will be potentially attracted by this approach across circuit-, architectural-, and application-levels [20].

References

[1] D. Ernst, S. Das, S. Lee, D. Blaauw, T. Austin, T. Mudge, N. S. Kim, and K. Flautner. Razor: circuit-level correction of timing errors for low-power operation. IEEE Micro, 24(6):10–20, 2004.
[2] M. Fojtik, D. Fick, Y. Kim, N. Pinckney, D. Harris, D. Blaauw, and D. Sylvester. Bubble Razor: An architecture-independent approach to timing-error detection and correction. International Solid-State Circuits Conference (ISSCC), 2012.
[3] B. Greskamp, L. Wan, U. R. Karpuzcu, J. J. Cook, J. Torrellas, D. Chen, and C. Zilles. Blueshift: Designing processors for timing speculation from the ground up. IEEE 15th International Symposium on High Performance Computer Architecture (HPCA), 2009.
[4] A. Yasin, J. Zhang, H. Chen, S. Garg, S. Roy, and K. Chakraborty. Synergistic timing speculation for multi-threaded programs. The 53rd Annual Design Automation Conference (DAC), ACM, 2016.
[5] J. Zhang and S. Garg. BandiTS: dynamic timing speculation using multi-armed bandit based optimization. IEEE Design, Automation & Test in Europe (DATE), 2017.
[6] M. De Kruijf, S. Nomura, and K. Sankaralingam. A unified model for timing speculation: Evaluating the impact of technology scaling, CMOS design style, and fault recovery mechanism. International Conference on Dependable Systems and Networks (DSN), IEEE, 2010.
[7] A. Gebregiorgis, S. Kiamehr, and M.B. Tahoori, Error propagation aware timing relaxation for approxi- mate near threshold computing. The 54th Annual Design Automation Conference (DAC), ACM, 2017.
[8] F. Nakhaee, M. Kamal, A. Afzali-Kusha, M. Pedram, S. M. Fakhraie, and H. Dorosti. Lifetime improvement by exploiting aggressive voltage scaling during runtime of error-resilient applications. Integration, 61:29–38, 2018.
[9] X. Jiao, M. Luo, J.-H. Lin, and R. K. Gupta. An assessment of vulnerability of hardware neural networks to dynamic voltage and temperature variations. The 36th International Conference on Computer-Aided Design (ICCAD), 2017.
[10] J. Zhang, K. Rangineni, Z. Ghodsi, and S. Garg. Thundervolt: enabling aggressive voltage underscaling and timing error resilience for energy efficient deep learning accelerators. The 55th Annual Design Automation Conference (DAC), ACM, 2018.
[11] P. N. Whatmough, S. K. Lee, H. Lee, S. Rama, D. Brooks, and G.-Y. Wei. A 28nm SoC with a 1.2 GHz 568nJ/prediction sparse deep-neural-network engine with > 0.1 timing error rate tolerance for IoT applications. IEEE International Solid-State Circuits Conference (ISSCC), 2017.
[12] J. Zhang, Z. Ghodsi, S. Garg, and K. Rangineni. Enabling timing error resilience for low-power systolic-array based deep learning accelerators. IEEE Design & Test, 37(2):93–102, 2020.
[13] L. Wan, M. Zeiler, S. Zhang, Y. L. Cun, and R. Fergus. Regularization of neural networks using dropconnect. International conference on machine learning (ICML), 2013.
[14] P. Pandey, P. Basu, K. Chakraborty, and S. Roy. GreenTPU: Improving timing error resilience of a near-threshold tensor processing unit. The 56th Design Automation Conference (DAC), ACM, 2019.
[15] W. Choi, D. Shin, J. Park, and S. Ghosh. Sensitivity based error resilient techniques for energy efficient deep neural network accelerators. The 56th Design Automation Conference (DAC), ACM, 2019.

[16] X. Wu, Y. Hung, Y. Chen, and S. Chang. Accuracy tolerant neural networks under aggressive power optimization. Design, Automation & Test in Europe Conference (DATE), IEEE, 2020.
[17] J. Wang, X. Fu, X. Wang, S. Liu, L. Gao, and W. Zhang. Enabling energy- efficient and reliable neural network via neuron-level voltage scaling. IEEE Transactions on Computers, 2020.
[18] Jiao, X., Ma, D., Chang, W. and Jiang, Y., TEVoT: Timing error modeling of functional units under dynamic voltage and temperature variations. The 57th Design Automation Conference (DAC), ACM, 2020.
[19] J. Zhang and S. Garg. FATE: Fast and accurate timing error prediction framework for low power DNN accelerator design. International Conference on Computer-Aided Design (ICCAD), IEEE, 2018.
[20] J. Zhang, K. Liu, F. Khalid, M.A. Hanif, S. Rehman, T. Theocharides, A. Artussi, M. Shafique, and S. Garg, 2019, Building robust machine learning systems: Current progress, research challenges, and opportunities. The 56th Design Automation Conference (DAC), ACM, 2019.

Back to Contents

Paper Submission Deadlines

FPT'20 - Int'l Conference on Field-Programmable Technology
Virtual
Deadline: Aug 5, 2020 (Abstracts due: July 29, 2020)
Dec 7-11, 2020
http://icfpt.org

SLIP^2 - System-Level Interconnect Problems and Pathfinding (co-located with ICCAD 2020)
San Diego, CA
Deadline: Aug 21, 2020 (Abstracts due: Aug 14, 2020)
Nov 5, 2020
http://sliponline.org

WOSET'20 - Workshop on Open-Source EDA Technology (co-located with ICCAD 2020)
San Diego, CA
Deadline: Sept 6, 2020
Nov 5, 2020
https://woset-workshop.github.io

ISED' 20 – Int' l Symposium on Electronic System Design
Kollam, India
Deadline: Sept 25, 2020
Dec 18-20, 2019
http://isedconf.org

Back to Contents

Upcoming Conferences and Symposia

IWBDA'20 - Int'l Workshop on Bio-Design Automation
Online
Aug 3-5, 2020
http://www.iwbdaconf.org/2020

ISLPED' 20 – ACM/IEEE Int' l Symposium on Low Power Electronics and Design
Boston, MA
Aug 10-12, 2020
http://www.islped.org

ESWEEK'20 - Embedded Systems Week (CASES, CODES+ISSS, and EMSOFT)
Hamburg, Germany
Deadline: Apr 17, 2020 (Abstracts due: Apr 3, 2020)
Sept 20-25, 2020
http://www.esweek.org

NOCS'20 – IEEE/ACM Int' l Symposium on Networks-on-Chip (co-located with ESWEEK'20)
Virtual Conference
Sept 24-25, 2020
http://nocs2020.engr.uky.edu/

GLSVLSI' 20 – ACM Great Lakes Symposium on VLSI
Beijing, China
Sept 7-9, 2020
http://www.glsvlsi.org

PACT'20 - Int'l Conference on Parallel Architectures and Compilation Techniques
Atlanta, GA
Oct 3-7, 2020
http://www.pactconf.org

VLSI-SoC' 20 – IFIP/IEEE Int' l Conference on Very Large Scale Integration
Salt Lake City, UT
Oct 5-7, 2020
http://www.vlsi-soc.com

ISCAS'20 – IEEE Int'l Symposium on Circuits and Systems
Seville, Spain
Oct 11-14, 2020
http://iscas2020.org

MICRO' 20 – IEEE/ACM Int'l Symposium on Microarchitecture
Athens, Greece
Oct 17-21, 2020
http://www.microarch.org/micro53

BodyNets'20 – Int' l Conference on Body Area Networks
Cyberspace
Oct 21-22, 2020
http://www.bodynets.org

ICCAD' 20 – IEEE/ACM Int' l Conference on Computer-Aided Design
San Diego, CA
Nov 2-5, 2020
http://www.iccad.com

HOST'20 – IEEE Int' l Symposium on Hardware-Oriented Security and Trust
San Jose, CA
Dec 6-9, 2020
http://www.hostsymposium.org

HiPC'20 – IEEE Int'l Conference on High Performance Computing, Data, And Analytics
Pune, India
Dec 16-19, 2020

http://www.hipc.org

iSES' 20 – IEEE Int' l Symposium on Smart Electronic Systems
Chennai, India
Dec 14-16, 2020
http://www.ieee-ises.org

Back to Contents

Best Paper Awards

Best Paper Awards at DAC 2020: The Design Automation Conference, https://www.dac.com

"Algorithm-Hardware Co-Design of Adaptive Floating-Point Encodings for Resilient Deep Learning Inference", by Thierry Tambe, En-Yu Yang, Zishen Wan, Yuntian Deng, Vijay Kanapa Reddi, Alexander Rush, David Brooks and Gu-Yeon Wei (Harvard Univ.)

Back to Contents

Researcher Spotlight

Hello readers,

Welcome to Researcher spotlight. In this edition we meet with Prof. Pierre-Emmanuel Gaillardon. He is an Associate Professor in the Department of Electrical and Computer Engineering with at The University of Utah, Salt Lake City, UT, where he leads the Laboratory for NanoIntegrated Systems (LNIS). He holds an Electrical Engineer degree from CPE-Lyon, a M.Sc. degree in Electrical Engineering from INSA Lyon, and a Ph.D. degree in Electrical Engineering from CEA-LETI, Grenoble, France and the University of Lyon, France.

''1. Can you share with us some of the research areas you are interested in?''

I like to describe the activities of my laboratory, the Laboratory for Nano-Integrated Systems at The University of Utah, as being related to computer engineering at large. Indeed, we cover all aspects related to computer engineering, starting as low as novel device fabrication and device physics all the way up to electronic design automation and embedded systems. We have 3 main activities currently running in my lab:
- Research related to reconfigurable transistors: We have developed a novel technology of transistors called TIGFET (for Three-Independent-Gate FETs) that can be reconfigured at runtime to behave either n- or p-type. They bring significant advantages from a system-level standpoint as they lead to the design of arithmetic logic with the same level of complexity (understand less # of transistors, reduced Energy Delay Product, etc.) than NAND gate logic, while being fully compatible with standard CMOS processes. We believe this can be quite transformative to many modern computing workloads.
- Research related to open-source IPs and EDA: Fueled by the DARPA Electronics Resurgence Initiative (ERI), the lab develops the OpenFPGA and LSOracle projects. I will focus on LSOracle as OpenFPGA will be discussed more below. LSOracle aims to be a logic synthesis framework that provides complete no-human-in-the-loop automation and high Quality of Results (QoR). It exploits machine learning and a large collection of underlying data structures and optimization recipes to autonomously partition a Boolean network, identifies the nature of its partitions (arithmetic vs. control) and select the proper associated recipes.
- Research related to sensor networks and environmental monitoring: We have developed and deployed a complete sensor network in the SLC valley that monitors in real-time environmental pollution (www.aqandu.org), essentially particulate matter. We have developed several new techniques for in-field calibration, Machine Learning (ML)-based data fusion, and are working on

novel generations of sensors with extremely high sensitivity and selectivity across a large array of chemicals of interest.
- Finally, we have also several additional research activities on for instance in-memory computing, RRAM-based FPGAs, novel ML based techniques, etc.

''2. OpenFPGA seems a very comprehensive suite providing end-to-end support for FPGA architecture exploration. To achieve a power and/or performance envelope, what is the fastest way to deploy OpenFPGA and converge on architecture?''

Indeed, while as you pointed out, the most "thorough" way is to bring the evaluation all the way down to GDS. OpenFPGA has analytical exploration capabilities that are built-in into it, providing you with XML to analytical performances in a MUCH quicker way than evaluating everything down to GDS. The proper approach would therefore be to port one of the "standard architectures" we provide to your tech node of interest to extract the performance results you would obtain with the actual target. You would use this data to calibrate the data in your XML files and start your architectural exploration sweeps (with a speed that is > 100x faster than the GDS flow of course). Because your analysis is calibrated on the results you might expect at the GDS level, your evaluation is very accurate. This whole analysis can be parallelized at a high degree.

''3. Fullchip simulations understandably provide maximum accuracy at the expense of simulation time and are thus the baseline for normalization. Is there any particular benefit to using one amongst Component or Grid levels particularly when these levels seem to have more-than-Full-chip and less-than-Full-chip power consumption in some cases? In what type of cases do component level numbers exceed grid level? Why?''

This is an excellent question, but unfortunately my answer will be less excellent! Indeed, it is all a matter of trade-off. As you pointed out, the full chip simulations provides the highest level of accuracy but largest runtime. We consider this analysis to be used essentially for calibration of accuracy and less meaningful in a full application context. The selection then boils down to picking up either fastest runtime or highest accuracy. Component level will exceed grid level when the complexity of running grid level simulations is too high (large grids or expected simulation loops to be faster, etc.). That being said our 'golden' flow is the grid-level simulation as we believe it provides the best trade-off between runtime and accuracy.

''4. The NWFETs (Nano Wire FETs) seem promising for a true 3D architecture, particularly with the improved performance and notably reduced area. With the freed-up area positioned for increased signal count delivering more current, would you have a simple estimate on the power overhead this approach could possibly incur? What class (CPUs, GPUs, FPGAs, ASICs etc) of applications are likely to benefit with the available features and potential constraints?''

Some of this work traces back to my PhD. At that time, the question of power overhead from 3D (parallel, sequential or more aggressive as this real 3D concept) was still nascent. However, while we never did a quantitative analysis, a simple qualitative analysis would tell us that for the same function, you end up having a more "distributed" power overhead that happens in your back-end, effectively making power distribution and power dissipation easier. Indeed, the main concerns with 3D is that the power is generated in "layers" and that those "layers" block power dissipation and therefore require the use of metal contacts to serve as thermal bridges.

This concept is currently revisited in my lab in a research is collaboration with IMEC and we found that most of the benefits are coming for GPU-classes of architectures and optimizations as you can exploit symmetries in the circuits and large level of stacking to create very high-performance SIMD systems.

''5. 3.5D technology as a hybrid of 3D TSV and 3D MI seems interesting. For readers outside of this domain, could you help explain the significance of testing costs, the parameters influencing the said costs and opportunities to repurpose contributing parts towards a product with an eye on costs, in

this collective space of 3D and 3.5D spaces.''

The 3.5D analysis was built around cost of test models published by my former colleagues at CEA-LETI, France. Test is a major concern in 3D assemblies and the general trade-off appears between a monolithic and parallel 3D chips. Monolithic chips are tested similarly to a regular 2D chip (but that potentially has more defect sources since it has more layers and complexity). Parallel 3D chips have more testing steps involved (know-good dies, known-good stack, etc.) and larger costs associated. Therefore, the trade-off appears between more testing costs vs. more defects, and we argue that 3.5D leads to a sweet spot.

Back to Contents

Call for Nomination

Call for Nomination for ACM/SIGDA Pioneering Achievement Award

Presented by the ACM Special Interest Group on Design Automation.

Deadline: Aug. 31, 2020

Description: To honor a person for lifetime, outstanding contributions within the scope of electronic design automation, as evidenced by ideas pioneered in publications, industrial products, or other relevant contributions. The award is based on the impact of the contributions throughout the nominee’ s lifetime.

Background: The ACM Special Interest Group on Design Automation sponsors or co-sponsors the ACM Transactions on Design Automation of Electronic Systems Best Paper Award, the William McCalla Award for best paper at the International Conference on Computer-Aided Design, and the ACM/IEEE A. Richard Newton Technical Impact in Electronic Design Automation Award which is given to authors of a publication authored at least ten years earlier and that has had an outstanding contribution to the field of EDA. In addition, SIGDA sponsors the ACM Outstanding Ph.D. Dissertation Award in Electronic Design Automation which is given each year to a graduating Ph.D. student in recognition of his/her thesis contributions to advancement in the EDA field. The Pioneering Achievement Award complements these awards and is intended for contributors whose impact is typically recognized over a lifetime of outstanding achievements.

Eligibility: Open to researchers in the field of electronic design automation who have had outstanding contributions in the field during their lifetime. Current members of the Board of the ACM Special Interest Group on Design Automation, or members of the Award Selection Committee are ineligible for the award. The awardee is invited to give a lecture at a SIGDA-sponsored event.

Award Items: A plaque for the awardee, a citation, and $1000 honorarium. The honorarium will be funded by the SIGDA annual budget.

Nominee Solicitation: The call for nominees will be published by email to members of SIGDA, on the web site of ACM’ s Special Interest Group on Design Automation, and in the SIGDA newsletter. The nomination should be proposed by someone other than the nominee. The nomination materials should be emailed to SIGDA-Award@acm.org (Subject: ACM/SIGDA Pioneering Achievement Award). Nominations for the award should include:

A nomination letter that gives: a 100-word description of the nominee’ s contribution and its impact; a 750-word detailed description of up to 10 of the nominee’ s major products (papers, patents, software, etc.), the contributions embodied in those products, and their impact; a list of at most 10 citations to the major products discussed in the description.
Three letters of recommendation (not including the nominator or nominee).
Contact information of the nominator.

In addition to the evidence of impact, the nomination form will include biographical information (including education and employment), professional activities, publications, and recognition. Three endorsements attesting to the impact of the work may be included.

Award Committee: Selection will be made by the ACM Special Interest Group in Design Automation Executive Committee based on the recommendation of a Pioneer Award committee. The Committee will meet to review nominations, review the recommendations of the Pioneer Award Committee, and make a selection. After selection, the committee will contact the recipient to ensure that the award will be accepted and he or she will be able to deliver the talk at the SIGDA Annual Member Meeting and Dinner at ICCAD.

All standard conflict of interest regulations as stated in ACM policy will be applied (see https://awards.acm.org/conflict-of-interest). Any awards committee members will recuse themselves from consideration of any candidates where a conflict of interest may exist.

Schedule: The call for nominees will be published annually. The nomination deadline is Aug 31st. The award will be announced at one or more subsequent SIGDA events and the awardee will be invited to give a talk on his/her work at the SIGDA Annual Member Meeting and Dinner at ICCAD.

Selection/Basis for Judging: This award honors an individual who has made an outstanding technical contribution in the scope of electronic design automation throughout his or her lifetime. The award is based on the impact of the contributions as indicated above. Nominees from universities, industry, and government worldwide will be considered and encouraged. The award is not a best paper or initial original contribution award. Instead, it is intended for lifetime, outstanding contributions within the scope of electronic design automation, throughout the nominee’ s lifetime.

Presentation: The award will be presented annually at the SIGDA Annual Member Meeting and Dinner at ICCAD.

Publicity: In ACM/SIGDA publications and at conferences sponsored by ACM/SIGDA.

For more information, please contact David Pan, ACM/SIGDA Award Chair, dpan@ece.utexas.edu

Back to Contents

Technical Activities

1, “Connected Car Industry: Big Picture”

It's been over 20 years since the first connected cars appeared in the United States. There are a lot more...

[https://www.eetasia.com/connected-car-industry-big-picture/]

2.” It's Time for Intel to Double Down on Foundry Business”

Intel Corp. dropped the ball. It should today be the world’ s biggest semiconductor foundry; not TSMC. This can still change...

[https://www.eetasia.com/its-time-for-intel-to-double-down-on-foundry-business/]

3. “Silicon 100: Emerging Startups to Watch”

This article originally appeared on EETimes.com Every year (or so) we publish a list of electronics and semiconductor startups that grabbed our...

[https://iot.eetimes.com/silicon-100-emerging-startups-to-watch/]

4. "Soaring Attendance at the 57th Design Automation Conference, as Premier Event for the Electronic Design Ecosystem Gets Even Bigger"

Virtual conference expands global reach with record attendance...

[https://www.edacafe.com/nbc/articles/1/1770304/Soaring-Attendance-57th-Design-Au...]

Job Openings:
---
1. University of Chicago Department of Computer Science United States

Job Title: Professor of Computer Science; Databases

Description: The Department of Computer Science at the University of Chicago invites applications for a position of Clinical Professor of Computer Science to teach Databases classes in its Masters Program in Computer Science (MPCS) and in its joint program with the Harris School of Public Policy, the Master of Science in Computational Analysis & Public Policy (CAPP). This full-time, benefit-eligible appointment is for an initial three-year term, with possibility of renewal. This is a teaching position with no research responsibilities, and a teaching load of six courses across three academic quarters of the year (Autumn, Winter, Spring). The person holding this position will teach at least two different courses: MPCS 53001 Databases and CAPP 30235 Databases for Public Policy. Syllabuses for the latest offerings of these classes can be found at https://mpcs-courses.cs.uchicago.edu/2019-20/spring/courses/53001 and https://www.classes.cs.uchicago.edu/archive/2019/spring/30235-1/syllabus.html. Depending on the applicant' s background and interests, the person holding this position may also be asked to teach classes covering advanced topics in Databases.
Applications must be submitted online through the University of Chicago's Academic Jobs website: apply.interfolio.com/76562. Review of applications will begin August 1, 2020, and continue until the position is filled.

2. University of Bonn, Germany

Job Title: W1-Professorship (Tenure Track W2) for Embedded Visual Computing

Description: The institute is looking for excellent candidates holding a doctoral degree who are about to establish their own research agenda, complementing our current activities in the areas of Visual Computing, Artificial Intelligence, Intelligent Systems, and Robotics. Applicants must have demonstrated scientific excellence in at least one of the following areas: Deep learning for visual computing, Parallel architectures for visual learning or inference, Embedded systems for visual perception or display, Energy-efficient computing, Edge computing, Novel computing paradigms, such as computing with stochastic elements or neuromorphic computing. The successful candidate has a strong publication record. It is expected that the successful candidate actively participates in collaborative research initiatives of the Computer Science Institute. Initial teaching duties are four hours per week during lecture periods, e.g., in our English Master Computer Science programme. Applicants are asked to send their applications, including a CV with list of ten most important publications, full list of publications, research statement, teaching statement, and copies of degree certificates, as one PDF file to fachgruppe@informatik.uni-bonn.de.

3. University of Northern British Columbia Canada

Title: Assistant Professor of Computer Science

Description: The University of Northern British Columbia invites applications for a term position in the Department of Computer Science at the rank of Assistant Professor, with a proposed starting date of

July 1, 2020. As an institution committed to fostering an inclusive and transformative learning environment, UNBC values high quality and growth in both teaching and scholarship. The duties of a regular full-time term appointment at the rank of Assistant Professor normally include teaching, scholarly activity, and service. Teaching responsibilities for this position involve teaching across the undergraduate curriculum, with teaching beginning in September. At least for the Fall Semester most teaching will be conducted online. The successful candidate will also have an opportunity to work on the development of course materials in collaboration with other faculty in the Computer Science Program, and to engage in scholarly activity. Applicants for this position should hold a PhD in Computer Science or closely related field (applicants nearing completion may be considered) and have a growing record of scholarly activity, as well as present evidence of a record of and commitment to teaching excellence. For this position, we are also interested in an individual with research interests that might enhance our undergraduate or graduate teaching offerings.
For more information about living and working with the University of Northern British Columbia please refer to http://www.unbc.ca/experience and http://www.unbc.ca/provost/new_faculty.html.
Make your mark with this leading post-secondary institution.

4. University of Queensland School of Information Technology and Electrical Engineering Australia

Title: Senior Lecturer/Associate Professor in Neuroengineering

Description: It is an exciting time to get involved with the School of Information Technology and Electrical Engineering, located on UQ' s St. Lucia campus. The School is ramping up its investment in teaching, research and engagement to create an inspiring, diverse and flexible workplace. The direction is backed by a bold, new strategic vision to ensure the School is at the forefront of meaningful research outcomes and pedagogy across its core impact areas of health, data, automation and energy. Boasting strong student enrolments in professionally accredited programs, combined with world-class researchers and facilities, the School is focused on strengthening its position in the global computer science and engineering communities. By attracting the brightest minds and fostering a truly innovative and collaborative work environment, the School will develop global solutions to contemporary issues and mentor the leaders of tomorrow.
As part of the Selection Process, applicants shortlisted for interview will be required to present a seminar and mock lecture to members of the school community prior to their scheduled interview. Research seminars and lecture topics will be provided with adequate notice to shortlisted candidates. The University of Queensland is committed to ensuring all candidates are provided with the opportunity to attend the panel interviews, however, for those candidates who are unable to attend in person, video interview options will be available. To discuss this role please contact Professor Aleksandar Rakić on a.rakic@uq.edu.au or Professor Pankaj Sah on pankaj.sah@uq.edu.au.

5. Duke Kunshan University

Title: Research Specialist on Software Development

Description: The Data Science Research Center (DSRC) at Duke Kunshan University (DKU) is seeking applicants with expertise in software development to engage in research collaboration between DKU and major industrial companies. A successful candidate will work with a group of faculty members on software design, implementation and validation, as well as providing consultation and assistance to industrial collaborators for technology transfer. As such, the position includes great opportunities for coinvestigation and co-authorship in dissemination activities.
Review of applications will begin immediately and continue until the position is filled. To be considered for the position, please email your curriculum vitae with the subject line of "Research Specialist on Software Development" to dkurecruitment@dukekunshan.edu.cn and xinli.ece@duke.edu.

Back to Contents

Notice to Authors

This newsletter is a free service for current SIGDA members and is added automatically with a new SIGDA membership.
Circulation: 2,700