sigda Special Interest Group
Design Automation
The resource for EDA Professionals

1 September 2019, Vol. 49, No. 9
Online archive: http://www.sigda.org/publications/newsletter



## Comments from the Editors

Dear ACM/SIGDA member,

We are excited to present to you the September 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 upcoming conference and funding deadlines, hot research topics to news 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!

Aida Todri-Sanial
Yu Wang
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:"

'"Aida Todri-Sanial"', E-Newsletter co Editor-in-Chief

'"Yu Wang"', 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

'"Yuan-Hao Chang"', E-Newsletter Associate Editor for SIGDA What is column

'"Wenchao Li"', E-Newsletter Associate Editor for SIGDA What is column

'"Debjit Sinha"', E-Newsletter Associate Editor for for SIGDA Paper submission deadline column

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

Back to Contents

## SIGDA News

(1) "Apple to Buy Intel's Modem Business for $1 Billion'"
[https://www.eetimes.com/document.asp?doc_id=1334964]

Apple will pay about $1 billion to acquire Intel's smartphone modem business, signaling that, despite a settlement reached with longtime supplier Qualcomm in April, Apple still has designs on its own silicon for 5G.

(2) "Alibaba Reveals 16-core RISC-V Chip"
[https://www.eetimes.com/document.asp?doc_id=1334966]

Alibaba Group's chip subsidiary, Pingtouge Semiconductor, this week announced what it claims is the most powerful RISC-V based processor, the Xuantie 910, targeting infrastructure for artificial intelligence (AI), 5G, and internet of things (IoT) as well as autonomous vehicles.

(3) "SoftBank Launches $108B Vision Fund 2 Targeting AI"
[https://www.eetimes.com/document.asp?doc_id=1334968]

Softbank has today launched its Vision Fund 2 to facilitate investment in late-stage artificial intelligence (AI) companies. The fund, for which $108 billion has been committed so far, will include participation from Apple, Foxconn, Microsoft, various banks from Japan and Kazakhstan, plus major participants from Taiwan.

(4) "As Trade Talks Resume, China Advances 5G"
[https://www.eetimes.com/document.asp?doc_id=1334879]

For the global semiconductor industry, the just-revealed truce in the “trade war” between U.S. President Donald Trump and China’s president, Xi Jinping is welcome news but hardly the end of a poorly scripted economic melodrama that continues to pose dire consequences for the American semiconductor market.

(5) "What the Japan-Korea Trade War Means to the World"
[https://www.eetimes.com/author.asp?section_id=36&doc_id=1334926]

By restricting exports of key chemical materials to South Korea, Japan has already done irreparable harm to the global electronics industry, at a price both incalculable and unnecessary.

(6) "TSMC Sees 5G Driving Strong Demand for 7nm"
[https://www.eetimes.com/document.asp?doc_id=1334951]

Taiwan Semiconductor Manufacturing Co. (TSMC) sees 5-nm and 7-nm demand improving from its earlier expectations, as the worldwide 5G development accelerates.

(7) "For Mass-Market Cars, Forget L3-L5 Autonomy"
[https://www.eetimes.com/author.asp?section_id=36&doc_id=1334845]

Let's apply the SAE automation levels to the mass-market cars. Nobody should be surprised how the future of autonomous vehicles looks a lot different in Detroit, Stuttgart and Tokyo.

(8) "Startup Puts AI Core in SSDs"
[https://www.eetimes.com/document.asp?doc_id=1334982]

Startup InnoGrit debuted a set of three controllers for solid-state drives (SSDs), including one for data centers that embeds a neural-network accelerator. They enter a crowded market with claims of power and performance advantages over rivals.

(9) "Battery Research Advances Quantum Computing Capabilities"
[https://www.eetimes.com/document.asp?doc_id=1334967]

Research in battery chemistry is getting a boost from quantum computing that in turn is furthering research to improve quantum computing.

(10) "STEM: It’s Time to Add an ‘A’"
[https://www.eetimes.com/document.asp?doc_id=1334730]

Yes, the world needs more people versed in STEM, but the growing imperative for STEM workers is hardening into dogma that isn’t good for anybody, including companies in the technology industry. As STEM skills are becoming increasingly exalted, non-technical skills are being not only minimized but denigrated. The problem is encapsulated in the tension between the two jokes above.

Back to Contents

## SIGDA Local Chapter News

We are pleased to announce that Design Automation Conference (DAC) has been promoted to Rank-A conference by China Computer Federation (CCF). The CCF webpage is https://www.ccf.org.cn/. Now DAC is CCF-A conference now! Thanks for the faculty members and researchers that have contributed to this.

Back to Contents

## SIGDA Award

Awards at ACM/SIGDA Sponsored Events

DAC 2019: The 56th Design Automation Conference, https://dac.com

1. Best Paper Award

"DREAMPLACE: Deep Learning Toolkit-enabled GPU Acceleration for Modern VLSI Placement" by Yibo Lin (Univ. of Texas at Austin), Shounak Dhar (Univ. of Texas at Austin), Wuxi Li (Univ. of Texas at Austin), Haoxing Ren (NVIDIA Corp., Austin), Brucek Khailany (NVIDIA Corp., Austin) and David Z. Pan (Univ. of Texas at Austin).

2. 2018 Phil Kaufman Award for Distinguished Contributions to Electronic Systems Design

Thomas W. Williams, Synopsys (retired, fellow): "For overall impact on electronic industry through contributions to scan design for testability, related test automation".

3. ACM IEEE A. Richard Newton Technical Impact Award in Electronic Design Automation

Edward B. Eichelberger (IBM retired fellow, IEEE fellow) and Thomas W. Williams (Synopsys retired fellow): for the paper "A Logic Design Structure for LSI Testability," in Proc. of the 14th Design Automation Conference, 1977.

4. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems Donald O. Pederson Best Paper Award

“Caffeine: Towards Uniformed Representation and Acceleration for Deep Convolutional Neural Networks,” by Chen Zhang (Microsoft Research Asia), Guangyu Sun (Peking University), Zhenman Fang (Simon Fraser Univ.), Peipei Zhou (University of California at Los Angeles), Peichen Pan (Falcon Computing), and Jason Cong (University of California at Los Angeles).

5. 2019 ACM TODAES Best Paper Award

"Security in Automotive Networks: Lightweight Authentication and Authorization," by Philipp Mundhenk (TUM CREATE Limited, Singapore), Andrew Paverd (Aalto University, Finland), Artur Mrowca (TUM CREATE Limited, Singapore), Sebastian Steinhorst (TUM CREATE Limited, Singapore), MArtin Lukasiewycz (TUM CREATE Limited, Singapore), Suhaib A. Fahmy(University of Warwick, United Kingdom) and Samarjit Chakraborty (Technische Universita¨t Mu¨nchen, Germany).

6. ACM SIGDA Outstanding Ph.D. Dissertation Award

"Distributed Timing Analysis," by Tsung-Wei Huang, Advisor: Martin D. F. Wong (Univ. of Illinois at Urbana-Champaign).

Back to Contents

## "What is" Column

What Is Processing-in-Memory (PIM)

Po-Chun Huang,
Assistant Professor,
po.chun.huang@mail.ntut.edu.tw
National Taipei University of Technology, Taiwan

Classical computer architecture used to adopt separate hardware components for computation (e.g., CPUs and GPGPUs) and those for data storage (e.g., DRAM-based main memory or disk-based secondary storage) [1]. As a result, the performance of modern data-intensive applications is often limited by the speeds of the memory/storage devices, as well as the bandwidth of the datapath between CPUs/GPUs and memory/storage, thereby creating the well-known von Neumann bottleneck [2]. To overcome this problem, the processing in memory (PIM) technology is proposed to incorporate the processing capability to random-access memory (RAM) in a single chip [2]. By moving some computation from processor to memory [3], the overheads to transmit data between processor and memory can be remarkably alleviated. As reported by prior literature [2], the PIM technology becomes more feasible as the main memory capacity increases, and are especially valuable for the emerging data-centric computing scenarios, such as data mining and machine learning applications.

As early proposals of PIM introduce fully programmable computation units (e.g., general-purpose processors or field programmable gate arrays) to the memory, the design efforts could be high, and the changes to the hardware/software stack might be inevitable [7]. Thus, the recent development of PIM technology is driven by the emergence of modern nonvolatile memories (NVMs), such as phase-change memory (PCM), metal-oxide resistive RAM (ReRAM), and spin-transfer torque RAM (STT-RAM), which can directly perform logical and arithmetic operations in memory [1]. Among these choices, ReRAM can support efficient matrix–vector multiply operations [1], and is widely used for neural computation [4], graphic algorithms [5], or performing bulk bitwise operations [6]. Nevertheless, the PIM technology is not preferred by all flavors of computations against CPUs/GPUs, and extra research efforts are needed to fully unleash its power.

Observing the heterogeneous computation capabilities of CPUs/GPGPUs and those of PIM, it is a brilliant approach to wisely determine whether to execute specific instructions, referred to as PIM-enabled instructions (PEIs), by PIM on the main memory, as suggested by Ahn et al. [7]. With PEIs, programmers can assign the instructions that should be executed by PIM, therefore optimizing the system performance. Representative examples of PEIs include the increment of integers, getting the minimum element, addition of floating-point numbers, computation of Euclidean distance, and computation of the dot product of vectors [7].

The PIM technology is especially suitable for applications in the artificial intelligence area, as observed by [1], [4], [5]. This is because that certain NVMs, such as ReRAM, can inherently support neural computation with its crossbar physical architecture. By allocating a full function (FF) subarray in the ReRAM space [1], a remarkable improvement of energy saving is observed at slight area overheads. Because the FF subarray is established dynamically in the ReRAM, the dynamic morphing of the ReRAM space for storing data and that for keeping the FF subarray becomes possible, allowing further performance enhancements.

In next-generation computing systems with explosive scales, the datapath between processing components and memory/storage components might become one of the major performance bottlenecks. In this case, PIM is promising to assist the establishing of active memory cubes (AMCs) [8], which provides not only data storage but also energy-efficient computation functionalities in a single component. A plural of AMCs can then be interconnected to construct a coherent and scalable main memory device for performance-demanding applications such as scientific computation [8]. Furthermore, the importance of the PIM technology is expected to be emphasized in the foreseeable future, due to the rapidly growing scale of computing systems.

As new applications keep emerging on the horizon, existing computer architecture is pushed further to deliver better performance and energy efficiency for the whole system. Among the potential technologies, PIM provides amazingly high performance and energy efficiency, and is a promising candidate for the key technologies in data-centric computing scenarios. We believe that the PIM technology is still worth more research attentions to reveal its value to a wider spectrum of applications.

[1] Chi, Ping, et al. “Prime: A novel processing-in-memory architecture for neural network computation in ReRAM-based main memory.” ACM SIGARCH Computer Architecture News. Vol. 44. No. 3. IEEE Press, 2016.

[2] Margaret Rouse, “What is processing in memory (PIM)?,” online available at: https://searchbusinessanalytics.techtarget.com/definition/processing-in-memory-P....

[3] Xu Yang, Yumin Hou, and Hu He, “A Processing-in-Memory Architecture Programming Paradigm for Wireless Internet-of-Things Applications,” MDPI Sensors, Vol. 19, No. 140, 2019.

[4] Song, Linghao, et al. “Pipelayer: A pipelined ReRAM-based accelerator for deep learning.” 2017 IEEE International Symposium on High Performance Computer Architecture (HPCA). IEEE, 2017.

[5] Long, Yun, Taesik Na, and Saibal Mukhopadhyay. “ReRAM-based processing-in-memory architecture for recurrent neural network acceleration.” IEEE Transactions on Very Large Scale Integration (VLSI) Systems 26.12 (2018): 2781-2794.

[6] Li, Shuangchen, et al. “Pinatubo: A processing-in-memory architecture for bulk bitwise operations in emerging non-volatile memories.” Proceedings of the 53rd Annual Design Automation Conference. ACM, 2016.

[7] J. Ahn, S. Yoo, O. Mutlu and K. Choi, “PIM-enabled instructions: A low-overhead, locality-aware processing-in-memory architecture,” 2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA), Portland, OR, 2015, pp. 336–348.

[8] R. Nair et al., “Active Memory Cube: A processing-in-memory architecture for exascale systems,” IBM Journal of Research and Development, vol. 59, no. 2/3, pp. 17:1-17:14, March–May 2015.

Back to Contents

# Researcher Spotlight

Hello readers,
In this edition of Researcher Spotlight we meet Professor Vaughn Betz. He is the NSERC/Intel Industrial Research Chair in Programmable Silicon in the Department of Electrical and Computer Engineering at the University of Toronto. He received his B.Sc in Electrical Engineering from the University of Manitoba followed by M.S. in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign and Ph.D. in Electrical and Computer Engineering from the University of Toronto. Excerpts from a recent conversation.

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

I'm very interested in how we can create novel FPGAs or other efficient spatial computing architectures. Some of my work in that area focuses on optimizing the architecture for deep learning, as that's clearly going to be a key part of many workloads in coming years, and it stresses the device in different ways: more math and less logic. I am also interested in how we can evolve programmable logic to be an easier and more efficient target for large systems and data center accelerators; a big focus in that area for my group has been how to integrate hardened Networks-on-Chip (NOCs) into the fabric to make system-level communication easier.

Another area I'm interested in is computer-aided design, as these devices are not very usable without good CAD algorithms. One focus is investigating how we can implement ever larger designs in a fixed compile time, using faster and more parallel algorithms, and new design styles that keep modules more timing-independent like latency-insensitive design and hard Networks-on-Chip. My group is doing a lot of work in the open-source VTR project, with many collaborators, to ensure there is a CAD framework that can be used to investigate novel FPGAs and algorithms. Some of the collaborations and work done by others in the VTR framework recently have allowed automatic layout of those novel FPGAs, and programming of commercial FPGAs from a fully open-source tool flow. It's been exciting to see how open-source CAD is gaining capabilities and momentum.

A third research area for me is a bit farther afield: it's a collaboration with a medical research team to create fast light simulators and custom CAD algorithms to model and optimize light-activated chemotherapy. That's been a very interesting project as it is so broad, yet it also has shown that knowledge of optimization algorithms and how to create fast software and hardware accelerators is useful in domains far from computer chip design.

2. Dynamic Voltage Scaling (DVS) as a power management technique has been successful in ASIC and microprocessor designs. How do you foresee its application in modern FPGAs, particularly since they have evolved from programmable fabrics to feature-rich systems containing gigabit-transceivers, embedded processor and memory instances?

Dynamic Voltage Scaling is definitely more difficult in FPGAs than in an ASIC or CPU, as the end use, and hence speed-limiting paths, aren't known when the FPGA is manufactured – they are only known when the designer programs the FPGA. That means you can't build speed-monitoring circuits into the device in the same way you do an ASIC and use them to decide what a safe voltage is. Despite that added difficulty, I think adaptive voltage and DVS is becoming more necessary for FPGAs, as without them voltage is not dropping much as we move to newer technologies, and that's leading to high power dissipations for the latest, very high capacity, FPGAs.

We've done a lot of work in this area recently. We created a CAD system that the end user runs (FRoC) to automatically create application-specific speed monitoring paths and determine how fast they are in each programmed FPGA, to find the lowest safe voltage setting for each chip. We think that combination of being automatic, but run by the end user who knows the application, is the key to making DVS work well for FPGAs.

You definitely raise a great point about modern FPGAs being much more than a programmable fabric: they include processors, gigabit transceivers, fast DDR interfaces and more. We think the key to applying DVS to such a complex system is voltage islands by feature type: the transceivers, DDR interface, processor, etc. should be on separate voltage islands so their voltage can be adjusted independently. A user running a complex transceiver at the max spec could choose to leave that at the nominal voltage while using DVS on the core logic that's consuming more of the power. Users with more timing margin on their transceivers and more stringent power requirements could extend the approach we take in FRoC: create a calibration design that tests the transceivers at lower voltages in each chip and in the required application configuration, to determine a safe set of transceiver voltages are at various temperatures.

3. Typical FPGAs, when implemented on silicon have extensive physical implementation of power grids shared by several sub-components with varying functionality, often operating at the same voltage domain. What do you think could

be major challenges in adopting DVS towards such scenarios?

I think you'll have to make some trade-offs. If you keep two sub-components that are quite distinct (such as an embedded processors and the programmable routing) on the same power grid, you have to take the worst-cast voltage that allows both to meet their timing requirements on your application on a specific chip. So that will generally reduce your DVS power savings, especially if those two sub-components behave quite differently as you scale voltage or tend to have very different amounts of timing margin.

To get more out of DVS, you can separate components that are likely to want different adaptive voltage onto separate power grids. You wouldn't want to make too many little grids, as it may make your power grid too sensitive to transients due to less capacitance, or increase area. However, there is already a trend to separate power grids to allow transceivers and PLLs to have a very stable supply, block RAMs to have a somewhat higher voltage, and the rest of the fabric a lower voltage. I think that trend will continue and we are looking at this question now in my group: should we separate FPGA circuitry into more voltage islands to make DVS more effective?

It's also interesting to look at where CPUs, which have been pioneers in DVS, have gone. Recent IBM Power CPUs have 30 voltage islands for example, indicating that CPU architects feel a reasonably large number of islands will pay for themselves with power savings.

4. Silicon interposers seem like a good opportunity to migrate from 2-dimension to 2.5D/3D FPGA designs. Moreover, interface bus-widths on interposers seem to help reduce required minimum channel width on FPGA fabric. While this may help free up valuable die-area towards logic and/or memory, the typical performance penalty incurred seems steep. How do you envision synthesis of large applications on such systems involving significant die-interposer-die connectivity?

Interposers certainly do open up some interesting new areas for FPGA architects. There are different ways to exploit them, and the two largest FPGA vendors, Intel and Xilinx, have taken fairly different approaches. Intel is using interposers to integrate a monolithic fabric die with various smaller "chiplet" dice that perform more specialized functions, particularly high speed I/O. That should keep the interconnect demand on the interposer more limited, and the partitioning of what logic goes where is essentially completed by the FPGA architect. The fact that the main die is programmable logic is interesting as it lets you communicate with a wide variety of chiplets just by programming parts of the fabric differently.

Xilinx has used a different approach; they use interposers to make very large programmable fabrics, where the actual fabric itself is divided into multiple dice. That puts more pressure on the interposer interconnect, as you can't get as many wires across an interposer as you can in a monolithic die, and the programmable routing of an FPGA is very wiring intensive. To deal with that interconnect choke point you need CAD tools that understand crossing an interposer is expensive, and it's also helpful to refine your architecture. We've done some work showing that if you have 20% of the usual programmable routing crossing an interposer, CAD tools that are aware of that choke point can do a good job of optimizing a design automatically. That 20% programmable routing crossing the interposer roughly matches what Xilinx's 28 nm FPGAs achieve, so that result lined up well with what they built. As interposer microbumps don't scale as well as on-die interconnect, this choke point is likely to get more challenging with device scaling. Xilinx has been refining their architecture, adding registers before and after interposer crossings for example, as well as their CAD tools, where they've published some interesting near-optimal place and route tools for small parts of the overall design that occur near the interposer boundary. Both those enhancements are helpful to avoid forcing end users to floorplan their devices or modify their designs to minimize routing across the interposer.

Embedded, or hard, networks-on-chip also have a role to play here. Since a hard NoC runs at a higher frequency than the programmable fabric and can be used by different traffic flows at different times, a hard NoC that crosses an interposer increases the interconnect bandwidth achievable with the limited interposer wiring. That's a nice synergy that Xilinx's upcoming Versal devices can exploit with their embedded hard NoC.

5. You have the distinction of having developed and worked on Versatile Place and Route (VPR) as a graduate student, possibly used variants of it in industry and continuing to develop it on your return to academia. What are your thoughts on this? Any advice to researchers and open source enthusiasts in similar pursuits?

VPR has definitely had far more impact and been a longer-lived project than I ever imagined when I started writing it during my PhD! It would be nice to say I had a grand plan that led to this outcome, but I really didn't.

During my PhD I knew I wanted to research FPGA architecture, and I liked writing algorithmic software, so developing VPR was a natural way to combine those two pursuits. I think there are two decisions that in retrospect were quite important to VPR's success. The first was that as I researched various FPGA architecture questions -- from how to make a good logic block, to how much interconnect should go in various parts of the chip, to the exact wire and switch patterns –

I added each new feature into the program and made it work with the prior features, rather than developing a new program or hacking in a solution for one problem that I would then throw away when moving on to the next question. That actually took longer, so it probably delayed my graduation some, but it meant that the capabilities of the tool were constantly growing. The second important decision was that I spent a reasonable amount of time on software engineering, often recoding an algorithm or set of data structures because I realized the initial implementation was becoming unmaintainable. Similarly, I added automatic data structure and algorithm checkers to the program so it could check its own output, and I commented reasonably extensively; both those things helped make the program more maintainable so both myself and others could keep extending it.

Taking VPR commercial, as Right Track CAD, which was ultimately acquired by Altera and now Intel, was a great experience. It was really driven by the fact that FPGA architecture and place and route were becoming much more complex at that time and we had some good research and a capable program in that area. So the timing was important – it's good to be in a place where a lot of change and disruption is happening to maximize the chances of commercializing your research. That experience was also very useful in terms of learning every detail of FPGA architectures and CAD tools; that's something that is easier to do in a commercial setting than in academia, and it's helped inform my research since becoming a professor.

There are now multiple commercial versions of VPR (inside Intel/Quartus and other companies), and also the open source version of VPR in the VTR project. In all those versions I'd say keeping the program flexible and data-driven so it can adapt to different chips automatically has been key. Another key challenge is to keep refactoring the code and emphasizing software engineering so the ever larger codebase doesn't become unmaintainable. Code reviews and good regression tests are probably the most important aspects of keeping the code maintainable and they're often overlooked since they don't lead to papers. That's probably the biggest challenge in an academic open-source project: how do you create enough expert developers to review the code of others, especially since that is not their main research task?

6. FPGA-2012 had a panel titled 'FPGAs in 2032'. FPGAs have come a long way from being vehicles for rapid prototyping to hardware emulation and now to hardware acceleration & deep learning. However, GPUs continue to dominate the deep learning space. In your view, what changes in FPGAs could tilt this in their favour?

I think FPGAs will play an important role in deep learning inference. Deep learning inference plays to their strengths: it is tolerant of low precision, has tight latency requirements, and often inference will be in edge or embedded devices with custom sensor interfaces and control logic that can also benefit from an FPGA's programmability. I think training is more difficult for FPGAs as it plays more to the strengths of GPUs: it requires higher precision, is not latency sensitive, has a higher memory footprint, and is performed in data centers with higher power budgets.

The ongoing research into lower precision neural network inference definitely helps FPGAs and I think FPGA architecture changes can help tilt the playing field more in their favour. My group has been working on low-cost changes to the logic block and DSP blocks that can improve FPGA performance on low precision deep learning by ~1.5x – 5x, for quite low hardware area costs. It's also possible that more extensive (and expensive!) hardware changes could pay for themselves by improving deep learning performance still more – it's a pretty wide open and exciting field right now.

Another big challenge for FPGAs has been the relative difficulty of creating an FPGA deep learning implementation vs. using a GPU. To overcome that, the work going on both commercially and in a variety of academic groups to create flows that go from machine learning frameworks to an FPGA implementation via a domain-specific compiler is very important.

7. In your view, what can academia learn from industry and vice-versa, particularly in the FPGA space? What should be the expectations for a mutually beneficial relationship? Is there scope for improvement in one or more areas?

In my experience, industry is great at scaling up and driving large projects. Academia's strengths are in exploring out-of-the-box ideas, and in building teams with a very diverse knowledge base to tackle unusual cross functional problems. The scale of the teams that I was part of in Altera to build Quartus and the Stratix series devices is hard to replicate in academia for example, and teams of that size really help you get all the details in a product complete. On the other hand, the work my group did on hard networks-on-chip has gone on for 8 years. While commercial devices are now incorporating hard networks-on-chips, when we started that project it was definitely considered radical and many commercial FPGA architects were sceptical; being in academia let us explore a higher risk idea. The medical collaboration to improve the effectiveness of light-activated chemotherapy is a good example of a project that suits academia well – it would be hard to assemble a group of medical and CAD experts in any one company, but in a university those researchers can be across the street from each other.
Overall I think academia's key contribution is in creating prototypes and researching ideas that industry hasn't explored, or won't explore. At some point if those ideas pan out, industry is best at scaling them into full products.
To get more synergy between academia and industry I think having some projects where academia and industry closely collaborate is very useful. That's happened on some projects I've been part of, and it can be a great experience as industry brings a deep knowledge of the real-world constraints, and grad students bring a can-do, optimistic attitude to try new

things. Those collaborations also help create a better training environment for graduate students, and in the end training students is the most important thing academia does – they go on to long careers in which they make contributions far beyond the research they do during their degrees.

Rajsaktish Sankaranarayanan
Associate Editor
Back to Contents

## Paper Submission Deadlines

WOSET'19 – Workshop on Open Source EDA Technology (co-located with ICCAD'19)
Westminster, CO
Deadline: Sep 6, 2019
Nov 7, 2019
http://woset.org

DATE'20 - Design Automation and Test in Europe
Grenoble, France
Deadline: Sep 8, 2019
Mar 9-13, 2020
http://www.date-conference.com

ISSCC'20 – IEEE Int'l Solid-State Circuits Conference
San Francisco, CA
Deadline: Sep 9, 2019
Feb 16-20, 2020
http://isscc.org

ISQED'20 - Int'l Symposium on Quality Electronic Design
Santa Clara, CA
Deadline: Sep 14, 2019
Mar 25-26, 2020
http://www.isqed.org

FPGA'20 – ACM/SIGDA Int'l Symposium on Field-Programmable Gate Arrays
Seaside, CA
Deadline: Sep 16, 2019
Feb 24-26, 2020
http://www.isfpga.org

Back to Contents

## Upcoming Symposia, Conferences and Workshops

PACT'19 - Int'l Conference on Parallel Architectures and Compilation Techniques
Seattle, WA
Sep 21-25, 2019
http://www.pactconf.org

BodyNets'19 – Int'l Conference on Body Area Networks
Florence, Italy
Oct 2-3, 2019
http://www.bodynets.org

VLSI-SoC'19 – IFIP/IEEE Int'l Conference on Very Large Scale Integration
Cuzco, Peru
Oct 6-9, 2019
www.vlsi-soc.com

MEMOCODE'19 – ACM/IEEE Int'l Conference on Formal Methods and Models for Codesign
San Diego, CA
Oct 9-11, 2019
https://memocode.github.io/2019

MICRO'19 – IEEE/ACM Int'l Symposium on Microarchitecture
Columbus, OH
Oct 12-16, 2019
http://www.microarch.org/micro52

ESWEEK'19 - Embedded Systems Week (CASES, CODES+ISSS, and EMSOFT)
New York, NY
Oct 13-18, 2019
http://www.esweek.org

NOCS'19 – IEEE/ACM Int'l Symposium on Networks-on-Chip
New York, NY
Oct 17-18, 2019
https://www.engr.colostate.edu/nocs2019

BIOCAS'19 – Biomedical Circuits and Systems Conference
Nara, Japan
Oct 17-19, 2019
http://www.biocas2018.org

ICCAD'19 – IEEE/ACM Int'l Conference on Computer-Aided Design
Westminster, CO
Nov 4-7, 2019
http://www.iccad.com

ICPADS'19 – IEEE Int'l Conference on Parallel and Distributed Systems
Tianjin, China
Dec 4-6, 2019
http://www.icpads2019.cn

FPT'19 - Int'l Conference on Field-Programmable Technology
Tianjin, China
Dec 9-13, 2019
http://icfpt.org

ISED'19 – Int'l Symposium on Electronic System Design
Kollam, India
Dec 13-15, 2019
http://isedconf.org

iSES'19 – IEEE Int'l Symposium on Smart Electronic Systems
Rourkela, India
Dec 16-18, 2019
http://www.ieee-ises.org

HiPC'19 – IEEE Int'l Conference on High Performance Computing
Hyderabad, India
Dec 17-20, 2019
http://www.hipc.org

VLSID'20 – Embedded and VLSI Design Conference
Bengaluru, India
Jan 4-8, 2020
http://www.vlsidesignconference.org

ASP-DAC'20 - Asia and South Pacific Design Automation Conference
Beijing, China
Jan 13-16, 2020

www.aspdac.com

HiPEAC'20: Int'l Conference on High Performance Embedded Architectures & Compilers
Balogna, Italy
Jan 20-22, 2020
https://www.hipeac.net

Back to Contents

## Funding Opportunities

'''North America'''

Keck Foundation
Deadline: various
http://www.wmkeck.org/grant-programs/research/

USDA Foundational Program
Deadline: various
http://www.grants.gov/web/grants/view-opportunity.html?oppId=283836

NATO: Science for Peace and Security
Deadline: Proposals accepted anytime; reviews take place February 1 and May 15
http://www.nato.int/cps/en/natolive/87260.htm

Air Force Research Laboratory: Research Collaboration Program (BAA-RQKM-2013-0005)
http://www07.grants.gov/search/search.do?&mode=VIEW&oppId=212295

Federal Aviation Administration Grants for Aviation Research (FAA-12-01)
Deadline: open to December 2019
http://www07.grants.gov/search/search.do?&mode=VIEW&oppId=134953

AFRL RD/RV University Cooperative Agreement
Deadline: open to Nov 23, 2020
http://www.grants.gov/web/grants/view-opportunity.html?oppId=280237

NASA Fellowship Programs
Deadline: Various
http://science.nasa.gov/researchers/sara/fellowship-programs/

Natural Sciences and Engineering Research Council of Canada
Deadline: Various
http://www.nserc-crsng.gc.ca/Students-Etudiants/PG-CS/index_eng.asp

Mitacs Accelerate PhD Fellowship: Ontario Business Grants Program
http://www.mentorworks.ca/what-we-offer/government-funding/research-development/...

Collaborative Research and Development Grants
(including DND/NSERC Research Partnership Grants)
http://www.nserc-crsng.gc.ca/Professors-Professeurs/RPP-PP/CRD-RDC_eng.asp

Natural Sciences and Engineering Research Council of Canada
Deadline: Various
http://www.nserc-crsng.gc.ca/Professors-Professeurs/CFS-PCP/index_eng.asp
http://www.nserc-crsng.gc.ca/Innovate-Innover/Engage-Mobiliser_eng.asp
http://www.nserc-crsng.gc.ca/Innovate-Innover/Collaborate-Collaborer_eng.asp
Research and Development Funding for Business Innovation (Multiple Organizations)
Deadline: Various
http://www.mentorworks.ca/what-we-offer/government-funding/research-development/

'''Europe'''

Horizon 2020
Deadline: Various
http://goo.gl/geBouC

German Academic Exchange Service (DAAD)
Deadline: Various
https://www.daad.org/scholarship

German Research Foundation (DFG)
Deadline: Various
http://www.dfg.de/en

Helmholtz Association
Deadline: Various
https://www.helmholtz.de/en

Leibniz Association
Deadline: Various
http://www.leibniz-gemeinschaft.de/en/home

Leopoldina
Deadline: Various
https://www.leopoldina.org/en/about-us

Max Planck Society
Deadline: Various
https://www.mpg.de/en

Swiss National Science Foundation
Deadline: Various
http://www.snf.ch/en/

'"Asia"'

Korea:

National Research Foundation of Korea
Deadline: Various
http://www.nrf.re.kr/nrf_eng_cms/show.jsp?show_no=90&check_no=89&c_relation=0&c_...

China:

National Natural Science Foundation of China
Deadline: Various
http://www.nsfc.gov.cn/publish/portal1/

Singapore:

National Research Foundation (NRF) Singapore
http://www.nrf.gov.sg
RIE 2020 plan
Deadline: Various
http://www.nrf.gov.sg/rie2020

India:

inistry of Electronics and Information Technology
Deadline: Various
http://meity.gov.in/content/research-development

Department of Science and Technology (Nano Mission)
Deadline: Various
http://nanomission.gov.in/

University Grants Commission
Deadline: Various
http://www.ugc.ac.in/

inistry of Education Academic Research Fund
Deadline: Various
https://www.olga.moe.gov.sg/default.aspx

Agency for Science Technology and Research (A*STAR) Science and Engineering Research Council (SERC)
Deadline: Various
https://www.a-star.edu.sg/Research/Funding-Opportunities/Grants-Sponsorship.aspx

Multiple Funding
Deadline: Various
http://www.computerscienceonline.org/cs-scholarships/

'"Oceania/Polynesia"'

Polynesian Cultural Center (Hawaii) - International Student Scholarship Program
For studying in BYU Hawaii
http://www.polynesia.com/students.html

New Zealand:

inistry of Business, Innovation and Employment
Deadline: Various
http://www.mbie.govt.nz/

Australia:

Premier's Research and Industry Fund
Deadline: Various
http://www.statedevelopment.sa.gov.au/science/premiers-research-and-industry-fun...

Australian Research Council
Deadline: Various
http://www.arc.gov.au

'"South America"'

Brazil:

Coordination for the Improvement of Higher Education Personnel (CAPES)
Deadline: Various
http://www.iie.org/programs/capes#.WAu2kJMrJPM

Ministry of Science, Technology, Innovation and Communications (CNPq)
Deadline: Various
http://www.cnpq.br/

'"Africa"'

Other scholarships for African Students (list of over 30 different scholarships)
Deadline: Various
http://www.afterschoolafrica.com/12264/list-annual-engineering-scholarships-afri...

Nigeria:

Bilateral Education Agreement (BEA) Awards (both Undergraduate and Post Graduate)
Deadline: Nomination interview - March'17
http://www.fsb.gov.ng/index.php/scholarship/menu-styles/bilateral-agreement

Federal Scholarship Board
Deadline: Various
http://www.fsb.gov.ng/index.php

Back to Contents

## Call for Participation: ESWEEK 2019

EMBEDDED SYSTEMS WEEK
Call for Participation

CASES * CODES+ISSS * EMSOFT * IoMT * NoCs * Workshops * Tutorials

New York City, USA, October 13 - 18, 2019
www.esweek.org

About Embedded Systems Week (ESWEEK)

Embedded Systems Week (ESWEEK) is the premier event covering all aspects of embedded systems and software. By bringing together three leading conferences (CASES, CODES+ISSS, EMSOFT), a special track on Internet of Medical Things (IoMT), a symposium (NOCS) and several workshops and tutorials, ESWEEK allows attendees to benefit from a wide range of topics covering the state of the art in embedded systems research and development.

Preliminary program: https://esweek.org/program
Registration: https://esweek.org/registration
Advanced Registration Deadline: Sep. 20

Registered attendees are entitled to attend sessions of all conferences CASES, CODES+ISSS, EMSOFT, and the IoMT Day. Symposium, workshops, and tutorials require separate registration.

Timeline
* September 9, 2019: "Child-Care Travel Support Program" extended application deadline
* Details: https://www.esweek.org/child-care-travel-support-program
* September 14, 2019: "ESWEEK Student Travel Grant" application deadline
* Details: https://www.esweek.org/esweek-student-travel-grant
* September 20, 2019: Advanced Registration Deadline
* Registration page: https://esweek.org/registration
* October 13 - 18, 2019: ESWEEK in New York City, USA

Keynotes
* Monday Keynote: "High Performance Computing in a World of Embedded Intelligence"
https://esweek.org/event-details?id=281--102-
Speaker: Steve Keckler - NVIDIA
* Tuesday Keynote: "Health Monitoring with Machine Learning and Wireless Sensors"
https://esweek.org/event-details?id=281--103-
Speaker: Dina Katabi - MIT
* Wednesday Keynote: "Cyber-Physical-Human Systems: Opportunities and Challenges"
https://esweek.org/event-details?id=281--104-
Speaker: Pramod Khargonekar - UC Irvine

Conferences and Special Track
* CASES: International Conference on Compilers, Architecture, and Synthesis for Embedded Systems
https://esweek.org/cases/about
Program Chairs: Akash Kumar, Technical University of Dresden, DE
Partha Pande, Washington State University, US
* CODES+ISSS: International Conference on Hardware/Software Codesign and System Synthesis

https://esweek.org/codes/about
Program Chairs: Sudeep Pasricha, Colorado State University, US
Roman Lysecky, University of Arizona, US
* EMSOFT: International Conference on Embedded Software
https://esweek.org/emsoft/about
Program Chairs: Sriram Sankaranarayanan, Univ. of Colorado Boulder, US
Timothy Bourke, Inria Paris, FR
* IoMT Day: Internet of Medical Things (IoMT)
https://esweek.org/iomt/about
IoMT Chair: Insup Lee, University of Pennsylvania, US
Paul Bogdan, University of Southern California, US

Symposium
* NOCS: International Symposium on Networks-on-Chip
https://esweek.org/nocs-about

Tutorials: https://esweek.org/events/2019-10-13
* Machine Learning for Design and Optimization of Embedded Systems
* Machine Learning Security
* Industry Tutorial: The Open Source ACRN Hypervisor on an Intel Embedded Platform
* Open-Source Hardware: Heterogeneous System Integration with Embedded Scalable Platforms
* HW/SW Modeling and Performance Analysis of Heterogeneous Safety-Critical Systems
* Industry Tutorial: PYNQ: Python Productivity for Zynq
* Hardware Design in the 21st Century with the Object Oriented and Functional Language Chisel

Workshops: https://esweek.org/workshops
* Accelerating AI for Embedded Autonomy (AAIEA)
* CyberCardia Workshop on Medical CPS (CyberCardia)
* Model-Based Design of Cyber Physical Systems (CyPhy)
* Embedded Operating Systems Workshop (EWiLi)
* International Workshop on Highly Efficient Neural Processing (HENP)
* INTelligent Embedded Systems Architectures and Applications (INTESA)
* International Workshop on Rapid System Prototyping (RSP)
* Workshops on Embedded Systems Education (WESE)

ESWEEK Student Travel Grant & Child-Care Travel Support Program
ESWEEK 2019 is happy to announce the ACM SIGBED, ACM SIGDA, IEEE CEDA, and NSF ESWEEK Opportunity Programs dedicated to supporting student participation at the event. The Programs offer travel support to a limited number of students. Preference will be given to students being an author or co-author of a paper accepted at one of the conferences of
ESWEEK 2019, and be in need for financial support. Details: https://esweek.org/esweek-student-travel-grant

Would you like to attend ESWEEK, but cannot because the cost of child-care is prohibitive? SIGBED provides funds for a limited number of grants that support child care for members that would like to participate in ESWEEK but are unable to do so without this support. SIGBED provides financial assistance to subsidize a variety of child-care options. Details: https://esweek.org/child-care-travel-support-program

Organization: https://esweek.org/committees/esweek/2019
Petru Eles, Linkˆping University, SE (General Chair)
Tulika Mitra, National University of Singapore, SG (Vice General Chair)
Soonhoi Ha, Seoul National University, KR (Past Chair)

ESWEEK Local Arrangement Chairs:
Ramesh Karri, New York University, US (Conference Chair)
Siddarth Garg, New York University, US
Back to Contents

## Call for Participation: NOCS 2019

13th IEEE/ACM International Symposium on Networks-on-Chip
Oct 17-18, 2019; co-located with ESWEEK 2019, New York, NY, USA
https://www.engr.colostate.edu/nocs2019/

The International Symposium on Networks-on-Chip (NOCS) is the premier event dedicated to interdisciplinary research on on-chip, package-scale, chip-to-chip, and datacenter rack-scale communication technology, architecture, design methods, applications and systems. NOCS brings together scientists and engineers working on NoC innovations and applications from inter-related research communities, including discrete optimization and algorithms, computer architecture, networking, circuits and systems, packaging, embedded systems, and design automation.

Registration for NOCS 2019 is open at: https://esweek.org/registration

The conference program includes several keynotes, tutorials, special sessions and regular paper session with participants from industry and academia. We hope you are able to attend!
Back to Contents

## Call for Participation: MLCAD

1st ACM/IEEE Workshop on Machine Learning for CAD (MLCAD)

The workshop focuses on Machine Learning (ML) methods for all aspects of CAD and electronic system design. The predecessor of this workshop was held at the Design, Automation and Test in Europe (DATE) Conference in March 2019. The workshop is sponsored by both IEEE Council on Electronic Design Automation (CEDA) and ACM Special Interest Group on Design Automation (SIGDA). Around one third of the workshop program will consist of invited and keynote speakers from major CAD and Industrial Companies, who will present their vision on machine learning for CAD.

Workshop: September 3-4, 2019
Register and attend: http://mlcad.itec.kit.edu

General Chairs
Marilyn Wolf, Georgia Institute of Technology
Jörg Henkel, Karlsruhe Institute of Technology

Industry Chairs
Ulf Schlichtmann, TU Munich
Paul Franzon, North Carolina State U.

Program Chairs
Hussam Amrouch, Karlsruhe Institute of Technology
Bei Yu, Chinese University of Hong Kong

Finance Chair
Hai Li, Duke University

Contact: henkel@kit.edu
http://mlcad.itec.kit.edu

SPONSORS: ACM SIGDA, IEEE CEDA

Back to Contents

## Call for Papers: WOSET 2019

Second Workshop on Open-Source EDA Technology (http://woset.org).
Co-located with ICCAD 2019, Nov 7,The Westin Westminster, Westminster CO.

This one-day workshop aims to galvanize the open-source EDA movement. The workshop will bring together EDA researchers who are committed to open-source principles to share their experiences and coordinate efforts towards developing a reliable, fully open- source EDA flow. The workshop will feature presentations that overview existing open-source tools, along with sessions and posters describing future planned EDA tools. The workshop will include a panel to brainstorm the current status and future challenges for open-source EDA, and to coordinate efforts and ensure quality and

interoperability across open-source tools.

Submissions (2-4 pages):
Overview of an existing or under-development open-source EDA tool.
Overview of support infrastructure (e.g. EDA databases and design benchmarks).
Open-source cloud-based EDA tools
Position statements (e.g. critical gaps, blockers/obstacles)

Important dates:
Sept 6th 2019: submission due date.
Sept 15th 2019: notification due date.

Submission URL:
https://easychair.org/conferences/?conf=woset19

For inquires: please contact sherief_reda@brown.edu
Back to Contents

## Call for Papers: ACM TECS Special Issue on LCTES 2019

ACM Transactions on Embedded Computing Systems
Special issue on Languages Compilers Tools and Theory of Embedded Systems

The special issue in the ACM Transactions on Embedded Computing Systems will consider peer-reviewed journal versions of top papers from LCTES 2019, as well as other papers received from the open call. This special issue solicits papers presenting original work on programming languages, compilers, tools, theory, and architectures that help in overcoming these challenges. Research papers on innovative techniques are welcome, as well as experience papers on insights obtained by experimenting with real-world systems and applications. We solicit original papers on the following topics of interest related to LCTES:
Programming languages
Compilers
Tools for analysis, specification, design, and implementation
Theory and foundations of embedded systems
Novel embedded architectures
Mobile systems and IoT
Industrial case studies

IMPORTANT DATES
Open for submissions in ScholarOne Manuscripts: August 15, 2019
Closed for submissions: October 15, 2019
Results of first round of reviews: January 01, 2020
Submission of revised manuscripts: March 01, 2020
Results of second round of reviews: May 01, 2020
Publication materials due: August 15, 2020

SUBMISSION GUIDELINES: Prospective authors are invited to submit their manuscripts electronically after the "open for submissions" date, adhering to the ACM Transactions on Embedded Computing Systems guidelines (tecs.acm.org/authors.cfm). Please submit your papers through the online system (mc.manuscriptcentral.com/tecs) and be sure to select the "SI:LCTES2019" option for the paper-type. Also, please indicate that you are submitting to the Special issue on LCTES 2019 in author's cover letter. Manuscripts should not be published or currently submitted for publication elsewhere. Extended versions of conference papers (including LCTES 2019) are acceptable with at least 30% new content. Any questions on this special issue should be addressed to Aviral Shrivastava (Aviral.Shrivastava@asu.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