

1 June 2019, Vol. 49, No. 6
Online archive: <http://www.sigda.org/newsletter>

1. [SIGDA News](#)
From: Xiang Chen <shawn.xiang.chen@gmail.com>
2. [SIGDA Local Chapter News](#)
From: Yanzhi Wang <yanzhiwang@northeastern.edu>
3. [SIGDA Award](#)
From: Pingqiang Zhou <zhoupq@shanghaitech.edu.cn>
4. ["What is" Column](#)
Contributing author: Prof. Wenchao Li, Boston University <wenchao@bu.edu>
From: Wenchao Li <wenchao@bu.edu>
5. [Paper Submission Deadlines](#)
From: [Debjit Sinha](#) <debjitsinha@yahoo.com>
6. [Upcoming Symposia, Conferences and Workshops](#)
From: [Debjit Sinha](#) <debjitsinha@yahoo.com>
7. [Funding Opportunities](#)
From: Jayita Das <jayita.365@gmail.com>
8. ["Researcher spotlight" Column](#)
From: Rajsaktish Sankaranarayanan <rajsaktish@gmail.com>
9. [Call for Papers: AAEAI 2019](#)
From: Luca Carloni <luca@cs.columbia.edu>
10. [Call for Papers: ACM TODAES Special Issue "Machine Learning for CAD"](#)
From: Joerg Henkel <henkel@kit.edu>
11. [Call for Papers: HiPC 2019](#)
From: Partha Pratim Pande <pande@wsu.edu>
12. [Call for Papers: ESWEEK 2019](#)
From: Lars Bauer <lars.bauer@kit.edu>
13. [Notice to Authors](#)

Comments from the Editors

Dear ACM/SIGDA member,

We are excited to present to you the June 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.

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

"[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) "MIPS R6 Architecture Now Available for Open Use"

[\[https://www.eetimes.com/document.asp?doc_id=1334489\]](https://www.eetimes.com/document.asp?doc_id=1334489)

MIPS 32-bit and 64-bit architecture – the most recent version, release 6 – will become available Thursday (March 28) for anyone to download at MIPS Open web page.

(2) "Samsung Ready with 5-nm EUV"

[\[https://www.eetimes.com/document.asp?doc_id=1334567\]](https://www.eetimes.com/document.asp?doc_id=1334567)

Samsung announced that it has completed work and is taking orders for a 5-nm foundry process using extreme ultraviolet lithography. It will offer 25% greater density and either 10% more performance or 20% less power consumption than its 7-nm node with EUV announced in October.

(3) "GPUs Holding Back AI Innovation"

[\[https://www.eetimes.com/document.asp?doc_id=1334579\]](https://www.eetimes.com/document.asp?doc_id=1334579)

GPUs are widely used to accelerate AI computing, but are the limitations of GPU technology slowing down innovation in the development of neural networks?

(4) "5G Needs More Memory to Compute"

[\[https://www.eetimes.com/document.asp?doc_id=1334512\]](https://www.eetimes.com/document.asp?doc_id=1334512)

Today's cellphone networks aren't your dad's cellphone networks. In fact, 5G not only represents a vast leap in communications compared to the flip phone days of 3G, it's also going to be more memory hungry.

(5) "Qualcomm Targets AI Inferencing in the Cloud"
[https://www.eetimes.com/document.asp?doc_id=1334548]

Qualcomm is taking another stab at the server, throwing its hat into what is rapidly becoming a crowded ring — AI inference processing for the data center.

(6) "AI Research Targets Nvidia, Mobile"
[https://www.eetimes.com/document.asp?doc_id=1334533]

A researcher from the University of Texas at Austin described a chip for training deep neural networks that he said can outperform an Nvidia V100 — even using low-cost mobile DRAM. At the same event, Arm discussed research on a chip that can significantly increase efficiency for computer vision jobs run on mobile systems.

(7) "AI Inferencing Chip Targets Edge Servers"
[https://www.eetimes.com/document.asp?doc_id=1334550]

Flex Logix — best known as a supplier of embedded FPGAs and other IP — is marketing its first chip, an edge inference co-processor targeting edge applications.

(8) "Intel & Lockheed Martin Collaborate to Launch New Hardened Security Solution"
[https://www.eetimes.com/document.asp?doc_id=1334582]

Lockheed Martin has collaborated with chip maker Intel to deliver a hardened security solution based on Intel's second-generation Intel Xeon Scalable processors that deliver more sophistication than software-only solutions. The solution, which is designed for use in the data center, aims to address the ever-increasing complexity of cyber threats while providing more consistent service performance.

(9) "Apple, Q'comm Deal Leaves Open Questions"
[https://www.eetimes.com/document.asp?doc_id=1334570]

Apple and Qualcomm announced that they have dropped all litigation in a multi-billion-dollar patent licensing dispute that has dragged on for more than two years across courts in China, Germany, and the U.S. However, the sketchy details of the settlement leave many unanswered questions.

(10) "AVs Need New Compute Platforms"
[https://www.eetimes.com/document.asp?doc_id=1334583]

While autonomous vehicle (AV) development dominates discussion in the electronics industry, firms are a long way off from having the right computing platforms, safety certification standards and regulations, says a new survey highlighting current thinking on exactly where we are on the path to autonomy.

[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

GLSVLSI 2019: The 29th the ACM Great Lakes Symposium on VLSI, <http://www.glsvlsi.org>

Best Paper Awards

Technical Track: "Crash Skipping: A Minimal-Cost Framework for Efficient Error Recovery in Approximate Computing

"What is" Column

What is Specification Mining?

Wenchao Li, Boston University

Specification is arguably the most important step in formal verification or assertion-based verification. It is a critical step in the design process with the goal to capture precisely the intended behavior of the design. However, manually creating quality specification can be a very time-consuming and laborious task. Similar to the question of "have I created enough tests?" in testing, verification in practice is often plagued by the question of "have I written enough specifications?" Many real-world experiences indicate that poor or the lack of sufficient specifications can easily lead to misses of critical bugs, and in turn design re-spins and time-to-market slips.

Specification mining is the idea that one can automate the process of specification generation either statically or by observing a design's execution. Examples of the former typically involve static analysis of source code, with some recent attempts that leverage natural language processing techniques to extract formal specification directly from design documents. One drawback of static approaches to specification mining is that it does not reflect that actual intended use of the program or the design. Dynamic approaches, on the other hand, mine specifications by observing simulation or execution traces. It is worth noting that these data-driven approaches often assume that what is commonly seen is likely correct. Thus, the mined specifications are only likely, and will need to be further validated by the designer or an automated verifier.

The type of specification plays a crucial role in the design of specification mining algorithms. Specifications can be largely categorized into automata, single-state predicates and temporal logic formulas. In the first case, the idea dates back to E. Mark Gold in 1978 on learning finite-state automata from traces. Logic is an important foundation for formal specification. Linear Temporal Logic (LTL) has been quite popular due to its expressiveness and being relatively user friendly, and forms the basis of Property Specification Language (now a IEEE standard). Many variants of temporal logic exist, and an important extension of LTL is Signal Temporal Logic, which allows for real-time and real-valued constraints. It also admits a non-Boolean semantics (the so-called robustness value) that can quantify how well a property is satisfied. This property turns out to be particularly useful for cyber-physical systems.

Paper Submission Deadlines

MEMOCODE'19 – ACM/IEEE Int'l Conference on Formal Methods and Models for Codesign

San Diego, CA

Deadline: Jun 7, 2019 (Abstracts due: May 31, 2019)

Oct 9-11, 2019

<https://memocode.github.io/2019>

HiPC'19 – IEEE Int'l Conference on High Performance Computing

Hyderabad, India

Deadline: Jun 14, 2019 (Abstracts due: Jun 7, 2019)

Dec 17-20, 2019

<http://www.hipc.org>

iSES'19 – IEEE Int'l Symposium on Smart Electronic Systems

Rourkela, India

Deadline: Jul 19, 2019

Dec 16-18, 2019

<http://www.ieee-ises.org>

FPT'19 - Int'l Conference on Field-Programmable Technology

Tianjin, China
Deadline: Jul 22, 2019 (Abstracts due: Jul 15, 2019)
Dec 9-13, 2019
<http://icfpt.org>

[Back to Contents](#)

Upcoming Symposia, Conferences and Workshops

DAC'19 – Design Automation Conference
Las Vegas, NV
Jun 2-6, 2019
<http://www.dac.com/>

SLIP'19 – ACM/IEEE International Workshop on System-Level Interconnect Prediction [co-located with DAC'19]
Las Vegas, NV
Jun 2, 2019
<http://sliponline.org>

LCTES'19 – Int'l Conference on Languages Compilers, Tools and Theory of Embedded Systems
Phoenix, AZ
Jun 22, 2019
<https://conf.researchr.org/home/LCTES-2019>

ISCA'19 – Int'l Symposium on Computer Architecture
Phoenix, AZ
Jun 22-26, 2019
<https://iscaconf.org>

ICDCS'19 – IEEE Int'l Conference on Distributed Computing Systems
Dallas, TX
Jul 7-10, 2019
<https://theory.utdallas.edu/ICDCS2019>

IWBDA'19 - Int'l Workshop on Bio-Design Automation
Cambridge, England
Jul 9-12, 2019
<http://www.iwbdaconf.org/2019>

ISVLSI'19 – IEEE Computer Society Annual Symposium on VLSI
Miami, FL
Jul 15-17, 2019
<http://www.isvlsi.org>

AHS'19 - NASA/ESA Conference on Adaptive Hardware and Systems
Colchester, UK
Jul 22-24, 2019
<http://www.ahs-conf.org>

ISLPED'19 – ACM/IEEE Int'l Symposium on Low Power Electronics and Design
Lausanne, Switzerland
Jul 29-31, 2019
<http://www.islped.org>

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

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>

AAIEA'19 - Accelerating AI for Embedded Autonomy

New York, NY

October 17, 2019

<https://alessandro-pinto.github.io/aaiea2019/>

[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

itacs 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.nsf.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:

Ministry 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/>

Ministry 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:

Ministry 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](#)

"Researcher spotlight" Column

Hello readers,

In this edition of the Researcher Spotlight column, we meet Prof. Yuanqing Cheng. He leads the Computer Architecture / Aided Design for Emerging Technologies (CADET) Lab at Beihang University, Beijing, China. He received the B.S. degree from Xidian University, Xi'an, China and M.S. degree from Harbin Institute of Technology, Harbin, China and Ph.D. degree from Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China. From 2012 to 2013, Dr. Cheng was a Post-doctoral Researcher in LIRMM, Montpellier, France. Then, Dr. Cheng joined the School of Electronic and Information Engineering, Beihang University as Assistant Professor. From 2015 to 2016, Dr. Cheng worked in University of California, Santa Barbara, CA, USA as a Visiting Scholar. Dr. Cheng is a CCF Senior Member of CCF and Member of IEEE and ACM. Excerpts from a recent conversation below:

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

Sure. My research topics include reliability aware physical design of three-dimensional integrated circuits, low power and reliable computer architecture design for emerging memory technologies (especially STT-RAMs based, and Carbon Nanotube based).

2. Thermal constrained task allocation for MPSoCs: Task allocation seems an interesting approach to reduce overall interconnect energy. Besides peak performance, if a range of performance is considered like say, Dynamic Frequency Scaling scenarios, how would task allocation have to be modified? Do you foresee a break-even point?

In fact, my focus is on the thermal aware task schedule synthesis for 3D MPSoCs. As you know, 3D IC has a severe thermal problem as more power-hungry dies are stacked vertically together. Considering the DVFS technique, it is

commonly used in low power processor design. In this scenario, a larger fraction of total power budget is expected to be consumed on interconnects due to data transfers. Therefore, in my opinion, it is more necessary to optimize the data communication energy with the high density vertical TSV interconnects available.

Furthermore, with the DVFS technique, some cores may switch between high voltage/frequency and low voltage/frequency modes frequently. It requires more adaptive task scheduling strategy to be aware of the dynamic thermal and power distribution and make a timely response on task reallocation to guarantee both the realtime and thermal constraints.

Nonetheless, as 3D IC technology evolves further towards monolithic 3D technique leveraging inter-layer dielectrics (ILDs), it would provide new opportunities for high performance 3D MPSoC design, and interconnect energy can be reduced further due to more vertical interconnects available on-chip.

3. TSV Test Wrapper Design: Tester time is expensive and is often a direct function of design complexity which needs to be tested. In methodologies that optimize the TSV, it appears there is a trade-off between reduced observability (of circuit functionality) and controllability (wrapper design). How do you envision optimality between these opposing goals and ensure functional bug escapes do not happen?

Testing and pre-testability is more important for 3D ICs. Considering the TSV-based 3D IC technology, different tiers are bonded together with TSVs for chip stacking. Except for testing the functionality of each tier, it also requires more thorough post-bond (3D assembly of dies) testing including block functionality, TSV and planar functionality, etc. It seems that high stacking integration density may reduce observability of critical signals. However, the observability can be improved through pre-bond (individual dies) testing with some pre-bond DFT support. The pre-bond DFT technique can make each tier partially functional, and each tier can be tested independently with better observability. Then, post-bond (3D assembly of dies) test can be used to test TSVs and the tier-partitioned functional blocks with test wrapper design. A good pre-bond DFT design can reduce the complexity of post-bond wrapper design effectively.

As the 3D IC moves to monolithic 3D era, it imposes more severe challenges on DFT and testing, and also calls for more research efforts to deal with ultra-high vertical integration density and different circuit partition methodologies.

4. STT-RAMs seem to offer distinctive levels of performance, power, area and reliability in comparison with other memory technologies. Can you share your thoughts on its architectural scope with non-CPU architectures, say GPUs, FPGAs?

STT-RAM has fast access speed, high integration density, ultra-low power and unlimited endurance. Thus, it is very suitable to replace SRAM as the on-chip last level cache. Moreover, it also has potential applications in FPGAs and GPUs. FPGAs are essentially based on SRAM technology, which is used to build up look up table, switching fabrics and on-chip memory. STT-RAMs can be used to construct these components with higher integration density. As the working speed of FPGAs mainly depends on the latency of the critical path, if the computing logic and storage can be put closer due to reduced area overhead, higher performance can be approached. Therefore, STT-RAMs may play an important role in improving working frequencies and power consumptions of FPGAs.

On the other hand, compared to the conventional general processor architecture, GPUs can approach much higher parallelism to achieve significant performance improvements. It has large volumes of register files and on-chip cache, which may also be suitable for STT-RAMs. Like Last Level Cache (LLC) design of CPU, we can use STT-RAM to build register files and on-chip caches of GPUs. However, considering the larger write latency of STT-RAM, it is crucial to prevent it from hurting performance. So, a combination of SRAM and STT-RAM to build register file and on-chip cache is very promising. To make it effective, it requires high level compiling technology to optimize the data layout among different memory regions.

In summary, the unique characteristics of STT-RAM provide brand new opportunities for GPUs and FPGAs to support a wide range of emerging applications.

5. Can you share your thoughts on reliable physical design in 3D integration technology for 3D ICs?

In the early stage of 3D integration technology, fabrication process was evolving. TSV-related defects such as TSV misalignment, TSV void and bonding interface contamination. Moreover, the power density increases remarkably due to the stacking structure. So, the current density may introduce severe electro-migration on defective TSVs. This forms the motivation of our research work. The good news is that with the progressive development of 3D IC fabrication process, TSV related yield loss has been reduced largely. As a result, many 3D based products emerge in the commercial market in recent years, such as 3D NAND Flash, HBM and HMC techniques.

Note that physical design and reliable design for 3D ICs are still indispensable technologies for its commercial success. Especially in the post-Moore era, 3D integration is a very promising technology trend enabling further scaling. As TSV-based 3D IC technique evolves into monolithic 3D IC (M3D IC), the physical design tools will play a more important role to improve IC design productivity. Considering the ultra-high density of vertical interconnects in M3D IC, it enables gate-level even transistor-level circuit partitioning within the 3D IC stack, which is impractical for TSV-based 3D ICs. New physical design methodologies are imperative to exploit huge volume of vertical MIV interconnects to improve performance and power consumption compared to 2D and TSV-based 3D ICs. The fine grain partitioning methodology

requires new P&R tools DFT, and reliable design techniques to pave the way for the success of M3D.

[Back to Contents](#)

Call for Papers: AAIAI 2019

Accelerating AI for Embedded Autonomy (AAIEA)

October 17, 2019, New York, NY

Co-located with Embedded Systems Week

The Workshop on Accelerating Artificial Intelligence for Embedded Autonomy aims at gathering researchers and practitioners in the fields of autonomy, automated reasoning, planning algorithms, and embedded systems to discuss the development of novel hardware architectures that can accelerate the wide variety of AI algorithms demanded by advanced autonomous and intelligent systems. Topics of interest include hardware architectures and design methodologies to accelerate:

Applications based on deep learning

Skill-level and instinctive autonomy based on deep reinforcement learning

Storage and retrieval of facts in knowledge bases

Logical reasoning methods such as deduction

Search for classical planning algorithms and Hierarchical Task Networks (HTN)

Inference in probabilistic models such as Bayesian networks and probabilistic logic

Planning algorithms for Markov Decision Processes (MDP) and Partial Observable Markov Decision Processes (POMDP)

The workshop is not limited to any particular hardware implementation platform. ASICs, FPGAs, and analog circuits, as well as heterogeneous platforms that combine hardware accelerators with CPUs and GPUs are all within the scope of the proposal. Similarly, emerging architectures such as those based on neuromorphic computing and so-called non Von Neumann machines are considered in scope.

For more information please refer to <https://alessandro-pinto.github.io/aaiea2019/>

Important Dates

Submission deadline: July 1, 2019

Author notification: July 31, 2019

Camera ready: August 30, 2019

Submissions

Submitted manuscripts should be written in English conforming to the IEEE conference proceedings format (8.5" x 11" page size, two-column) and not longer than 8 pages. All paper submissions must represent original and unpublished work. Papers must be submitted as PDF files through EasyChair (<https://easychair.org/conferences/?conf=aaiea19>).

Venue

The AAIEA Workshop is part of the Embedded Systems Week, which will be hosted at the Kimmel Center for University Life, in New York City. For detailed information on this venue, please refer to the conference website (<https://www.esweek.org/>).

Technical Program Committee

Luca Carloni, Columbia University

Luis Ceze, University of Washington

Aleksandra Faust, Google

Alessandro Pinto, United Technologies Research Center

Brandon Reagen, Facebook

Vijay Reddi, Harvard

Gerald Wang, United Technologies Research Center

Zhiru Zhang, Cornell

Yuhao Zhu, University of Rochester

Organizers

Alessandro Pinto, United Technologies Research Center

Gerald Wang, United Technologies Research Center

Luca Carloni, Columbia University

Call for Papers: ACM TODAES Special Issue "Machine Learning for CAD"

ACM TODAES:

*** Special Issue on Machine Learning for CAD (ML-CAD)***

This Special Issue focuses on machine learning methods for all aspects of CAD for VLSI and electronic system design. Advances in machine learning (ML) over the past half-dozen years have revolutionized the effectiveness of ML for a variety of applications. However, design processes present challenges that require parallel advances in ML and CAD as compared to traditional ML applications such as image classification. This Special Issue focuses on both design time and run-time techniques.

Topics of interest to this Special Issue include but are not limited to:

- * ML approaches to logic design.
- * ML for physical design.
- * ML for analog design.
- * ML methods for design space exploration in CAD
- * ML methods for design constraints like low power, reliability etc..
- * Labeled and unlabeled data in ML for CAD.
- * ML for low power, thermal and other management techniques.
- * ML for multicore and manycore architectural design.
- * ML for Design Technology Co-Optimization (DTCO)

Deadline for submission: *** June 15th 2019 ***

Notification first review round: July 31st 2019

Revisions due: September 15th 2019

Final Notification: Oct 31st, 2019

Submission instructions: <https://todaes.acm.org>

Upon submission: On the first page of the submitted work please mark clearly: "Submitted to the Special Issue on Machine Learning for CAD (ML-CAD)". (All other author guidelines for this special issue are identical to the author guideline for regular submissions to ACM TODAES).

Submit to: <https://mc.manuscriptcentral.com/todaes>

When submitting, please check "Yes" in response to the question "Is this manuscript a candidate for a special issue?".

Guest Editors:

Marilyn Wolf, Georgia Tech

Jörg Henkel, Karlsruhe Institute of Technology

Hussam Amrouch, Karlsruhe Institute of Technology

Contact: henkel@kit.edu

Call for Papers: HiPC 2019

26th IEEE International Conference on High Performance Computing, Data, and Analytics

December 17-20, 2019, Hyderabad, India

<https://hipc.org/>

HiPC 2019 will be the 26th edition of the IEEE International Conference on High Performance Computing, Data, Analytics and Data Science. HiPC serves as a forum to present current work by researchers from around the world as well as highlight activities in Asia in the areas of high performance computing and data science. The meeting focuses on all aspects of high performance computing systems, and data science and analytics, and their scientific, engineering, and

commercial applications.

Authors are invited to submit original unpublished research manuscripts that demonstrate current research in all areas of high performance computing, and data science and analytics, covering all traditional areas and emerging topics including from machine learning, big data analytics and blockchain. Each submission should be submitted to one of the tracks listed under the two broad themes of High Performance Computing and Data Science.

High Performance Computing tracks:

Algorithms: This track invites papers that describe original research on developing new parallel and distributed computing algorithms, and related advances. Examples of topics that are of interest include (but not limited to):

- * New parallel and distributed algorithms and design techniques;
- * Advances in enhancing algorithmic properties or providing guarantees (e.g., fault tolerance, resilience, concurrency, data locality, communication-avoiding);
- * Classical and emerging computation models (e.g., parallel/distributed models, quantum computing, neuromorphic and other bioinspired models);
- * Provably efficient parallel and distributed algorithms for advanced scientific computing and irregular applications (e.g., numerical linear algebra, graph algorithms, computational biology); and
- * Algorithmic techniques for resource allocation and optimization (e.g., scheduling, load balancing, resource management);

Architectures: This track invites papers that describe original research on the design and evaluation of high performance computing architectures, and related advances. Examples of topics of interest include (but not limited to):

- * Design and evaluation of high performance processing architectures (e.g., reconfigurable, system-on-chip, manycores, vector processors);
- * Design and evaluation of networks for high performance computing platforms (e.g., interconnect topologies, network-on-chip);
- * Design and evaluation of memory, cache and storage architectures (e.g., 3D, photonic, Processing-In-Memory, NVRAM, burst buffers, parallel I/O);
- * Approaches to improve architectural properties (e.g., energy/power efficiency, reconfigurable, resilience/fault tolerance, security/privacy); and
- * Emerging computational architectures (e.g., quantum computing, neuromorphic and other bioinspired architectures).

Applications: This track invites papers that describe original research on the design and implementation of scalable applications for execution on parallel and distributed platforms, and related advances. Examples of topics of interest include (but not limited to):

- * Design and implementation of shared and distributed memory parallel applications (e.g., scientific computing and industry applications, emerging applications in IoT and life sciences - biology, medicine, chemistry, etc.);
- * Design and simulation methodologies for scaling applications on peta- and exascale platforms (e.g., co-design approaches, hardware/software co-design, heterogeneous and hybrid programming);
- * Hardware acceleration of parallel applications (e.g., CPU/GPUs, multi-GPU clusters, FPGA, vector processors, manycore); and
- * Design of application benchmarks for parallel and distributed platforms.

Systems Software: This track invites papers that describe original research on the design, implementation and evaluation of systems software for high performance computing platforms, and related advances. Examples of topics of interest include (but not limited to):

- * Scalable systems and software architectures for high performance computing (e.g., middleware, operating systems, I/O services);
- * Techniques to enhance parallel performance (e.g., compiler/runtime optimization, learning from application traces, profiling);
- * Techniques to enhance parallel application development and productivity (e.g., Domain-Specific Languages, programming environments, performance/correctness checking and debugging);
- * Techniques to deal with uncertainties, hardware/software resilience, and fault tolerance;
- * Software for cloud, data center, and exascale platforms (e.g., middleware tools, schedulers, resource allocation, data migration, load balancing); and
- * Software and programming paradigms for heterogeneous platforms (e.g., libraries for CPU/GPU, multi-GPU clusters, and other accelerator platforms);

Data Science tracks:

Scalable Algorithms and Analytics: This track invites papers that describe original research on developing scalable algorithms for data analysis at scale, and related advances. Examples of topics of interest include (but not limited to):

- * New scalable algorithms for fundamental data analysis tasks (supervised, unsupervised learning, and pattern discovery);
- * Scalable algorithms that are designed to address the characteristics of different data sources and settings (e.g., graphs, social networks, sequences, data streams);
- * Scalable algorithms and techniques to reduce complexity of large-scale data (e.g., streaming, sublinear data structures, summarization, compressive analytics);
- * Scalable algorithms that are designed to address requirements in different data-driven application domains (e.g., life sciences, business, agriculture); and
- * Scalable algorithms that ensure the transparency and fairness of the analysis.

Scalable Systems and Software: This track invites papers that describe original research on developing scalable systems and software for handling data at scale, and related advances. Examples of topics of interest include (but not limited to):

- * Design of scalable system software to support various applications (e.g., recommendation systems, web search, crowdsourcing applications, streaming applications)
- * Design of scalable system software for various architectures (e.g., OpenPower, GPUs, FPGAs).
- * Architectures and systems software to support various operations in large data frameworks (e.g., storage, retrieval, automated workflows, data organization, visualization, visual analytics, human-in-the-loop);
- * Design and implementation of systems software for distributed data frameworks (e.g., distributed file system, virtualization, cloud services, resource optimization, scheduling); and
- * Standards and protocols for enhancing various aspects of data analytics (e.g., open data standards, privacy preserving and secure schemes).

One or more best paper awards will be given for outstanding contributed papers.

IMPORTANT DATES

Abstract Submission: Friday, June 7, 2019

Paper Submission: Friday, June 14, 2019

Reviews for Rebuttals: Monday, July 29, 2019

Rebuttals due: Monday, August 5, 2019

Initial Submission Decision: Monday, August 19, 2019

Revisions Due: Friday, September 20, 2019

Author Notification: Monday, September 30, 2019

Camera Ready: Monday, October 14, 2019

EasyChair Submission link: <https://easychair.org/conferences/?conf=hipc2019>

Manuscript Guidelines

Submitted manuscripts should be structured as technical papers and may not exceed ten (10) single-spaced double-column pages using 10-point size font on 8.5x11 inch pages (IEEE conference style), including figures, tables, and references. See IEEE style templates at this page for details.

Journal Special Issue: Authors of selected high quality papers in HiPC 2019 will be invited to submit extended version of their papers for possible publication in a special issue of Journal of Parallel and Distributed Computing.

PROGRAM CHAIRS

Ananth Kalyanaraman, Washington State University, USA

George Karypis, University of Minnesota, USA

PROGRAM VICE-CHAIRS

HPC Tracks:

Algorithms: Bora Uçar, CNRS and École normale supérieure de Lyon, France

Applications: Alba Cristina M.A. de Melo, University of Brasilia, Brazil

Architecture: Smruti Ranjan Sarangi, Indian Institute of Technology Delhi, India

System Software: Sriram Krishnamoorthy, Pacific Northwest National Laboratory, USA

Data Science Tracks:

Scalable Algorithms and Analytics: Srinivasan Parthasarathy, Ohio State University, USA

Scalable Systems and Software: Gagan Agrawal, Ohio State University, USA

Contact Information

* High Performance Computing tracks:

Ananth Kalyanaraman, Washington State University, USA, ananth@wsu.edu

* Data Science tracks:

Call for Papers: ESWEEK 2019

EMBEDDED SYSTEMS WEEK

Call for Papers, Workshops, Tutorials, and Special Sessions

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

New York City, USA, October 13 - 18, 2019

www.esweek.org

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 IoMT Day, a symposium (NOCS), and hot-topic workshops and tutorials, ESWEEK presents attendee a wide range of topics unveiling the state of the art in embedded systems design and HW/SW architectures.

https://esweek.org/sites/default/files/ESWEEK2019_CfP_1.pdf

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

Journal Track:

- Abstract Submission: April 5, 2019
- Full Paper Submission: April 12, 2019 (firm)
- Notification of Acceptance: July 10, 2019

Work-in-Progress Track:

- Paper Submission: June 7, 2019 (firm)
- Notification of Acceptance: July 10, 2019

Workshop Proposals: April 18, 2019

Tutorial Proposals: April 30, 2019

Special Session Proposals: April 30, 2019

Paper Process

ESWEEK 2019 continues a dual publication model comprising the Journal track and the Work-in-Progress (WiP) track. Journal track papers, which are full-length (10-page) papers describing mature work, will be published in the ACM Transactions on Embedded Computing Systems (TECS). The WiP track papers, which are short (2-page) papers representing not- yet-mature but promising research, will be published in the ESWEEK proceedings and will be listed as regular publications within the IEEE and/or ACM digital libraries. Authors of WiP papers have the opportunity to publish the extended form of their work in any conference or journal they prefer. Journal and WiP papers are mutually exclusive, i.e., a work can only be in submission in one of the two tracks. For more information, please refer to:

<https://www.esweek.org/author-information>

CASES: International Conference on Compilers, Architectures, and Synthesis for Embedded Systems

CASES is a premier forum where researchers, developers and practitioners exchange information on the latest advances in compilers and architectures for high-performance, low-power embedded systems. The conference has a long tradition of showcasing leading edge research in embedded processor, memory, interconnect, storage architectures and related compiler techniques targeting performance, power, predictability, security, reliability issues for traditional and emerging applications. In addition, we invite innovative papers that address design, synthesis, and optimization in heterogeneous and accelerator-rich architectures.

https://esweek.org/sites/default/files/2019-CASES-cfp_0.pdf

CASES 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

The International Conference on Hardware/Software Codesign and System Synthesis is the premier event in system-level design, modeling, analysis, and implementation of modern embedded and cyber-physical systems, from system-level

specification and optimization down to system synthesis of multi-processor hardware/software implementations. The conference is a forum bringing together academic research and industrial practice for all aspects related to system-level and hardware/software co-design. High-quality original papers will be accepted for oral presentation followed by interactive poster sessions. <https://esweek.org/sites/default/files/CODES%2BBISS-2019-CFP.pdf>

CODES+ISSS Program Chairs:

[Sudeep Pasricha](#), Colorado State University, US

Roman Lysecky, Arizona State University, US

EMSOFT: International Conference on Embedded Software

The ACM SIGBED International Conference on Embedded Software (EMSOFT) brings together researchers and developers from academia, industry, and government to advance the science, engineering, and technology of embedded software development. Since 2001, EMSOFT has been the premier venue for cutting-edge research in the design and analysis of software that interacts with physical processes, with a long-standing tradition for results on cyber-physical systems, which compose computation, networking, and physical dynamics.

https://esweek.org/sites/default/files/EMSOFT2019_CfP_0.pdf

EMSOFT Program Chairs:

Sriram Sankaranarayanan, University of Colorado Boulder, US

Timothy Bourke, Inria Paris, FR

IoMT: Internet of Medical Things

The Internet of Medical Things (IoMT) paves the foundations for intelligent and reliable personalized precision medicine. Grounded in the mathematical and physical modeling of human anatomy and physiology, it offers accurate multiscale medical monitoring through smart sensing, enabling continuous diagnosis via on-fly communication with medical experts. It provides hyperspectral and hyperdimensional processing and restores health through patient-specific actuation. The IoMT special day provides a forum for academic and industry representatives from areas such as medical and bio-engineering and embedded systems to discuss innovative ideas and solutions for precise personalized medicine. Sub-missions to the IoMT day are via the three conferences. <https://esweek.org/iomt/about>

IoMT Chairs:

Insup Lee, University of Pennsylvania, US

Paul Bogdan, University of Southern California, US

Call for Workshop Proposals

ESWEEK 2019 will host several workshops on Oct. 17/18th. ESWEEK workshops are excellent opportunities to bring together researchers and practitioners from different communities to share their experiences in an interactive atmosphere. We invite you to submit workshop proposals on any topic related to the broad set of research, education, and application areas in embedded systems before the deadline of April 18, 2019.

https://esweek.org/sites/default/files/ESWEEK2019_CfP_1.pdf

Workshop Chair:

Laura Pozzi, USI Lugano, CH

Call for Tutorial Proposals

ESWEEK 2019 is looking for high-quality tutorials that will take place on Oct. 13th. Tutorials offer a unique opportunity where presenters can interact with attendees and attendees can gain in-depth knowledge on specific topics. Tutorials on all topics related to embedded system design, analysis and development are welcome and can be either half or full day, lecture style or hands on. We invite you to submit tutorial proposals before the deadline of April 30, 2019.

Call for Special Session Proposals

ESWEEK 2019 will host several special sessions. They should cover hot, contemporary topics that are complementary to regular sessions and can constitute individual presentations, panels or other formats. Participants of each accepted special session will have the opportunity to co-author an overview paper (maximum 10 pages) of the session, published in the ESWEEK proceedings. We invite you to submit special session proposals on any topic related to the broad areas of interest of the conference or beyond before the deadline of April 30, 2019.

https://esweek.org/sites/default/files/ESWEEK2019_CfP_1.pdf

Tutorials and Special Sessions Chair:

Andreas Gerstlauer, University of Texas Austin, US

Organization

ESWEEK 2019 General Chairs:

Petru Eles, Link[^]ping University, SE

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

Siddarth Garg, New York University, US

[Back to Contents](#)

Notice to Authors

By submitting your article for distribution in this Special Interest Group publication, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights: to publish in print on condition of acceptance by the editor; to digitize and post your article in the electronic version of this publication; to include the article in the ACM Digital Library and in any Digital Library related services; and to allow users to make a personal copy of the article for noncommercial, educational or research purposes. However, as a contributing author, you retain copyright to your article and ACM will refer requests for republication directly to you.

This newsletter is a free service for current SIGDA members and is added automatically with a new SIGDA membership.

Circulation: 2,700

This ACM/SIGDA E-NEWSLETTER is being sent to all persons on the ACM/SIGDA mailing list. To unsubscribe, send an email to listserv@listserv.acm.org with "signoff sigda-announce" (no quotes) in the body of the message. Please make sure to send your request from the same email as the one by which you are subscribed to the list.