Special Interest Group on Design Automation
ACM/SIGDA E-NEWSLETTER, Vol. 55, No. 8

SIGDA - The Resource for EDA Professionals

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

Online archive: https://www.sigda.org/publications/newsletter

SIGDA News

1. **TSMC to Top $100bn as it Plans 11 Fabs in Taiwan**
TSMC to build 11 fabs in Taiwan to meet demand for AI chips as it ramps up production in Arizona and Japan. Demand for AI chips is driving a 30% increase in revenue this year, enabling the company to break through the US$100 billion.

2. **World's First AI Photonic Processor Fires Up**
German photonic processor developer Q.ANT has delivered its Native Processing Server (NPS) to the Leibniz Supercomputing Centre (LRZ), marking the first integration of an analog photonic co-processor into an operational high-performance computing (HPC) environment.

3. **Europe Prepares its Quantum Act**
The European Commission is preparing a Quantum Act to boost the development of quantum technologies, although not in the same way as the EU CHIPS Act for semiconductors. The Act is being prepared for 2026 and follows the publication of the Quantum Strategy to make Europe a leader in technology by 2030. This will see six pilot lines and a Quantum design centre.

4. **UK Switches on Its AI Supercomputer**
The UK Secretary of State for Science, Innovation and Technology, Peter Kyle, has switched on the second phase of the UK's AI supercomputer, Isambard AI. The machine now has over 5400 GPUs and is by far the largest machine in the UK and ranks sixth in Europe (and 11th in the world).

5. **Europe's RISC-V Processor Developer Up for Sale**
The board of directors at Codasip in Germany has put the RISC-V processor developer up for sale. The company, under CEO Ron Black, is looking to sell itself in the next three months. It has developed tools to produce processor cores using the open RISC-V instruction set and has funding of up to €380m with various

Message from the EiC

Dear Readers,

In this edition, we bring you the latest news and activities in our community, upcoming conferences, paper deadlines, an insightful article on What is Processing-in-Memory (PiM) for Efficient AI, and job openings worldwide.

Please do not hesitate to write to us if you want to contribute articles and announcements or share your thoughts and feedback.

Sandeep Chandran,
Editor-in-Chief,
SIGDA e-Newsletter

equity and grants that include follow-on project funding. There are currently 250 staff members, with 57% in hardware and 30% in software.

6. Korea Joins Horizon Europe R&D Program

The Republic of Korea is the first Asian country to join and help fund the Horizon Europe research and innovation programme. After a trial period that started in January 2025, Korean researchers and companies can now participate in several areas of the programme on equal terms with those for EU Member States.

7. China Offers Olive Branch on Rare Earth Elements

The Chinese government is offering a 'fast track' approval process to European companies for rare earth elements used in magnets for electric vehicles and renewable energy. The move would help reduce supply chain issues that have hit European companies amid the US-China trade war.

8. Mitsubishi Invests in Things for AI-Assisted PLM Systems

Mitsubishi Electric Corporation, through its ME Innovation Fund, has invested in Things, Inc., a Japan-based startup that develops and provides AI-assisted product lifecycle management (PLM) systems for manufacturing, specializing in comprehensive document management from product planning to development to disposal.

9. Quantum Dot Lasers for Silicon Chiplets

Researchers in the US have successfully integrated indium arsenide quantum dot (QD) lasers monolithically on silicon photonics chiplets. Integrating lasers onto silicon is a key technique for reducing the power consumption of AI chips in data centers as interconnect speeds rise. Nvidia, for example, is working with TSMC on a co-packaged optics (CPO) silicon photonics system with 1.6 Tbits/s performance based on micro ring resonators for the chiplets.

10. First LPDDR6/5X 14.4-Gbps Memory IP Targets AI Infrastructure

The new Cadence LPDDR6/5X memory IP system is a key enabler for scaling up AI infrastructure to accommodate the memory bandwidth and capacity demands of next-generation AI LLMs, agentic AI, and other compute-intensive workloads across various verticals. Multiple engagements are currently underway with leading AI, high-performance computing (HPC), and data centre customers.

SIGDA Awards

2025 BEST PAPER NOMINEES @ GLSVLSI 2025
http://glsvlsi.org/program.html

- ***MILS: Modality Interaction Driven Learning for Logic Synthesis***
  Mingyu Zhao, Jiawei Liu, Jianwang Zhai and Chuan Shi

- ***Titanus: Enabling KV Cache Pruning and Quantization On-the-Fly for LLM Acceleration***
  Peilin Chen and Xiaoxuan Yang

SIGDA E-News Editorial Board

- ***Aphelios: A Selective Lock-step Neural Processing Unit Design***
  Wenhao Sun, Yiming Gan, Yuhui Hao and Yinhe Han

- ***Modeling, Design and In-situ Demonstration of Bio-inspired Central Pattern Generator and Neuromorphic Computing Circuits for Complex Kinematic Control of Quadruped Robots***
  Qiankai Cao, Yuhao Ju, Zhiwei Zhong, Zhengyu Chen and Jie Gu

- ***Impact of Error Rate Misreporting on Resource Allocation in Multi-tenant Quantum Computing and Defense***
  Subrata Das and Swaroop Ghosh

- ***Microarchitecture Evaluation Framework for Transient Execution Attack Vulnerability: Metrics, Fuzzing, and Sensitivity Analysis***
  Jordan McGhee, Nayra Lujano, Aiden Peterson, Henry Duwe, Akhilesh Tyagi and Berk Gulmezoglu

What is Processing-in-Memory (PiM) for Efficient AI?

Contributing author: Gabriel Falcao <gff@deec.uc.pt>
Associate Professor, Instituto de Telecomunicações & Electrical and Computer Engineering Department, University of Coimbra, Portugal, and
Visiting Associate Professor at the Computer Science Department, New York University, Abu Dhabi, UAE

AE: Alberto Marchisio <alberto.marchisio@nyu.edu>

As Artificial Intelligence (AI) continues to expand across edge devices [1]—such as wearables, healthcare systems, autonomous vehicles, and smart sensors—**there is a growing demand for computing architectures that can meet stringent energy and performance constraints**. A central challenge lies in the **data movement bottleneck**: transferring data between memory and processing units consumes significantly more energy than computation itself [2], particularly in memory-bound workloads like today's neural network models. This issue is especially pronounced in sparse models, or when activations, weights, and gradients must be frequently fetched from off-chip memory (e.g., DRAM), or in federated learning models used in edge deployments with limited memory bandwidth [3].

Processing-in-Memory (PiM) emerges as a disruptive architectural paradigm that addresses this bottleneck by relocating computation closer to—or directly within—the memory components [4]. By minimizing the physical and temporal

Paper Deadlines

HOST'26 – IEEE Int'l Symposium on Hardware-Oriented Security and Trust
Washington DC, USA
Abstracts due (fall): Aug. 19, 2025
Deadline (fall): Aug 26, 2025
Abstracts due (Winter): Dec. 1, 2025
Deadline (winter): Dec. 8, 2025
May 4-7, 2026
http://www.hostsymposium.org

ISSCC'26 – IEEE Int'l Solid-State Circuits Conference
San Francisco, CA, USA
Deadline: Sept. 3, 2025
Feb. 16-20, 2026
http://isscc.org

DATE'26 - Design Automation and Test in Europe
Verona, Italy
Abstracts due: Sept 7, 2025
Deadline: Sept. 14, 2025
April 20-22, 2026
http://www.date-conference.com

ISQED'26 - Int'l Symposium on Quality Electronic Design
San Francisco, CA, USA
Deadline: Sept. 17, 2025
Apr. 8-10, 2026
http://www.isqed.org

FPGA'26 – ACM/SIGDA Int'l Symposium on Field-Programmable Gate Arrays
Seaside, CA, USA
Abstracts due: Sept. 24, 2025
Deadline: Oct. 1, 2025
Feb. 22-24, 2026
http://www.isfpga.org

ISPD'26 – ACM Int'l Symposium on Physical Design
Bonn, Germany
Abstracts due: Sept. 21, 2025

distance between data storage and processing, PiM offers substantial improvements in throughput, latency, energy efficiency, and autonomy—key requirements for both edge AI and high-performance computing. PiM architectures are generally categorized into two classes: **Processing-near-Memory (PnM)**, where computation is performed in dedicated processing elements (e.g., accelerators, cores, or reconfigurable logic) placed adjacent to memory arrays; and **Processing-using-Memory (PuM)**, which performs logic operations directly within memory arrays by exploiting the analog properties of memory cells. While PuM is highly efficient for basic operations, traditional implementations face challenges with complex functions such as multiplication or exponentiation due to logic overhead. Nevertheless, recent studies have demonstrated PuM's potential as a key enabler of efficient AI inference at the edge [5, 6], particularly for quantized Convolutional Neural Networks (CNNs). As CNNs dominate image inference tasks, aggressive quantization (e.g., using 8-bit or less to represent weights, activations or gradients) significantly reduces memory footprint and energy consumption with minimal accuracy degradation, making them well-suited for PiM deployment [7].

Architectures such as Ambit, SIMDRAM, pLUTo, PRIME, ISAAC and others [4] exemplify the performance and energy gains achievable over conventional CPU/GPU platforms. Building on top of previous technology proposed by these works, **pLUTo** (Processing-using-Memory with Lookup Table Operations) stands out as a **novel DRAM-based PuM architecture that replaces complex arithmetic with high-throughput lookup table (LUT) queries** [6]. By transforming computation into parallel memory reads, pLUTo achieves notable improvements in performance and energy efficiency while offering design flexibility through tunable area overheads in custom DRAM layouts.

Beyond CNNs, PiM also shows strong promise for accelerating **Large Language Models (LLMs)** [8, 9], which involve massive tensor matrix operations and memory-intensive sparse attention mechanisms [10]. In fact, as the first chips for running LLMs in edge devices are expected to gain more relevance, PiM architectures show potential to alleviate pressure on memory hierarchies (during both training and inference) by improving data locality and reducing bandwidth demands—particularly in mitigating the effects of irregular memory access patterns common in sparse matrix and graph-based computations.

To fully realize PiM's potential, further advancements are needed in compiler support, algorithm redesign, software/hardware co-design, and benchmarking ecosystems. Rather than replacing traditional computing, PiM is expected to augment it—reducing data movement and enhancing efficiency, especially in edge AI scenarios. As edge intelligence and large-scale AI models continue to evolve, data-centric architectures like PiM—and particularly PuM designs such as pLUTo—are poised to become foundational technologies for scalable, efficient, and sustainable high-performance AI systems.

References:

[1] Satyanarayanan, M. Edge computing: A new disruptive force. Keynote address at the 13th Annual ACM Intern. Systems and Storage Conf., 2020.
[2] Horowitz, Mark. "1.1 computing's energy problem (and what we can do about it)." In 2014 IEEE international solid-state circuits conference digest of technical papers (ISSCC), pp. 10-14. IEEE, 2014.
[3] Bonawitz, Kallista, Peter Kairouz, Brendan McMahan, and Daniel Ramage. "Federated learning and privacy: Building privacy-preserving systems for machine

Deadline: Sept. 28, 2025
Mar 15-18, 2026
http://www.ispd.cc/

ISCAS'26 – IEEE Int'l Symposium on Circuits and Systems
Shanghai, China
Deadline: Oct. 12, 2025
May 24-27, 2026
https://2026.ieee-iscas.org/

RTAS'26 - IEEE Real-Time and Embedded Technology and Applications Symposium
Saint Malo, France
Deadline: Nov. 13, 2025
May 12-14, 2026
http://2026.rtas.org

learning and data science on decentralized data." Queue 19, no. 5 (2021): 87-114.
[4] Mutlu, Onur, Saugata Ghose, Juan Gómez-Luna, and Rachata Ausavarungnirun. "A modern primer on processing in memory." In Emerging computing: from devices to systems: looking beyond Moore and Von Neumann, pp. 171-243. Singapore: Springer Nature Singapore, 2022.
[5] Oliveira, Geraldo F., Juan Gómez-Luna, Saugata Ghose, Amirali Boroumand, and Onur Mutlu. "Accelerating neural network inference with processing-in-DRAM: from the edge to the cloud." IEEE Micro 42, no. 6 (2022): 25-38.
[6] Ferreira, João Dinis, Gabriel Falcao, Juan Gómez-Luna, Mohammed Alser, Lois Orosa, Mohammad Sadrosadati, Jeremie S. Kim et al. "pLUTo: Enabling massively parallel computation in DRAM via lookup tables." In 2022 55th IEEE/ACM International Symposium on Microarchitecture (MICRO), pp. 900-919. IEEE, 2022.
[7] Falcao, Gabriel, and João Dinis Ferreira. "To PiM or not to PiM." Communications of the ACM 66, no. 6 (2023): 48-55.
[8] He, Yintao, Haiyu Mao, Christina Giannoula, Mohammad Sadrosadati, Juan Gómez-Luna, Huawei Li, Xiaowei Li, Ying Wang, and Onur Mutlu. "Papi: Exploiting dynamic parallelism in large language model decoding with a processing-in-memory-enabled computing system." In Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2, pp. 766-782. 2025.
[9] Gu, Yufeng, Alireza Khadem, Sumanth Umesh, Ning Liang, Xavier Servot, Onur Mutlu, Ravi Iyer, and Reetuparna Das. "PIM is all you need: A CXL-enabled GPU-free system for large language model inference." In Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2, pp. 862-881. 2025.
[10] Gonçalves, Nuno, Marcos V. Treviso, and André F. T. Martins. ”AdaSplash: Adaptive Sparse Flash Attention”. To appear in Proceedings of the 42nd International Conference on Machine Learning (ICML), 2025. (Spotlight) or from ArXiv: Gonçalves, Nuno, Marcos Treviso, and André FT Martins. "AdaSplash: Adaptive Sparse Flash Attention." arXiv preprint arXiv:2502.12082 (2025).

SIGDA Partner Journal

ACM Transactions on Design Automation of Electronic Systems (TODAES) features groundbreaking research and development in the specification, design, analysis, simulation, testing, and evaluation of electronic systems, with a focus on computer science and engineering. The journal's impact factor increased to 2.2 in 2023, more than doubling its value from 2020. Additionally, each issue highlights a notable contribution as the Editor's Pick for special recognition.

TODAES also recognizes papers and outstanding junior researchers through the best paper and rookie of the year awards. Authors can send their paper submissions to the manuscript portal.

TODAES welcomes special issue proposals from leading researchers and practitioners. Such proposals should be emailed to Prabhat Mishra, Senior Associate Editor, at prabhat@ufl.edu

Upcoming Conferences

ISLPED’25 – ACM/IEEE Int’l Symposium on Low Power Electronics and Design
University of Iceland, Iceland
Aug. 6-8, 2025
http://www.islped.org

MLCAD'25 - ACM/IEEE Workshop on Machine Learning for CAD
Santa Cruz, CA, USA
Sep. 8-10, 2025
https://mlcad.org/symposium

ESWEEK'25 - Embedded Systems Week
Taipei, Taiwan
Sept. 28 - Oct. 3, 2025
http://www.esweek.org

VLSI-SoC’25 – IFIP/IEEE Int’l Conference on Very Large Scale Integration
Puerto Varas, Chile
Oct. 12-15, 2025
http://www.vlsi-soc.com

MICRO’25 – IEEE/ACM Int'l Symposium on Microarchitecture
Seoul, Korea
Oct. 18-22, 2025
http://www.microarch.org/micro58

ICCAD’25 – IEEE/ACM Int’l Conference on Computer-Aided Design
Munich, Germany
Oct 26-30, 2025
https://iccad.com/

PACT'25 - Int'l Conference on Parallel Architectures and Compilation Techniques
Irvine, CA, USA
Nov. 3-6, 2025
http://www.pactconf.org

TODAES Special Issue Call for Papers

Special Issue on Co-Design and Design Automation for Optical/Photonic Computing Systems

This special issue seeks original submissions on pioneering research aimed at advancing co-design and EDA methodologies to support the modeling, simulation, design optimization, and physical implementation toward hybrid integration of optical computing/interconnect and electronic systems with high reliability, scalability, and efficiency. All these topics, as well as further potential topics mentioned below, are of interest to this special issue.

Important Dates

- Submissions deadline: Oct 15, 2025
- First-round review decisions: Dec 15, 2025
- Deadline for revision submissions: Jan 15, 2026
- Notification of final decisions: Feb 15, 2026
- Tentative publication: Spring 2026

Submissions should be made through the ACM TODAES submission site (http://mc.manuscriptcentral.com/todaes)

For questions and further information, please contact guest editors at:

- Jiaqi Gu, Arizona State University, jiaqigu@asu.edu
- Cunxi Yu, University of Maryland, cunxiyu@umd.edu
- Sudeep Pasricha, Colorado State University, sudeep@colostate.edu
- Xu Wang, Cadence Design Systems, xubc@cadence.com

More information can be found in this call for papers.

Technical Activities

1. **Optical AI Just Redefined Precision**

Artificial intelligence is driving meaningful improvements in optical technologies, such as improving optical measurement for increased precision in critical applications or enhancing the deployment of fiber-optic communication networks. These strategically chosen use cases reduce errors and elevate efficiency, making decision-makers across industries interested in using them to optimize outcomes. Knowing about the options could help you decide whether to pursue them…

2. **Volvo Taking Platform Route to Scale Autonomous Driving Beyond Trucks**

Volvo Autonomous Solutions has been crafting a platform for autonomous trucking that can be deployed across all Volvo brands once higher levels of autonomy are achieved and business cases identified…

3. **Key Considerations for Automotive Body Control Network Design**

The rapid transformation of the automotive industry has led to a complex web of body control electronics, comprising a dense network of ECUs, actuators, and interconnects…

ICCD'25 – IEEE Int'l Conference on Computer Design
Dallas, Texas, USA
Nov. 10-12, 2025
http://www.iccd-conf.com

FPT'25 - Int'l Conference on Field-Programmable Technology
Shanghai, China
Dec. 2-5, 2025
http://icfpt.org

ISED'25 – Int'l Conference on Intelligent Systems and Embedded Design
Chhattisgarh, India
Dec. 17-19, 2025
http://isedconf.org

iSES'25 – IEEE Int'l Symposium on Smart Electronic Systems
Hyderabad, India
Dec. 17-20, 2025
http://www.ieee-ises.org

HiPC'25 – IEEE Int'l Conference on High Performance Computing, Data, And Analytics
Hyderabad, India
Dec. 17-20, 2025
http://www.hipc.org

VLSID'26 – International Conference on VLSI Design & International Conference on Embedded Systems
Pune, Maharashtra, India
Jan. 3 - 7, 2026
https://vlsid.org/

ASP-DAC'26 - Asia and South Pacific Design Automation Conference
Hong Kong, China
Jan. 19-22, 2026
http://www.aspdac.com

HiPEAC'26: Int'l Conference on High Performance Embedded Architectures & Compilers
Krakow, Poland
Jan. 26-28, 2026
https://www.hipeac.net/2026/krakow/

4. **As Data Center Growth Soars, Startup Uses AI to Cut Power Binge**
Startup Bay Compute can help data centers cut power consumption by as much as 20% with the company's AI-based operating system…

Job Positions

University of Limassol, Cyprus

Job Title: Faculty Positions in Computer Science

Description: The Department of Information Technology at the School of Technology and Innovation, of the University of Limassol, invites applications for open-rank faculty positions in Computer Science. We seek outstanding individuals with a proven record of research excellence and a strong commitment to teaching, who will contribute to the continued growth and academic distinction of the department. We welcome applications from scholars in all areas of Computer Science. Particular interest will be given to candidates whose research focuses on applied or systems-oriented domains, as well as those working in data science and artificial intelligence. Areas of special interest include, but are not limited to: Software Engineering / Operating Systems; Internet Technologies; Cloud Computing; Digital Transformation; Data Mining and Machine Learning; Symbolic AI; Deep Learning; Reinforcement Learning; Natural Language Processing and Large Language Models. For more information, please refer to https://facultyvacancies.com/faculty-positions-in-computer-science,i42529.html

University of Doha, Qatar

Job Title: Professor of Data and Cyber Security

Description: The College of Computing and Information Technology (CCIT) invites applications for the position of NCSA Professorship Chair at the rank of Full Professor in Data and Cyber Security. This 10-month renewable full-time position will entail teaching (1 course per semester), applied research (with focus on critical infrastructure security), and service. Reporting to the Head of Data and Cyber-Security Department, the successful candidate will collaborate with the faculty in the CCIT and staff members in the NCSA to achieve their respective missions. She/He will liaise with NCSA and key industry stakeholders in the State of Qatar to advance critical infrastructure security posture by working on research projects, developing strategic initiatives, and conducting workshops relevant to the field. The current data and cyber-security degree program in the CCIT encompasses a broad range of courses from applied cryptography to industrial control systems security, through network security, ethical hacking, security monitoring, and incident response. The degree program is also complemented with a master program in artificial intelligence and cognitive

cybersecurity. For more information, please refer to https://facultyvacancies.com/professor-of-data-and-cyber-security,i42465.html.

Laval University, Canada

Job Title: Faculty Position in Electronic

Description: The Department of Electrical and Computer Engineering invites applications for a full-time tenure-track professorship. The successful candidate will conduct high-impact research in electronics to solve complex challenges in wireless or optical communications technologies, signal processing, and/or receiver/transmitter devices, as well as the hybrid integration of these technologies. Expertise in high-frequency electronics is an asset. The candidate will be expected to conduct research and publish non-exclusively in the following areas: High-frequency electronics; Power electronics; Microelectronics. For more information, please refer to https://polytechnicpositions.com/faculty-position-in-electronic,i13855.html.

Notice to authors